# The Journal of Economic Perspectives

A journal of the
American Economic Association

Fall 2018

The Journal of
# Economic Perspectives

*A journal of the American Economic Association*



---

*Editorial offices:*
Journal of Economic Perspectives
*American Economic Association Publications*
*2403 Sidney St., #260*
*Pittsburgh, PA 15203*
*email: jep@jepjournal.org*

---

The *Journal of Economic Perspectives* gratefully acknowledges the support of Macalester College.
Registered in the US Patent and Trademark Office (®).

Composed by American Economic Association Publications, Pittsburgh, Pennsylvania, USA.
Printed by LSC Communications, Owensville, Missouri, 65066, USA.

*THE JOURNAL OF ECONOMIC PERSPECTIVES* (ISSN 0895-3309), Fall 2018, Vol. 32, No. 4. The *JEP* is published quarterly (February, May, August, November) by the American Economic Association, 2014 Broadway, Suite 305, Nashville, TN 37203-2418. Annual dues for regular membership are $20.00, $30.00, or $40.00 depending on income; for an additional $15.00, you can receive this journal in print. The journal is freely available online. For details and further information on the AEA go to https://www.aeaweb.org/. Periodicals postage paid at Nashville, TN, and at additional mailing offices.
POSTMASTER: Send address changes to the *Journal of Economic Perspectives*, 2014 Broadway, Suite 305, Nashville, TN 37203. Printed in the U.S.A.

*The Journal of*
# *Economic Perspectives*

## Contents

*Volume 32 • Number 4 • Fall 2018*

**Symposia**

## Statement of Purpose

The *Journal of Economic Perspectives* attempts to fill a gap between the general interest press and most other academic economics journals. The journal aims to publish articles that will serve several goals: to synthesize and integrate lessons learned from active lines of economic research; to provide economic analysis of public policy issues; to encourage cross-fertilization of ideas among the fields of economics; to offer readers an accessible source for state-of-the-art economic thinking; to suggest directions for future research; to provide insights and readings for classroom use; and to address issues relating to the economics profession. Articles appearing in the journal are normally solicited by the editors and associate editors. Proposals for topics and authors should be directed to the journal office, at the address inside the front cover.

## Policy on Data Availability

It is the policy of the *Journal of Economic Perspectives* to publish papers only if the data used in the analysis are clearly and precisely documented and are readily available to any researcher for purposes of replication. Details of the computations sufficient to permit replication must be provided. The Editor should be notified at the time of submission if the data used in a paper are proprietary or if, for some other reason, the above requirements cannot be met.

## Policy on Disclosure

Authors of articles appearing in the *Journal of Economic Perspectives* are expected to disclose any potential conflicts of interest that may arise from their consulting activities, financial interests, or other nonacademic activities.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 3–32*

# An Economist's Guide to Climate Change Science

Solomon Hsiang and Robert E. Kopp

Humans have engaged in large-scale transformation of natural systems for millennia. Stone Age hunting technologies led to extinctions of large mammals; agricultural revolutions transformed forests into farmlands; pursuit of minerals has carved the earth's surface; dams and reservoirs now manipulate the flow of almost all rivers; and synthetic fertilizers now flood the nitrogen cycle. But among these transformations, the restructuring of the global carbon cycle and the accompanying alteration of the climate stands apart in its sheer scale, complexity, and economic significance. Essentially all humans that have ever lived contributed, in their own small ways, to reshaping this planetary-scale system. Thousands of years of forest clearance may have added hundreds of billions of tons of carbon to the atmosphere. In the industrial era, every home lit by a coal or natural gas-fired power plant and every petroleum-powered train, plane, and motor vehicle has contributed to the net accumulation of carbon dioxide in the atmosphere. The average human contributes about 5 tonnes of carbon dioxide ($CO_2$) every year (Le Quéré et al. 2018), about a quarter of which will remain in the atmosphere for well over a millennium (Archer et al. 2009).

■ *Solomon Hsiang is Chancellor's Associate Professor of Public Policy, University of California, Berkeley, California, and is a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. Robert E. Kopp is Director of the Institute of Earth, Ocean & Atmospheric Sciences and Professor of Earth & Planetary Sciences, both at Rutgers University, New Brunswick, New Jersey. Hsiang and Kopp are both Principal Investigators, Climate Impact Lab, http://www.impactlab.org. Their email addresses are shsiang@berkeley.edu and robert.kopp@rutgers.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.3

doi=10.1257/jep.32.4.3

Those emissions of $CO_2$, together with other greenhouse gases, distort the planet's energy balance. In steady state, the sunlight that makes it to the Earth's surface is absorbed and then re-radiated to space as an equal quantity of heat (technically, infrared light). The accumulation of greenhouse gases in the atmosphere blocks some of this re-radiation, redirecting energy back toward the Earth's surface: about 27 trillion watts (0.05 watts per square meter) per 1 percent increase in atmospheric $CO_2$ concentrations, equivalent to the energy of one Hiroshima-scale atomic bomb spread over the surface of the Earth every 2.3 seconds. The resulting climatic distortion affects not just temperatures around the world, but also where clouds form, when it floods, how cyclones move, and the volume of water in the ocean. Thus, while fossil-fueled human industriousness has raised unprecedented multitudes out of poverty, the scale of the climate change externality it has produced is similarly extraordinary.

At least since Nordhaus's (1977) presentation at an American Economic Association annual meeting, the analysis and management of climate change has been recognized as an important economic problem, and a growing number of economists are lending the world their expertise in understanding the problem and developing solutions. However, conversations with colleagues indicate to us that a general discomfort with physical sciences—a subject sometimes not studied since high school—prevents many economic minds from engaging more deeply with the problem of climate change.

The goal of this article is to provide a brief introduction to the physical science of climate change, aimed towards economists. We begin by describing the physics that controls global climate, how scientists measure and model the climate system, and the magnitude of human-caused emissions of carbon dioxide. We then summarize many of the climatic changes of interest to economists that have been documented and that are projected in the future. We conclude by highlighting some key areas in which economists are in a unique position to help climate science advance. An important message from this final section, which we believe is deeply underappreciated among economists and thus highlight here, is that all climate change forecasts rely heavily and directly on economic forecasts for the world. On timescales of a half-century or longer, the largest source of uncertainty in climate science is not physics, but economics (Hawkins and Sutton 2009).

## Basics of Climate Change Science

For most economic and scientific purposes, climate can be defined as the joint probability distribution describing the state of the atmosphere, ocean, and freshwater systems (including ice). Each of these systems is itself an extraordinarily high-dimensional system, so it is appealing to work with summary statistics such as *global mean surface temperature* or temperature distributions for major cities. Indeed, global mean surface temperature is intimately tied to the fundamental physics of planetary energy balance that explain global warming. However, consumers of

climate science should recognize that such simplifications, while sometimes useful, do not capture the entire picture.

The idea that human activity could alter the climate has a long history, going back almost two centuries (for an overview, see Weart 2018). However, it took focused research during the second half of the 20th century to achieve the level of confidence we now possess that human activity is altering the climate (Stocker et al. 2013; US Global Change Research Program 2017). This confidence comes from many lines of evidence based on observations at Earth's surface and throughout different layers of the atmosphere and oceans, geological reconstructions of historical climates, and two centuries of physical theory. The null hypothesis that humans have had no influence on global climate is now easily rejected given available data (for example, Hegerl et al. 2007).

**Planetary Energy Balance and Greenhouse Gases**

Sunlight continuously enters our planet's atmosphere from space. In order for the earth to maintain a stable surface temperature, this flow of incoming energy must be balanced by a flow of energy leaving the atmosphere. For the Earth, about 30 percent of incident sunlight is immediately reflected back out to space from the surface or from clouds. The remaining 70 percent is absorbed by the Earth's surface and atmosphere, and must be balanced by the planet's own emission of infrared radiation to space, which intensifies with higher temperatures. Without greenhouse gases, the equilibrium global mean surface temperature would be –18°C (about 0°F)[1], fully determined by the Sun's temperature, the Earth's distance from the Sun, and the reflectivity (also known as "albedo") of the Earth. If a larger flow of energy somehow were to reach the Earth's surface—for example, if the Sun were to grow in brightness, or the Earth to decline in albedo—the planet would heat up until the additional outgoing flow of infrared radiation exactly offset this new source of energy.

Greenhouse gases distort Earth's energy balance because they are transparent to incoming visible and ultraviolet sunlight but absorb infrared radiation, hindering the return flow of this energy from the surface and the lower atmosphere into space. When a greenhouse gas molecule intercepts infrared radiation headed from the surface to space, the absorbed energy is re-emitted in all directions, sending some energy that might otherwise have escaped to space back down to the surface of the Earth. This causes the surface and lower atmosphere to warm, increasing their emission of infrared radiation slightly. Equilibrium is re-established when the intensified outgoing infrared radiation is sufficient to offset the trapping effects of the greenhouse gases.

Because of the presence of greenhouse gases, the average height in the atmosphere from which infrared radiation can escape to space and contribute to balancing the planet's energy budget is not the Earth's surface; it is a level of the atmosphere known as the "effective radiating level." At present, Earth's effective

[1] To convert any temperature *change* from Celsius to Fahrenheit, multiply by nine-fifths—so 2°C of warming is $2 \times 9/5 = 3.6$°F of warming. To convert a temperature in *levels* from Celsius to Fahrenheit, multiply by nine-fifths and then add 32—thus a day with level temperature of $30°C = 30 \times 9/5 + 32 = 86$°F.

radiating level occurs at about 5.5 km altitude; on average, this level has the necessary temperature of about –18°C—the same that the Earth's surface would have in the absence of greenhouse gases. The relationship between temperature and altitude in Earth's lower atmosphere—on average about 6°C/km—makes the surface nearly 33°C (about 59°F) warmer than this level.

When greenhouse gases are added to the atmosphere, the first reaction is that the height of the effective radiating level moves upward. This temporarily leads to a decrease in the amount of radiation escaping from the Earth to space; but conservation of energy implies the surface and lower atmosphere must then warm up, so the higher (and originally cooler) effective radiating level would warm to the equilibrium temperature of –18°C. In the absence of additional feedbacks, doubling carbon dioxide concentrations would lead to the effective radiating level being about 200 meters higher, which in turn would lead to an equilibrium surface warming of about 1.2°C (Hansen et al. 1981).

However, the warming surface and atmosphere trigger feedbacks, which change the shift in effective radiating level and surface temperature associated with a given change in greenhouse gas concentration. Estimates of equilibrium climate sensitivity (the long-term, equilibrium response to a doubling of $CO_2$ concentrations) that include atmospheric and sea ice feedbacks are generally 2.0–4.5°C (3.6–8.1°F) (Collins, Knutti et al. 2013). The most important feedback involves water vapor: a warmer atmosphere is a more humid atmosphere, and water vapor is the most powerful natural absorber of longwave infrared radiation. Other important feedbacks involve sea ice (which reflects incoming solar energy), clouds (which can both trap heat and reflect incoming solar energy), and the response to warming of the ocean and land ecosystems (which drive most of the flow of $CO_2$ out of the atmosphere and can also affect albedo).

Because greenhouse gases alter the climate by changing the radiative properties of the atmosphere, their influence is measured in units of "radiative forcing," defined as the extent to which the human-generated stock of gas distorts the net flow of radiation into the atmosphere on average (incoming minus outgoing), relative to a preindustrial baseline. For example, a rise in atmospheric $CO_2$ concentrations from the historical baseline of 278 parts per million to the current (as of 2018) level of 409 parts per million exerts about 2.1 $W/m^2$ of radiative forcing. For reference, the energy from the sun reaching the top of Earth's atmosphere is 342 $W/m^2$, and central estimates of the equilibrium warming associated with a change in radiative forcing are about 0.8°C per $W/m^2$; thus the equilibrium warming associated with the current level of $CO_2$ forcing is about 1.6°C above the preindustrial baseline.

Radiative forcing by greenhouse gas emissions does not translate immediately into surface warming, in part because the deep ocean takes centuries to warm and, through exchange of heat with the surface ocean, slows overall warming. Nonetheless, modeling experiments indicate that—because of the relative timescales over which the planet warms and $CO_2$ is naturally removed from the atmosphere—most of the warming associated with a marginal emission of $CO_2$ occurs within a couple decades and persists for millennia (Joos et al. 2013). Thus, climatic changes experienced today are a result of both relatively recent emissions

and also cumulative emissions during the past centuries of fossil fuel combustion and past millennia of deforestation.

**Establishing Baseline Climates**

Within climate science, *paleoclimatology* is a well-developed subfield that focuses on the reconstruction of historical climates, thus setting a baseline for explaining climate changes. For examples, gases trapped in air bubbles of ice contain information on atmospheric chemistry at the moment they froze (for example, Luthi et al. 2008); the width of tree rings reflect growing-season temperatures and rainfall (for example, Jones et al. 2009); microscopic fossils in salt-marsh sediments reflect changes in salinity, and thus in local sea level (for example, Edwards and Horton 2000); and the relative abundance of different isotopes of oxygen in ocean sediments reflect the extent of "ice ages" because polar ice sheets lock up lighter isotopes, thereby restricting their supply to the deep ocean (for example, Cramer, Miller, Barrett, and Wright 2011). In some cases, physical data can be corroborated with observations in historical records, such as records of cyclone-caused shipwrecks maintained by insurers (Trouet, Harley, and Domínguez-Delmás 2016). While most proxies and historical observations are inherently local, spatio-temporal statistical methods and comparison to physical models can be used to estimate global mean values of quantities such as surface temperature and sea level from local data.

Figure 1 presents reconstructions of atmospheric $CO_2$ concentrations, global mean surface temperature, and global mean sea level over two different timescales. In the context of the last 11,600 years, known as the Holocene Epoch, the recent sharp jump in atmospheric $CO_2$ concentrations is quite striking and is unequivocally explained by human-caused emissions (Luthi et al. 2008; MacFarling Meure et al. 2006). The higher-resolution, post-1958 observational record from Mauna Loa in Hawaii also reflects higher-frequency patterns of largely natural variability, like the seasonal cycle and inter-annual variability in the strength of the land and ocean carbon sink (Keeling et al. 2001). The Holocene temperature record reveals a long-term decline, caused by slow variations in Earth's orbit, that began around 5,500 years ago (Marcott, Shakun, Clark, and Mix 2013).[2] The post-1850 reconstruction from direct observations reveals that this decline was interrupted in the 20th century, by a rise totaling about 1.0°C from the late-19th to the early-21st century (Rohde et al. 2013). This rise, which is superimposed by a spectrum of

[2]The exact timing of the Holocene decline is currently contested. Marsicek, Shuman, Bartlein, Shafer, and Brewer (2018) suggest that the global analysis underlying Figure 1 is seasonally biased and conceals a more complex pattern, at least in North America and Europe, where their analysis suggests that summer temperatures declined starting around 5,500 years ago, but that winter temperatures did not cool until about 2,000 years ago. Some other, very long term cycles include Milanković cycles, which are global periodic climate changes driven by variations in the orientation of Earth's axis of rotation (19,000- and 23,000-year periods), the tilt of Earth's axis of rotation (41,000-year periods), and the shape of Earth's orbit around the Sun (roughly 100,000- and 400,000-year periods) (Berger 2012). Changes in incoming solar radiation caused by these cycles, amplified by natural feedbacks, serve as the pacemaker for ice ages over the last 2.6 million years.

*Figure 1*
**Atmospheric $CO_2$ Concentrations, Global-Mean Surface Temperature, and Global-Mean Sea Level**

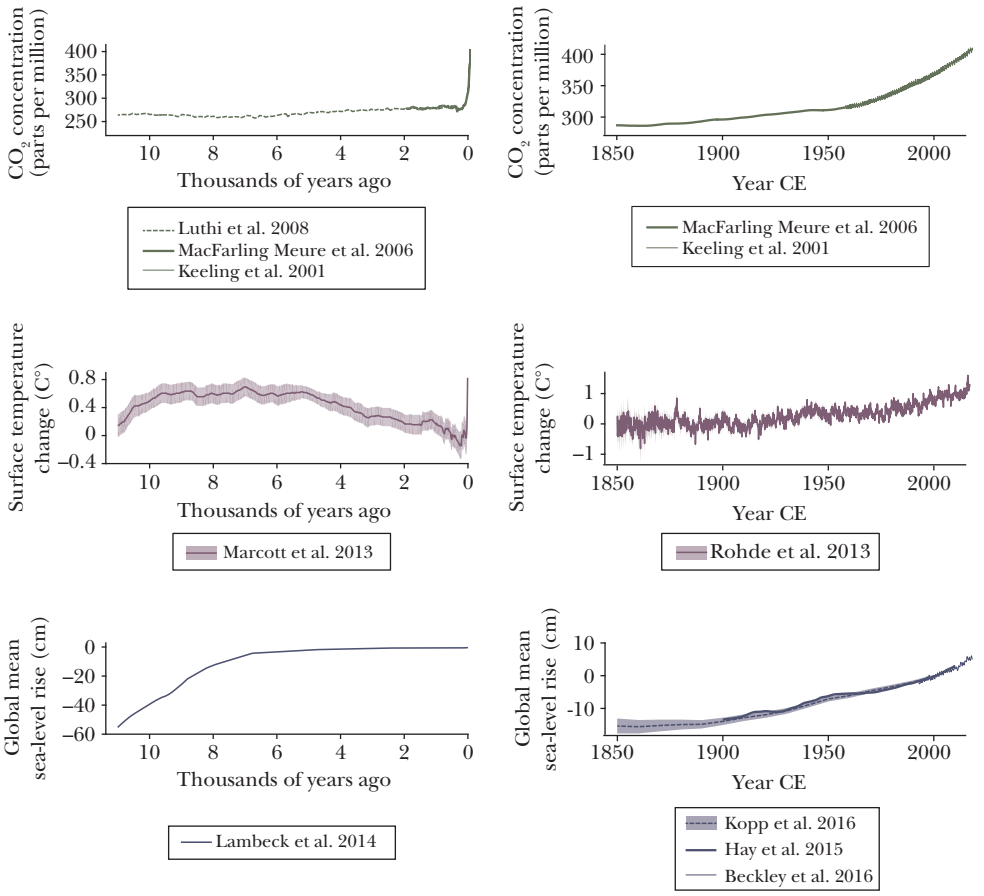


*Data:* Luthi et al. (2008); MacFarling Meure et al. (2006); Keeling et al. (2001); Marcott, Shakun, Clark, and Mix (2013); Rohde et al. (2013); Lambeck, Rouby, Purcell, Sun, and Sambridge (2014); Kopp, Kemp et al. (2016); Hay, Morrow, Kopp, and Mitrovica (2015); Beckley et al. (2016).
*Note:* The figure shows historical atmospheric $CO_2$ concentrations from ice cores and direct measurements (top), reconstructed historical global mean surface temperatures relative to the 1850–1900 average (middle), and reconstructed global mean sea level relative to the 1991–2009 average (bottom), over the last 11,000 years (left) and since 1850 CE (right). Shaded areas are 95 percent confidence intervals.

higher-frequency variability, some internal to the climate system and some driven by changes in forcing, is well-explained by the response to human-caused emissions.

Sea level responds more sluggishly than temperature to changes in forcing, because both the oceans (which expand when they absorb heat) and ice sheets (which can shrink in response to warming temperature) are large systems with an ability to absorb tremendous quantities of heat while warming only modestly. The first half of the Holocene is characterized by relatively rapid sea-level rise, which

ended with the final disappearance of Laurentide Ice Sheet in North America. This rise was a delayed response to about 5°C warming since the thermal nadir of the last ice age (about 21,000 years). The twentieth-century sea-level rise was the fastest in at least 2,800 years, and the last quarter-century was characterized by a rate about twice as fast as the 20th century average (Sweet, Horton, Kopp, LeGrande, and Romanou 2017).

In general, a core challenge to determining whether humans are changing the climate is assessing whether systematic changes in the behavior of the climate system are explained or confounded by the sources of natural variation. As one example of such natural variation, El Niño–Southern Oscillation is the dominant pattern in the global climate at annual frequencies and has occasionally been studied by economists (for example, Hsiang and Meng 2015). Other longer-term ocean-related oscillations include the North Atlantic Oscillation, which varies on seasonal, annual, decadal, and centennial timescales (Hurrell 1995; Trouet, Scourse, and Raible 2012); the Pacific Decadal Oscillation (Mantua and Hare 2002); and the Atlantic Multidecadal Oscillation (Clement et al. 2015). Climate models form the basis for inference in this setting, seeking to separate a secular trend signal from these oscillating sources of noise.

**Climate Models**

Climate models mathematically represent physical understanding of the climate system. They fall along a hierarchy of complexity from simple models that capture key aspects of the longer-term, global-scale response to detailed, full-complexity Earth system models that provide greater insight into processes at finer temporal and spatial scales (Hayhoe et al. 2017).

The simplest climate change models, called energy balance models, can simulate millennia of global mean climate change in a single second on a laptop. These models are based on a budgeting of sunlight and thermal energy in the atmosphere, as well as the role of key feedbacks. Early pen-and-paper versions of such models date back to the work of Svante Arrhenius in the 1890s; by the 1960s, they had also been adapted to include a single spatial dimension representing the vertical structure of the atmosphere, which allowed the models to describe vertical motions of air (Manabe and Strickler 1964).

In the 1960s, the equations of fluid dynamics were incorporated to produce early atmospheric "general circulation models" that capture both the three-dimensional structure and dynamical evolution of the global atmosphere (for example, Manabe and Smagorinsky 1965). Later generations of models elaborated their representation of the ocean, as well as of sea ice, land surfaces, and atmospheric chemistry. These general circulation models[3] were the ancestors of today's full-complexity Earth system models, which also endogenize vegetation dynamics and the carbon cycle. Full-complexity Earth system models represent the best tools available for simulating spatial patterns of the climate response, but they have several drawbacks.

[3]As general circulation models evolved to include more than just the fluid dynamics of the ocean and the atmosphere, the acronym GCM was sometimes adapted to stand for "global climate model."

First, they are computationally expensive, taking several hours on a high-performance computing cluster just to simulate one year of climate. Second, although such models provide fairly high spatial resolution—with grid cells that are roughly 100 km along a side in the generation of models used in the Fifth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC)—this resolution may still be inadequate for capturing details relevant to many economic impacts. Third, detailed models may produce baseline climate projections that differ from observed historical patterns. To address these last two issues, the climate science community has developed post-processing techniques for bias-correction and spatial "downscaling," thus increasing the spatial resolution of the final output. Such techniques include both statistical approaches (like using quantile regressions to mimic historical variability around the mean) and physical modeling approaches that embed higher-resolution regional climate models within boundary conditions set by a global model (as in Wood, Leung, Sridhar, and Lettenmaier 2004). In addition, cutting-edge climate models are run at increasingly higher resolutions; some of the most recent models have resolutions below 50 km × 50 km, and in some cases can achieve local resolutions as high as 10 km × 10 km.

Within the context of a single climate model simulation, uncertainty arises from the imperfect representation of physical processes—that is, from structural and parametric uncertainty—as well as from the imperfectly known initial conditions of a model run. As famously discovered by Lorenz (1963) in early numerical weather models, tiny errors in initial conditions can produce dramatically different forecasts within the same model, chaotic behavior known colloquially as "the butterfly effect." This endogenous chaotic behavior turns out to be more difficult to predict than global average conditions, which are tightly constrained by energy budgets. As a result, climate modeling teams usually run their model multiple times with perturbed initial conditions, creating a collection of results known as an *initial-conditions ensemble*. Individual realizations of the model are never interpreted as literal forecasts; rather, the ensemble as a whole is thought to capture statistical properties of the climate system. Indeed, most climate scientists generally avoid the terms "forecast" and "prediction," preferring instead the term "projection" to describe a simulation of future climate under an assumed emission scenario. Producing decadal projections with global climate models is a frontier research area, with one of the key challenges being aligning the internal variability of a climate model with the internal variability of the real climate (Meehl et al. 2009).

*Emissions of Radiative Pollutants*

Not all the greenhouse gases emitted by humans remain in the atmosphere today; a substantial fraction has been absorbed by carbon sinks on land (like plants and soil) and in the ocean (for example, by phytoplankton and chemical dissolution). If all 1434 Gt of fossil $CO_2$ emitted since 1750 had stayed in the atmosphere, the current atmospheric $CO_2$ concentration would be about 475 ppm rather than the observed 409 ppm, even without considering emissions from deforestation.[4]

[4] One part-per-million $CO_2$ in the atmosphere is equal to about 7.8 Gt $CO_2$ in physical mass.

*Table 1*

**Historical and Top 15 Current Emissions of Carbon Dioxide from Fossil Fuel Combustion and Cement Production**

| *Country* | *Cumulative 1751–2014 (gigatonnes $CO_2$)* | *% of Global* | *Emissions 2014 (gigatonnes $CO_2$)* | *% of Global* | *Emissions per capita (tonnes $CO_2$), 2014* |
|---|---|---|---|---|---|
| China | 174.7 | 12% | 10.3 | 30% | 7.5 |
| United States | 375.9 | 26% | 5.3 | 15% | 16.2 |
| India | 41.7 | 3% | 2.2 | 7% | 1.7 |
| Russia / USSR | 151.3 | 11% | 1.7 | 5% | 11.9 |
| Japan | 53.5 | 4% | 1.2 | 4% | 9.6 |
| Germany | 86.5 | 6% | 0.7 | 2% | 8.9 |
| Iran | 14.8 | 1% | 0.6 | 2% | 8.3 |
| Saudi Arabia | 12.0 | 1% | 0.6 | 2% | 19.5 |
| South Korea | 14.0 | 1% | 0.6 | 2% | 11.7 |
| Canada | 29.5 | 2% | 0.5 | 2% | 15.1 |
| Brazil | 12.9 | 1% | 0.5 | 2% | 2.6 |
| South Africa | 18.4 | 1% | 0.5 | 1% | 9.1 |
| Mexico | 17.5 | 1% | 0.5 | 1% | 3.8 |
| Indonesia | 11.0 | 1% | 0.5 | 1% | 1.8 |
| United Kingdom | 75.2 | 5% | 0.4 | 1% | 6.5 |
| ***World*** | **1,434.0** | **100%** | **34.1** | **100%** | **4.7** |

*Source:* Boden, Marland, and Andres (2017).

However, cumulative emissions of $CO_2$ are nonetheless a useful metric, as the $CO_2$-caused warming is approximately proportional to cumulative emissions (Allen et al. 2009), with every trillion tons of $CO_2$ causing about 0.2–0.7°C of warming.

Table 1 presents the estimated cumulative emissions of $CO_2$ from fossil fuels and cement production during 1751–2014, as well as the flow of emissions in 2014 (Boden, Marland, and Andres 2017). The United States is responsible for over one-fourth of historical emissions, followed by China (12 percent) and Russia (11 percent, including the former Soviet Union); together with Germany (6 percent) and the United Kingdom (5 percent), these five countries account for 60 percent of historical emissions. However, if one examines flows today rather than the stock of historical emissions, the picture is changing; China (30 percent) dominated emissions in 2014, followed by the United States (15 percent), India (7 percent), Russia (5 percent), and Japan (4 percent). Germany is the largest emitter in the European Union (2.1 percent), with the EU28 collectively ranking third in global $CO_2$ emissions, responsible for about 10 percent (Janssens-Maenhout et al. 2017). High national emissions reflect high carbon intensity per capita (per-capita emissions are 16.2 tonnes/year in the United States, 3.4 times the global average), high population levels (per capita emissions in India, the third-leading emitter, are about one-third the global average), or a mix of both factors (per-capita emissions in China are about 60 percent more than the global average).

These metrics do not include $CO_2$ emissions from deforestation, which are significant: Pongratz and Caldeira (2012) estimate that these accounted for about

230 gigatonnes of $CO_2$ from 800–1850, and 425 gigatonnes of $CO_2$ from 800–2006, compared to about 1,175 gigatonnes of $CO_2$ from fossil fuels over this latter time period (Boden, Marland, and Andres 2017). At present, the ratio of fossil fuel to land use emissions is about 7.6 (Le Quéré et al. 2018).

These metrics also do not include emissions of non-$CO_2$ greenhouse gases and other climate-altering pollutants. The climatic impact of an emission depends on both its radiative forcing of the molecules emitted and their lifetime in the atmosphere. For example, methane survives for only 12 years on average in the atmosphere before breaking down into $CO_2$ and water, whereas a substantial fraction of emitted $CO_2$ lasts for millennia. Thus, while methane has large radiative impact per molecule per year, the integrated lifetime impact of a marginal molecule of methane emissions is partially offset by its short lifetime.[5] Blanco et al. (2014) provide a discussion of non-$CO_2$ emissions.

Emissions of particulate matter and aerosol precursors (like sulfur dioxide) also influence the radiative balance of the atmosphere. Both pollutants lead to the formation of aerosols—particles that are solid or liquid, not gases, but which are small enough to remain aloft in the atmosphere for substantial periods of time (days in the lower atmosphere; years in the stratosphere). Most aerosols reflect incoming sunlight, leading to surface cooling (negative radiative forcing), but some, notably black carbon, absorb solar energy and increase warming. Through their effects on cloud physics, aerosol emissions have complex regional consequences for precipitation that are distinct from the effects of greenhouse gases (Rosenfeld et al. 2008). Because the spatial distribution and net radiative effects of aerosols are difficult to monitor, and change more quickly than gases, the overall radiative impact of aerosols is highly uncertain and remains an important open question in climate change science. The global average effective radiative forcing of aerosols is estimated to be between –1.9 and –0.1 $W/m^2$—opposite in sign and between about 5 percent and 90 percent of the forcing from $CO_2$ (Boucher et al. 2013).

As one more level of complexity, coal combustion emits both $CO_2$ and aerosol pollution, which leads to a tradeoff of timescales: burning less coal reduces particulate matter and sulfur dioxide emissions which is directly beneficial to human health, but also leads to a short-term increase in warming due to the reduction in aerosol emissions, even though the long-term effect of reduced $CO_2$ emissions is a substantial reduction in warming (Wigley 2011). Similarly, efforts to target reductions in particulate pollution from coal power plants without tackling $CO_2$ emissions will lead to climate warming (Westervelt et al. 2015).

**Emissions Scenarios**

There are many climate modeling research programs, each of which develop, maintain, and run global climate models whose outputs are compared against one

[5] Methane, with an atmospheric concentration of about 1.8 ppm, currently exerts a forcing of about 0.5 $W/m^2$; nitrous oxide, at 0.3 ppm, exerts a forcing of about 0.2 $W/m^2$, and fluorinated gases like chlorofluorocarbons and hydrofluorocarbons, with concentrations less than 1 part per billion, exert forcing of about 0.3 $W/m^2$.

another. The Coupled Model Intercomparison Project (CMIP) (Taylor, Stouffer, and Meehl 2012) is the largest comparative effort, and plays a major role in informing the assessment reports of the Intergovernmental Panel on Climate Change. To ensure that model outputs are comparable across groups, standardized emissions scenarios are used as inputs to all models. The latest effort, CMIP Phase 5 (CMIP5), used a range of emission scenarios, known as the Representative Concentration Pathways (RCPs), that exogenously prescribe the flow of human-caused emissions over the coming decades. These emissions scenarios, which begin in 2005, are labeled by the overall radiative forcing (in $W/m^2$) that occurs in 2100 in each scenario. RCP 8.5 has the strongest forcing, with $CO_2$ emissions nearly doubling from their current levels by 2050 and continuing to rise thereafter; RCP 4.5 has a moderate forcing, with $CO_2$ emissions stabilizing at close to their current levels through the middle of the century and declining thereafter, reaching about 40% of their current levels by 2080; and RCP 2.6 has the weakest forcing, with $CO_2$ emissions declining immediately, to less than a third of the current levels by 2050, and becoming net-negative during the 2080s. In RCP 8.5, atmospheric $CO_2$ concentration climbs to 541 ppm by 2050 and 936 ppm by 2100; in RCP 4.5, to 487 ppm by 2050 and 538 ppm by 2100; and in RCP 2.6, to 443 by 2050, declining to 421 ppm by 2100. Below, when we discuss "high-", "moderate-" and "low-" emissions scenarios, we are referring to RCP 8.5, 4.5, and 2.6, respectively.

## Observed and Projected Climate Changes in the Modern World

In this section, we describe how historical changes in the climate are identified and attributed to human activity, as well as climate changes that are projected to occur. Interested readers should consult the IPCC Fifth Assessment Report (Stocker et al. 2013), USGCRP (2017), and the readings cited below for additional details.

### Detection and Attribution of Climate Change

Over the last several decades, a core objective of climate science has been to detect changes in the climate and to determine whether these changes can be attributed to human activity. *Detection* refers to the empirical problem of determining whether there has been an actual shift in the joint distribution of environmental variables that we refer to as the climate. *Attribution* refers to the inferential problem of assigning a cause to the observed changes (Bindoff et al. 2013). Attribution studies generally simulate what counterfactual climates would look like in the absence of human activity, altering the model parameters that describe human inputs to the climate. Thus, for example, human emissions of greenhouse gases and aerosols might be eliminated in a model's "control" simulation. If it is not possible, or sufficiently unlikely, that these human-free simulations can account for observed changes in the climate, then scientists attribute these changes to human activity.

The scientific community is in broad and strong agreement that overall, human activity has already substantially altered the global climate and that continued changes should be expected as emissions of greenhouse gases continue (Stocker et al. 2013; USGCRP 2017). The vast majority of actively publishing researchers now acknowledge

*Table 2*
**Statements of the Intergovernmental Panel on Climate Change (IPCC) on Detection and Attribution of Global Climate Change**

| | |
|---|---|
| *First Assessment Report* (1990) | "Unequivocal detection of the enhanced greenhouse effect from observations is not likely for a decade or more." |
| *Second Assessment Report* (1995) | "The balance of evidence suggests a discernible human influence on global climate." |
| *Third Assessment Report* (2001) | "Most of the observed warming over the last 50 years is *likely** to have been due to the increase in greenhouse gas concentration." |
| *Fourth Assessment Report* (2007) | "Most of the observed increase in global average temperatures since the mid-20th century is *very likely* due to the observed increase in anthropogenic greenhouse gas concentrations." |
| *Fifth Assessment Report* (2013) | "It is *extremely likely* that human influence has been the dominant cause of the observed warming since the mid-20th century." |

*Source:* The IPCC Assessment Reports can be found at https://www.ipcc.ch/publications_and_data/publications_and_data_reports.shtml.
* The uncertainty language used by the IPCC is precisely defined: *likely* refers to an assessed probability of at least 66 percent, *very likely* implies at least 90 percent, and *extremely likely* means at least 95 percent.

the strength of the evidence implicating a human-caused signal in climate change (Cook et al. 2016). The agreement in the scientific community has grown stronger over the last quarter century, reflected in the IPCC's increasingly strong statements regarding the detection and attribution of global warming shown in Table 2. Some of the public confusion regarding the strength of scientific evidence appears to have been sown intentionally. For example, a textual analysis of ExxonMobil documents from 1977–2014 indicates that internal documents generally acknowledged that climate change is real and human caused while public-facing documents did not (Supran and Oreskes 2017).

Figure 2 (which is best viewed in the color version of this article available at the JEP website) shows some of the most important evidence in support of the conclusion that human emissions are causing global temperatures to rise. In the left panel, red [or upper light grey] bands indicate the range of global mean surface temperature simulated in 90 percent of climate models that exogenously impose observed human emissions. Blue [or lower light grey] bands indicate the analogous range for the same models but in a "control" simulation that imposes only natural forces. Observed temperatures, indicated by the black line, began to separate from the envelope of control simulations in the 1980s and now lie far outside this range. In contrast, observed temperatures are fully consistent with the range of temperatures simulated when human emissions are included. We note that this consistency extends not just to global mean surface temperature, but also to changes in stratospheric temperature and ocean heat content. Thus, it is extremely difficult to explain current temperatures in the absence of human activity. The

*Figure 2*

**Average Annual Global Mean Surface Temperature, Compared to Distributions of Climate Model Simulations**

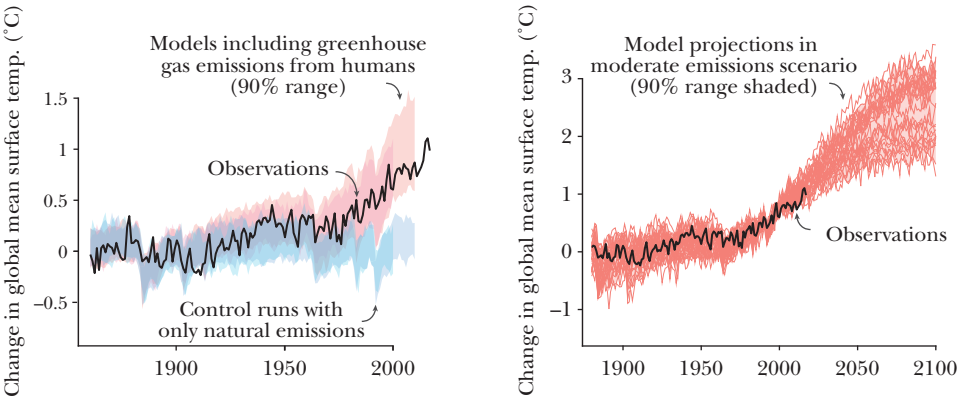


*Sources:* Data comes from Jones, Stott, and Christidis (2013), Morice, Kennedy, Rayner, and Jones (2012), and Taylor, Stouffer, and Meehl (2012).

*Note:* This graph is best viewed in color; the electronic version of this article available at the JEP website is in color. The heavy black line shows observed average annual global mean surface temperature. The red [or light grey] distributions are exogenously "treated" with anthropogenic greenhouse gas emissions, while the blue [or light grey] distributions (shown only in the left panel) are "control" runs that only contain natural forcings. In the left panel, climate model distributions are from the Third Coupled Model Intercomparison Project (CMIP3) published in 2007 and displayed until 2000, and CMIP5 published in 2013 and displayed until 2010. In the right panel, all climate model projections come from CMIP5 in the moderate emissions scenario (RCP 4.5). Temperatures shown are relative to the 1880–1900 average.

gradually increasing confidence of the scientific community can be understood by noting the envelope of model results published in association with the 2007 IPCC report (displayed ending in simulation year 2000) were less cleanly separated than those published in association with the 2013 IPCC report (displayed ending in 2010), although the separation visible through 2000 was already reflected in the IPCC's 2007 statement that temperatures were "*very likely* due to anthropogenic greenhouse gas concentrations" (Table 2).

It is now virtually certain (at least 99 percent probability) that the observed modern warming trend exceeds the bounds of natural variability (Bindoff et al. 2013). Furthermore, humans are *likely* (with at least 66 percent probability) responsible for 0.6°C–0.8°C of the observed 0.6°C of warming over 1951–2010. Values greater than 0.6°C are possible for the anthropogenic contribution because of the possibility that natural forcing and variability could otherwise impose a slightly negative baseline trend (for example, as a result of volcanic eruptions), a pattern which is visible in the control runs of Figure 2.

**Temperature Changes**

Since the late 19th century, global mean surface temperature has increased by about 1.0°C, with the trend accelerating after 1980. Almost every location on the planet has exhibited an upward temperature trend over this period (Wuebbles et al. 2017). Warming has also been substantially faster over land than the ocean—between 1880–1900 and 1997–2017, the land has warmed 1.4°C (2.5°F) on average while the oceans warmed roughly 0.6°C (1.1°F) (GISTEMP Team 2018).

As one would expect, given the array of factors affecting temperatures, the overall rise in temperatures has not been smooth over time or homogenous across space. For example, warming was dampened in the 1950s–1970s, most likely as a result of both aerosol emissions, which reflected sunlight away from the planet (Maher, Gupta, and England 2014), and natural variability. Since 1980, the most rapid warming has occurred in the far north, where the replacement of highly reflective summer sea ice with dark, open ocean rapidly increases the absorption of sunlight and local warming (Serreze and Barry 2011).

A heavily discussed period of slowed average warming over 1998–2013, the so-called "hiatus," now appears fully consistent with the natural variance of the climate system (Cahill, Rahmstorf, and Parnell 2015), as can be readily seen in the overlay of simulated and observed temperature time series in the right panel of Figure 2. Relative to the distribution of simulations for the first decade of the 21st century, the observed values fall toward the low end of projections but never leave the envelope of expected variations. However, in addition to natural variability, it is thought that some model simulations warmed too quickly because the emissions scenarios in the RCPs underrepresented volcanic and human aerosol emissions after 2005 (Medhaug, Stolpe, Fischer, and Knutti 2017).

Based on the assessment of the Intergovernmental Panel on Climate Change of CMIP5 simulations, projected global mean surface temperature is likely to rise 0.9–2.3°C (1.6–4.1°F) above preindustrial levels (defined as the 1850–1900 average) by 2080–2100 under a low-emissions scenario, 1.7–3.3°C (3.1–5.9°F) under a moderate-emissions scenario (shown in Figure 2), and 3.2–5.4°C (5.8–9.7°F) under a high-emissions scenario (Collins et al. 2013).

As in the past, warming will be more rapid over land, where most economic activity occurs, compared to over the ocean. The only location on the surface that is projected by some models to cool is a very small portion of the North Atlantic Ocean just south of Greenland, where changing ocean circulation may induce cooling. Although warming will continue to occur fastest over the Arctic, average summer temperature will diverge from the historical range soonest in low-latitude regions, which experience lower historical variance. Figure 3A illustrates regional heterogeneities in the rate of warming (in °C) that are otherwise masked by globally averaged summary statistics. The map depicts the average warming at each location associated with a 1°C increase in global mean temperature; values greater than 1°C indicate rates of warming faster than the global mean, while values below 1°C indicate warming that is slower than the global mean.

To help grasp the potentially transformative scale of these thermal changes, Figure 4, Panel A, plots the average summertime temperatures for the lower 48 US

*Figure 3*

**Projected Change in Local Average Temperatures and Local Average Rainfall per 1°C of Warming in Global Mean Temperatures**

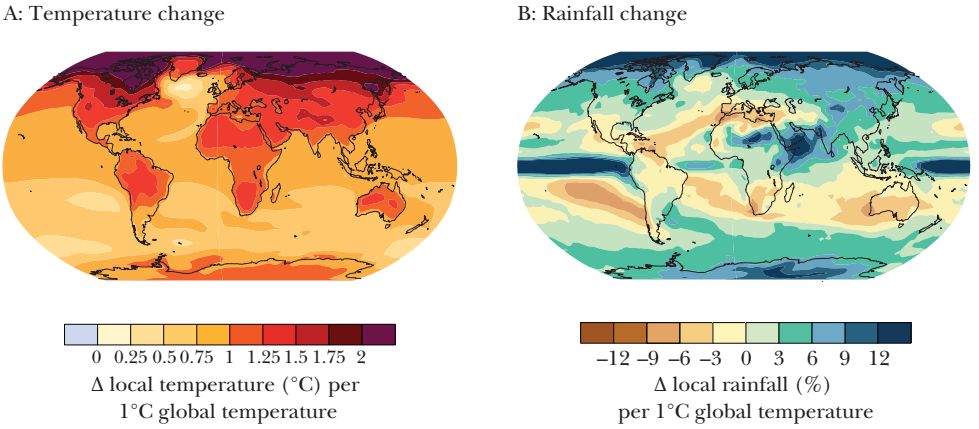


*Source:* Collins, Knutti, et al. (2013).
*Note:* Changes are differences in means between 1986–2005 and 2081–2100 in CMIP5 simulations of RCP 4.5, scaled by the overall change in global mean temperature. These heatmaps should be viewed in color. See the electronic versions on the JEP website.

states, adapted from Houser et al. (2015). The cluster in the bottom left of the figure in blue text indicates historically observed temperatures, while the cluster in the upper right of the figure in red text indicates average projected mean temperatures for 2080–2099 across models simulating a high-emissions scenario. This layout allows for projected future temperatures to be matched to historical analogs. For example, future summers in Vermont will be similar to historical summers in Maryland, summers in Connecticut will be similar to past summers in Arizona, future summer in New Jersey will be slightly hotter than historical Louisiana summers, and future summers in Georgia and Florida will be much hotter than anything previously experienced in the United States. As shown in Panel B, a similar analysis at the level of countries shows that future temperatures in Norway are projected to be similar to historical temperatures in Germany, future Mexico will be slightly hotter than historical Iraq, future Indonesia will be similar to historical Mali, and India and Thailand are projected to be hotter than any country presently on Earth.[6]

**Precipitation Changes**

A warmer atmosphere is capable of holding more water vapor, leading to an increase in overall average precipitation (rainfall and snowfall). Observed precipitation in the mid-latitude Northern Hemisphere has increased since the 1950s. Heavy precipitation events in particular have increased, most clearly in North America

[6] The Appendix available with this paper at http://e-jep.org shows current and projected annual average temperature for 166 countries.

*Figure 4*

**Average Temperatures for Lower 48 US States Observed during 1981–2010 and Projected for 2080–2099 in a High Emission (RCP 8.5) Scenario.**

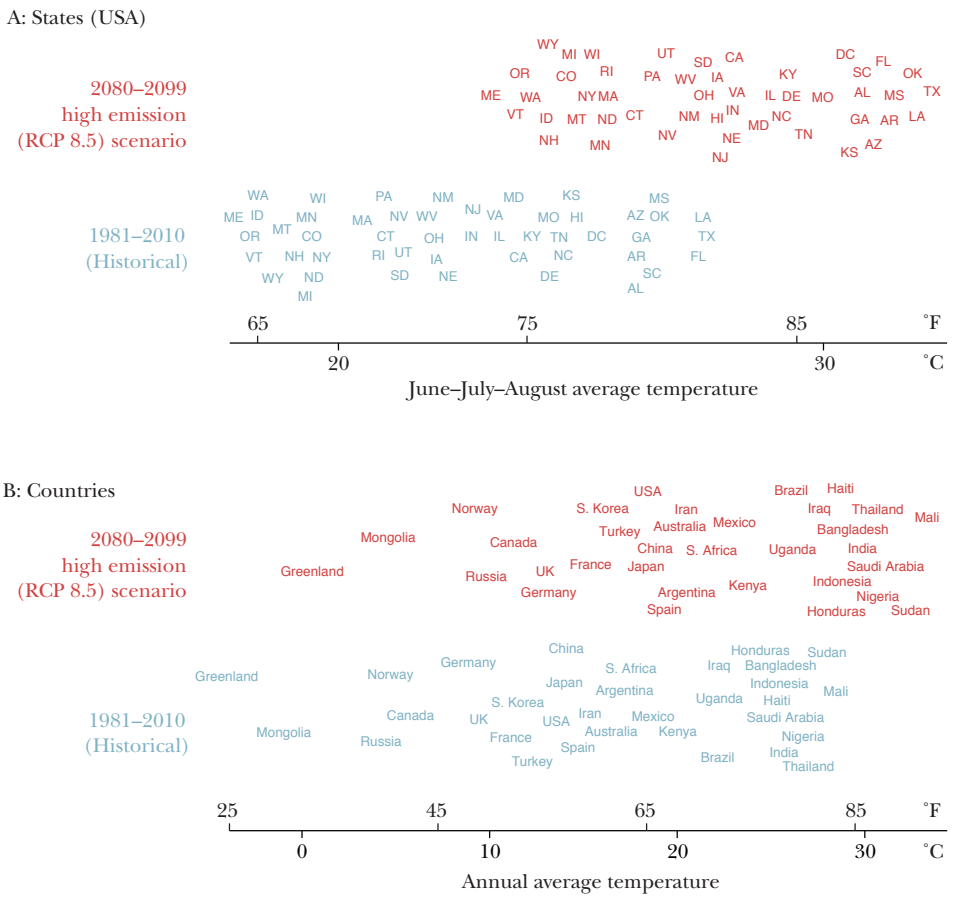


*Note:* Panel A displays summertime area-average temperatures adapted from Houser et al. (2015). Panel B displays population-weighted annual average temperatures, using data from Burke, Hsiang, and Miguel (2015). Markers are vertically jittered for readability.

and Europe, where the most data is collected (Hartmann et al. 2013). Both of these changes are consistent with those expected on a warming planet. However, because the atmospheric dynamics that govern precipitation involve both large motions of air masses and processes that occur at scales below the spatial resolution of many climate models, precipitation changes are considerably more challenging to model numerically than temperature changes. This difficulty, combined with the array of changes in temperature, wind, humidity and other factors that all affect when and where precipitation falls, have rendered projections of precipitation changes more complex and more uncertain than projections of temperature.

There is large heterogeneity in the sign of projected precipitation change, with many locations getting wetter while many others get drier. Precipitation dynamics are also strongly affected by internal variability—such as the El Niño–Southern

Oscillation—and projections for specific locations depend upon changes in large-scale patterns of atmospheric circulation (Collins et al. 2013). Figure 3B illustrates average changes in local rainfall for each 1°C increase in global mean temperature. At many locations, there is a large range of uncertainty for projected changes, with plausible projections allowing for no change. Simple summary statements like "dry regions are likely become generally drier and wet regions are likely to become generally wetter" hold well over the ocean, but are coarse descriptions of the complex precipitation changes that may occur over land (Greve et al. 2014). In the United States, the most robust projections are for a springtime drying of the Southwest, summertime drying of the Northwest, and increase in winter and spring precipitation in the Northeast, upper Midwest, and northern Great Plains (Houser et al. 2015).

**Humidity Changes**

*Specific humidity* is the total moisture content of air. *Relative humidity* is the ratio of specific humidity to a theoretical maximum moisture capacity, which rises exponentially as temperature increases. Since the 1970s, global mean specific humidity has increased; however, there is little evidence of an increase in global mean relative humidity and some evidence for a decline, possibly reflecting faster warming of the land than of the oceans, which are the primary source of atmospheric moisture (Sherwood and Fu 2014). Models that project that the largest increases in temperature on land also tend to predict the largest decreases in relative humidity (Fischer and Knutti 2013).

One reason that humidity is thought to be economically important is that it affects human health, since higher humidity levels make it more difficult for the human body to cool itself in hot conditions through sweating. One physical metric closely related to the combined effect of heat and humidity is *wet-bulb temperature*. Wet-bulb temperatures are measured using a ventilated thermometer wrapped in a wet cloth, and are strongly related to the experienced conditions of a sweating person.

Dangerously hot and humid conditions are projected to become dramatically more likely in several regions around the world (Sherwood and Huber 2010). For example, Houser et al. (2015) defined "dangerously hot and humid days" as those characterized by peak wet-bulb temperatures over 80°F. By this definition, dangerously hot and humid days are "typical of the most humid parts of Texas and Louisiana in the hottest summer month, and the most humid summer days in Washington and Chicago." Their analysis found that, in the southeastern United States, the population-weighted frequency of dangerously hot and humid days are projected to rise from 8 per year on average in 1981–2010 to 17–28 days per year over 2040–2059 in a moderate emissions scenario and to 40–70 days/year on average over 2080–2099 in a high emissions scenario (Houser et al. 2015).

**Tropical Cyclones**

Tropical cyclones are the class of phenomena that includes tropical storms, typhoons, hurricanes, and cyclones; these categories are distinguished by wind speed and the ocean basin where the storm occurs. Tropical cyclones are driven by the temperature difference between the warm ocean surface and cooler temperatures higher in the atmosphere. The warm ocean moistens overlying air, which rises and

cools, releasing energy and rain. Thus, climate change is thought to have countervailing effects on storms: warming sea surface temperatures fuel storms, but warming temperatures higher in the atmosphere may suppress them (Knutson et al. 2010).

Tropical cyclone formation and storm trajectory depend on myriad additional factors, especially wind patterns, which introduce additional complexity into projections of their future changes. Furthermore, inconsistent historical data on storms in the open ocean prior to the satellite era make inferences difficult. However, there is evidence that the frequency and intensity of the strongest storms in the Atlantic have been increasing since the 1970s, and some evidence that humans contributed to this change (Walsh et al. 2016).

Efforts to model all of these factors together broadly agree that the frequency of intense tropical cyclones (such as category 4 and 5 hurricanes), as well as the average intensity of their associated rainfall, is projected to increase with warming (Kossin et al. 2017). The effect on total number of storms remains less certain, though most studies suggest a stable or decreasing quantity of lower-intensity storms. The effect of climate change on *storm tracks* (the paths that storms take toward land) is uncertain and may offset or enhance the effect of increased storm intensity in some regions.

The spatial distribution of these changing risks is heterogeneous. For example, systematic changes in the spatial distribution of storm tracks within an ocean basin may reallocate cyclone risk between populations, even if the overall frequency of storms does not change. Across ocean basins, models generally agree in projecting substantial increases of storm intensity in the West Pacific, affecting East Asia and Oceania, with some decreases in activity occurring in the Indian Ocean, affecting Southern Asia and East Africa. Projecting changes in the North Atlantic, which affects Central and North America, has been more challenging; the greatest scientific uncertainty persists for this area (Knutson et al. 2010; Christensen et al. 2013).

Humans may also affect the genesis and growth of tropical cyclones through the regional effects of aerosol pollution. Aerosols aloft may cool local sea surface temperatures, by reflecting sunlight before it reaches the sea surface, as well as heat higher levels in the atmosphere, by absorbing sunlight when the particles are dark-colored. Both of these effects generally work to weaken storms, and storm activity in recent decades may have been greater had greenhouse gas co-pollutants been absent (Walsh et al. 2016).

**Sea-Level Rise**

Global mean sea-level rise is driven by two processes: an increase in the volume of the water already in the ocean, which occurs as the water warms and expands, and an increase in the mass of water in the ocean, primarily from the melting of ice on land. Since 1900, global mean sea level has increased by about 18–21 cm, with the rate of rise since about 1990 being 2–2.5 times faster than during the preceding nine decades. A substantial fraction of this rise is attributable to human-caused climate change (Sweet et al. 2017). Regional sea-level changes can differ substantively from this global trend, modulated by changes in currents and winds; changes in Earth's gravitational field, rotation, crust, and mantle that occur as land ice changes; and changes in the height of land that result from compaction of sediments, plate

tectonics, the ongoing mantle response to historical land ice changes, and other factors (Kopp, Hay, Little, and Mitrovica 2015). Historic sea-level rise has led to a detectable increase in the frequency of coastal flooding, in some cases by more than an order of magnitude (Sweet and Park 2014).

Due to the slow response time of the oceans and ice sheets, sea-level rise is fairly insensitive to alternative emissions scenarios for the first half of this century. Across studies, median projections of future global mean sea-level rise are 20–30 cm during 2000–2050, with less than a 5 percent chance of exceeding 50 cm (for a review see Horton et al. 2018).

After 2050, projections become more deeply uncertain, due to both uncertain human emissions and the uncertain response of the polar ice sheets (Kopp, DeConto, et al. 2017). Median projections for 2000–2100 range from 40–80 cm for a low-emissions scenario to 70–150 cm for a high-emissions scenario. However, global mean sea-level rise of as much as 250 cm by 2100 cannot be ruled out. For reference, the last time global temperatures were about as high as they currently are (about 125,000 years ago), global mean sea-level was about 6–9 meters higher than today. Some coastal areas will be inundated permanently; others will be protected by additional investments to be incurred. The resulting increase in the frequency of tide- and wave-driven flooding is expected to render some low-lying island states uninhabitable (Storlazzi et al. 2018).

In addition to increasing average high-tide levels, a major economic consequence of sea-level rise results from its interaction with tropical cyclones and extratropical cyclones (like the "nor'easter" storms that sometimes hit the northeastern United States). Storm surges that occur during these storms can impose major costs and sea-level rise adds roughly linearly to peak storm surge height. Figure 5 illustrates the joint effects of projected sea-level rise and changing tropical cyclone activity on the flood risk of Miami, FL, and New York, NY (Hsiang et al. 2017). The extent of areas expected to flood with a 1 percent annual probability increases substantially after 2050 for Miami but much earlier in the century for many regions of New York.

**Droughts and Floods**

By altering temperature and precipitation patterns, climate change alters the frequency and intensity of extreme moisture conditions, such as droughts and floods. There is a limited but increasing ability to attribute intensifying extreme floods and drought to human activity. For example, Emanuel (2017) estimated that climate forcing by humans amplified the probability of rainfall experienced by Texans during Hurricane Harvey six-fold.

Since 1950, the likelihood of drought has increased in the Mediterranean and West Africa and decreased in central North America and northwestern Australia (Hartmann et al. 2013); again, assessment of the human role in this trend is challenging (Bindoff et al. 2013). Some measures of drought in the United States have increased due to warming, which increases evaporation and exhibits an anthropogenic signal (Wehner, Arnold, Knutson, Kunkel, and Legrande 2017). In model projections, the frequency of droughts tends to increase in dry regions (Collins et al. 2013), which are also projected to expand. Prolonged hot and dry periods are

*Figure 5*
**Areas Projected to Experience Floods at Least Once every 100 Years on Average (1% annual risk) in Miami, FL, and New York, NY**

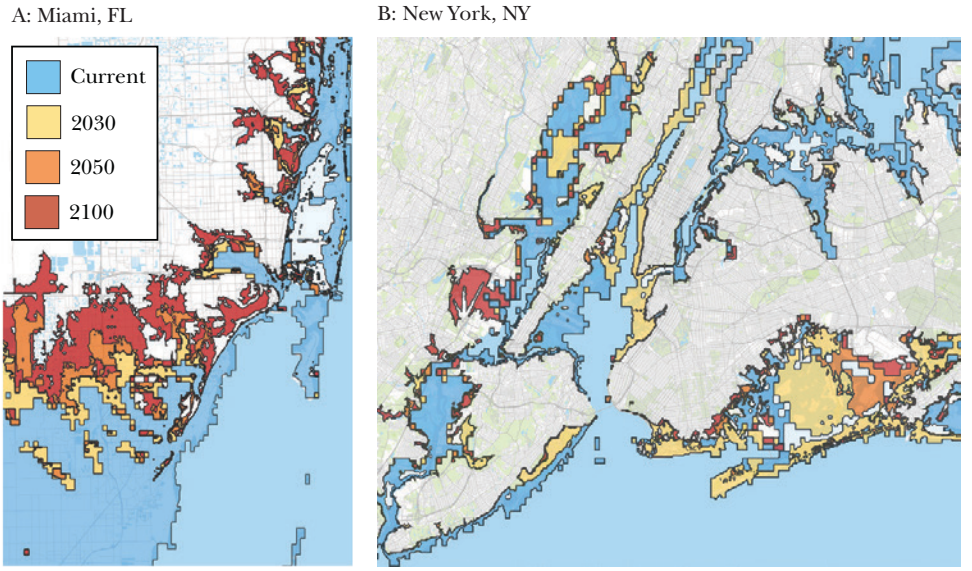


*Source:* Hsiang, Kopp, Jina, Rising, et al. (2017).
*Note:* These projections account for median projected sea-level rise and for projected changes in tropical cyclone intensity in a high-emission (RCP 8.5) scenario.

projected to become substantially more frequent in many grassland areas with low agricultural productivity, regions that today often depend on livestock production (Bell, Sum, Longmate, Tseng, and Hsiang 2018). In regions with more vegetation, amplifying oscillations between heavy rainfall and drought is projected to increase the frequency of wildfires (Abatzoglou and Williams 2016) because vegetative fuel grows rapidly during wet periods then becomes flammable during dry periods.

The observed increase in heavy precipitation suggests that climate change is contributing to increasing rain-driven flood damages (Jiménez Cisneros, Oki, et al. 2014). Projected increases in heavy precipitation and shorter-lived snowpack are likely to further increase the frequency of inland flooding. In addition to changing rainfall patterns, rising sea levels amplify the frequency of coastal flooding (Buchanan, Oppenheimer, and Kopp 2017). Changing patterns of flood risk are of particular economic importance, as flooding is recognized as one of the costliest classes of disaster globally (Swiss Re Institute 2018). However, in many cases, the dominant drivers of increased flood damages are related to the number of people and extent of development affected by the flood rather than the physical size of the flood.

**Clouds**

Understanding clouds is scientifically important, because they generate competing feedbacks in the climate system. However, cloud physics are complex, and many important dynamics occur at a spatial resolution finer than those used

in most global climate models, generating substantial uncertainty in the projected changes in cloud cover for many regions of the world.

On the one hand, clouds reflect visible light, so increases in cloud cover, particularly low-altitude cloud cover, can increase the fraction of incoming sunlight that is reflected before it warms the Earth's surface. On the other hand, clouds absorb outgoing infrared radiation leaving the Earth's surface and thus contribute to the greenhouse effect, so increases in cloud cover can amplify warming (Boucher et al. 2013).

To date, global-scale changes in cloudiness remain unclear. Looking forward, some analyses suggest the potential for circulation changes leading to large-scale, nonlinear reductions in low-latitude cloudiness with warming that could substantially increase the sensitivity of temperature to $CO_2$ forcing (Caballero and Huber 2013).

One way in which cloud science may become important to economists is in the rapidly growing research field of "geoengineering" or "climate engineering," which considers various proposals to intentionally alter the climate so as to counteract some effects of greenhouse gas emissions (Caldeira, Bala, and Cao 2013). The most widely proposed intervention is "solar radiation management," which involves increasing the reflectivity of the atmosphere in order to shade and cool the surface. One proposed mechanism for increasing reflectivity is to spray aerosol precursors into the upper atmosphere, mimicking the mechanism through which historical volcanic eruptions have cooled the surface. Another proposal, with more localized and shorter lasting effects, is "cloud brightening," achieved by manipulating cloud-droplet size by spraying particles into lower portions of the atmosphere. A theoretically appealing feature of cloud-brightening proposals is that cloud brightening might be used to temporarily cool a city or ecosystem during particularly damaging heat waves, although the various economic costs and unintended effects of such policies remain poorly studied (Proctor et al. 2018).

**Ocean Acidification**

As the $CO_2$ concentration in the atmosphere increases, some of this $CO_2$ will dissolve into ocean water, where it will form carbonic acid and increase ocean acidity. Currently, the ocean absorbs roughly one-quarter of global $CO_2$ emissions through this process.

The rate of acidification depends on local chemistry, temperature, circulation patterns, and freshwater inputs. For example, in the North Pacific during the last three decades, surface ocean acidity has increased by about 12 percent. Globally, the current rate of acidification is unparalleled in at least the last 55 million years, as reflected in a variety of chemical indicators from ocean sediments (Hönisch et al. 2012). In a high-emissions scenario, global mean surface ocean acidity is projected to increase 100–150 percent (Jewett and Romanou 2017).

Ocean acidification is thought to alter marine ecosystems substantially, although the magnitude of these effects is not well understood. The acidity of ocean waters is known to alter the ability of organisms, such as clams or corals, to create the hard shells and reefs that they depend on for survival, effects with largely unknown consequences for the various fish stocks and other marine products consumed around the world.

**Ecosystems**

Numerous ecosystem changes that can be directly related to climate change have been observed. In many locations around the world, a broad suite of terrestrial organisms is migrating toward higher altitudes and latitudes (Chen, Hill, Ohlemüller, Roy, and Thomas 2011). In the oceans, fish are migrating to stay within their preferred water temperatures (Pinsky, Worm, Fogarty, Sarmiento, and Levin 2013). Under moderate- and high-emissions scenarios, many slow-moving terrestrial species like the coastal redwoods (Roberts and Hamann 2016) may be unable to track the northward movement of climate zones—roughly 0.1–1.3 km per year (Loarie et al. 2009)—quickly enough to stay within their thermal tolerances. In many instances, the ability of species to migrate is further aggravated by fragmentation of habitat. Overall, high-latitude ecosystems are likely to be transformed by invasions from lower latitudes, and extinctions may be common at lower latitudes (Pörtner et al. 2014).

Coral reefs—home to more than a million species—are threatened by both high temperatures and ocean acidification (Hoegh-Guldberg et al. 2007). Bleaching events associated with high temperatures have become more frequent and extensive, with widespread events spanning the tropics in 1998, 2010, and 2015–2016 (Hughes et al. 2017), and a majority of coral reefs around the world are projected to be at risk of degradation even under a low-emissions scenario (Frieler et al. 2013).

The relationships between climate change and ecosystem change can be difficult to untangle, because climate is only one of many human-caused factors affecting ecosystems. For example, land-use change, overexploitation, species introductions, nitrogen deposition, and water resource development also play major roles and may exhibit trends that are correlated with climate change (Chapin et al. 2000). But it is worth noting that although the causes of mass extinction events in Earth's geological history are complex and difficult to pin down, a growing body of evidence suggests that a number of the largest mass extinctions coincided with large-scale climate changes (for example, Payne and Clapham 2012).

**Tipping Elements and Critical Thresholds**

Nonlinearities and feedbacks in the Earth systems give rise to the potential for multiple stable states of different parts of the Earth system, with potentially rapid lock-in of a state shift once critical thresholds are crossed. These parts of the Earth system are often called "tipping elements," and the their thresholds called "tipping points." However, the "tipping point" language can create confusion about the speed with which state shifts can occur. In popular discourse and much of the economic analysis of climate change, changes associated with a "tipping point" are described as occurring "rapidly." This description is accurate in a geological context, insofar as these changes are rapid relative to other comparable drivers of similar changes in the Earth system. But this description can be misleading in some economic contexts; while some state shifts may be rapid on a human timescale, others may play out over millennia. Below, we summarize a few examples; see Kopp, Shwom, Wagner, and Yuan (2016) for a detailed review.

Tipping elements can exist in the atmosphere/ocean circulation. For example, many of the climate oscillations mentioned earlier, such as El Niño–Southern

Oscillation, occur due to tipping elements, and these patterns may undergo substantial shifts in frequency and amplitude in a warmer climate. Large-scale patterns of ocean circulation, such as the Atlantic Meridional Overturning Circulation—an important component of global ocean circulation that plays a major role in setting temperature, precipitation, and sea level in the North Atlantic—are also potential tipping elements. These atmosphere/ocean tipping elements are among those most likely to undergo rapid shifts.

Tipping elements also occur in ice sheets. For example, positive feedbacks involving ocean-ice sheet interactions might cause sustained ice sheet loss in the Antarctic that would eventually raise global mean sea level by multiple meters or tens of meters. Indeed, some evidence suggests that multiple meters of future sea-level rise from the Antarctic may already be locked in, although depending on the pace of regional warming this rise may take many centuries to manifest.

Tipping elements can also exist in the carbon cycle and in ecosystems. For example, warming of previously frozen soils (permafrost) is allowing microbes to decompose freshly unfrozen organic material into $CO_2$ and methane. These releases may be an important positive feedback on warming, which could potentially amplify warming by several tenths of a degree in the 21st century. Ecosystems are also well known to undergo rapid regime shifts; coral reefs, whose bleaching is discussed above, are a notable example.

## How Economists Can Help Climate Science

We close with a few thoughts on how economists can provide support to climate science.

### Improving Emissions Forecasts

Forecasts for greenhouse gas emissions in the coming decades and centuries are a key ingredient to physical simulations of climate change. Emissions clearly depend on global economic activity, but there is not a one-to-one mapping of economic forecasts and the standardized RCP emissions scenarios discussed earlier. Because global emissions are a single time series, there are many possible future configurations of the global economy, technology, and policy that could produce each emissions trajectory.

The coordinated standardization exercise that produced the RCP emissions scenarios also constructed a set of five Shared Socioeconomic Pathways (SSPs), which represent standardized population projections, forecasts of economic growth and convergence, and forecasts of technological change in both the energy sector and adaptation technologies. SSPs can be loosely thought of as potential "states of the world" that might be realized in the future and which no single country can unilaterally change through policy. The narrative for SSP 1 is "Sustainability," representing a world with low barriers to both mitigation and adaptation; SSP 2 represents a "Middle of the Road" scenario; SSP 3 is "Regional Rivalry," representing a world with high barriers to both mitigation and adaptation; SSP 4 is "Inequality," representing a world with high barriers to adaptation but low barriers to mitigation (due to slow economic growth); and SSP 5 is "Fossil-fueled development," representing

a world with high barriers to mitigation but low barriers to adaptation. Different combinations of SSPs and policy choices give rise to different global emission trajectories. For example, SSP 5 in the absence of policy measures can give rise to emissions consistent with RCP 8.5, but other SSPs would require carbon subsidies or similar policies to give rise to such high emissions. In all SSPs, emissions low enough to be consistent with RCPs 4.5 and 2.6 require carbon mitigation policy.

The construction of these Shared Socioeconomic Pathways and the corresponding sets of RCP emissions scenarios represent the output of a modeling program coordinated across numerous research groups (Moss et al. 2010). The energy/agriculture/economic/climate integrated assessment models used to construct the scenarios are elaborate process models that have been assembled by interdisciplinary teams of engineers and economic modelers, mostly from energy economics. At the heart of most models are assumptions about exogenous population growth and about the rate and convergence of technical change. Researchers using these models have addressed many issues over the years: for example, the tradeoffs among different technologies and their roles in meeting different emissions targets (for example, Clarke et al. 2009).

However, many economists with expertise that would be useful to these modeling exercises have remained unengaged with (or unaware of) this enterprise. In our view, deepening engagements with economists in subfields outside of energy economics, such as macroeconomics, development economics, and political economy, will help strengthen and accelerate this research program (see also Barron 2018). Further, there should be a stronger emphasis on using empirical results and hindcasting experiments to constrain the behavior of these models (for example, Calvin, Wise, Kyle, Clarke, and Edmonds 2017).

**Focusing Climate Research to Support Investigation of Economic and Social Questions**

Much of climate science research is focused on answering key research questions formulated by physical scientists about the nature of the global climate system. These questions are scientifically important and of substantial consequence, but in many cases, key questions or measurements about the climate system that are *economically* or *socially* important remain unanswered. We see three general areas where the potential gains from intellectual exchanges between climate scientists and economists seem large.

First, empirically disentangling the economic consequences of climate change is a large research enterprise (discussed by Auffhammer in this symposium; for a summary, see Carleton and Hsiang 2016). Such analyses universally require some "data engineering" to map physical observations appropriately onto social systems (Auffhammer, Hsiang, Schlenker, and Sobel 2013; Hsiang 2016). Several advances in this literature have arisen from methodological innovations in how physical information from climate science has been summarized and integrated into theoretical and econometric models. For example, Schlenker and Roberts (2009) developed an approach for accounting for the accumulating effects of exposure to extreme heat; Hsiang (2010) developed a method for reconstructing continuous human

exposure to tropical cyclones using standard, albeit limited, meteorological data; Hsiang, Meng, and Cane (2011) introduced a technique for identifying populations heavily impacted by the El Niño–Southern Oscillation; and Proctor et al. (2018) developed an approach to isolate the optical effects of overhead volcanic aerosols. These innovations required both insight into the physics of the climate system coupled to insights from economists regarding the construction of economically meaningful measures. Continued empirical progress will require deepening engagement between researchers in these two fields.

Second, global intercomparison programs for climate change models specify outputs that modeling groups must record in order to participate, but in the past these outputs have been geared towards scientific questions rather than toward calculation of economic outcomes that result from climate change. For example, the expected future correlation of drought events across the major agricultural regions is likely to be important for future food prices, but statistics about such correlated extremes are not commonly computed (Kopp, Easterling, et al. 2017). By supplying the climate modeling community with information on what variables, patterns, and scales are of key economic interest, economists can help the climate modeling community synthesize and output their findings in new and useful ways.

Finally, economists can aid climate scientists in identifying research questions that would be the most valuable to address from a socioeconomic perspective. For example, economists can help distinguish between those economic impacts for which it is most valuable to improve climate model resolution, and those for which is it more important to explore structural or parametric uncertainties, even if at lower resolution. Economists can help distinguish between economic outcomes for which it is more important to constrain long-term climate sensitivity and those for which better characterization of short-term responses, natural variability, or spatial patterns is a priority. Economists can also identify new climate observations that would contribute the greatest value to improved risk-management strategies. In short, economists are well suited to support climate scientists by valuing the different types of information that these scientists could potentially provide to the world.

Climate change management is an urgent and pervasive societal challenge. Natural scientists have had a century's head start over economists in studying the topic, but at this point, input from social scientists can be especially productive. We hope this article can spur more economists to engage in this challenge.

■ *We thank Tamma Carleton, Megan Lickley, Daniel Gilford, and DJ Rasmussen for helpful comments.*

## References

**Abatzoglou, John T., and A. Park Williams.** 2016. "Impact of Anthropogenic Climate Change on Wildfire across Western US Forests." *PNAS* 113 (42): 11770–75.

**Allen, Myles R., David J. Frame, Chris Huntingford, Chris D. Jones, Jason A. Lowe, Malte Meinshausen, and Nicolai Meinshausen.** 2009. "Warming Caused by Cumulative Carbon Emissions towards the Trillionth Tonne." *Nature* 458(7242): 1163–66.

**Archer, David, Michael Eby, Victor Brovkin, Andy Ridgwell, Long Cao, Uwe Mikolajewicz, Ken Caldeira, Katsumi Matsumoto, Guy Munhoven, Alvaro Montenegro, and Kathy Tokos.** 2009. "Atmospheric Lifetime of Fossil Fuel Carbon Dioxide." *Annual Review of Earth and Planetary Sciences* 37(1): 117–34.

**Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel.** 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy* 7(2): 181–98.

**Barron, Alexander R.** 2018. "Time to Refine Key Climate Policy Models." *Nature Climate Change* 8(5): 350–52.

**Beckley, B., N.P. Zelensky, S.A. Holmes, F.G. Lemoine, R.D. Ray, G.T. Mitchum, S. Desai, S.T. Brown.** 2016. Global Mean Sea Level Trend from Integrated Multi-Mission Ocean Altimeters TOPEX/Poseidon Jason-1 and OSTM/Jason-2 Version 4.2. Ver. 4.2. PO.DAAC, CA, USA, http://dx.doi.org/10.5067/GMSLM-TJ142.

**Bell, K., S. Sum, J. Longmate, J. Tseng, and S. Hsiang.** 2018. "Empirically Valuing Factors of Natural Capital in Pasture Systems." Working paper.

**Berger, André.** 2012. "A Brief History of the Astronomical Theories of Paleoclimates." In *Climate Change: Inferences from Paleoclimate and Regional Aspects*, edited by André Berger, Fedor Mesinger, and Djordje Sijacki. 107–129. Springer.

**Bindoff, N. L., P. A. Stott, K. M. AchutaRao, M. R. Allen, N. Gillett, D. Gutzler, K. Hansingo, G. Hegerl, Y. Hu, S. Jain, I. I. Mokhov, J. Overland, J. Perlwitz, R. Sebbari and X. Zhang.** 2013. "Detection and Attribution of Climate Change: From Global to Regional." Chap. 10 in *Climate Change 2013: The Physical Science Basis*, Working Group 1 Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, edited by Thomas F. Stocker et al. Cambridge, UK: Cambridge University Press.

**Blanco G., R. Gerlagh, S. Suh, J. Barrett, H. C. de Coninck, C. F. Diaz Morejon, R. Mathur, N. Nakicenovic, A. Ofosu Ahenkora, J. Pan, H. Pathak, J. Rice, R. Richels, S. J. Smith, D. I. Stern, F. L. Toth, and P. Zhou.** 2014. "Drivers, Trends and Mitigation." Chap. 5 in *Climate Change 2014: Mitigation of Climate Change*, Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, edited by O. Edenhofer et al. Cambridge University Press..

**Boden, T., G. Marland, and R. J. Andres.** 2017. "Global, Regional, and National Fossil-Fuel $CO_2$ Emissions (1751–2014) (V. 2017)." Oak Ridge, TN: Carbon Dioxide Information Analysis Center (CDIAC), Oak Ridge National Laboratory (ORNL). http://dx.doi.org/10.3334/CDIAC/00001_V2017.

**Boucher, O., D. Randall, P. Artaxo, C. Bretherton, G. Feingold, P. Forster, V.-M. Kerminen, Y. Kondo, H. Liao, U. Lohmann, P. Rasch, S.K. Satheesh, S. Sherwood, B. Stevens, and X.Y. Zhang.** 2013. "Clouds and Aerosols." Chap. 7 in *Climate Change 2013: The Physical Science Basis*. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, edited by T. F. Stocker et al. Cambridge University Press.

**Buchanan, Maya K., Michael Oppenheimer, and Robert E. Kopp.** 2017. "Amplification of Flood Frequencies with Local Sea Level Rise and Emerging Flood Regimes." *Environmental Research Letters* 12(6): 064009.

**Burke, Marshall, Solomon M. Hsiang, and Edward Miguel.** 2015. "Global Non-Linear Effect of Temperature on Economic Production." *Nature* 527 (7577): 235–39.

**Caballero, Rodrigo, and Matthew Huber.** 2013. "State-Dependent Climate Sensitivity in Past Warm Climates and Its Implications for Future Climate Projections." *PNAS* 110(35): 14162–67.

**Cahill, Niamh, Stefan Rahmstorf, and Andrew C. Parnell.** 2015. "Change Points of Global Temperature." *Environmental Research Letters* 10(8): 084002.

**Caldeira, Ken, Govindasamy Bala, and Long Cao.** 2013. "The Science of Geoengineering." *Annual Review of Earth and Planetary Sciences* 41(1): 231–56.

**Calvin, Katherine, Marshall Wise, Page Kyle, Leon Clarke, and Jae Edmonds.** 2017. "A Hindcast Experiment Using the Gcam 3.0 Agriculture and Land-Use Module." *Climate Change Economics* 08(01): 1750005.

**Carleton, Tamma A., and Solomon M. Hsiang.** 2016. "Social and Economic Impacts of Climate." *Science* 353(6304): aad9837.

**Chapin, F. Stuart, Erika S. Zavaleta, Valerie T. Eviner, Rosamond L. Naylor, Peter M. Vitousek, Heather L. Reynolds, David U. Hooper, Sandra Lavorel, Osvaldo E. Sala, Sarah E. Hobbie, Michelle C. Mack, and Sandra Díaz.** 2000. "Consequences of Changing Biodiversity." *Nature,* 405: 234–242.

**Chen, I.-Ching, Jane K. Hill, Ralf Ohlemüller,**

**David B. Roy, and Chris D. Thomas.** 2011. "Rapid Range Shifts of Species Associated with High Levels of Climate Warming." *Science* 333 (6045): 1024–26.

**Christensen, J. H., K. Kanikicharla, E. Aldrian, S.-I. An, I. F. A. Cavalcanti, M. de Castro, W. Dong, P. Goswami, A. Hall, J. K. Kanyanga, A. Kitoh, J. Kossin, N.-C. Lau, J. Renwick, D. B. Stephenson, S.-P. Xie and T. Zhou.** 2013. "Climate Phenomena and Their Relevance for Future Regional Climate Change." Chap. 14 in *Climate Change 2013: The Physical Science Basis*, edited by Thomas F. Stocker, et al. Cambridge, UK: Cambridge University Press. Available at: http://www.ipcc.ch/report/ar5/wg1/.

**Clarke, Leon, Jae Edmonds, Volker Krey, Richard Richels, Steven Rose, and Massimo Tavoni.** 2009. "International Climate Policy Architectures: Overview of the EMF 22 International Scenarios." *Energy Economics* 31(S2): S64–S81.

**Clement, Amy, Katinka Bellomo, Lisa N. Murphy, Mark A. Cane, Thorsten Mauritsen, Gaby Rädel, and Bjorn Stevens.** 2015. "The Atlantic Multidecadal Oscillation without a Role for Ocean Circulation." *Science* 350 (6258): 320–24.

**Collins, M., R. Knutti, J. Arblaster, J.-L. Dufresne, T. Fichefet, P. Friedlingstein, X. Gao, W.J. Gutowski, T. Johns, G. Krinner, M. Shongwe, C. Tebaldi, A.J. Weaver, and M. Wehner.** 2013. "Long-Term Climate Change: Projections, Commitments and Irreversibility." Chap. 12 in *Climate Change 2013: The Physical Science Basis*, edited by Thomas F. Stocker, et al. Cambridge University Press. Available at: http://www.ipcc.ch/report/ar5/wg1/.

**Cook, John, Naomi Oreskes, Peter T. Doran, William R. L. Anderegg, Bart Verheggen, Ed W. Maibach, J. Stuart Carlton, Stephan Lewandowsky, Andrew G., Skuce, Sarah A. Green, Dana Nuccitelli, Peter Jacobs, Mark Richardson, Bärbel Winkler, Rob Painting, and Ken Rice.** 2016. "Consensus on Consensus: A Synthesis of Consensus Estimates on Human-Caused Global Warming." *Environmental Research Letters* 11(4): 048002.

**Cramer, B. S., K. G. Miller, P. J Barrett, and J. D. Wright.** 2011. "Late Cretaceous–Neogene Trends in Deep Ocean Temperature and Continental Ice Volume: Reconciling Records of Benthic Foraminiferal Geochemistry (Δ18O and Mg/Ca) with Sea Level History." *Journal of Geophysical Research: Oceans* 116 (C12): C12023.

**Edwards, R. J., and B. P. Horton.** 2000. "Reconstructing Relative Sea-Level Change Using UK Salt-Marsh Foraminifera." *Marine Geology* 169(1–2): 41–56.

**Emanuel, Kerry.** 2017. "Assessing the Present and Future Probability of Hurricane Harvey's Rainfall." *PNAS*, 114 (48): 12681–12684. https://doi.org/10.1073/pnas.1716222114.

**Fischer, E. M., and R. Knutti.** 2013. "Robust Projections of Combined Humidity and Temperature Extremes." *Nature Climate Change* 3(2): 126–30.

**Frieler, K., M. Meinshausen, A. Golly, M. Mengel, K. Lebek, S. D. Donner, and O. Hoegh-Guldberg.** 2013. "Limiting Global Warming to 2 °C Is Unlikely to Save Most Coral Reefs." *Nature Climate Change* 3(2): 165–70.

**GISTEMP Team.** 2018. GISS Surface Temperature Analysis (GISTEMP). Dataset. NASA Goddard Institute for Space Studies. https://data.giss.nasa.gov/gistemp/.

**Greve, Peter, Boris Orlowsky, Brigitte Mueller, Justin Sheffield, Markus Reichstein, and Sonia I. Seneviratne.** 2014. "Global Assessment of Trends in Wetting and Drying over Land." *Nature Geoscience* 7(10): 716–21.

**Hansen, J., D. Johnson, A. Lacis, S. Lebedeff, P. Lee, D. Rind, and G. Russell.** 1981. "Climate Impact of Increasing Atmospheric Carbon Dioxide." *Science* 213(4511): 957–66.

**Hartmann, D. L., A. M. G. Klein Tank, M. Rusticucci, L. V. Alexander, S. Brönnimann, Y. Charabi, F.J. Dentener, E. J. Dlugokencky, D. R. Easterling, A. Kaplan, B. J. Soden, P. W. Thorne, M. Wild, and P.M. Zhai.** 2013. "Observations: Atmosphere and Surface." Chap. 2 in *Climate Change 2013: The Physical Science Basis*, edited by Thomas F. Stocker, et al. Cambridge University Press. Available at: http://www.ipcc.ch/report/ar5/wg1/.

**Hawkins, Ed, and Rowan Sutton.** 2009. "The Potential to Narrow Uncertainty in Regional Climate Predictions." *Bulletin of the American Meteorological Society* 90(8): 1095–1107.

**Hay, Carling C., Eric Morrow, Robert E. Kopp, and Jerry X. Mitrovica.** 2015. "Probabilistic Reanalysis of Twentieth-Century Sea-Level Rise." *Nature* 517(7535): 481–84.

**Hayhoe, K., J. Edmonds, R. E. Kopp, A. N. LeGrande, B. M. Sanderson, M. F. Wehner, and D. J. Wuebbles.** 2017. "Climate Models, Scenarios, and Projections." Chap. 4 in *Climate Science Special Report: Fourth National Climate Assessment,* Vol. I, edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock, 411–29. Washington, DC, USA: U.S. Global Change Research Program.

**Hegerl, G. C., F. W. Zwiers, P. Braconnot, N. P. Gillett, Y. Luo, J. A. Marengo Orsini, N. Nicholls, J. E. Penner, and P. A. Stott.** 2007. "Understanding and Attributing Climate Change." Chap. 9 in *Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, edited by S. Solomon, D. Qin, M. Manning, Z. Chen, M. Marquis, K. B. Averyt, M. Tignor, and H. L. Miller. Cambridge University Press.

**Hoegh-Guldberg, O., et al.** 2007. "Coral Reefs Under Rapid Climate Change and Ocean Acidification." *Science* 318(5857): 1737–42.

**Hönisch, Bärbel, et al.** 2012. "The Geological Record of Ocean Acidification." *Science* 335(6072): 1058–63.

**Horton, Benjamin P., Robert E. Kopp, Andra J. Garner, Carling C. Hay, Nicole S. Khan, Keven Roy, and Timothy A. Shaw.** 2018. "Mapping Sea-Level Change in Time, Space and Probability." *Annual Review of Environment and Resources.* https://doi.org/10.1146/annurev-environ-102017-025826.

**Houser, Trevor, Solomon Hsiang, Robert Kopp, Kate Larsen, Michael Delgado, Amir Jina, Michael Mastrandrea, Shashank Mohan, Robert Muir-Wood, D. J. Rasmussen, James Rising, and Paul Wilson.** 2015. *Economic Risks of Climate Change: An American Prospectus.* New York: Columbia University Press.

**Hsiang, Solomon M.** 2010. "Temperatures and Cyclones Strongly Associated with Economic Production in the Caribbean and Central America." *PNAS* 107(35): 15367–72.

**Hsiang, Solomon.** 2016. "Climate Econometrics." *Annual Review of Resource Economics* 8(1): 43–75.

**Hsiang, Solomon, Robert Kopp, Amir Jina, James Rising, Michael Delgado, Shashank Mohan, D. J. Rasmussen, Robert Muir-Wood, Paul Wilson, Michael Oppenheimer, Kate Larsen, and Trevor Houser.** 2017. "Estimating Economic Damage from Climate Change in the United States." *Science* 356(6345): 1362–1369.

**Hsiang, Solomon M., and Kyle C. Meng.** 2015. "Tropical Economics." *American Economic Review* 105(5): 257–61.

**Hsiang, Solomon M., Kyle C. Meng, and Mark A. Cane.** 2011. "Civil Conflicts Are Associated with the Global Climate." *Nature* 476(7361): 438–41.

**Hughes, Terry P. et al.** 2017. "Global Warming and Recurrent Mass Bleaching of Corals." *Nature* 543(7645): 373–77.

**Hurrell, James W.** 1995. "Decadal Trends in the North Atlantic Oscillation: Regional Temperatures and Precipitation." *Science* 269(5224): 676–79.

**Janssens-Maenhout, G., M. Crippa, D. Guizzardi, M. Muntean, E. Schaaf, J. G. J. Olivier, J. A. H. W. Peters, and K. M. Schure.** 2017. *Fossil CO2 & GHG Emissions of All World Countries.* JRC Science for Policy Report, EUR 28766 EN. Luxembourg: Publications Office of the European Union.

**Jewett, L., and A. Romanou.** 2017. "Ocean Acidification and Other Ocean Changes." Chap. 13 in *Climate Science Special Report: Fourth National Climate Assessment,* Vol. I, edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock, 411–29. Washington, DC, USA: U.S. Global Change Research Program.

**Jimenez Cisneros, Blanca E., Taikan Oki, Nigel W. Arnell, Gerardo Benito, J. Graham Cogley, Petra Döll, Tong Jiang, and Shadrack S. Mwakalila.** 2014. "Freshwater Resources." In *Climate Change 2014: Impacts, Adaptation and Vulnerability,* edited by C. B. Field et al., 229–269. Cambridge University Press.

**Jones, Gareth S., Peter A. Stott, and Nikolaos Christidis.** 2013. "Attribution of Observed Historical Near-Surface Temperature Variations to Anthropogenic and Natural Causes Using CMIP5 Simulations." *Journal of Geophysical Research: Atmospheres* 118(10): 4001–24.

**Jones, P. D., et al.** 2009. "High-Resolution Palaeoclimatology of the Last Millennium: A Review of Current Status and Future Prospects." *The Holocene* 19(1): 3–49.

**Joos, F., et al.** 2013. "Carbon Dioxide and Climate Impulse Response Functions for the Computation of Greenhouse Gas Metrics: A Multi-Model Analysis." *Atmospheric Chemistry and Physics* 13(5): 2793–2825.

**Keeling, Charles D., Stephen C. Piper, Robert B. Bacastow, Martin Wahlen, Timothy P. Whorf, Martin Heimann, and Harro A. Meijer.** 2001. "Exchanges of Atmospheric $CO_2$ and $13CO_2$ with the Terrestrial Biosphere and Oceans from 1978 to 2000. I. Global Aspects," June. SIO Reference no 01-06. https://escholarship.org/uc/item/09v319r9.

**Knutson, Thomas R., John L. McBride, Johnny Chan, Kerry Emanuel, Greg Holland, Chris Landsea, Isaac Held, James P. Kossin, A. K. Srivastava, and Masato Sugi.** 2010. "Tropical Cyclones and Climate Change." *Nature Geoscience* 3(3): 157.

**Kopp, Robert E., Robert M. DeConto, Daniel A. Bader, Radley M. Horton, Carling C. Hay, Scott Kulp, Michael Oppenheimer, David Pollard, and Benjamin H. Strauss.** 2017. "Evolving Understanding of Antarctic Ice-Sheet Physics and Ambiguity in Probabilistic Sea-Level Projections." *Earth's Future* 5(12): 1217–33.

**Kopp, Robert E., D. R. Easterling, T. Hall, K. Hayhoe, R. Horton, K. E. Kunkel, and A. N. LeGrande.** 2017. "Potential Surprises—Compound Extremes and Tipping Elements." Chap 15 in *Climate Science Special Report: Fourth National Climate Assessment, Volume I,* edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock. Washington, DC, USA: U.S. Global Change Research Program.

**Kopp, Robert E., Carling C. Hay, Christopher M. Little, and Jerry X. Mitrovica.** 2015. "Geographic Variability of Sea-Level Change." *Current Climate Change Reports* 1(3): 192–204.

**Kopp, Robert E., Andrew C. Kemp, Klaus Bittermann, Benjamin P. Horton, Jeffrey P. Donnelly, W. Roland Gehrels, Carling C. Hay, Jerry X. Mitrovica, Eric D. Morrow, and Stefan Rahmstorf.** 2016. "Temperature-Driven Global Sea-Level Variability in the Common Era." *PNAS* 113(11): E1434–41.

**Kopp, Robert E., Rachael Shwom, Gernot Wagner, and Jiacan Yuan.** 2016. "Tipping Elements and Climate-Economic Shocks: Pathways toward Integrated Assessment." *Earth's Future* 4(8): 346–72.

**Kossin, J. P., T. Hall, T. Knutson, K. E. Kunkel,**

**R. J. Trapp, D. E. Waliser, and M. F. Wehner.** 2017. "Extreme Storms." Chap. 9 in *Climate Science Special Report: Fourth National Climate Assessment, Volume I*, edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock. Washington, DC, USA: U.S. Global Change Research Program.

**Lambeck, Kurt, Hélène Rouby, Anthony Purcell, Yiying Sun, and Malcolm Sambridge.** 2014. "Sea Level and Global Ice Volumes from the Last Glacial Maximum to the Holocene." *PNAS*: 111(43): 15296–303.

**Le Quéré, Corinne, et al.** 2018. "Global Carbon Budget 2017." *Earth System Science Data* 10(1): 405–48.

**Loarie, Scott R., Philip B. Duffy, Healy Hamilton, Gregory P. Asner, Christopher B. Field, and David D. Ackerly.** 2009. "The Velocity of Climate Change." *Nature* 462(7276): 1052–55.

**Lorenz, Edward N.** 1963. "Deterministic Nonperiodic Flow." *Journal of the Atmospheric Sciences* 20(2): 130–41.

**Lüthi, Dieter, Martine Le Floch, Bernhard Bereiter, Thomas Blunier, Jean-Marc Barnola, Urs Siegenthaler, Dominique Raynaud, Jean Jouzel, Hubertus Fischer, Kenji Kawamura, and Thomas F. Stocker.** 2008. "High-Resolution Carbon Dioxide Concentration Record 650,000–800,000 Years before Present." *Nature* 453(7193): 379–82.

**MacFarling Meure, C., D. Etheridge, C. Trudinger, P. Steele, R. Langenfelds, T. Van Ommen, A. Smith, and J. Elkins.** 2006. "Law Dome $CO_2$, $CH_4$ and $N_2O$ Ice Core Records Extended to 2000 Years BP." *Geophysical Research Letters* 33(14): L14810.

**Maher, Nicola, Alexander Sen Gupta, and Matthew H. England.** 2014. "Drivers of Decadal Hiatus Periods in the 20th and 21st Centuries." *Geophysical Research Letters* 41(16): 5978–86.

**Manabe, Syukuro, and Joseph Smagorinsky.** 1965. "Simulated Climatology of a General Circulation Model with a Hydrological Cycle." *Monthly Weather Review* 95(4): 769–98.

**Manabe, Syukuro, and Robert F. Strickler.** 1964. "Thermal Equilibrium of the Atmosphere with a Convective Adjustment." *Journal of the Atmospheric Sciences* 21(4): 361–85.

**Mantua, Nathan J., and Steven R. Hare.** 2002. "The Pacific Decadal Oscillation." *Journal of Oceanography* 58(1): 35–44.

**Marcott, Shaun A., Jeremy D. Shakun, Peter U. Clark, and Alan C. Mix.** 2013. "A Reconstruction of Regional and Global Temperature for the Past 11,300 Years." *Science* 339(6124): 1198–1201.

**Marsicek, Jeremiah, Bryan N. Shuman, Patrick J. Bartlein, Sarah L. Shafer, and Simon Brewer.** 2018. "Reconciling Divergent Trends and Millennial Variations in Holocene Temperatures." *Nature* 554(7690): 92–96.

**Medhaug, Iselin, Martin B. Stolpe, Erich M. Fischer, and Reto Knutti.** 2017. "Reconciling Controversies about the 'Global Warming Hiatus.'" *Nature* 545(7652): 41–47.

**Meehl, Gerald A., et al.** 2009. "Decadal Prediction: Can It Be Skillful?" *Bulletin of the American Meteorological Society* 90(10): 1467–86.

**Morice, Colin P., John J. Kennedy, Nick A. Rayner, and Phil D. Jones.** 2012. "Quantifying Uncertainties in Global and Regional Temperature Change Using an Ensemble of Observational Estimates: The HadCRUT4 Data Set." *Journal of Geophysical Research: Atmospheres* 117(D8): D08101.

**Moss, Richard H., et al.** 2010. "The next Generation of Scenarios for Climate Change Research and Assessment." *Nature* 463(7282): 747–56.

**Nordhaus, William D.** 1977. "Economic Growth and Climate: The Carbon Dioxide Problem." *American Economic Review* 67(1): 341–46.

**Payne, Jonathan L., and Matthew E. Clapham.** 2012. "End-Permian Mass Extinction in the Oceans: An Ancient Analog for the Twenty-First Century?" *Annual Review of Earth and Planetary Sciences* 40(1): 89–111.

**Pinsky, Malin L., Boris Worm, Michael J. Fogarty, Jorge L. Sarmiento, and Simon A. Levin.** 2013. "Marine Taxa Track Local Climate Velocities." *Science* 341(6151): 1239–42.

**Pongratz, Julia, and Ken Caldeira.** 2012. "Attribution of Atmospheric $CO_2$ and Temperature Increases to Regions: Importance of Preindustrial Land Use Change." *Environmental Research Letters* 7(3): 034001.

**Pörtner, Hans-Otto, David M. Karl, Philip W. Boyd, William Cheung, Salvador E. Lluch-Cota, Yukihiro Nojiri, Daniela N. Schmidt, Peter O. Zavialov, Jürgen Alheit, and Javier Aristegui.** 2014. "Ocean Systems." Chap. 6 in *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Part A: Global and Sectoral Aspects. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.* Cambridge University Press.

**Proctor, Jonathan, Solomon Hsiang, Jennifer Burney, Marshall Burke, and Wolfram Schlenker.** 2018. "Estimating Global Agricultural Effects of Geoengineering Using Volcanic Eruptions." *Nature* 560(7719): 480–83.

**Roberts, David R., and Andreas Hamann.** 2016. "Climate Refugia and Migration Requirements in Complex Landscapes." *Ecography* 39(12): 1238–46.

**Rohde, Robert, Richard A. Muller, Robert Jacobsen, Elizabeth Muller, Saul Perlmutter, Arthur Rosenfeld, Jonathan Wurtele, Donald Groom, and Charlotte Wickham.** 2013. "A New Estimate of the Average Earth Surface Land Temperature Spanning 1753 to 2011." *Geoinformatics & Geostatistics: An Overview* 1(1).

**Rosenfeld, Daniel, Ulrike Lohmann, Graciela B.**

**Raga, Colin D. O'Dowd, Markku Kulmala, Sandro Fuzzi, Anni Reissell, and Meinrat O. Andreae.** 2008. "Flood or Drought: How Do Aerosols Affect Precipitation?" *Science* 321(5894): 1309–13.

**Schlenker, Wolfram, and Michael J. Roberts.** 2009. "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields Under Climate Change." *PNAS* 106(37): 15594–15598.

**Serreze, Mark C., and Roger G. Barry.** 2011. "Processes and Impacts of Arctic Amplification: A Research Synthesis." *Global and Planetary Change* 77(1): 85–96.

**Sherwood, Steven C., and Matthew Huber.** 2010. "An Adaptability Limit to Climate Change Due to Heat Stress." *PNAS* 107(21): 9552–55.

**Sherwood, Steven, and Qiang Fu.** 2014. "A Drier Future?" *Science* 343(6172): 737–39.

**Stocker, T. F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P.M. Midgley (eds).** 2013. *Climate Change 2013: The Physical Science Basis.* Contribution of Working Group 1 to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press.

**Storlazzi, Curt D., Stephen B. Gingerich, Ap van Dongeren, Olivia M. Cheriton, Peter W. Swarzenski, Ellen Quataert, Clifford I. Voss, Donald W. Field, Hariharasubramanian Annamalai, Greg A. Piniak, and Robert McCall.** 2018. "Most Atolls Will Be Uninhabitable by the Mid-21st Century Because of Sea-Level Rise Exacerbating Wave-Driven Flooding." *Science Advances* 4(4): eaap9741.

**Supran, Geoffrey, and Naomi Oreskes.** 2017. "Assessing ExxonMobil's Climate Change Communications (1977–2014)." *Environmental Research Letters* 12(8): 084019.

**Sweet, William V., R. Horton, R. E. Kopp, A. N. LeGrande, and A. Romanou.** 2017. "Sea Level Rise." Chap. 12 in *Climate Science Special Report: Fourth National Climate Assessment,* Vol. I, edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock, 333–63. Washington, DC, USA: U.S. Global Change Research Program.

**Sweet, William V., and Joseph Park.** 2014. "From the Extreme to the Mean: Acceleration and Tipping Points of Coastal Inundation from Sea Level Rise." *Earth's Future* 2(12): 579–600.

**Swiss Re Institute.** 2018. "Sigma 2/2017: Natural Catastrophes and Man-Made Disasters in 2016: A Year of Widespread Damages." http://institute.swissre.com/research/overview/sigma/2_2017.html.

**Taylor, Karl E., Ronald J. Stouffer, and Gerald A. Meehl.** 2012. "An Overview of CMIP5 and the Experiment Design." *Bulletin of the American Meteorological Society* 93(4): 485–98.

**Trouet, Valerie, Grant L. Harley, and Marta Domínguez-Delmás.** 2016. "Shipwreck Rates Reveal Caribbean Tropical Cyclone Response to Past Radiative Forcing." *PNAS* 113(12): 3169–74.

**Trouet, V., J. D. Scourse, and C. C. Raible.** 2012. "North Atlantic Storminess and Atlantic Meridional Overturning Circulation during the Last Millennium: Reconciling Contradictory Proxy Records of NAO Variability." *Global and Planetary Change* 84–85: 48–55.

**U.S. Global Change Research Program (USGCRP).** 2017. *Climate Science Special Report: Fourth National Climate Assessment,* Vol I. Edited by Donald J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock. Washington, DC: U.S. Global Change Research Program.

**Walsh, Kevin J. E., John L. McBride, Philip J. Klotzbach, Sethurathinam Balachandran, Suzana J. Camargo, Greg Holland, Thomas R. Knutson, James P. Kossin, Tsz-cheung Lee, Adam Sobel, and Masato Sugi.** 2016. "Tropical Cyclones and Climate Change." *Wiley Interdisciplinary Reviews: Climate Change* 7(1): 65–89.

**Weart, Spencer R.** 2018. *The Discovery of Global Warming.* A hypertext history. February. https://history.aip.org/climate/index.htm.

**Wehner, M. F., J. R. Arnold, T. Knutson, K. E. Kunkel, and A. N. Legrande.** 2017. "Droughts, Floods, and Wildfires." Chap. 8 in *Climate Science Special Report: Fourth National Climate Assessment,* Vol. I, edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock. Washington, DC: U.S. Global Change Research Program.

**Westervelt, D. M., L. W. Horowitz, V. Naik, J.-C. Golaz, and D. L. Mauzerall.** 2015. "Radiative Forcing and Climate Response to Projected 21st Century Aerosol Decreases." *Atmospheric Chemistry and Physics* 15(22): 12681–703.

**Wigley, T. M. L.** 2011. "Coal to Gas: The Influence of Methane Leakage." *Climatic Change* 108(3): 601–608.

**Wood, A. W., L. R. Leung, V. Sridhar, and D. P. Lettenmaier.** 2004. "Hydrologic Implications of Dynamical and Statistical Approaches to Downscaling Climate Model Outputs." *Climatic Change* 62 (1–3): 189–216.

**Wuebbles, D. J., D. R. Easterling, K. Hayhoe, T. Knutson, R. E. Kopp, J. P. Kossin, K. E. Kunkel, et al.** 2017. "Our Globally Changing Climate." Chap. 1 in *Climate Science Special Report: Fourth National Climate Assessment,* Vol. I, edited by D. J. Wuebbles, D. W. Fahey, K. A. Hibbard, D. J. Dokken, B. C. Stewart, and T. K. Maycock. Washington, DC, USA: U.S. Global Change Research Program.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 33–52*

# Quantifying Economic Damages from Climate Change

Maximilian Auffhammer

Climate scientists have spent billions of dollars and eons of supercomputer time studying how increased concentrations of greenhouse gases and changes in the reflectivity of the earth's surface affect dimensions of the climate system relevant to human society: surface temperature, precipitation, humidity, and sea levels. Recent incarnations of physical climate models have become sophisticated enough to be able to simulate intensities and frequencies of some extreme events, like tropical storms, under different warming scenarios. The current consensus estimates from what may be the most heavily peer-reviewed scientific publication in human history, the 5th Assessment Report of the Intergovernmental Panel on Climate Change, are that the average global surface temperature has increased by 0.85° Celsius (1.5° F) since the industrial revolution. Estimates of future warming by the end of the current century range from 0.9 to 5.4°C (1.6 - 9.7°F) (IPCC 2013; Hsiang and Kopp in this issue of the journal).

In a stark juxtaposition, the efforts involved in and the public resources targeted at understanding how these physical changes translate into economic impacts are disproportionately smaller, with most of the major models being developed and maintained with little to no public funding support. This is concerning, because optimal policy design in the context of addressing the biggest environmental market failure in human history requires an understanding of the external cost imposed

■ *Maximilian Auffhammer is the George Pardee Jr. Professor of International Sustainable Development, University of California, Berkeley, California, and Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. His email address is auffhammer@berkeley.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.33

doi=10.1257/jep.32.4.33

by additional emissions of greenhouse gases. Estimating this number is far from straightforward for two main reasons: First, climate change is a global phenomenon and hence local emissions result in global damages, the quantification of which is challenging as damages vary across space and time. Second, greenhouse gases are long-lived, which means that today's emissions affect generations hundreds of years from now. Hence if one would like to calculate the external cost of one more ton of $CO_2$ equivalent emitted—which is about what you would emit if you drove a Ford Mustang GT from San Francisco to Chicago—you would need to calculate the discounted stream of global damages from that additional ton over the next 300 years or so relative to a baseline with one less ton of $CO_2$.

The goal of this paper is first to shed light on how (mostly) economists have gone about calculating this "social cost of carbon" for regulatory purposes and to provide an overview of the past and currently used estimates. In the second part, I will focus on where in this literature empirical economists may have the highest value added: specifically, the calibration and estimation of economic damage functions, which map weather patterns transformed by climate change into economic benefits and damages. A broad variety of econometric methods have recently been used to parameterize the dose (climate) response (economic outcome) functions. The paper seeks to provide both an accessible and comprehensive overview of how economists think about parameterizing damage functions and quantifying the economic damages of climate change. There are a number of more technical surveys, which I invite the interested reader to consult (useful starting points include Carleton and Hsiang 2016; Dell, Jones, and Olken 2014; Diaz and Moore 2017).

## The Social Cost of Carbon

The social cost of carbon is an estimate of the discounted present value of damages from one additional ton of $CO_2$ equivalent emitted at a certain point in time. This social cost of carbon is increasing over time, as later emissions result in larger damages due to the elevated stock of greenhouse gases in the atmosphere, and because GDP grows over time and some damage categories are modeled as proportional to GDP (EPA 2016). Calculations of the social cost of carbon are obtained through so-called Integrated Assessment Models. The most well-known of these models are DICE (Dynamic Integrated Climate–Economy model by 2018 Nobel Laureate William Nordhaus), FUND (Climate Framework for Uncertainty, Negotiation and Distribution model by David Anthoff and Richard Tol), and PAGE (Policy Analysis of the Greenhouse Effect model by Chris Hope), although there are a number of more recent and ambitious modeling efforts. These models "integrate" simple socioeconomic scenarios that produce future emissions trajectories, which are fed into a simple climate model that translates emissions paths into concentrations and produces scenarios for future temperatures, precipitation, and sea levels. These climatic outcomes are then fed into a set of damage functions, which map the climate model output into economic damages at the regional or global level. The discounted difference in damages between a baseline future and a future with one

*Figure 1*
**Sample of Social Cost of Carbon Estimates Used in Federal Rulemakings**

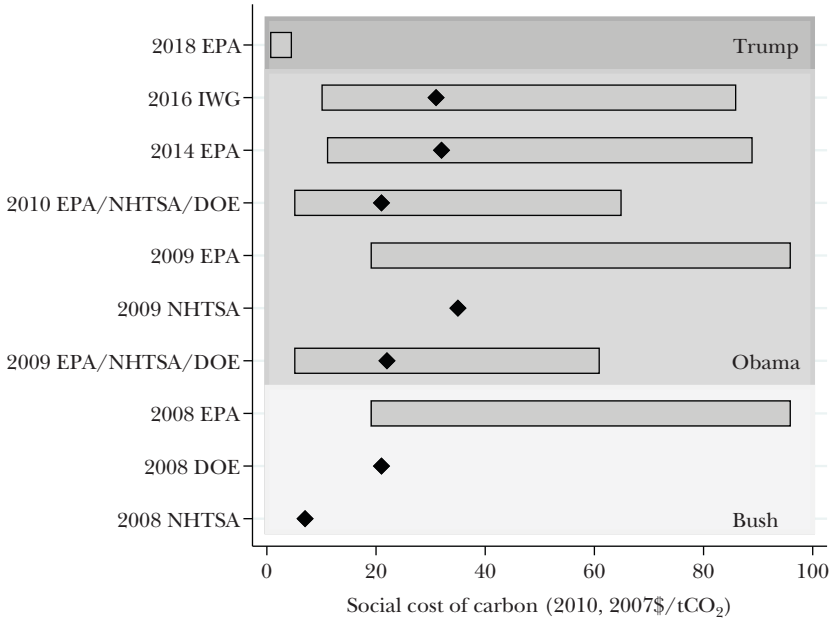


*Sources:* Rose (2012); Rose et al. (2014); IWG (2016); EPA (2018).
*Note:* Estimates for the social cost of carbon are for emissions of a ton of $CO_2$ in 2010 in 2007 dollars. NHTSA is National Highway Traffic Safety Administration; IWG is Interagency Working Group; EPA is Environmental Protection Agency; DOE is Department of Energy. The black diamond indicates the "central estimate," if one was identified. The grey bars indicate selected upper and lower bounds used in regulatory analyses.

more ton of emissions then becomes the social cost of carbon—essentially the external cost of one ton of additional $CO_2$ emissions at a point in time. There is a nascent literature calculating social costs of other greenhouse gases (for example, methane is a more potent greenhouse gas, but with a shorter atmospheric lifetime).

A tremendous number of modeling assumptions need to be made to calculate the social cost of carbon for use in rulemaking. The modeler needs to decide on the time horizon to be considered, the approach to discounting and the rate to be used, the reflection of uncertainties, the changes to risks, which impacts can be included, the choice of reference conditions, whether one should equity weight across countries, and what recent literature should be incorporated (Rose 2012). Among these, the three factors of possibly biggest consequence are the choice of discount rate, which sectors are omitted (for example, ecosystem services), and whether one should consider only domestic or global damages. The latter decision is really a legal question, as the externality is global and hence, from an economic point of view, the global number is the correct estimate of the externality. Figure 1 shows the evolution of the social cost of carbon for a ton emitted in 2010 (measured in 2007 US dollars) in federal rulemaking for a sample of rules.

The first official estimates of the social cost of carbon in 2008 were made under the Bush administration. The 2008 National Highway and Traffic Safety Administration (NHTSA) number was an estimate of global damages used for setting fuel economy standards. The 2008 Department of Energy (DOE) number was a global social cost of carbon used for setting air conditioner equipment and gas range standards. The 2008 Environmental Protection Agency (EPA) estimates were used in the proposed rulemaking for regulating greenhouse gas emissions under the Clean Air Act. The bar here indicates the distribution of the central number used. The actual analysis also considered an additional range from –$7 to $781. It is noteworthy that this first round of proposed rulemaking under the Bush administration stated that $CO_2$ is a global pollutant and that "economic principles suggest that the full costs to society of emissions should be considered in order to identify the policy that maximizes the net benefits to society, i.e., achieves an efficient outcome (Nordhaus, 2006)." The document further acknowledges that "domestic estimates omit potential impacts on the United States (for example, economic or national security impacts) resulting from climate change impacts in other countries" (US EPA 2008).

President Obama convened an Interagency Working Group, which was charged with calculating an official social cost of carbon to be used across the board in federal rulemaking (Greenstone, Kopits, and Wolverton 2013). Three prominent Integrated Assessment Models—Nordhaus's DICE model,[1] Anthoff and Tol's FUND model,[2] and Hope's proprietary PAGE model—were used to calculate a distribution of the social cost of carbon across time and scenarios for a set of common socioeconomic assumptions, discount rates, and uncertainty over a number of parameters. The central and often-cited estimate of the social cost of carbon, which is the mean number across 50,000 simulations for each model at a 3 percent discount rate, is $42 (in 2007 dollars) for one ton of emissions made in the year 2020.[3] If one uses a 5 percent discount rate, this value drops to $12; if one uses a 2.5 percent discount rate, it increases to $62. The Interagency Working Group also ran a so-called "high-impact scenario," which is the 95th percentile number at a 3 percent discount rate and valued at $123. The central estimate of the social cost of carbon was projected to rise to $50/ton in 2030 and $69 in 2050.

The Obama administration later commissioned the National Academies of Sciences to assess the Interagency Working Group modeling exercise and suggest improvements. The National Academies of Sciences (2017) recommended substantial revisions to the way the social cost of carbon is estimated. President Trump, however, disbanded the Interagency Working Group, which could have implemented these changes. Two current proposed rulemakings under the Trump administration use a social cost of carbon that only considers *domestic damages* and discount rates of *3 percent* and *7 percent*.

[1] The DICE model is at https://sites.google.com/site/williamdnordhaus/dice-rice.
[2] The FUND model is at http://www.fund-model.org.
[3] Of course, 42 is also the Answer to the Ultimate Question of Life, the Universe, and Everything, according to the *The Hitchkiker's Guide to the Galaxy*.

The top bar in Figure 1 indicates the range of the domestic social cost of carbon using the 3 percent and 7 percent discount rates currently proposed by the National Highway Traffic Safety Administration for the "revision" of the Corporate Average Fuel Economy (CAFE) standards for fuel economy of cars and light trucks, which clearly represent a drastic decrease in the estimated externality to between $1 and $7 for a ton emitted in 2020.

The estimates also do not incorporate any of the major updates suggested by the National Academies of Sciences (2017) report, which implies that the 2018 estimates do not represent best available science. For example, the National Academies of Sciences made suggestions relating to how one constructs a baseline future economy out to the year 2300, assumptions made in the climate modeling, and the discounting approach taken. Maybe most importantly for the purposes of this paper, the National Academies of Sciences report points a stern finger at the damage functions used in all three Integrated Assessment Models.

The damage functions in the Integrated Assessment Models, which are used to calculate the social cost of carbon, are outdated. Greenstone (2016) points out that the most recent studies in the FUND model stem from 2009, with the majority of the literature cited stemming from the early and mid-1990s. For example, the damage function for agriculture in the FUND model implies that warming up to roughly 5°C produces benefits for the sector (Rose et al. 2014). This is not consistent with the recent literature on agricultural impacts, which for example, points at the significant negative impact of extreme heat days. Moore, Baldos, Hertel, and Diaz (2017) updated the FUND damage function by incorporating the most recent empirical estimates for agriculture and find a doubling of the social cost of carbon by simply updating this sector alone. The literature underlying the DICE damage function also mostly comes from studies conducted in the 1990s. None of the cites for the PAGE model are from after 2010. As Greenstone (2016) shows, this ignores more than 100 studies published since 2010—which use more up-to-date econometric techniques and exploit the explosion in data availability.

## Damage Functions, Weather, and Climate

In the context of climate change studies, a "damage function" refers to a mapping of climate into economic outcomes—essentially what is broadly called a "dose response function." One question arises immediately: What is "climate"? When we leave our homes in the morning, weather is what we encounter. Weather outcomes are draws from an underlying distribution. For the purposes of this paper, I consider the moments of this distribution the *climate*. This approach is consistent with the often-used definition that climate is a 30-year average of (for example) surface temperature, although thinking about climate as a set of statistical moments is broader than just an average. Climate change is hence a slow shift in some moments of the weather distribution over time. The changes could be variance-preserving mean shifts or higher-order changes to the distribution. It is important to remember that even the simple case of a variance-preserving rise in mean temperature—think

of a bell-shaped curve of daily temperature outcomes shifting upward—will lead to a higher frequency of "extreme events"—the incidence of what would have been 95th or 99th percentile temperature outcomes under the old climate regime.

To estimate economic effects of such changes, we need to take into account how economic actors respond to weather generated by a new climate regime. For example, individuals in San Francisco have historically recognized that extreme warm-weather outcomes were rare, and so almost no one had air conditioners installed in their homes. However, if San Franciscans learn that climate is changing and their summers will resemble Fresno's much hotter summers in most future years, many will go ahead and install room air conditioners or central air units in new construction. Hence, a hotter climate will result in higher electricity consumption due to the presence of more air conditioners, which consumers incurred costs to install. In terms familiar to the economist, there is an extensive margin response in many sectors (the installation of air conditioning, irrigation equipment, sea walls) as well as an intensive margin response (the more frequent operation of air conditioners and irrigation equipment).

In order to provide estimates of damages from climate change, one needs to estimate damage functions that take both extensive and intensive margin adaptation into account—and to do this for *all* climate-sensitive sectors across the globe for a number of dimensions of climate. Some key climate-related changes would include changes in temperature, humidity, precipitation, sea level, and the occurrence of extreme events like storms.

With this perspective in mind, what are the properties that damage functions used in policy analysis of the economic impacts of climate change should posses? First, we would like to parameterize damage functions between the distribution of pre- and post-climate-change weather and economic outcomes of interest. Second, we would like these functions to identify and estimate parameters that carry a plausibly causal interpretation. Third, we would like the damage function to account for adaptation and measure the full costs of adaptation. Fourth, we would like the damage function to allow an estimation of economic welfare impacts.

This sounds as difficult as it is in practice. Figure 2 helps to explain why. The top left panel shows the weather pattern of temperature generated in two climate regimes. The light gray time series depicts a pre-climate-change world and the dark series shows a post-climate-change world, displaying higher mean and variance of the temperature series. The top right panel displays two damage functions (the parabolas) which map weather into an outcome, in this case temperature into household electricity consumption (measured in kilowatt-hours). The damage function, as has been confirmed in many empirical settings, is highly nonlinear. When it is cold outside and temperatures rise, electricity consumption falls, as people heat less. When it is hot outside and temperatures rise, electricity consumption increases as people air condition the indoor environment. In the pre-climate-change San Francisco, this response is relatively shallow, as few people have air conditioners as indicated by the solid damage function. If climate changes and produces the warmer more variable weather, we assume that people eventually will adapt by buying and operating air conditioners, which changes the damage function to the

*Figure 2*
**Mapping Weather into Impacts—The Importance of Accounting for Adaptation**

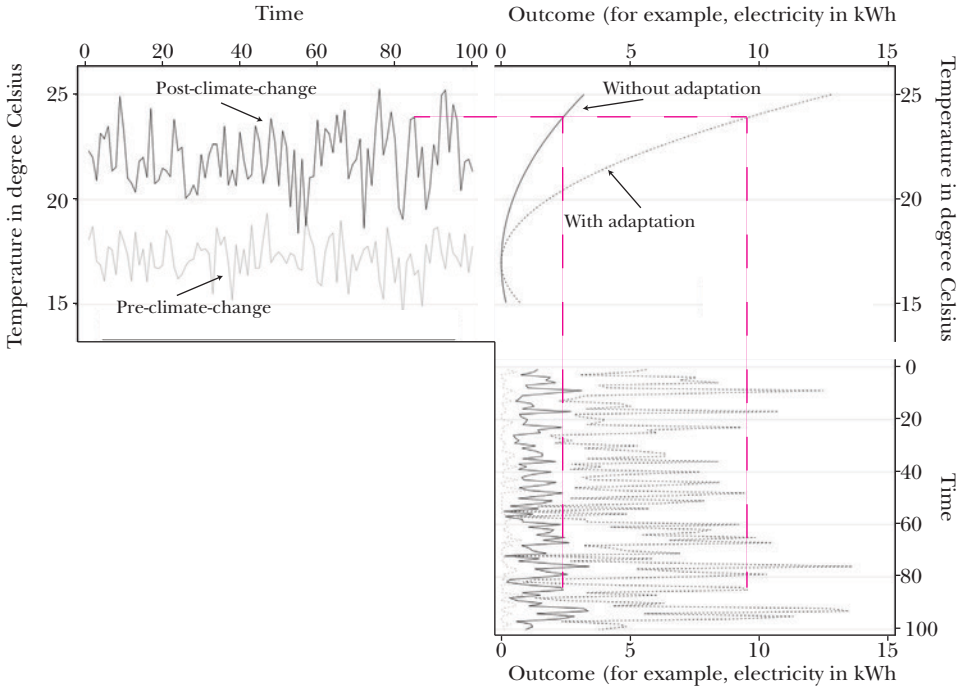


*Source:* Author.
*Note:* The top left panel shows the weather pattern of temperature generated in two climate regimes. The light gray time series depicts a pre-climate-change world and the dark series shows a post-climate-change world, with a temperature series displaying higher mean and variance. The top right panel displays two damage functions (the parabolas) which map weather into an outcome, in this case temperature into household electricity consumption (measured in kilowatt-hours). The effect can be seen in the bottom panel.

dotted parabola (labeled "With adaptation"). The response, especially at higher temperatures, is now steeper—resulting in stronger post-adaptation increases in electricity consumption on a one-degree warmer day when it is hot outside.

The effect can be seen in the bottom panel. If climate changes and we use the flatter (and wrong) pre-climate-change response function, which ignores the extensive margin adaptation, projected electricity consumption is the black solid line. This is clearly incorrect, as one is using the right weather but the wrong damage response function. The correct response function is the dotted parabola, which results in the dotted time series of electricity consumption in the bottom panel. It is much higher and much more variable compared to the no adaptation prediction. In the literature, this distinction is often referred to as the "weather versus climate response." I think it a better way to phrase this is "the impacts of weather simulated with versus without an extensive margin adaptation response." In a world changed by climate, we will still face weather when we walk out of our front door. As I will

discuss below, a rapidly growing empirical literature uses weather variation to identify response functions that partially or fully allow for adaptation.

So how does one go about calibrating these damage functions and using them to project damages? The question asked of any empirical economist these days is "what would the perfect counterfactual be?" In this context, a researcher actually needs to be concerned about two counterfactuals: 1) the counterfactual future climate; and 2) the counterfactual for identifying the appropriate damage function.

The first counterfactual, the climatic one, asks the question: What level of climate change will occur? Given our metaphysical inability to experiment by randomly imposing different levels of greenhouse gases on a large sample of otherwise identical Planet Earths, researchers instead resort to computational counterfactuals of the climate system, which are referred to as "global circulation models" (GCMs). These models use different scenarios of greenhouse gas emissions and physical representations of the climate system to predict changes in the climate system (IPCC 2013; Auffhammer, Hsiang, Schlenker, and Sobel 2013). They provide projections of, for example, surface temperatures, precipitation, and sea-level rise at a reasonable level of disaggregation and make these freely available through public depositories (Climate Impact Lab 2018; NASA 2018). A companion paper in this symposium by Hsiang and Kopp discusses these models and their limitations in more detail.

For the second counterfactual, we need to identify how agents in a given location respond to weather generated from a different climate regime. As a thought experiment, what is the right counterfactual for climate change in the United States by end of century? The US average historical (1986–2005) June/July/August temperature is 74°F. By end of century (under the aggressive RCP8.5 scenario), this temperature is projected to be 84° (Climate Impact Lab 2018).

One could contemplate a number of counterfactuals that might be used. If one has a set of units that are similar on observables and unobservables, but with different weather due to different local climate regimes, one might use a cross-sectional comparison. If one has long time series over a period of time where climate has changed, one might exploit time-series variation, possibly across units, to get econometric identification. But these approaches become questionable when we are comparing places that are far apart in characteristics space. Neighboring counties in California might possibly serve as counterfactuals for each other. However, using the economies of Pakistan, India, Mali, and Thailand as "hotter counterfactuals" for the United States or Europe, on the grounds that current mean temperatures in these countries are close to 84 degrees, is a stretch.

The econometric approaches discussed below all suffer from this issue of a fundamental lack of comparability, and I am afraid that there is no perfect way to overcome it. Indeed, the problem is even more severe than thinking about counterfactuals based on geographic and time-series variation would suggest. Comparing any current day or preindustrial society to a climate-changed world 100 years from now will be an imperfect comparison.

Many of the econometric studies I will describe below, including ones I have authored, use a counterfactual where we impose end-of-century climate on today's economy, which is a suboptimal way to circumvent the challenge of characterizing

an end-of-century economy as attempted by the Integrated Assessment Models used to calculate the social cost of carbon. As I will discuss below, the current state of knowledge predicts that climate change will affect economic growth, the distribution of population and wealth across space, and also significantly affect technology—both through mitigation and adaptation channels. An ideal counterfactual for several decades into the future would need to compare how these demographic and economic factors would change in the absence and presence of climate change as well.

## Estimating Economic Damages from Climate Change

One of the first known reflections on an association between human/economic activity and climate goes back to Parmenides, a disciple of Pythagoras writing in the fifth century BCE, who divided the world into five zones: one torrid, two temperate, and two frigid (Sanderson 1999). The torrid zones (which we call the tropics today) he thought were too hot and the frigid zones too cold for human habitation. Aristotle later agreed with this view. He believed that the only areas on earth habitable by humans were located between the tropics and the Arctic and Antarctic circles—the area where he lived.

The emergence of climate change as field of study in the physical sciences in the late 1970s led social scientists to think about estimating the possible consequences of a *changing* climate on economic sectors such as agriculture (D'Arge 1975; Kokoski and Smith 1987; Adams, McCarl, Dudek, and Glyer 1988; Adams 1989).

### Ricardian Cross-Sectional Approaches

Thousands of econometric papers control for weather in regressions, but Mendelsohn, Nordhaus, and Shaw (1994) offered the first attempt at estimating a damage function econometrically with the purpose of simulating the impacts of climate change on an important economic sector. They proposed a cross-sectional Ricardian framework, which is maybe the most widely used approach in climate impact estimation to this day. The intuition underlying this approach is that in a stationary climate, farmers optimize their production technology and crop choice according to the environment they face. This includes soil quality, slope of the land, agro-ecological zone, and of course climate, as captured by a set of statistical moments of the weather distribution over a substantial period of time. If land markets function perfectly, the land value should reflect the discounted present value of expected profits for a given parcel of land. In a regression framework, one can then decompose land values into their different components, one of which is long-run (for example, 30-year) averages of weather. In standard practice, one regresses farmer self-reported land values on polynomials of climate, which are often broken out by season. The marginal effects on the climate variables then indicate the marginal value of a one-unit change in a measure of climate.

Figure 3 helps cement the economic intuition behind this approach. Imagine a single farmer, who is currently growing crop 1 and earning profits corresponding

to the *y* value at point *A*. If faced with a significantly hotter climate, the farmer becomes indifferent between growing crop 1 and crop 2 at point *B*. If climate warms further still, the farmer would be much better off at point *C*, that is, switching to crop 2, rather than at point D where the farmer continues to grow crop 1. Because the cross-sectional regression observes optimizing farmers across the climate spectrum, this approach estimates the envelope of the individual crop-specific payoff functions and allows for climate adaptation. As a result, this approach both estimates a response that allows for adaptation to climate change and relies on data that are readily available in many regions in both the developed and developing world. It uses hotter locations as a counterfactual for the response of cooler location to climate change.

Three main criticisms of this method have been raised. First, this cross-sectional approach to damage function estimation is vulnerable to omitted variables bias, hence putting in question whether the estimates are *plausibly causal*. Any drivers of land values (or net profits) that are correlated with the climate indicators and outcome and are excluded from the model will confound the estimates of the marginal value of climate. As one vivid illustration, Schlenker, Hanemann, and Fisher (2005) reexamined the analysis of Mendelsohn, Nordhaus, and Shaw (1994), and point out that irrigation is an important driver of farm profits. This was omitted from the original regression model. When correcting for this by limiting the analysis to agricultural land east of the 100th meridian (the 100th meridian runs down through the middle of North and South Dakota and down through the middle of Texas) where agriculture is mostly non-irrigated, the marginal value of climate changed significantly. The estimated impacts of climate change went from being slightly beneficial to robustly negative.

Second, this Ricardian approach essentially assumes *costless adaptation* to climate change. But switching crops is not costless (Quiggin and Horowitz 1999). The fixed costs to switching from growing one crop to another may include investment in new harvesting equipment, irrigation infrastructure, and the acquisition of technical know-how. If these costs are big enough, it may be optimal for the farmer to delay or avoid change—in Figure 3, to continue farming crop 1 at point *D* rather than changing to crop 2 at point *C*. Hence, this method may provide biased estimates of the effect of climate change depending on how costly it is for farmers to switch from one crop to the next.

Third, this framework is applied retrospectively under the assumption that only historical climate matters. This assumption may no longer be tenable, as the climate has been changing since the 1960s. If agents know this, they should base their actions on expected rather than historical climate. Severen, Costello, and Deschenes (2016) provide an interesting extension of the Ricardian method by incorporating climate expectations. They show evidence that farmers already incorporate this information, suggesting that failing to incorporate expectations leads to a significant underestimation of projected impacts of climate change.

This cross-sectional framework has been applied in a number of other sectors. For example, Albouy, Graff, Kellogg, and Wolff (2016) back out the marginal value of climate in a cross-sectional study looking at residential home values. Mansur,

*Figure 3*
**Crop Choice and Profits in the Long and Short Run**

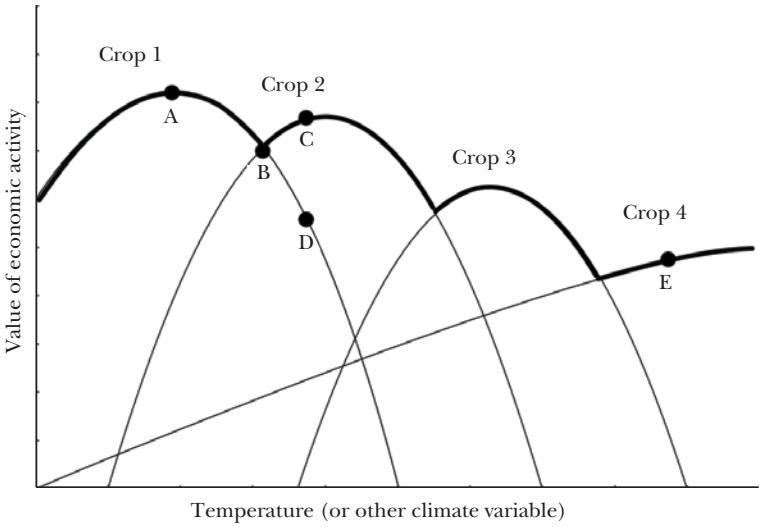


*Source:* Figure inspired by Mendelsohn, Nordhaus, and Shaw (1994).
*Note:* Imagine a single farmer, who is currently growing crop 1 and earning profits corresponding to the *y* value at point *A*. If faced with a significantly hotter climate, the farmer becomes indifferent between growing crop 1 and crop 2 at point *B*. If climate warms further still, the farmer would be much better off at point *C* (switching to crop 2) rather than at point *D* (continuing to grow crop 1).

Mendelsohn, and Morrison (2008) use this approach to study the effects of impacts of climate change on energy consumption, where the adaptation is not crop-switching, but rather fuel-switching.

**Panel Data Approaches**

Motivated by concern over the possibility of omitted variables in the Ricardian approach, Auffhammer, Ramanathan, and Vincent (2006) and Deschênes and Greenstone (2007) proposed using year-to-year variation in agricultural outcomes, temperature, and precipitation to estimate damage functions. Observing longitudinal panels of India's state-level rice output and US corn/soy and wheat yields, respectively, these papers can control for unit-specific and time-period fixed effects, which does away with some of the concerns over omitted variables bias. The regression equation in this approach regresses outcomes of interest (say, crop yields) on measures of contemporaneous weather (instead of the long-run averages of historical weather). If the right-hand-side weather variable enters the regression linearly, the estimated response has often been characterized as a short-run/weather/no-adaptation response—which is of course different from the weather response after a future persistent change in climate which accounts for adaptation. In this simplest version of the framework, econometric identification arises from within-unit year-to-year fluctuations in weather and the outcome of interest.

From the standpoint of analyzing the economic effects of climate change, an obvious concern with this approach is that it may capture short-run (intensive margin) adaptation to weather fluctuations, but not long-run (extensive margin) adaptation. For example, this approach captures farmer responses to bad weather draws in the short-run (like lower fertilizer application in a drier year) rather than in the long-run (like installation of irrigation infrastructure).

It is generally true that agents have more adaptation choices in the long run, especially along the extensive margin, and thus estimates that do not take this adaptation into account may overstate impacts. For example, farmers in the long run can switch crops, change the cropping calendar, or move their operations north, all of which would dampen the estimated impacts of climate change. However, there are also examples of adaptation options that are available in the short run and not in the long run. One example is a farmer with very limited groundwater resources and a slow refilling aquifer, who can smooth bad rainfall outcomes in the short run, yet continued water withdrawals would deplete the aquifer. As a result, this kind of adaptation would be only available in the short run, not the long run. Hence the bias may work in either direction depending on the nature of the adaptation options available to economic agents.

The critique that it is difficult to infer long-run adaptations based on short-run changes has some validity, but as I discuss later, several methods have been proposed for deriving long-run adaptation to climate change from panel data. Moreover, while the criticism of the lack of long-run adaptation in this approach may seem intuitive, it does not apply to all panel studies using weather as a right-hand-side variable. McIntosh and Schlenker (2006) consider the case in which the weather variable on the right-hand-side enters as a second-degree polynomial. Because the response function is calibrated by two parameters, the coefficient on the higher-order term uses both variation from *within units* as well as *across units*. Econometric identification arises from both within-unit time series variation as well as cross-sectional variation across units. Hence, it has been argued, that studies using this nonlinear specification allow for plausibly causal estimates that incorporate adaptation.

The papers leaning on this approach most strongly are panel studies of GDP *growth rates* across countries as a function of annual temperature fluctuations (Dell, Jones, and Olken 2012; Burke, Hsiang, and Miguel 2015a). The most recent of these papers find impacts of climate change on global GDP around 20 percent by end of the century, which is an *order of magnitude* larger than what is found by most Integrated Assessment Models. There is broad enthusiasm for this approach, especially in the interdisciplinary climate literature. Aside from the fact that the choice of growth rate as the dependent variable implies that temperature shocks have persistent effects on economic growth, it is important to remember that this approach introduces cross-sectional variation and all that comes with it in the identification of the higher-order term.

Another critique of the panel data approach is that if weather is measured with error, then as more fixed effects are included in the regression, concerns over measurement error loom larger (Fisher, Hanemann, Roberts, and Schlenker 2012). In the vast majority of locations, weather is measured with error, and the bigger

the distance between weather stations, the bigger measurement error concerns become. The United States and Europe have tens of thousands of weather stations, but many locations in in sub-Saharan Africa do not have a weather station within hundreds of miles. If the measurement error is classical, this is likely to attenuate the response towards zero.

**Long Difference Estimation**

Motivated by omitted variables bias issues in Ricardian models and the possible issues relating to capturing long-run adaptation in panel data models, Burke and Emerick (2016) proposed an alternate approach, which seeks to provide plausibly causal estimates of damages that fully account for observable adaptation. Climate has already changed in the United States over the previous half-century; in particular, they show that warming and precipitation trends are quite heterogeneous across US counties east of the 100th meridian. Hence, one can use differential climate trends as a source of econometric identification. The beauty of this approach is that the distribution of observed trends includes changes similar in magnitude to those expected over the next century, which creates some overlap between the temperature and precipitation variation used for identification and out-of-sample projection.

In their estimation, they use the difference between five-year moving averages of crop yields two decades apart and regress these on five-year moving averages of weather also two decades apart for all agricultural counties east of the 100th meridian. The differencing is equivalent to the inclusion of county fixed effects and the variation used to identify a climate effect incorporates adaptation. The marginal effects from this estimation show that the long-run estimates are at best half of those estimated from panel data models using short-run variation in weather. However, given the range of statistical significance, one cannot rule out that the two are equivalent. The authors interpret this finding as evidence of only limited long-run adaptation, which is one interpretation. Those working with panel data approaches might argue that the comparison here is flawed, because the baseline used for comparison incorporates some degree of adaptation.

This long difference approach is appealing because it provides plausibly causal estimates of climate impacts that account for adaptation. However, the data requirements are significant. One needs broad spatial coverage of data over long periods of time. The other application where this long difference approach has been applied is in measuring the impacts of climate change on aggregate GDP across countries (Dell, Jones, and Olken 2012). However, other than in the cross-country sense, there are no applications of this estimator in nonagricultural sectors or in the developing country context. There should be more applications of this method in settings where data are sufficient.

**Ricardo Meets Panels: Climate Adaptive Response Estimation**

A small but rapidly expanding literature attempts to estimate how the dose response function between weather and outcomes of interest changes *as a function of a changing climate*. There are two approaches. The first is similar to a "split sample approach," where one splits a long panel of observable outcomes and weather

into two periods and estimates the response function separately. One can then use statistical tests to search for evidence of adaptation between the two periods. For example, Barreca, Clay, Dechenes, Greenstone, and Shapiro (2016) examine the mortality response to weather over time in the United States and show a massive decrease in the effect of a hot day on mortality over time, which is due to the significant rollout of air conditioning in the hot and often humid areas of the United States. One example of this approach is Roberts and Schlenker (2011).

A second approach along these lines represents a marriage of the panel data estimation approach using short-run weather fluctuations and the Ricardian approach. The concept here is that if one observes a large number of units (like counties, households, or firms) over a significant number of periods covering a spatial area with large heterogeneity in climate, one can estimate separate response functions for subgroups of the individual units using observed short-run weather fluctuations (for example, use within-household variation to identify a short-run response function by zip code). By controlling for unit- and time-fixed-effects, it is possible to obtain plausibly causal estimates of local short-run dose response functions. One can then either in a second step regress the slopes of the dose response on climate (for example, long run average summer temperature) across subgroups, or, through an interaction term in a single regression, estimate how the slope of the dose response function varies across areas with different climates, incomes, and other observables that vary across space. Sightings of this approach include Bigano, Hamilton, and Tol (2007), Auffhammer and Aroonruengsawat (2012), Hsiang and Narita (2012), Butler and Huybers (2013), Davis and Gertler (2015); Heutel, Miller, and Molitor (2017), and Carleton et al. (2018).

This approach offers two important forward steps beyond the panel studies discussed above. First, it explicitly models climate adaptation by exploiting cross-sectional differences in the slopes of dose response functions. Second, it allows us to model explicitly the effects of income and population on the damage functions.

While this approach has significant appeal, it does not overcome some of the shortfalls of the Ricardian and panel methods. The econometrician is always limited by using historical observations in order to parameterize equations. The best we can do is simulate how income, population, and climate have affected short-run dose response functions historically and to assume that this relationship remains stable. We can approximate a future San Francisco with the climate of Fresno by assigning the appropriate climate, income, and population, but none of these approaches properly address the fact that Fresno may be structurally very different from a future San Francisco—even if we assign the right income and population. We simply lack the crystal ball that lets us look to 2100 and beyond. But this issue has plagued social science broadly, because predicting what the world looks like 100 years out is, well, rather difficult.

## Room for Expert Elicitation?

This literature on estimating the economic damages of climate change has been criticized on four grounds, which have been well-enunciated in Pindyck

(2013, 2016, 2017). Ultimately, these criticisms raise the possibility that for studying climate change, conventional econometric studies may need to be supplemented with a healthy dose of "expert elicitation."

Pindyck's first criticism is that in Integrated Assessment Models, the functional form of relationships and their parameterization—including those in damage functions—are "arbitrary." Second, he expresses concern that many of the studies cited above "are limited to short time periods and small fluctuations in temperature and other weather variables," which is effectively the same as pointing out that econometricians rely on observed data and technology to parameterize their dose response functions. In whichever way one phrases this concern, the bottom line is that existing studies may not account well for long-term adaptation and in particular for the possibility of very significant changes in technology. Third, the biggest impacts of climate change may result from extreme and catastrophic events, which can be thought of as low-probability events with possibly massive economic consequences. Examples would include the shutdown of the Thermohaline Circulation that gives Europe its lovely climate, the melting of the West Antarctic ice sheet, and the possible rapid release of significant amounts of methane from the tundra. We have (fortunately!) not observed these events in the measured historical record and hence econometric estimation cannot provide estimates of the economic damages from such events. A final concern is that there is little agreement over the correct approach to discounting and which discount rate to apply in placing a value on future damages from climate change.

In response to these concerns, Pindyck has strongly argued for "expert elicitation." For example, in response to estimating the risks and costs of extreme climate events, one can imagine that teams of scientists with an understanding of the physical and economic consequences might be able to provide coarse estimates of the damages resulting from such large events. There are well-established procedures for such expert elicitation, and this may be a fruitful avenue forward to make progress on this topic. However, experts in this arena have to rely on "process understanding," as there are no data here to help. Similarly, one can imagine a group of experts who might tackle the question of what discount rate is most appropriate to use, which is what Drupp, Freeman, Groom, and Nesje (forthcoming) did. The median answer for the risk-free social discount rate is 2 percent in their study, which is quite different from the 3 percent and 7 percent rates applied in the most recently used social cost of carbon in proposed US government rulemakings for automotive fuel economy (CAFE) standards.

However, expert elicitation seems less useful in coming up with better estimates of damages in order to overcome the first two of Pindyck's critiques. I would argue that the recent literature has made significant headway in estimating plausibly causal damage functions incorporating adaptive response from partially cross-sectional variation. The formulations doubtless can be critiqued and questioned, but they are not arbitrary. I question whether experts would come up with "better" estimates than the cutting-edge papers in this literature. Maybe more fundamentally, a group of experts called upon to participate in an expert elicitation exercise concerning the functional form of damage models and

extrapolations to larger climate changes or time periods would begin with—of course—a review of the existing recent models in this area, which brings us back to the importance of better econometric models.

## What We Know and What We Don't

Cline (1992) put forth a list of important sectors for which we require a better understanding of their climate sensitivity. Table 1 below replicates his table and I have subjectively filled in where this literature currently stands in terms of published and ongoing efforts. A glance shows that there is a lot of work to do.

Yet it is clear that the literature on the econometric estimation of damage functions of climate change is rapidly expanding—both in terms of methods as well as sectoral and spatial coverage. The previously stagnant state of affairs where most of the damage functions in Integrated Assessment Models had not been updated significantly in over a decade has changed dramatically. Economists need to push forward in improving sectoral and spatial coverage of the damage functions provided to modelers, using methods that allow us to parameterize plausibly causal damage functions, which account for adaptation and allow us to estimate welfare impacts of climate change. The current frontier is probably best described by work using the "Ricardo meets panel data" approach.

Moore et al. (2017) is one published attempt to incorporate the most recent estimates of damage functions for the agricultural sector into an Integrated Assessment Model (the FUND model) and this one-sector exercise doubles the social cost of carbon (SCC), which underlines the importance of these efforts.

Those interested in this area will want to keep an eye on two major efforts that involve ambitious ongoing collaborations between climate scientists and economists. The Climate Impacts Lab, managed jointly by researchers at the University of Chicago, UC Berkeley, Rutgers University, and the Rhodium Group, produces damage functions for mortality, migration, energy consumption, agricultural yields, and conflict which satisfy the characteristics laid out above and have global coverage. At the same time, a group at Resources for the Future has undertaken the task of implementing the changes suggested by the National Academies of Sciences in the modeling of the social cost of carbon. The governments of Mexico and Canada have pledged their support of these efforts, as all US federal government development of the modeling behind the social cost of carbon has been halted—a fact which is deeply concerning.

As these and other researchers dig deeper, three key areas require especially deep thinking. First, we need to improve how we incorporate damages from catastrophic events, which may well require abandoning the econometric toolkit and relying on cross-disciplinary expert solicitation. Second, we need to think about general equilibrium effects across space and spillover effects across sectors in our models. Collaborations between trade and climate economists (Dingel, Meng, and Hsiang 2018), as well as academics working on supply chains (for example, Seetharam 2018), will likely yield fruitful insights. Finally, it is shocking how little work has been done on the effects of climate change on nonmarket goods other than mortality. It is

*Table 1*
**Coverage of the Damage Function Literature**

| *Sector* | *Plausibly causal estimates* | *Adaptation addressed* | *Global coverage* | *Examples* |
|---|---|---|---|---|
| Agriculture | Yes | Yes | Yes | Schlenker and Roberts (2009); Moore, Baldos, Hertel, and Diaz (2017) |
| Forestry | No | No | No | |
| Species loss | No | No | No | |
| Sea-level rise | Yes | Yes | No | Houser, Hsiang, Kopp, and Larsen (2015) |
| Energy | Yes | Yes | No | Auffhammer (2018) |
| Human amenity | Yes | ~ Yes | No | Albouy, Graf, Kellogg, and Wolff (2016); Baylis (2015) |
| Morbidity and mortality | Yes | Yes | Yes | Deschênes and Greenstone (2011); Carleton et al. (2018) |
| Migration | Yes | No | No | Bohra-Mishra, Oppenheimer, Hsiang (2014); Missiran and Schlenker (2017) |
| Crime and conflict | Yes | No | Maybe | Burke, Hsiang, and Miguel (2015b) |
| Productivity | Yes | No | No | Peng, Deschênes, Meng, and Zhang (2018) |
| Water consumption | No | No | No | |
| Pollution | Yes | Maybe | No | Bento, Mookerjee, and Severenini (2017) |
| Storms | Yes | Yes | No | Hsiang and Narita (2012); Deryugina, Kawano, and Levitt (2018) |

*Source:* Cline (1992) put forth a list of important sectors for which we require a better understanding of their climate sensitivity. Table 1 below replicates this list and I have subjectively filled in where this literature currently stands in terms of published and ongoing efforts.

paramount that we begin developing approaches that will allow us to quantify damages from species loss, ecosystem services—as well as effects on human morbidity—and incorporate these into the models that estimate costs of climate change.

■ *I thank the Berkeley Climate Economics Lunch group, the members of the National Academy of Sciences panel on the Social Cost of Carbon, David Anthoff, Lint Barrage, Marshall Burke, Tamma Carleton, Chris Costello, Olivier Deschênes, Tony Fisher, Michael Greenstone, Michael Hanemann, Sol Hsiang, David Lobell, Rob Mendelsohn, Pierre Mérel, Frances Moore, Michael Roberts, Wolfram Schlenker, Joe Shapiro, Andy Solow, Richard Tol, Fiona Wilkes, Gary Yohe, and many others for numerous conversations that have informed my thinking on the topics discussed in this survey over the years. All misguided thinking is solely mine.*

## References


**Adams, Richard M.** 1989. "Global Climate Change and Agriculture: An Economic Perspective." *American Journal of Agricultural Economics* 71(5): 1272–79.

**Adams, Richard M., Bruce A. McCarl, Daniel J. Dudek, and J. David Glyer.** 1988. "Implications of Global Climate Change for Western Agriculture." *Western Journal of Agricultural Economics* 13(2): 348–56.

**Albouy, David, Walter Graf, Ryan Kellogg, and Hendrik Wolff.** 2016. "Climate Amenities, Climate Change, and American Quality of Life." *Journal of the Association of Environmental and Resource Economists* 3(1): 205–46.

**Auffhammer, Maximilian.** 2018. "Climate Adaptive Response Estimation: Short and Long Run Impacts of Climate Change on Residential Electricity and Natural Gas Consumption Using Big Data." NBER Working Paper 24397.

**Auffhammer, Maximilian, and Anin Aroonruengsawat.** 2012. "Hotspots of Climate-Driven Increases in Residential Electricity Demand: A Simulation Exercise Based on Household Level Billing Data for California." White Paper from the California Energy Commission's Climate Change Center.

**Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel.** 2013. "Using Weather Data and Climate Model Output in Economic Analyses of Climate Change." *Review of Environmental Economics and Policy* 7(2): 181–98.

**Auffhammer, Maximilian, V. Ramanathan, and Jeffrey R. Vincent.** 2006. "Integrated Model Shows that Atmospheric Brown Clouds and Greenhouse Gases Have Reduced Rice Harvests in India." *PNAS* 103(52): 19668–672.

**Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro.** 2016. "Adapting to Climate Change: The Remarkable Decline in the US Temperature–Mortality Relationship over the Twentieth Century." *Journal of Political Economy* 124(1): 105–159.

**Baylis, Patrick.** 2015. "Temperature and Temperament: Evidence from a Billion Tweets." Energy Institute at Haas Working Paper 265. November.

**Bento, Antonio, Mehreen Mookerjee, and Edson Severnini.** 2017. "A New Approach to Measuring Climate Change Impacts and Adaptation." Unpublished paper. http://econweb.ucsd.edu/cee/papers/BentoetaI_ClimatePenalty.pdf.

**Bigano, Andrea, Jacqueline Hamilton, and Richard S. J. Tol.** 2007. "The Impact of Climate Change on Domestic and International Tourism: A Simulation Study." *Integrated Assessment Journa*l 7(1).

**Bohra-Mishra, Pratshya, Michael Oppenheimer, and Solomon M. Hsiang.** 2014. "Nonlinear Permanent Migration Response to Climatic Variations But Minimal Response to Disasters." *PNAS* 111(27): 9780–85.

**Burke, Marshall, and Kyle Emerick.** 2016. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy* 8(3): 106–40.

**Burke, Marshall, Solomon M. Hsiang, and Edward Miguel.** 2015a. "Global Non-Linear Effect of Temperature on Economic Production." *Nature* 527(7577): 235–39.

**Burke, Marshall, Solomon M. Hsiang, and Edward Miguel.** 2015b. "Climate and Conflict." *Annual Review of Economics* 7(1): 577–617.

**Butler, Ethan E., and Peter Huybers.** 2013. "Adaptation of US Maize to Temperature Variations." *Nature Climate Change* 3(1): 68–72.

**Carleton, Tamma, Michael Delgado, Michael Greenstone, Solomon M. Hsiang, Andrew Hultgren, Amir Jina, Robert Kopp, Ishan Nath, James Rising, Ashwin Rode, Samuel Seo, Justin Simcock, Arvid Viaene, Jiacan Yuan, and Alice Zhang.** 2018. "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits." In *The Economic Consequences of Environmental Change: Essays on Climate and Water.* University of California Berkeley, Dissertation.

**Carleton, Tamma A., and Solomon M. Hsiang.** 2016. "Social and Economic Impacts of Climate." *Science* 353(6304): aad9837.

**Climate Impact Lab.** 2018. Climate Impact Map. https://www.impactlab.org/map/#usmeas=absolute&usyear=-1981-2010&gmeas=absolute&gyear=1986-2005&tab=global (accessed April 15, 2018).

**Cline, William R.** 1992. *The Economics of Global Warming.* Institute for International Economics.

**d'Arge, Ralph C.** 1975. *Economic and Social Measures of Biologic and Climatic Change,* vol. 6. Institute for Defense Analyses.

**Davis, Lucas W., and Paul J. Gertler.** 2015. "Contribution of Air Conditioning Adoption to Future Energy Use under Global Warming." PNAS 112(19): 5962–67.

**Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken.** 2012. "Temperature Shocks and Economic Growth: Evidence from the Last Half

Century." *American Economic Journal: Macroeconomics* 4(3): 66–95.

**Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken.** 2014. "What Do We Learn from the Weather? The New Climate–Economy Literature." *Journal of Economic Literature* 52(3): 740–98.

**Deryugina, Tatyana, Laura Kawano, and Steven Levitt.** 2018. "The Economic Impact of Hurricane Katrina on Its Victims: Evidence from Individual Tax Returns. *American Economic Journal: Applied Economics* 10(2): 202–33.

**Deschênes, Olivier, and Michael Greenstone.** 2007. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather." *American Economic Review* 97(1): 354–85.

**Deschênes, Olivier, and Michael Greenstone.** 2011. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* 3(4): 152–85.

**Diaz, Delavane, and Frances C. Moore.** 2017. "Quantifying the Economic Risks of Climate Change." *Nature Climate Change* 7: 774–82.

**Dingel, Jonathan I., Kyle C. Meng, and Solomon Hsiang.** 2018. "The Spatial Structure of Productivity, Trade, and Inequality: Evidence from the Global Climate." https://www.dropbox.com/s/2dibmgy94kn7any/DingelMengHsiang.pdf?dl=1.

**Drupp, Moritz A., Mark C. Freeman, Ben Groom, and Frikk Nesje.** Forthcoming. "Discounting Disentangled." *American Economic Journal: Economic Policy.*

**Environmental Protection Agency (EPA).** 2018. Social cost of carbon spreadsheet provided by EPA to OMB, July, 16, 2018. https://www.regulations.gov/contentStreamer?documentId=EPA-HQ-OAR-2018-0283-0453&attachmentNumber=16&contentType=excel12book.

**Fisher, Anthony C., W. Michael Hanemann, Michael J. Roberts, and Wolfram Schlenker.** 2012. "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather: Comment." *American Economic Review* 102(7): 3749–60.

**Greenstone, Michael.** 2016. "A New Path Forward for an Empirical Social Cost of Carbon." Presentation to the National Academies of Sciences, https://sites.nationalacademies.org/cs/groups/dbassesite/documents/webpage/dbasse_172599.pdf.

**Greenstone, Michael, Elizabeth Kopits, and Ann Wolverton.** 2013. "Developing a Social Cost of Carbon for US Regulatory Analysis: A Methodology and Interpretation." *Review of Environmental Economics and Policy* 7(1): 23–46.

**Heutel, Garth, Nolan H. Miller, and David Molitor.** 2017. "Adaptation and the Mortality Effects of Temperature Across U.S. Climate Regions." NBER Working Paper 23271.

**Houser, Trevor, Solomon Hsiang, Robert Kopp, and Kate Larsen.** 2015. *Economic Risks of Climate Change: An American Prospectus.* Columbia University Press.

**Hsiang, Solomon.** 2016. "Climate Econometrics." *Annual Review of Resource Economics* 8: 43–75.

**Hsiang, Solomon M., Marshall Burke, and Edward Miguel.** 2013. "Quantifying the Influence of Climate on Human Conflict." *Science* 341(6151): 1235367

**Hsiang, Solomon M., and Daiju Narita.** 2012. "Adaptation to Cyclone Risk: Evidence from the Global Cross-Section." *Climate Change Economics* 3(2): 1250011.

**Interagency Working Group (IWG) on Social Cost of Greenhouse Gases, United States Goverment.** 2016. Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866. https://www.epa.gov/sites/production/files/2016-12/documents/sc_co2_tsd_august_2016.pdf.

**Intergovernmental Panel on Climate Change (IPCC).** 2013: "Summary for Policymakers." In: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change,* edited by T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P.M. Midgley. Cambridge University Press.

**Kokoski, Mary F., and V. Kerry Smith.** 1987. "A General Equilibrium Analysis of Partial-Equilibrium Welfare Measures: The Case of Climate Change." *American Economic Review* 77(3): 331–41.

**Livingstone, David N.** 2011. "Environmental Determinism." In *The SAGE Handbook of Geographical Knowledge*, edited by John A. Agnew, 368–380. SAGE publishing.

**Mansur, Erin T., Robert Mendelsohn, and Wendy Morrison.** 2008. "Climate Change Adaptation: A Study of Fuel Choice and Consumption in the US Energy Sector." *Journal of Environmental Economics and Management* 55(2): 175–93.

**McIntosh, Craig T., and Wolfram Schlenker.** 2006. "Identifying Non-Linearities in Fixed Effects Models." Unpublished paper, https://gps.ucsd.edu/_files/faculty/mcintosh/mcintosh_research_identifying.pdf.

**Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw.** 1994. "The Impact of Global Warming on Agriculture: A Ricardian Analysis." *American Economic Review* 84(4): 753–71.

**Missirian, Anouch, and Wolfram Schlenker.** 2017. "Asylum Applications and Migration Flows." *American Economic Review* 107(5): 436–40.

**Moore, Frances C., Uris Baldos, Thomas Hertel, and Delavane Diaz.** 2017. "New Science of Climate Change Impacts on Agriculture Implies Higher Social Cost of Carbon." *Nature Communications* 8: 1607.

**NASA.** 2018. "NASA Earth Exchange Global Daily Downscaled Projections (NEX-GDDP)." https://cds.nccs.nasa.gov/nex-gddp/ (accessed June 2018).

**National Academies of Sciences, Engineering, and Medicine.** 2017. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*. Washington, DC: The National Academies Press. https://doi.org/10.17226/24651.

**Nordhaus, William.** 2006. "Paul Samuelson and Global Public Goods." Chap. 5 in *Samuelsonian Economics and the Twenty-First Century*, edited by Michael Szenberg, Lall Ramrattan, and Aron A. Gottesman. Oxford University Press.

**Peng, Zhang, Olivier Deschênes, Kyle C. Meng, and Junjie Zhang.** 2018. "Temperature Effects on Productivity and Factor Reallocation: Evidence from a Half Million Chinese Manufacturing Plants." *Journal of Environmental Economics and Management* 88(3): 1–17.

**Pindyck, Robert S.** 2013. "Climate Change Policy: What Do the Models Tell Us?" *Journal of Economic Literature* 51(3): 860–72.

**Pindyck, Robert S.** 2016. "The Social Cost of Carbon Revisited." NBER Working Paper 22807.

**Pindyck, Robert S.** 2017. "The Use and Misuse of Models for Climate Policy." *Review of Environmental Economics and Policy* 11(1): 100–114.

**Quiggin, John and Horowitz, John K.** 1999. "The Impact of Global Warming on Agriculture: A Ricardian Analysis: Comment." *American Economic Review*, 1999, 89(4), pp. 1044–45.

**Roberts, Michael J., and Wolfram Schlenker.** 2011. "The Evolution of Heat Tolerance of Corn: Implications for Climate Change." In *The Economics of Climate Change: Adaptations Past and Present*, pp. 225–51, edited by Gary D. Libecap and Richard H. Steckel. University of Chicago Press.

**Roberts, Malcolm J., Pier Luigi Vidale, Matthew S. Mizielinski, Marie-Estelle Demory, Reinhard Schiemann, Jane Strachan, Kevin Hodges, Ray Bell, and Joanne Camp.** 2015. "Tropical Cyclones in the UPSCALE Ensemble of High-Resolution Global Climate Models." *Journal of Climate* 28(2): 574–96.

**Rose, Steven K.** 2012. "The Role of the Social Cost of Carbon in Policy." *Wiley Interdisciplinary Reviews: Climate Change* 3(2): 195–212.

**Rose, S., D. Turner, G. Blanford, J. Bistline, F. de la Chesnaye, and T. Wilson.** 2014. *Understanding the Social Cost of Carbon: A Technical Assessment*. EPRI Report 3002004657.

**Sanderson, Marie.** 1999. "The Classification of Climates from Pythagoras to Koeppen." *Bulletin of the American Meteorological Society* 80(4): 669–73.

**Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher.** 2005. "Will US agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach." *American Economic Review* 95:1, 395–406.

**Schlenker, Wolfram, and Michael J. Roberts.** 2009. "Nonlinear Temperature Effects Indicate Severe Damages to US Crop Yields under Climate Change." *PNAS* 106(37): 15594–98.

**Seetharam, Ishuwar.** 2018. "The Indirect Effects of Hurricanes: Evidence from Firm Internal Networks." https://web.stanford.edu/~ishuwar/JMP.pdf.

**Severen, Christopher, Christopher Costello, and Olivier Deschenes.** 2016. "A Forward Looking Ricardian Approach: Do Land Markets Capitalize Climate Change Forecasts?" NBER Working Paper 22413.

**US Environmental Protection Agency (EPA).** 2008. "Regulating Greenhouse Gas Emissions Under the Clean Air Act; Proposed Rule." *Federal Register*, Wednesday, July 30, 73(147).

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 53–72*

# The Cost of Reducing Greenhouse Gas Emissions

Kenneth Gillingham and James H. Stock

What is the most economically efficient way to reduce greenhouse gas emissions? The principles of economics deliver a crisp answer: reduce emissions to the point that the marginal benefits of the reduction equal its marginal costs. This answer can be implemented by a Pigouvian tax, for example a carbon tax where the tax rate is the marginal benefit of the emissions reduction or, equivalently, the monetized damages from emitting an additional ton of carbon dioxide ($CO_2$). The carbon externality will then be internalized and the market will find cost-effective ways to reduce emissions up to the amount of the carbon tax.

However, most countries, including the United States, do not place an economy-wide tax on carbon, and instead have an array of greenhouse gas mitigation policies that provide subsidies or restrictions typically aimed at specific technologies or sectors. Such climate policies range from automobile fuel economy standards, to gasoline taxes, to mandating that a certain amount of electricity in a state comes from renewables, to subsidizing solar and wind electrical generation, to mandates requiring the blending of biofuels into the surface transportation fuel supply, to supply-side restrictions on fossil fuel extraction. In the world of a Pigouvian tax, markets sort out the most cost-effective ways to reduce emissions, but in the world

■ *Kenneth Gillingham is Associate Professor of Economics, Yale University, New Haven, Connecticut. James H. Stock is the Harold Hitchings Burbank Professor of Political Economy, Harvard University, Cambridge Massachusetts. Gillingham is a Faculty Research Fellow and Stock is a Research Associate, both at the National Bureau of Economic Research, Cambridge Massachusetts. Their email addresses are kenneth.gillingham@yale.edu and james_stock@harvard.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.53

doi=10.1257/jep.32.4.53

we live in, economists need to weigh in on the costs of specific technologies or narrow interventions.

This paper reviews the costs of various technologies and actions aimed at reducing greenhouse gas emissions. Our aim is twofold. First, we seek to provide an up-to-date summary of costs of actions that can be taken now using currently available technology. These costs focus on expenditures and emissions reductions over the life of a project compared to some business-as-usual benchmark—for example, replacing coal-fired electricity generation with wind or weatherizing a home. We refer to these costs as static because they are costs over the life of a specific project undertaken now, and they ignore spillovers. In the environmental economics literature, these static costs are an element in creating what is called a marginal abatement cost (MAC) curve, which plots out the marginal costs of achieving a cumulative level of emissions abatement in order from the lowest- to highest-cost technology or measure.

To economists not in the energy-environment field, these marginal abatement costs might contain some surprises. Although we are skeptical of most "free lunch" static estimates, for some technologies the cost of emissions reductions is remarkably low. For example, blending corn ethanol into gasoline up to a 10 percent ratio provides essentially costless emissions reductions (our point estimate is in the "free lunch" range) in the United States because ethanol is a less-expensive octane booster than alternatives derived from petroleum. Another low or negative static cost source of emissions reductions is replacing coal-fired electricity generation with natural gas, a switch that has been widely adopted by power generators located where gas prices are low because of the fracking revolution. On the other hand, some actions that might seem green are, from a static perspective, anything but. For example, driving a Ford Focus electric vehicle in a region in which electricity is generated by coal has approximately the same $CO_2$ footprint as a Ford Explorer sport utility vehicle that averages 25 miles per gallon, and costs nearly as much. We find a wide range of costs for interventions currently being employed, both across and within different types of interventions. This heterogeneity in costs implies that we could achieve the same amount of greenhouse gas emissions reductions that we are achieving now at a much lower static cost, or greater emissions reductions for the same cost. Possible reasons for the use of more expensive policies include the chosen policies having less transparent costs, individual policies having justifications beyond just climate policy, differences in the marginal costs across locations, and lobbying by businesses that could potentially be affected by lower-cost policies. In some cases, especially policies aimed at developing nascent technologies, the policies are developed with a longer-term vision in mind.

These estimates of static costs help to inform discussions about climate policy, but they miss the critical consideration that climate change is a long-term problem. As a result, the proper answer to our opening question is not necessarily what is the least expensive mitigation strategy among options available today, but what are the actions if, taken today, will minimize the cost of mitigation both today *and* into the future, recognizing that actions taken today can influence future costs. We refer to such costs as dynamic, because they outlive the life of a specific project.

Our second aim is to distinguish between dynamic and static costs and to argue that some actions taken today with seemingly high static costs can have low dynamic costs, and vice versa. We make this argument at a general level and through two case studies, of solar panels and of electric vehicles. The cost of both technologies has fallen sharply, arguably driven in part by demand-side incentives that in turn stimulated learning-by-doing and technological improvements, the benefits of which are only partially captured by the manufacturing firm. In addition, purchasing an electric vehicle today drives the demand for charging stations, which in effect reduces the cost (here, the cost of time and worry) to potential future purchasers. Under the right circumstances, such dynamic effects can offer a justification for policies that a myopic calculation suggests have high costs.

## Estimates of Static Abatement Costs

Before we begin, we briefly digress on units. The standard units of emissions costs and benefits are dollars per metric ton (1,000 kilograms) of $CO_2$ emissions avoided. As a point of comparison, the social cost of carbon is an estimate of the net present value of monetized social damages from emission of an additional metric ton of $CO_2$; under the Obama administration, the US government estimated the social cost of carbon to be approximately $46 in 2017 dollars for a ton of emissions in 2017 (IWG 2016).[1] Burning one gallon of petroleum gasoline produces roughly nine kilograms of $CO_2$, so a social cost of carbon value of $46/metric ton $CO_2$ corresponds to $0.41 per gallon. Also, carbon dioxide is only one of many greenhouse gases; others include methane, nitrous oxide, and hydrofluorocarbons. To facilitate comparisons, it is conventional to convert costs for reducing non-$CO_2$ greenhouse gases into $CO_2$-equivalent units, and we adopt that convention here.[2]

### Brief Background on Marginal Abatement Cost Curves

The marginal abatement cost (MAC) curve plots measures to abate emissions in order from the least to most expensive. For each, there is a cost per ton of emissions reduced and a quantity of emissions reductions available at that cost. The

[1] The Trump administration withdrew this estimate by executive order and forbid agencies from using the underlying research for regulatory purposes; as of this writing, the Environmental Protection Agency is using two estimates, $1 and $6 per ton, depending on the discount rate (3 or 7 percent) (Newell 2017). The estimate of $46/ton is in the range of the academic literature, although some estimates are much higher (as one example, see Gillingham et al. 2018). There is currently a cross-institutional interdisciplinary effort to provide a comprehensive update to the social cost of carbon based on recommendations made by the National Academy of Sciences (2017), which is discussed on the Resources for the Future website at http://www.rff.org/research/collection/rffs-social-cost-carbon-initiative.

[2] A complication in developing $CO_2$-equivalent estimates is that the atmospheric residence time of greenhouse gases varies. The most common approach, the global warming potential approach, is only an approximation when used to calculate the social cost of non-$CO_2$ greenhouse gases. See Marten and Newbold (2012) for a more comprehensive approach to calculating the social cost of non-$CO_2$ greenhouse gases.

*Figure 1*
**The McKinsey (2009) Marginal Abatement Cost Curve: "Global GHG Abatement Cost Curve Beyond Business-As-Usual-2030"**

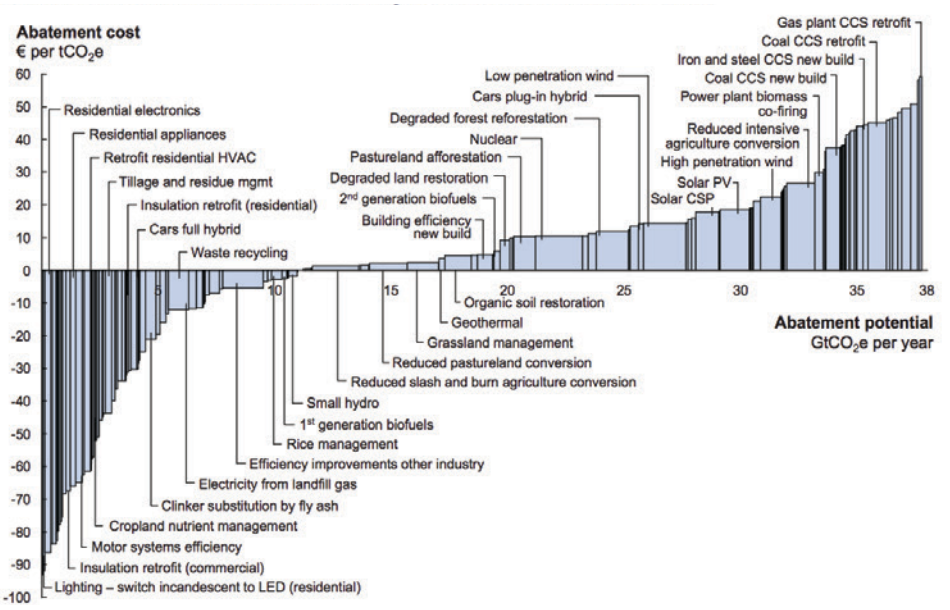


*Source:* Global GHG Abatement Cost Curve v2.0. Figure and notes reproduced with permission from McKinsey (2009).
*Note:* The curve presents an estimate of the maximum potential of all technical GHG abatement measures below €60 per $tCO_2e$ if each lever was pursued aggressively. It is not a forecast of what role different abatement measures and technologies will play.

use of MAC curves to support climate policy analysis dates back at least a quarter century (for an early review, see Grubb, Edmonds, ten Brink, and Morrison 1993). All models that estimate the mitigation costs of climate policy either implicitly or explicitly use a MAC curve.

The most prominent attempt at developing a comprehensive marginal abatement cost curve is the well-known McKinsey curve, which is constructed using engineering estimates of the cost of implementing new technologies or other measures.

Figure 1 displays the global version of the McKinsey curve (McKinsey & Company 2009). A striking feature of the McKinsey curve, which is shared by MAC curves more generally (for example, see figure 2 in Grubb et al. 1993), is that some interventions have negative abatement costs: that is, emissions can be reduced, and money saved, at the same time. Economists, including ourselves, are often skeptical of these "free lunch" estimates, unless they are supported by convincing evidence and explanations. Negative costs require institutional entities, such as firms, not to be optimizing, or require the existence of behavioral failures in consumer decision-making (like consumers acting myopically). In some cases, entities such as

*Table 1*
**New Source Generation Costs when Compared to Existing Coal Generation**
*(ordered from lowest to highest)*

| *Technology* | *Cost estimate ($2017/ton $CO_2$)* |
|---|---|
| Onshore wind | 24 |
| Natural gas combined cycle | 24 |
| Utility-scale solar photovoltaic | 28 |
| Natural gas with carbon capture and storage | 42 |
| Advanced nuclear | 58 |
| Coal retrofit with carbon capture and storage | 84 |
| New coal with carbon capture and storage | 95 |
| Offshore wind | 105 |
| Solar thermal | 132 |

*Source:* Author's calculations updating methodology from Clean Air Task Force (2013) based on Energy Information Administration estimates from the 2018 Annual Energy Outlook. Costs are projected for facilities that come online in 2022. Costs do not incorporate federal renewable tax credits.

governments are institutionally complex and/or not minimizing costs, so these free-lunch savings are potentially valid but institutionally difficult to realize. When these negative costs are for energy efficiency programs, this is often called the "energy efficiency gap" and there is a continued debate in the literature on whether there is a real gap or whether the gap can be explained by unaccounted-for costs (Gerarden, Newell, and Stavins 2017; Gillingham and Palmer 2014; in this journal, Allcott and Greenstone 2012).

The concern over negative costs highlights a limitation of marginal abatement curves like the McKinsey curve in Figure 1: specifically, that they are based on engineering estimates, which have their own assumptions and typically do not include behavioral considerations. An example of such a behavioral effect is turning the heat up because the cost of doing so has declined because of weatherization. Economists are typically interested in the combined effect of behavioral responses and the engineering costs.

**Static Cost Comparisons**

In addition to these and other methodological concerns, the cost estimates in the McKinsey curve in Figure 1 are out of date. We therefore turn to more current estimates of marginal costs. These estimates are drawn from the economics and trade literatures, supplemented by our own calculations.

To fix orders of magnitude, we begin with some "bottom-up" or engineering cost estimates for the power sector, presented in Table 1. These estimates compare the cost per ton of $CO_2$ abated by replacing electricity generated by an existing coal-fired power plant with electricity generated by a cleaner alternative. The estimates are based on the US Energy Information Administration's (2018) so-called

"levelized" cost of electricity for the different sources, which combines discounted capital, operating, and maintenance expenses to produce a cost of energy per megawatt-hour, given the typical utilization rate or capacity factor for each generation type. These estimates are similar to private sector estimates, such as those by Lazard (2017).

According to these estimates, the least expensive technologies to reduce emissions relative to existing coal are onshore wind, natural gas combined cycle, utility-scale solar photovoltaics, and natural gas with carbon capture and storage technology. Advanced nuclear technologies are more expensive, followed by other carbon capture and storage technologies, offshore wind, and solar thermal. The technologies in this set of estimates that are less expensive (when replacing existing coal) than the Obama administration's social cost of carbon estimate of $46 per ton of $CO_2$ are onshore wind, natural gas combined cycle, utility scale photovoltaic, and natural gas with 90 percent carbon capture and storage. In comparison, offshore wind and solar thermal are currently quite expensive ways to reduce emissions (although offshore wind costs are falling). These estimates only consider climate benefits of switching from coal, not any other health co-benefits arising from reductions in local air pollutants.

From a policy perspective, engineering cost estimates such as those in Table 1 have important limitations. Some of these technologies are in wide current use, so cost estimates are reasonably reliable (onshore wind, natural gas combined cycle), whereas other technologies have demonstrated technical feasibility but current projects are subject to large cost overruns, so the engineering costs could be underestimates (for example, advanced nuclear, carbon capture and storage). Another limitation is that these are national averages, and costs vary regionally depending on local conditions (for example, local fuel prices, wind conditions, and insolation). In addition, these are costs of switching technologies, which differ from the costs of a policy designed to encourage technology switching. These engineering estimates do not incorporate behavioral responses or any indirect emissions such as fugitive methane emissions from the production and transport of natural gas.

We therefore turn to a systematic review of costs of interventions—typically policies—aimed at reducing greenhouse gas emissions. This review draws on more than 50 recent articles in the economics literature. We selected papers based on a few criteria. First, the paper must be an economic analysis, so we draw most heavily from papers published in economics journals and economics working paper series. Second, the paper must either have enough information so that we can calculate a cost per ton of emissions reduction or include an explicit estimate of this cost. Most papers we review have an explicit estimate in dollars per ton $CO_2$. Third, we focus on papers published in the past decade, and nearly all of the papers included in our review are published after 2006. In some cases, we have supplemented the estimates from the economics literature with studies from the trade literature and/or our own calculations.

The results are summarized in Table 2. The table presents ranges of estimates whenever there are multiple estimates from either the same study or multiple studies; the online Appendix available with this paper at http://e-jep.org provides

*Table 2*
**Static Costs of Policies based on a Compilation of Economic Studies**
*(ordered from lowest to highest cost)*

| *Policy* | *Estimate ($2017/ton $CO_{2e}$)* |
|---|---|
| Behavioral energy efficiency | –190 |
| Corn starch ethanol (US) | –18 to +310 |
| Renewable Portfolio Standards | 0–190 |
| Reforestation | 1–10 |
| Wind energy subsidies | 2–260 |
| Clean Power Plan | 11 |
| Gasoline tax | 18–47 |
| Methane flaring regulation | 20 |
| Reducing federal coal leasing | 33–68 |
| CAFE Standards | 48–310 |
| Agricultural emissions policies | 50–65 |
| National Clean Energy Standard | 51–110 |
| Soil management | 57 |
| Livestock management policies | 71 |
| Concentrating solar power expansion (China & India) | 100 |
| Renewable fuel subsidies | 100 |
| Low carbon fuel standard | 100–2,900 |
| Solar photovoltaics subsidies | 140–2,100 |
| Biodiesel | 150–250 |
| Energy efficiency programs (China) | 250–300 |
| Cash for Clunkers | 270–420 |
| Weatherization assistance program | 350 |
| Dedicated battery electric vehicle subsidy | 350–640 |

*Note:* Figures are rounded to two significant digits. We have converted all estimates to 2017 dollars for comparability. See Appendix Table A-1 for sources and methods. $CO_{2e}$ denotes conversion of tons of non-$CO_2$ greenhouse gases to their $CO_2$ equivalent based on their global warming potential.

an expanded version of Table 2 with sources and methods. As in a marginal abatement cost curve, we have ordered the estimates in Table 2 from lowest to highest cost.

We highlight seven features of Table 2.

First, the range of costs of these interventions is extremely wide, from less than $10 per ton to over $1,000 per ton. What is striking about this range is that all the interventions in Table 2 are either policy steps that have been implemented, at least in some jurisdiction, or have been actively proposed and considered. Most of the costs are relatively expensive, in the sense that they exceed $46/ton. Evidently, static cost is only one consideration when a policy is proposed or considered. This heterogeneity likely stems from multiple sources, including the carbon intensity of the displaced fuel (for example, is the electricity on the grid coming from coal or hydropower?) and the other policies in place.

Second, there is a wide range of costs *within* a type of intervention. For example, subsidies to wind generation, such as the wind production tax credit in the United States, have estimated carbon abatement costs ranging from $2 to more than $260

per ton of reduced $CO_2$. For wind power, one reason for the large range is that there is large variation across sites in wind potential. The range is even wider for subsidies for solar photovoltaics, in part because there is wide variation in solar potential across locations (the solar power potential in southwestern Arizona is roughly twice that in upstate New York[3]), in part because of the timing of the programs (for example, earlier programs faced higher solar panel costs than later programs), and in part because of differences in scale (utility-scale arrays cost much less to install per kilowatt than rooftop arrays) (Baker, Fowlie, Lemoine, and Reynolds 2013). The wide ranges of estimates in Table 2 underscore that policies may have very different costs per ton of $CO_2$ depending on the empirical setting and/or the methodology of the study. The ranges of the estimates should not necessarily be taken as a proxy for uncertainty, for they simply show how estimates vary across studies. Due to within-study uncertainty, values above and below the ranges are likely to occur with some probability.

Third, some of the interventions that have negative economic costs in the McKinsey curve (and in other marginal abatement cost curves) have positive costs here. For example, engineering estimates of weatherization programs often suggest that they have negative costs. So why have such changes not already been undertaken? This is the energy efficiency paradox. In a randomized controlled trial, however, Fowlie, Greenstone, and Wolfram (2018) found that the actual costs of the weatherization exceeded the savings, leading to the $350/ton estimate of the mitigation cost reported in Table 2. They attribute the difference between the negative engineering costs and the actual positive costs for the homes in their study primarily to flaws in the engineering models.

Fourth, some of the costs in Table 2 *are* negative. A striking estimate arises from behavioral economics studies of how small nudges can get consumers to reduce their energy consumption, thereby saving money while reducing emissions; the estimate in Table 2 is taken from Allcott and Mullainathan's (2010) meta-analysis of behavioral interventions. An example of such a nudge is the OPOWER program, in which an insert in the residential electricity bill compares the homeowner's usage to that of neighbors, costing the utility very little and leading to consumer savings. One concern, which we share, is that while the cost of such reductions is negative, the total emissions reductions from such nudges are likely to be relatively small and partially transitory. The other negative estimate in Table 2 is for corn ethanol, which some might find surprising.

In the United States, petroleum gasoline blend stock must be blended with an octane booster to bring it up to the 87 octane standard of regular unleaded gasoline. Ethanol is a lower-cost octane booster than its petroleum alternatives (Irwin and Good 2017). In 2012—a year in which there were no direct federal subsidies and the federal ethanol mandate under the Renewable Fuel Standard was not binding—ethanol comprised just under 10 percent of the US retail gasoline supply.

[3] See the National Renewable Energy Laboratory (NREL) National Solar Radiation Database (NSRDB) Data Vewer at https://maps.nrel.gov/nsrdb-viewer/.

The California Air Resources Board (2018) estimates that ethanol from new corn ethanol plants has roughly 70 percent of the life-cycle $CO_2$ emissions of petroleum, including the carbon effects of induced land use change. Thus, for blends up to 10 percent, ethanol has negative greenhouse gas emissions reductions costs, and indeed is the market choice. Blending ethanol up to approximately 30 percent continues to enhance octane. The US fueling infrastructure, however, generally cannot handle blends above 10 percent, nor are engines designed to harness those octane advantages to improve energy efficiency, a situation known as the "E10 blend wall." As a result, subsidies are needed to incentivize ethanol consumption in blends higher than E10, and those costs increase quickly when measured in dollars per ton of $CO_2$ avoided.

Fifth, a few of the interventions have very low costs. Some, like the Clean Power Plan—the Obama administration's rulemaking for $CO_2$ emissions standards in the power sector—and regulations to reduce methane flaring from fracked oil wells that coproduce natural gas, are examples in which the regulation intensity was chosen with cost in mind. The Clean Power Plan is notable for its low cost per ton of emissions reductions (this estimate is taken from the *Regulatory Impact Analysis for the Clean Power Plan*, US Environmental Protection Agency 2015). This cost per ton is less than any of the engineering costs in Table 1, for two reasons. First, some of the emissions reduction comes from switching generation from existing coal-fired plants to existing gas-fired plants, and so does not require building a new plant as in Table 1. Second, because the Clean Power Plan allowed interstate trading of emissions permits, new low-greenhouse gas generating facilities would be built where it is most economically efficient to do so, yielding lower costs than the generic plant replacement costs in Table 1. The Clean Power Plan is also notable because its projected $CO_2$ emissions reductions are the largest, or nearly so, among the interventions in Table 2.

Sixth, some of the interventions have very high static costs. The United States and Europe have programs that require blending biodiesel into the diesel fuel supply. Biodiesel can be made from many oil feedstocks, including waste grease, but on the margin it is made from food-competing feedstocks such as soybean oil. These food oils are expensive and production of soy biodiesel requires a large subsidy, which is provided in the United States primarily through a tax credit and through the Renewable Fuel Standard. In other cases, the high costs are a result of inefficiencies in program design. For example, the temporary Cash for Clunkers program was installed at the depth of the recession in 2009 to provide an infusion of demand for new cars to support the auto industry and to provide countercyclical fiscal stimulus. Because the program exchanged old vehicles for more efficient new ones, it boosted fleet fuel economy. However, it had substantial temporary inframarginal transfers that were not a problem for its primary purpose—to pull forward auto demand—but made it a costly way to reduce emissions.

Seventh, the literature suggests that the cost of reducing carbon is low for some land use policies (see "Reforestation" on Table 2). In a randomized controlled experiment that lasted two years, Jayachandran et al. (2017) found that cash payments

for forest conservation in Uganda substantially reduced deforestation and cost $1 per ton of carbon sequestered. They do not, however, provide evidence on what happened after the payments stopped, and a natural concern is that there would be a reversion to the deforestation baseline. If so, the emissions reduction would be temporary, that is, the emissions would simply be postponed, not eliminated.[4] This distinction between permanent and temporary sequestration, along with the difficulty of ascertaining whether the payments actually induce incremental carbon retention in practice (something that was in fact found in Jayachandran et al.'s experiment), are at the heart of the controversy over the use of carbon offsets (for example, van Benthem and Kerr 2013; Bento, Kanbur, and Leard 2016).

One sobering insight from the estimates in Table 2 is that many of the least-expensive interventions cover a small amount of $CO_2$ reductions, whereas the scalable technologies that that are at the center of discussions about a transformation to a low-carbon economy—electric vehicles, solar photovoltaic panels, and offshore wind turbines—are among the most expensive on the list. Behavioral nudges are a very small step towards deep decarbonization. In contrast, the more expensive scalable technologies have a much greater potential for substantial emissions reductions. For these technologies, what matters most are not the static costs today, but the costs and consequences of these interventions over time, that is, the dynamic costs of the intervention. It is informative to know what are the cheapest interventions to do today, but we would argue that it is even more important to know what interventions might most effectively drive down the price of large-scale reductions in emissions in the future.

## Dynamic Costs

The long residence time of $CO_2$ in the atmosphere makes climate change a long-term problem, in which (to a first approximation) what matters is the total number of tons emitted over some long horizon. As a result, the key to reducing emissions in the future is to have low-cost alternatives to fossil fuels that are zero- or low-carbon. The true total cost of investments or interventions today therefore must include both their static or face-value cost, and any spillovers those investments have for future costs of emissions reduction. The importance of a dynamic perspective is hardly new—see Popp, Newell, and Jaffe (2010) for a review—but it is often neglected both in the public debate and in the literature on costs of abatement. Yet, the welfare benefits of even small growth rates in the efficiency of clean technologies may be large, as suggested by simulations in Hassler, Krussel, Olovsson, and Reiter (2018).

[4]The distinction between temporary and permanent forest sequestration is important. Temporary rainforest sequestration is equivalent to storing emissions then releasing them later. In a manner analogous to how generating electricity from wind displaces retired coal-fired electricity, permanent sequestration permanently keeps the $CO_2$ in question out of the atmosphere.

**Conceptual Framework**

The static cost estimates of the previous section focus on direct reductions in emissions in the relatively short-run. However, expenditures on certain kinds of short-run reductions in emissions today can also affect emissions in the future, above and beyond direct emissions from the project. There are at least four reasons why this second component of emissions reduction could be nonzero and possibly large for some green technologies. Three of these stem from externalities, while the fourth is the difference between myopic and dynamic cost minimization.

First, many of these low-carbon technologies are nascent, and there could be substantial gains in production efficiency as more units are produced. Such gains can arise from engineering and managerial improvements made as production increases, a channel referred to as learning by doing, and from scale economies. To the extent that such gains are only partially appropriable by the firm, an expenditure today provides a positive externality that reduces costs in the future. The first case study that we discuss in the next subsection—solar panels—focuses on this learning-by-doing effect.

Second, a related externality arises from research and development spillovers because research results are only partially appropriable. These spillovers also represent a market failure, and economists have argued that the spillovers are likely to be particularly large for emerging clean technologies (Nordhaus 2011). To the extent that purchases today spur additional research, which then reduces costs, expenditures today reduce emissions tomorrow. It can often be difficult to separate the effects of research and development spillovers from learning-by-doing spillovers, for as a firm ramps up production, it also may ramp up research. For this reason, economists have often encouraged caution in relying too heavily on learning-by-doing to model technological change (Nordhaus 2014).

Third, a separate externality that is present for some technologies is a network or "chicken and egg" externality, in which an expenditure today influences the options that are available to others in the future. For example, purchases of electric vehicles today will, on the margin, stimulate demand for charging stations, which once installed will lower the effective cost for future potential purchasers of electric vehicles. Our second case study, of electric vehicles, in principle includes both learning-by-doing and network externalities.

Fourth, energy investments typically have substantial irreversible components, which in general implies state dependence so that the dynamically optimal path may differ from a sequence of myopic optimizations each chosen at a point in time. This potential for lock-in is at the heart of the debate about the merits of natural gas as a bridge fuel towards decarbonizing the power sector, in which renewable proponents argue that natural gas is cheaper only if one optimizes myopically and fails to recognize that the power sector will eventually need to be decarbonized. This intuition underlies Vogt-Schilb, Meunier, and Hallegatte (2018), who show that if abatement is achieved through investment in long-lived capital, it can be optimal to begin emissions abatement with expensive abatement investments that have large emissions reduction potential because they crowd-out dirtier long-lived investments. Irreversibility (state

dependence) also underlies the results of Fischer and Newall (2008), Acemoglu, Aghion, Bursztyn, and Hemous (2012), and Acemoglu, Akcigit, Hanley, and Kerr (2016), who show that a carbon price combined with research subsidies for low-greenhouse-gas technologies may be desirable to attain dynamically efficient outcomes.

Of course, long-term considerations may not always lower the cost of emissions reductions. For example, nuclear power has long had major federal research subsidies but its cost has gone up, not down (Davis and Hausman 2016). Additionally, as the marginal ton of displaced electricity becomes cleaner (for example, displacing natural gas instead of coal), the cost per ton abated by low-carbon renewables will tend to increase. One major reason why dynamic considerations are often ignored is that they tend to be highly uncertain. But that uncertainty should be viewed as a research challenge rather than an excuse to ignore dynamic considerations. And there is some evidence from the recent literature.

**Dynamic Cost Case Study 1: Solar Panels**

From 2010 to 2015, the price of solar photovoltaic panels fell by two-thirds, while annual global panel installations grew by 250 percent, as shown in Figure 2. The fact that panel sales increased when their price fell is hardly surprising, but more intriguing is that the steepest decline in panel prices after about 2007 post-dated the initial growth in panel sales, which began around 2002. The growth in sales in the mid-2000s was associated with policies that provided aggressive financial support for installing rooftop photovoltaic arrays through the German Energiewende, which provided a substantial feed-in tariff that allowed solar installations to be compensated at a very high rate for electricity fed into the grid, and the California Solar Initiative, which provided generous upfront subsidies for solar installations.[5] These early panel purchases were very expensive and account for some of the high photovoltaic cost estimates in Table 2. As stressed by a number of researchers (for example, Borenstein 2017), the static cost per ton of $CO_2$ reduced from policies to encourage solar installations tends to be high. Our literature review finds costs ranging from more than \$100 per ton of $CO_2$ to in the thousands per ton of $CO_2$. On the lower end, Hughes and Podolefsky (2015) estimate costs of the California Solar Initiative at between \$130 and \$196 per ton. On the high end, Abrell, Kosch, and Rausch (2017) find a static cost per ton of €500–1300 (roughly \$574–\$1,492 in 2017 dollars) for solar feed-in tariffs in Germany and Spain (a solar feed-in tariff is a long-term fixed price contract for purchasing electricity from a solar array).

However, both the timing shown in Figure 2 and recent research suggests that the early push in demand, stimulated by deep government subsidies, did in fact help to drive down the price of solar panels. One channel is that current subsidies may

[5] Many US states have generous net metering policies that act as implicit subsidies by compensating solar fed into the grid at the retail rate. Some states have Renewable Portfolio Standards that require utilities to procure certain amounts of renewable power (sometimes with a solar carve-out) by certain dates. At the federal level, since 2008, there has been a 30 percent investment tax credit for the installation of a residential solar system.

*Figure 2*
**Solar Panel Price Indexes Excluding Subsidies and Cumulative Worldwide Installed Capacity, 1983–2015**

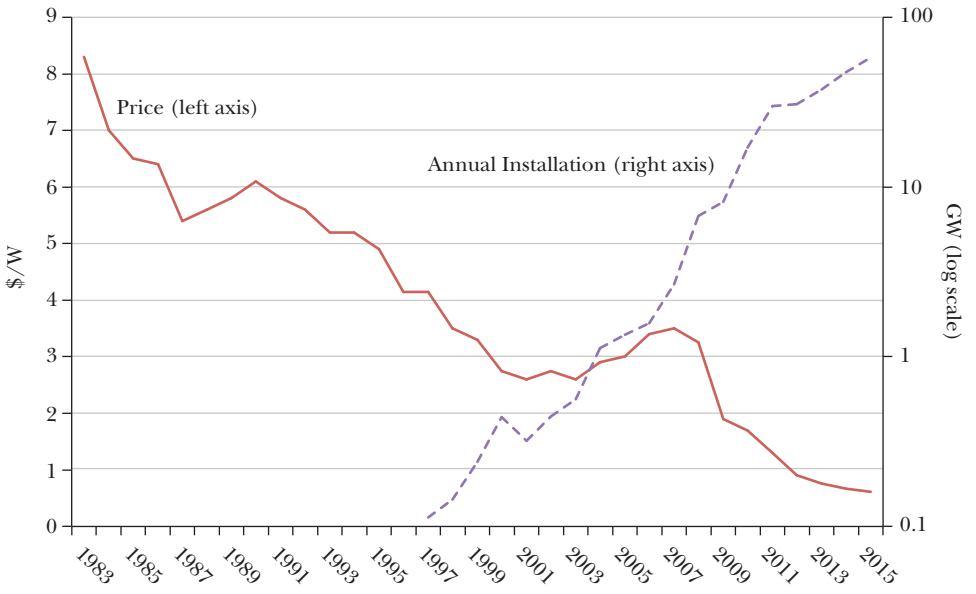


*Source:* International Energy Agency (2017), Navigant Consulting (2009), and Gerarden (2018).

encourage firms to innovate to reduce their future costs. Gerarden (2018) estimates that this induced innovation effect, which does not include learning-by-doing, contributed to the decline in solar array prices and increased the long-run external social benefits from global government subsidies to solar adoption by at least 22 percent. His results further suggest an important spillover from any single country that subsidizes solar to the rest of the world due to the investment in innovation by international firms. In this sense, the German Energiewende subsidized lower-cost solar for the rest of the world.

Other channels for cost reduction in the production of solar panels include learning-by-doing and economies of scale. Nemet (2006) decomposes the reduction in cost into the manufacturing plant size, module efficiency, and cost of silicon, finding that between 1980 and 2001, economies of scale from larger manufacturing plant sizes accounted for 43 percent of the cost reduction. Most of the remaining cost reduction could be attributed to improvements in module efficiency due to research and development investment. The substantial cost declines in solar module prices over the past decade are often attributed to economies of scale (Carvalho, Dechezleprêtre, and Glachant 2017). Economies of scale and learning-by-doing can in many cases be appropriable by the firms making decisions to scale up (this appears to be the case for learning-by-doing among rooftop solar installers, as Bollinger and Gillingham 2018 explain), so that learning-by-doing and scale economies do not by

themselves necessarily constitute reasons for policy intervention. Absent a carbon price, however, the demand for solar panels will be less than it would be were there a carbon price. As a result, second-best policies that are initially expensive (like the German Energiewende) can in principle stimulate production that would not normally happen because fossil fuels are cheaper than they would be, were their externality priced. For solar panels, at least, all this seems to have been the case.

Going forward, we might continue to see policy-induced cost reductions for solar technology. As the penetration of solar rises, and as the rest of the electricity system decarbonizes, such cost reductions will have to continue to be substantial to offset the higher potential costs of additional storage needed because of solar intermittency.

**Dynamic Cost Case Study 2: Electric Vehicles**

Like solar panels, the static costs of $CO_2$ reductions obtained by using electric vehicles is high in Table 2 (the last row). Today, many electric vehicles in the United States are charged using electricity that on the margin is produced by fossil fuels. Holland, Mansur, Muller, and Yates (2016) use the method of Graff Zivin, Kotchen, and Mansur (2014) for computing marginal emissions to examine the static optimal second-best purchase subsidy on electric vehicles accounting for both greenhouse gases and local air pollution. Holland, Mansur, Muller, and Yates (2016) find that the subsidy ranges from a subsidy of $2,785 in California (with relatively clean electricity on the margin) to a *penalty* of $4,964 in North Dakota, where electricity is generated from coal. Archsmith, Kendall, and Rapson (2015) perform similar calculations that additionally include life-cycle considerations, and find that on average electric vehicles currently only slightly reduce greenhouse gases relative to gasoline-powered vehicles.

From a dynamic perspective, however, the case against programs to support electric vehicle purchases is far less clear. The static calculations ignore the fact that the grid is evolving and becoming cleaner. Moreover, the general issues raised for solar panels—induced innovation, learning by doing, and economies of scale that would not otherwise be achieved because carbon is not priced—apply to electric vehicles as well. Unlike the case of solar panels, however, we are not aware of any research that investigates drivers of price dynamics for electric vehicles, perhaps because the cost declines and sales growth are so recent. The available data are, however, suggestive that the analogy to demand-pull effects for solar panels also applies to electric vehicles.

Figure 3 plots electric vehicles that entered the market from model years 2011 to 2018 based on their suggested retail price (y-axis) and battery range (x-axis). The price-range frontier has strikingly shifted out: more recent market entrants have greater battery range at lower cost, underscoring this rapid improvement in technology.

The large declines in price for vehicles with the same range is mainly due to the ongoing decline in battery prices. From 2009 to 2015, the price of batteries for electric vehicles fell by 75 percent (US Department of Energy 2016). Like solar

*Figure 3*
**Electric Vehicle Manufacturers Suggested Retail Price (MSRP) Plotted against the Battery Range Shows Impressive Technology Improvements within a Short Time**

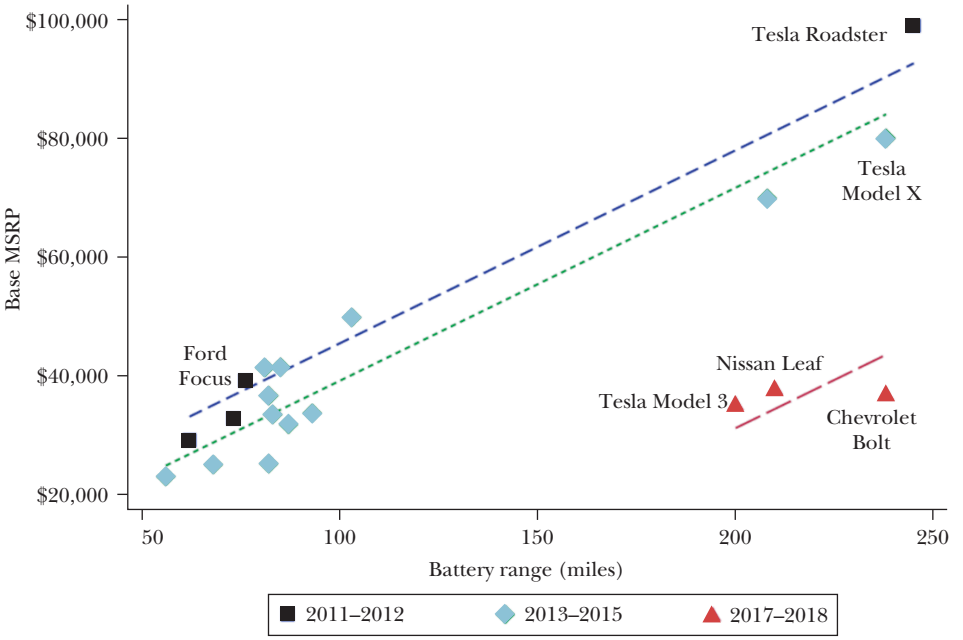


*Source:* J. Li (2017) and authors' calculations.
*Note:* Dates indicate year the model is introduced. Regression lines are fit with a common slope and different intercept for each group of model years.

photovoltaic arrays, electric vehicles have been a target of demand-pull subsidy programs. Since the Energy Policy Act of 2005, there has been a US federal income tax credit of $7,500 (which phases out with production by any given manufacturer). Many states have additional incentives, such as a $3,000 rebate in Connecticut, eligibility for driving in a high-occupancy lane with only a single occupant in California, and a zero-emissions vehicle mandate in 10 states that requires automakers to sell a certain number of zero tailpipe emission vehicles (including electric vehicles) for every non-zero-emissions vehicle sold. Numerous papers in the transportation literature have provided evidence suggesting that electric vehicle subsidies increase demand for electric vehicles, as one would expect (reviewed in Zhou, Levin, and Plotkin 2016). The general pattern of demand-pull policies combined with subsequent sharp declines in costs is similar to that found for solar panels. We note that it is consistent with learning-by-doing and scale economy effects, and that confirming or refuting this hypothesis is an important area for future research by economists.

Electric vehicles also exhibit network effects, whereby the purchase of an additional electric vehicles makes the installation of a charging station more profitable. Thus, a positive feedback can exist, leading to multiple equilibria. For example, there may be one equilibrium with few charging stations and few or no electric vehicles, and another with many charging stations and electric vehicles. There is a growing literature on electric vehicles and network effects. Zhou and S. Li (2017) point out the possibility of multiple equilibria in electric vehicles and argue that a subsidy targeted at the marginal electric vehicle purchaser can be much more efficient than a policy that provides large inframarginal gains to those who would purchase an electric vehicle anyway. Yu, S. Li, and Tong (2016) discuss how network effects can lead the market solution to underinvest in electric vehicles compared to what is socially optimal. J. Li (2017) develops a structural model of two-sided market estimated with vehicle registration data from the United States and finds that mandating compatibility in charging stations would benefit consumers, enhance network effects, and increase the size of the electric vehicle market. Springel (2018) uses vehicle registration data from Norway—the country with the highest penetration of electric vehicles—to estimate a structural model showing that subsidies for charging stations are more effective for increasing electric vehicles uptake than are purchase subsidies for electric vehicles, but their effectiveness tapers off with increased subsidy.

The findings of these papers on network effects point to how a static perspective on policies to encourage technologies such as electric vehicles miss important aspects germane to the long-term cost-effectiveness of different policy approaches.

**Static versus Dynamic Costs: Other Examples**

Our two case studies present the sanguine view that seemingly expensive investments today result in lower costs in the future, a finding broadly akin to the theoretical work of Vogt-Schilb, Meunier, and Hallegatte (2018), Newbery (2018), Acemoglu et al. (2012), and Acemoglu et al. (2016). This happy result, however, is not preordained. For example, taking the dynamic approach may lead one to invest *less* in a carbon abatement technology if costs are expected to increase, rather than decrease, over time. Nuclear technology may fall into this category as construction costs of nuclear energy have risen, not fallen (Davis and Hausman 2016). Increasing costs of integrating renewable electricity into the electric grid can also work in this direction. In other cases, the static approach is perfectly appropriate. Consider policies to reduce methane leaks from the natural gas distribution system: the costs of sealing these leaks is likely to be similar in the near future as it is today because the process of sealing leaks is well understood but costly (digging up pavement and replacing pipes). Still other cases are less clear. Policies that would promote fuel switching to natural gas may reduce emissions in the short-run, but have potential to lead to investments in long-lived capital assets, and possibly even technological lock-in (Gillingham and Huang 2018).

## Challenges

The costs of reducing carbon emissions discussed in this paper pose several challenges. One of these challenges is that some politically appealing programs, such as support for biodiesel or subsidies for energy efficiency programs, can be quite costly either for technological reasons or because of behavioral responses. Because the costs for these programs are often masked and only apparent upon scrutiny by economists, they appear low-cost—but are not.

A second challenge is the reverse, where highly visible programs are perceived as high-cost, but are not. A prominent example is the Clean Power Plan, which would have resulted in large emissions reductions for a cost far below that of many other programs already in place.

A third challenge is that the static costs provide at best an incomplete picture of the true costs of a particular action, which must include the dynamic consequences. The sign of those dynamic consequences in general depends on the intervention. If the intervention is replacing coal electricity generation by natural gas, low short-term costs might lead to higher longer-term costs if the result is long-lived natural gas infrastructure that is locked in and costly to abandon as the price of renewables drops. In contrast, if the intervention is providing subsidies for purchasing electric vehicles, the demand-pull effects of induced learning by doing and economies of scale can make dynamic costs much lower than a myopic static calculation would suggest. Because climate change is a long-term problem and the changes ultimately needed to reduce emissions are vast, the dynamic costs are far more important than the static ones.

A fourth challenge is to the economic research community, and it stems from the previous observation. As is clear from our review, most of the empirical studies of costs by economists focus on static costs, typically static costs of programs that have already been in place. This is natural because there is data on these programs, and understanding the costs of previous programs is a helpful guide to designing future programs. But particularly in the field of climate change research, more attention is needed on the determinants of dynamic costs. This exciting field of research merges environmental and energy economics with the extant literature on productivity, diffusion, and learning-by-doing. We have highlighted two areas—solar photovoltaics and electric vehicles—in which demand-pull policies appear to have induced cost reductions; however, that need not always happen and magnitudes surely vary from one case to the next.

Climate change is a long-term problem, and the focus of policy must be on long-term solutions. To make major progress on climate goals, like 80 percent decarbonization by 2050 in the United States, will require new technology deployed on a vast scale. Even if each technological step is evolutionary—cheaper electric vehicle batteries, connecting the grid to harness the wind potential in the Midwest, reducing the cost of offshore wind, developing and commercializing low-carbon fuels for air transport—the overall change will be revolutionary. If a price on carbon is not politically feasible—and arguably even if it is—these long-term considerations

need to be incorporated into our short-term policy tradeoffs. From the perspective of the cost calculations in this paper, one clear implication is that choosing low-cost interventions without a future, including ones that lock in fossil fuel infrastructure, can result in too much emphasis being placed on what is cheapest to do today. We are always surprised by the specifics of technological progress, but as economists, we are not surprised that it is more likely to occur when the right incentives are in place.

■ *The authors thank Todd Gerarden and Jing Li for providing data and comments, and Tim Bialecki for research assistance. The authors also thank Joe Aldy, Rick Duke, Matt Kotchen, Derek Lemoine, Will Rafey, and Gernot Wagner for helpful discussions and/or comments on an early draft of this paper.*

## References

**Abrell, Jan, Mirjam Kosch, and Sebastian Rausch.** 2017. "The Economic Cost of Carbon Abatement with Renewable Energy Policies." CER-ETH (Center of Economic Research at ETH Zurich) Working Paper 17/273.

**Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous.** 2012. "The Environment and Directed Technical Change." *American Economic Review* 102(1): 131–66.

**Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr.** 2016. "Transition to Clean Technology." *Journal of Political Economy* 124(1): 52–104.

**Allcott, Hunt, and Michael Greenstone.** 2012. "Is There an Energy Efficiency Gap?" *Journal of Economic Perspectives* 26(1): 3–28.

**Allcott, Hunt, and Sendhil Mullainathan.** 2010. "Behavior and Energy Policy." *Science* 327(5970): 1204–1205.

**Archsmith, James, Alissa Kendall, and David Rapson.** 2015. "From Cradle to Junkyard: Assessing the Life Cycle Greenhouse Gas Benefits of Electric Vehicles." *Research in Transportation Economics* 52: 72–90.

**Baker, Erin, Meredith Fowlie, Derek Lemoine, and Stanley S. Reynolds.** 2013. "The Economics of Solar Electricity." *Annual Review of Resource Economics* 5: 387–426.

**Bento, Antonio, Ravi Kanbur, and Benjamin Leard.** 2016. "On the Importance of Baseline Setting in Carbon Offsets Markets." *Climatic Change* 137(3–4): 625–37.

**Bollinger, Bryan, and Kenneth Gillingham.** 2018. "Learning-by-Doing in Solar Photovoltaic Installations." Unpublished paper, Yale University.

**Borenstein, Severin.** 2017. "The Private Net Benefits of Residential Solar PV: The Role of Electricity Tariffs, Tax Incentives and Rebates." *Journal of the Association of Environmental and Resource Economists* 4(S1): S85–S122.

**California Air Resources Board.** 2018. "LCFS Pathway Certified Carbon Intensities." at https://www.arb.ca.gov/fuels/lcfs/fuelpathways/pathwaytable.htm and https://www.arb.ca.gov/fuels/lcfs/fuelpathways/current-pathways_all.xlsx, accessed August 2, 2018

**Carvalho, Maria, Antoine Dechezleprêtre, and Matthieu Glachant.** 2017. "Understanding the Dynamics of Global Value Chains for Solar Photovoltaic Technologies." WIPO Economic Research Working Paper 40, World Intellectual Property Organization. http://www.wipo.int/edocs/pubdocs/en/wipo_pub_econstat_wp_40.pdf.

**Clean Air Task Force.** 2013. "Comparison of $CO_2$ Abatement Costs in the United States for Various Low and No Carbon Resources." Factsheet, September. Available at: http://www.catf.us/resources/factsheets/

**Davis, Lucas W., and Catherine Hausman.** 2016. "Market Impacts of a Nuclear Power Plant

Closure." *American Economic Journal: Applied Economics* 8(2): 92–122.

**Fischer, Carolyn, and Richard G. Newell.** 2008. "Environmental and Technology Policies for Climate Mitigation." *Journal of Environmental Economics and Management* 55(2): 142–62.

**Fowlie, Meredith, Michael Greenstone, and Catherine Wolfram.** 2018. "Do Energy Efficiency Investments Deliver? Evidence from the Weatherization Assistance Program." *Quarterly Journal of Economics* 133(3): 1597–1644.

**Gerarden, Todd.** 2018. "Demanding Innovation: The Impact of Consumer Subsidies on Solar Panel Production Costs." Unpublished paper. Available at: https://scholar.harvard.edu/gerarden/publications/innovation-solar-panel-manufacturing.

**Gerarden, Todd D., Richard G. Newell, and Robert N. Stavins.** 2017. "Assessing the Energy-Efficiency Gap." *Journal of Economic Literature* 55(4): 1486–1525.

**Gillingham, Kenneth, and P. Huang.** 2018. "The Environmental and Welfare Implications of Abundant Natural Gas." Unpublished Paper, Yale University.

**Gillingham, Kenneth, William Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, and John Reilly.** 2018. "Modeling Uncertainty in Integrated Assessment of Climate Change: A Multimodel Comparison." *Journal of the Association of Environmental and Resource Economists* 5(4): 791–826.

**Gillingham, Kenneth, and Karen Palmer.** 2014. "Bridging the Energy Efficiency Gap: Policy Insights from Economic Theory and Empirical Evidence." *Review of Environmental Economics and Policy* 8(1): 18–38.

**Graff Zivin, Joshua S., Matthew J. Kotchen, and Erin T. Mansur.** 2014. "Spatial and Temporal Heterogeneity of Marginal Emissions: Implications for Electric Cars and Other Electricity-Shifting Policies." *Journal of Economic Behavior & Organization* 107(Part A): 248–68.

**Grubb, Michael, Jae Edmonds, Patrick ten Brink, and Michael Morrison.** 1993. "The Costs of Limiting Fossil-Fuel $CO_2$ Emissions: A Survey and Analysis." *Annual Review of Energy and the Environment* 18: 397–478.

**Hassler, John, Per Krusell, Conny Olovsson, and Michael Reiter.** 2018. "Integrated Assessment in a Multi-region World with Multiple Energy Sources and Endogenous Technical Change." Unpublished paper.

**Holland, Stephen P., Erin T. Mansur, Nicholas Z. Muller, and Andrew J. Yates.** 2016. "Are there Environmental Benefits from Driving Electric Vehicles? The Importance of Local Factors." *American Economic Review* 106(12): 3700–29.

**Hughes, Jonathan E., and Molly Podolefsky.** 2015. "Getting Green with Solar Subsidies: Evidence from the California Solar Initiative." *Journal of the Association of Environmental and Resource Economists* 2(2): 235–75.

**International Energy Agency.** 2017. "Renewables 2017." https://www.iea.org/publications/renewables2017/.

**Intragency Working Group on Social Cost of Greenhouse Gases (IWG), United States Government.** 2016. "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866." Interagency Working Group on the Social Cost of Greenhouse Gases, United States Government. At https://www.epa.gov/sites/production/files/2016-12/documents/sc_co2_tsd_august_2016.pdf.

**Irwin, Scott, and Darrel Good.** 2017. "On the Value of Ethanol in the Gasoline Blend." *farmdoc daily,* March 15.

**Jayachandran, Seema, Joost de Laat, Eric F. Lambin, Charlotte Y. Stanton, Robin Audy, and Nancy E. Thomas.** 2017. "Cash for Carbon: A Randomized Trial of Payments for Ecosystem Services to Reduce Deforestation." *Science,* July, 21, 367(6348): 267–73.

**Lazard**. 2017. "Lazard's Levelized Cost of Energy Analysis Version 11.0." Report. https://www.lazard.com/perspective/levelized-cost-of-energy-2017/.

**Li, Jing.** 2017. "Compatibility and Investment in the U.S. Electric Vehicle Market." Unpublished Paper. Available at: http://www.mit.edu/~lijing/.

**Marten, Alex L., and Stephen C. Newbold.** 2012."Estimating the Social Cost of Non-$CO_2$ GHG Emissions: Methane and Nitrous Oxide." *Energy Policy* 51: 957–72.

**McKinsey & Company.** 2009. "Pathways to a Low-Carbon Economy: Version 2 of the Global Greenhouse Gas Abatement Cost Curve." Report.

**National Academy of Sciences, Committee on Assessing Approaches to Updating the Social Cost of Carbon.** 2017. *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide.* Washington, DC: National Academies Press.

**Navigant Consulting.** 2009. Navigant Consulting PV Module Index (a proprietary data source).

**Nemet, Gregory F.** 2006. "Beyond the Learning Curve: Factors Influencing Cost Reductions in Photovoltaics." *Energy Policy* 34(17): 3218–32.

**Newbery, David.** 2018. "Evaluating the Case for Supporting Renewable Electricity." *Energy Policy* 120: 684–96.

**Newell, Richard G.** 2017. "Unpacking the Administration's Revised Social Cost of Carbon."

October, 10. http://www.rff.org/blog/2017/unpacking-administration-s-revised-social-cost-carbon.

**Nordhaus, William.** 2011. "Designing a Friendly Space for Technological Change to Slow Global Warming." *Energy Economics* 33(4): 665–73.

**Nordhaus, William D.** 2014. "The Perils of the Learning Model for Modeling Endogenous Technological Change." *Energy Journal* 35(1): 1–13.

**Popp, David, Richard G. Newell, and Adam B. Jaffe.** 2010. "Energy, the Environment, and Technological Change." Chap. 21 in *Handbook of the Economics of Innovation,* vol. 2, edited by Bronwyn H. Hall and Nathan Rosenberg. Elsevier.

**Springel, Katalin.** 2018. "Network Externality and Subsidy Structure in Two-Sided Markets: Evidence from Electric Vehicle Incentives." Unpublished paper.

**US Department of Energy.** 2016. "Revolution … Now: The Future Arrives for Five Clean Energy Technologies—2016 Update." Primary author: Paul Donohoo-Vallett. Contributing authors: Patrick Gilman, David Feldman, James Brodrick, David Gohlke, Roland Gravel, Amy Jiron, Carol Schutte, Sunita Satyapal, Tien Nguyen, Paul Scheihing, Blake Marshall, Sarah Harman. Available at: https://www.energy.gov/eere/downloads/revolution-now-future-arrives-five-clean-energy-technologies-2016-update.

**US Energy Information Administration.** 2018. "Levelized Cost and Levelized Avoided Cost of New Generation Resources in the Annual Energy Outlook 2018." March.

**US Environmental Protection Agency**. 2015. *Regulatory Impact Analysis for the Clean Power Plan Final Rule.* October 13. https://www3.epa.gov/ttnecas1/docs/ria/utilities_ria_final-clean-power-plan-existing-units_2015-08.pdf.

**van Benthem, Arthur, and Suzi Kerr**. 2013. "Scale and Transfers in International Emissions Offset Programs." *Journal of Public Economics* 107: 31–46.

**Vogt-Schilb, Adrien, Guy Meunier, and Stéphane Hallegatte.** 2018. "When Starting with the Most Expensive Option Makes Sense: Optimal Timing, Cost and Sectoral Allocation of Abatement Investment." *Journal of Environmental Economics and Management* 88: 210–33.

**Yu, Zhe, Shanjun Li, and Lang Tong.** 2016. "Market Dynamics and Indirect Network Effects in Electric Vehicle Diffusion." *Transportation Research Part D: Transport and Environment* 47: 336–56.

**Zhou, Yan, Todd Levin, and Steven E. Plotkin.** 2016. *Plug-in Electric Vehicle Policy Effectiveness: Literature Review*, Argonne National Laboratory ANL/ESD-16/8, May. https://www.energy.gov/sites/prod/files/2017/01/f34/Plug-In%20Electric%20Vehicle%20Policy%20Effectiveness%20Literature%20Review.pdf.

**Zhou, Yiyi, and Shanjun Li.** 2017. "Technology Adoption and Critical Mass: The Case of the U.S. Electric Vehicle Market." NET Institute Working Paper 15-10.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 73–96*

# Is This Tax Reform, or Just Confusion?

Joel Slemrod

Based on the experience of recent decades, the United States apparently musters the political will to change its tax system comprehensively about every 30 years, so it seems especially important to get it right when the chance arises. But at the annual meeting of the National Tax Association in November 2017, one of the keynote speakers began his remarks by asserting that if one set a monkey in front of a typewriter and, after some period, saved only those parts of what had been typed having to do with tax reform, the resulting tax law would be better than what the Republicans were poised to legislate. Ouch. The keynote speaker was Harvard economist and former Secretary of the Treasury Lawrence Summers. Fareed Zakaria (2017), the author and journalist, called the Tax Cuts and Jobs Act just passed "possibly the worst piece of major legislation in a generation . . ." and said, "Those who vote for this tax bill … will live in infamy . . ." Not just the worst tax legislation, but the worst legislation of any kind! The headline of a *Wall Street Journal* op-ed written by Alan Blinder (2017) was "Almost Everything Is Wrong with the New Tax Law."

Defenders of the tax legislation were, naturally, more positive. One supporting petition that was signed by 137 economists, including many academics, asserted that "[e]conomic growth will accelerate if the Tax Cuts and Jobs Act passes, leading to more jobs, higher wages, and a better standard of living for the American people" (as reported by CNBC 2017). The hyperbole of champions was often quantitative.

■ *Joel Slemrod is a Professor of Economics and the Paul W. McCracken Collegiate Professor of Business Economics and Public Policy at the Stephen M. Ross School of Business, both at the University of Michigan, Ann Arbor, Michigan. His email address is jslemrod@umich.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.73

doi=10.1257/jep.32.4.73

A report released by the Council of Economic Advisers (2017) asserted that a tax law like the one eventually passed—featuring a big cut in the corporate income tax rate—could lead to as much as a $9,000 increase in average wages. President Donald Trump remarked that he saw no reason why the tax bill could not deliver a growth rate as high as 6 percent (as reported in Bach 2017).

All in all, a casual observer of the public discussion of tax reform in 2017 could reasonably conclude that economists agree about almost nothing, including whether a specific tax law should even be categorized as tax reform or as mere confusion.[1] In this paper, I address that question and, more importantly, offer an assessment of the Tax Cuts and Jobs Act.[2] The law is clearly not "tax reform" as economists usually use that term: that is, it does not seek to broaden the tax base and reduce marginal rates in a roughly revenue-neutral manner. However, the law is not just a muddle. It seeks to address some widely acknowledged issues with corporate taxation, and takes some steps toward broadening the tax base, in part by reducing the incentive to itemize deductions.

## The Last Big Tax Reform, in 1986

I begin by revisiting the Tax Reform Act of 1986, which is the last time the US income tax was substantially overhauled. In the lead-up to that legislation, dissatisfaction with the income tax law had been simmering. President Jimmy Carter (1976) had called the tax code "a disgrace to the human race." Serious political debate can be traced to a bill introduced by Senator Bill Bradley (D-NJ) and Representative Richard Gephardt (D-MO) in the summer of 1983 called "The Fair Tax Act," which featured dramatically lower top individual and corporate tax rates, a broadening of the tax base at both the corporate and individual levels, and a structure designed to preserve both the existing distribution of the tax burden and the level of revenue collected. The Bradley–Gephardt bill was followed by a very similar proposal co-sponsored by two Republican senators, Jack Kemp of New York and Bob Kasten of Wisconsin.

President Reagan fueled the process in early 1984 by calling for a tax reform study. The US Department of the Treasury (1984) study—a multivolume report and set of proposals—was formulated essentially without political interference. It contained recommendations for a quite radical base-broadening, rate-reducing tax reform, including repealing the deduction of state and local taxes, introducing a deduction for 50 percent of dividends paid, as well as inflation indexing of depreciation, inventories, capital gains, and interest. Simplification was stressed, and the

[1] The title for this article is inspired by a line in the 1967 song "Love and Confusion" by Jimi Hendrix: "My heart burns with feeling, but my mind, it's cold and reeling/Is this love, baby, or is it just confusion?"

[2] Although the tax bill is commonly referred to as the Tax Cuts and Jobs Act, that is not its official name. The short name violated a Senate rule, so it was renamed: "To provide for reconciliation pursuant to titles II and V of the concurrent resolution on the budget for fiscal year 2018." A thoughtful preliminary analysis of the bill is offered in Gale, Gelfond, Krupkin, Mazur, and Toder (2018).

report spoke positively of a tax system in which two-thirds of individual taxpayers could opt for "return-free" filing in which the IRS would calculate tax liability based on information reports from third parties, with taxpayers needing only to verify or correct the "pre-filled" return. Politics subsequently intervened, and all of the just-mentioned aspects were abandoned when a reformulated proposal was released in May 1985. But the comprehensive Tax Reform Act of 1986 did survive the political process, and was signed by President Reagan on October 22, 1986.[3] It had wide bipartisan support: among Democrats, only 12 senators and 74 House members voted against the bill; among Republicans, only 11 senators and 62 House members voted against the bill—almost the same level of opposition as from the Democrats.

At least at first glance, the Tax Reform Act of 1986 and the Tax Cuts and Jobs Act of 2017 may appear to have much in common. Both bills substantially lowered the headline corporate rate, from 46 to 34 percent in 1986 and from 35 to 21 percent in 2017. Both substantially raised the standard deduction. Both reduced individual tax rates, although the 1986 law reduced them much more dramatically at the top, as the statutory rate on the highest incomes fell from 50 to 28 percent over a two-year period, capping a stunning fall from 70 percent to 28 percent between 1981 and 1986. Both bills featured some broadening of the tax base. The Tax Reform Act of 1986 introduced full taxation of realized capital gains, eliminated the sales tax deduction, and put severe limits on passive losses. One key difference is that the 1986 corporate tax rate cut was accompanied by the elimination of the investment tax credit and a slight slowing of depreciation allowances, such that most observers concluded that the cost of capital would increase, rather than fall (for example, Jorgenson and Yun 1990). Another key difference is that the Tax Cuts and Jobs Act was neither revenue neutral nor distributionally neutral, while the Tax Reform Act of 1986 was designed to be both, with a projected increase in corporate tax revenue offsetting projected declines in individual tax revenue.

## How Has the Environment for Tax Law Changed since 1986?

Tax law in the United States and in the world did not stay put between 1986 and 2017. The top US personal income tax rate crept up from 28 percent, reaching 39.6 percent by 1993, then falling to 35 percent for 2003 until 2012, and then returning to 39.6 percent. A preferential tax rate for long-term capital gains was reintroduced in the 1990 budget act (by raising the tax rate on other income while leaving unchanged that on capital gains), and a lower preferential rate was enacted

[3] Jeffrey Birnbaum and Alan Murray, the two Boswells of the Tax Reform Act of 1986, brilliantly recounted the political ups and downs in their 1987 book *Showdown at Gucci Gulch*, managing to make tax policy formulation *exciting*. Steuerle (2008) offers a valuable history of postwar US tax policy, focusing on developments after 1980. Two papers, McLure and Zodrow (1987) and Auerbach and Slemrod (1997), address the economics and politics of the Treasury proposals and the 1986 tax act.

in 1997; subsequently, an even lower rate was adopted and extended to dividends in 2003. To increase incentives for investment, "bonus" depreciation rules (allowing a fraction of capital expenditure to be written off immediately) were introduced in 2002 initially at a rate of 30 percent; they lapsed for a few years, were then revived by the Economic Stimulus Act of 2008, eventually increased to 100 percent for assets acquired after September 8, 2010, and then fell back to 50 percent.

Because the world has changed, we wouldn't expect that the best tax system in 1986 would also be ideal in 2018. (Nor, of course, did the Tax Reform Act of 1986 attain the ideal at the time.) Several aspects of change in the economic environment especially matter for thinking about tax policy. First, the US economy has become more integrated into the world economy. The sum of US imports and exports as a percentage of GDP nearly doubled from 1986, when this ratio was 16.9 percent, to 30.9 percent in 2013, only to decline slightly to 26.5 percent in 2016 (based on data from the World Development Indicators of the World Bank at https://data.worldbank.org/indicator/NE.TRD.GNFS.ZS?locations=US). As discussed further below, US corporations held $2.6 trillion abroad by 2017. As a result, the effect of taxes on cross-border stocks and flows has become even more important than in 1986.

Second, corporate tax rates in other countries around the world have declined steadily, largely driven by the pressure of global competition for investment and revenue, so that by 2017 the statutory US corporate income tax rate stood near the top of the league table. Between 1986 and 2017, the average statutory corporate income tax rate of OECD countries fell from 47.2 percent to 24.5 percent.[4] Of course, having an economic policy out of step with the rest of the world is hardly a sufficient argument for change; indeed, I imagine most supporters of the Tax Cuts and Jobs Act would not want to use this line of reasoning to justify, for example, US adoption of a national health insurance system or of a 15 percent value-added tax.

Third, US income inequality has increased steadily in the past three decades. By one calculation, the share of income received by the top 0.1 percent of earners has more than doubled from about 5 percent in the mid-1980s to over 10 percent by 2015 (Alvaredo, Atkinson, Piketty, Saez, and Zucman, No date).

Fourth, intangible capital such as research and development has gained in importance. While investment in intangible capital was about 80 percent of investment in tangible capital in the mid-1980s, by the early 2010s it was about 140 percent (Branstetter and Sichel 2017, figure 2). The income generated by intangible capital is easier to shift across countries, and spending on intangible capital was generally immediately deductible even before the Tax Cuts and Jobs Act.

Finally, the US fiscal picture looks very different. In 1986, federal debt held by the public amounted to 38.4 percent of GDP. In 2017, it stood at 76.5 percent,

[4]These figures are author's calculations, based on data in Historical Table II.1 in http://www.oecd.org/tax/tax-policy/tax-database.htm#C_CorporateCaptial and Table II.1 in http://stats.oecd.org/index.aspx?DataSetCode=TABLE_II1. The tax rates are unweighted averages and include both federal and subfederal rates.

in part because of the stimulus programs undertaken during the Great Recession, and partly because of the inexorable demographic changes that have been pushing promised entitlements for the elderly above the taxes in place to fund them.

## The Tax Cuts and Jobs Act

Tax reform was not a major preoccupation of the Obama administration, although some notable changes did occur. For example, tax increases accompanied the Patient Protection and Affordable Care Act of 2010, including a 3.8 percent surtax on investment income of households with over $250,000 ($200,000 for single filers) of income and a 0.9 percent Medicare payroll surtax on households with over $250,000 of wages and self-employment earnings ($200,000 for single filers). The American Taxpayer Relief Act of 2012 made many of the Bush-era tax cuts permanent, and created an additional tax bracket at the top with a marginal tax rate of 39.6 percent.

The Obama administration did issue reports in 2012 and then again in 2016 about business taxation (White House and US Department of the Treasury 2012, 2016). The 2016 report called for broadening the corporate tax base by eliminating dozens of tax expenditures, reducing the maximum corporate tax rate from 35 percent to 28 percent, eliminating the corporate alternative minimum tax, limiting interest paid deductions, and eliminating oil and gas tax preferences. All but the last-mentioned aspect of this Obama-era report were echoed in the Tax Cuts and Jobs Act of 2017. Also, the Obama-era report rejected the move to a pure territorial system, under which the income earned abroad by foreign subsidiaries of US multinational corporations is exempt from US tax,[5] and instead proposed a 19 percent minimum tax on foreign earnings when they are earned rather than—as was current law—upon repatriation, and a one-time tax on unrepatriated earnings.

Even before gaining control of Congress and the presidency in 2016, Republicans were formulating changes in the tax system. The "Path to Prosperity" plan issued by the House Budget Committee in March 2012 (and renewed in 2013) called for reducing the corporate tax rate to 25 percent, repealing the alternative minimum tax, and moving to a territorial system. In 2014, the then-chairman of the House Ways and Means Committee, Dave Camp (R-MI), released a revenue-neutral tax reform plan that would reduce the maximum corporate income tax rate to 25 percent and levy a top rate of 25 percent on the income of "pass-through" business entities (whose income is not subject to the corporate income tax but "passes through" to the owners who are taxed), but require slower methods and longer periods for depreciation and introduce a 95 percent exclusion of dividends

[5] In contrast, a pure worldwide system would subject the foreign earnings of US corporations to US tax. Before the Tax Cuts and Jobs Act of 2017, the US corporate system was a hybrid system under which, in principle, foreign earnings were subject to tax but were offset by a limited foreign tax credit and taxed only upon repatriation of the earnings.

from foreign subsidiaries plus a tax on their accumulated earnings; on the individual side, it would lower rates and reduce the number of brackets, fold the personal exemption into a higher standard deduction and child tax credit, and put a floor on deductible charitable contributions of 2 percent of adjusted gross income. In June 2016, a Tax Reform Task Force (2016) set up by the House Republicans issued a report that offered a "blueprint" for a new tax code that cut the top individual tax rates from 39.6 percent to 33 percent, abolished the individual alternative minimum tax, exempted half of interest receipts, and repealed the estate tax. On the business side, it provided full expensing (that is, immediate write-off of the costs of investment rather than depreciating over time) for tangible assets, capped the tax rate on pass-through entities at 25 percent, eliminated deductions for net interest expense, cut the corporate rate to a flat 20 percent and, notably, adopted a destination-based business tax system under which exports are exempted from the tax base and imports are taxed.

The tax legislative process started to take shape in July 2017, when key congressional and administration officials issued a one-page statement stressing tax "relief" but backing away from the controversial destination-based business tax plan. In September, a nine-page plan with more details was released. These documents provided a starting point for the Congressional debate. Eventually, a bill containing related but generally different provisions was passed in Congress on December 20 and signed into law by President Trump on December 22, 2017.

**Changes in the Corporate Income Tax**

Changes in the corporate income tax are the focal point of the Tax Cuts and Jobs Act. The issues here are discussed more fully in Auerbach's article in this symposium so I will not concentrate on corporate income tax changes; but they are too important to any discussion of the new tax law to ignore entirely.

Many nonrate aspects of the Tax Cuts and Jobs Act will affect decisions for certain firms: for example, the new law disallows the carryback of net operating losses (to reduce taxes paid in previous years), but allows for their indefinite carryforward (potentially to reduce taxes paid in future years). But I will focus on the central change, the cut in the corporate tax rate from 35 percent to 21 percent.[6] Two issues with reducing the corporate income tax rate are of primary importance: the effects on investment and on income-shifting.

The new tax law affects investment incentives in a variety of ways, but two avenues are especially relevant. The first is how it affects the cost of capital. Although normally a corporate income tax cut reduces the cost of capital, the effects of the Tax Cuts and Jobs Act rate cut are less clear. Lowering the corporate income tax rate always provides a lump-sum benefit to *past* investments, so some of the revenue loss does not map into improved incentives for future investment. In this case, with the enactment of expensing of most tangible capital expenditures (until the end of

[6] Prior to 2017, the corporate tax rate schedule was graduated, but the great majority of income was taxed at a 35 percent rate.

2022) and the continued expensing of intangible investments,[7] the rate cut does not as substantially reduce the cost of capital; with expensing, the government is in effect a silent partner to private investment, bearing the same fraction of costs that it takes in revenue, and the extent of its partnership (that is, the tax rate) doesn't matter for the cost of capital.[8] One attempt to quantify the impact of the Tax Cuts and Jobs Act on the cost of capital concluded that the 2017 law would reduce the cost of capital for equity-financed corporate (and noncorporate) investment, increase it for debt-financed investment, and decrease it overall; the estimated decrease was negligible for intellectual property (DeBacker and Kasher 2018).

Business tax cuts could also affect investment by increasing cash flow. This controversial notion has been around since Fazzari, Hubbard, and Petersen (1988) found empirical evidence suggesting that the investment of financially constrained firms appears to be more sensitive to cash flow, which is consistent with the notion that more cash flow relaxes financing constraints on investment. But this finding has not been consistently replicated. For example, Chen and Chen (2012, p. 394) conclude that "investment-cash flow sensitivity has completely disappeared in recent years," while other papers argue that cash flow is still a significant determinant of investment even when there is an exogenous shock to cash flow without a corresponding change in firm growth opportunities, as is arguably the case for the Tax Cuts and Jobs Act.

The statutory rate of corporate income tax also influences the direction and volume of taxable income shifting across national borders. With a 21 percent corporate tax rate, the incentive for outward shifting will decline, although 0 percent tax rate havens will still be attractive destinations for taxable income, and I expect to see some nonhaven countries react to the Tax Cuts and Jobs Act by lowering their corporate tax rate.

A shift from a worldwide toward a territorial corporate tax system by itself *increases* the reward to shifting taxable income abroad because no residual tax is due upon repatriation. However, the Tax Cuts and Jobs Act does not enact a pure territorial system (some have even called it a hybrid worldwide system), because it contains provisions designed to restrain income shifting. In addition, it contains a significant expansion of the base of cross-border income to which *current* US taxation would apply via the GILTI provision, which stands for "global intangible low-taxed income," that imposes a 10.5 percent minimum tax without deferral on profits earned abroad that exceed a firm's "normal" return.

The Tax Cuts and Jobs Act did eliminate US tax due upon the act of repatriating the future earnings of foreign subsidiaries of US multinational companies.

[7] Under the Tax Cuts and Jobs Act, after 2021 companies will not be able to write off research and development expenses immediately, but rather will be able to deduct them over a five-year period. Whether this provision will survive expected intense lobbying to overturn it is unclear.

[8] With a tax rate of $\tau$, expensing of all costs and full deductibility of losses, the present value of equity-financed investments is $\Sigma((1-\tau)R_t - (1-\tau)C_t)/(1+r)^t = (1-\tau)\Sigma(R_t - C_t)/(1+r)^t$. Those projects that have a positive present value (and thus are value enhancing) will still have it regardless of the value of $\tau$; those that do not, still do not regardless of $\tau$.

But what about the as-yet-unrepatriated past earnings of foreign subsidiaries, which by 2017 had grown to $2.6 trillion (ITEP 2017), in part because some companies were waiting for the next tax holiday—which has now arrived? The Tax Cuts and Jobs Act taxed them retroactively, at a rate of 15.5 percent rate on cash and 8 percent on noncash, payable in installments over eight years, whether they were repatriated or not. Some claimed that, by eliminating the tax cost to the repatriation of these funds, this would stimulate investment in the United States due to the parent companies' newly "available" funds, but this outcome is unlikely. The fact that companies had significant resources in the form of cash held by their foreign subsidiaries presumably had already made it much easier to obtain domestic loans when funds were needed, as lenders knew that, if necessary, the foreign profits could be used to repay the loans. Furthermore, on average, US companies are hardly cash-constrained, holding $1.84 trillion at the end of 2016 (Moody's Investor Service 2017). For empirical evidence, we can look back to the American Jobs Creation Act of 2004, which offered a temporary tax holiday for repatriations in 2004 and 2005, allowing funds to be brought back subject to a 5.25 percent tax rate. One careful study found that those companies that repatriated were ones with limited investment opportunities, and that much of the money went to share repurchases; there was no evidence of a stimulus to business investment, employment, or research and development expenditure (Blouin and Krull 2009; for a dissenting view, see Faulkender and Petersen 2012).

After the Tax Reform Act of 1986, as the top personal income tax rate fell below the corporate rate, there was a massive surge in businesses classifying themselves as pass-through entities—like partnerships or S corporations—whose income is not subject to the corporate income tax but instead becomes part of the individual taxable income for the owners. One might expect the reverse to occur starting in 2018, when the corporate income tax rate fell to well below the top personal tax rate (even considering the 20 percent deduction of income allowed to certain types of pass-through business income). In fact, the financial management firms KKR and Ares Management are two examples of companies that have already announced that they will convert from a partnership to a corporation.

Such business reclassifications have implications not only for the actual distribution of income, but also for how it is measured using tax return data. The reported taxable income of pass-through business entities appears on personal tax returns, but the income generated by businesses subject to the corporate income tax appears on personal returns only if and when it is paid out as dividends or shows up as capital gains. Thus, when, as after the 1986 tax reform, business income shifts to pass-through entities, inequality based on personal tax returns will seem to have risen when it has not really done so (Slemrod 1996). Now that the corporate rate has fallen substantially below the top personal rate, students of income distribution need to be on the lookout for the reverse phenomenon: some business income that pre-2018 was showing up on personal tax returns might "disappear," with only traces seen as dividends and capital gains. The more general issue here is that, whenever the definition of taxable income changes, there will be a mechanical change

in measured income and its distribution. This can matter a lot for some important issues. Auten and Splinter (2017) revisit the trends in US income inequality since 1960 adjusting for tax base changes, including the shifting of business income between C corporations and pass-through entities, and find that these adjustments reduce the increase in the top 1 percent's share of income since 1960 by as much as two-thirds.

**Changes in the Individual Income Tax**

The Tax Cuts and Jobs Act included cuts in individual tax rates (and changes in the bracket breakpoints), but the cuts were not nearly as big as in 1986. In 2017, the highest individual income tax rate fell from 39.6 percent to 37 percent; in 1986, the rate for the highest income earners fell from 50 percent to 28 percent over two years. In principle, the individual rate cuts in the Tax Cuts and Jobs Act expire after 2025, although almost no one thinks that will really happen. In spring 2018, there was already talk of a second tax bill that would extend those rate cuts past their scheduled expiration in 2025 (as reported in York 2018).

Just as the Tax Cuts and Jobs Act lowered individual rates a bit, it broadened the individual tax base a bit, although not as extensively as did the Tax Reform Act of 1986. Much of this happened indirectly, via the near doubling of the standard deduction, which reduces the value of claiming itemized deductions relative to claiming the standard deduction. One substantive and controversial base-broadener was a $10,000 cap on the itemized deduction of state and local taxes. This was one of the most progressive changes in the 2017 tax act; the Joint Committee on Taxation (2018, p. 8) estimated that over 45 percent of the lost tax benefit would have gone to taxpayers with more than $500,000 of annual income. To be sure, there are principled arguments for cutting back on this deduction—for example, that it bluntly subsidizes subfederal spending regardless of any externality involved—but one suspects that red-state, blue-state politics were also involved. After all, among states, the median of the deduction for state and local taxes taken as a percentage of adjusted gross income in 2014 was 4.5 percent, but it was 9.1 percent in New York, 8.7 percent in New Jersey, 8.3 percent in Connecticut, and 7.9 percent in California (Tax Foundation 2017).[9]

As in 1986, several end-of-year tax planning strategies became attractive in 2017 as the shape of the tax bill became clear. For example, those who likely will not itemize in 2018 had an incentive to move their deductible expenses such as charitable contributions up to 2017; similarly, some itemizers should have tried to accelerate payments of state and local taxes. Conversely, pass-through enterprises

[9] The states of Connecticut, New Jersey, and New York have filed a lawsuit challenging the cap. The state of New York has also enacted a "workaround" to the cap, creating a new optional payroll tax that shifts the state and local tax deduction from individuals who can no longer fully take it to businesses that can; the state will count these remittances toward state income taxes owed by the workers (for discussion, see Rubin and Vilensky 2018). Some states have also enacted laws to designate funds to solicit private contributions, not subject to the $10,000 cap, to support public services. Whether any of this will fly with the IRS is not yet known.

had an incentive to postpone income into 2018 when the new 20 percent deduction kicked in. In 1986, the anticipatory effects were in some cases stunning, especially in the case of capital gains, whose tax rate was scheduled to increase beginning on January 1, 1987. Burman, Clausing, and O'Hare (1994) determined that long-term capital gains on corporate stock realized and reported to tax authorities in December 1986 were seven times their level in 1985. The possibility of anticipatory behavior is important to bear in mind when the historical record for 2017 and 2018 is analyzed, because comparing behavior in 2017 to behavior in 2018 might reveal more about a short-term timing elasticity than a long-run elasticity, which will generally be smaller.

**Implications for Growth**

To the extent that economics played a role in the debate leading up to the Tax Cuts and Jobs Act, what mattered most were projections about its impact on the aggregate economy, in large part because this affected the estimate of the revenue loss from the legislation—the bigger the stimulus to economic activity, the smaller the revenue loss.

The Congressional Budget Office (2018, Box B-2, p. 117) summarizes and provides sources for ten different estimates of the impact of the Tax Cuts and Jobs Act. The official arbiter of the revenue cost, the Joint Committee on Taxation, put the "static"[10] ten-year revenue loss at $1.456 trillion from 2018 to 2027. It also estimated that the Tax Cuts and Jobs Act would increase the level of GDP by 0.8 percent on average over the ten-year budget window due to increases in labor supply and investment. These effects were estimated to reduce the revenue loss by $458 billion, or about one-third of the static revenue loss; the projected increase in interest rates would add another $51 billion to deficits. As the basis for these estimates, the Joint Committee on Taxation used a weighted average of the results of three models: the "macroeconomic equilibrium growth model" of its own staff, plus an overlapping generations and a dynamic stochastic general equilibrium model. As a comparison, the Congressional Budget Office put the offset due to macroeconomic effects on total deficits from 2018–2028 at 20 percent, reducing the sum of the deficits from $2.314 trillion to $1.854 trillion.

Other prominent groups came out with growth estimates on either side of those offered by the Joint Committee on Taxation. For example, the Penn Wharton model estimated that GDP by 2027 would be between 0.6 percent and 1.1 percent larger than otherwise, offsetting the static revenue estimate by between just 8 and 19 percent. On the high side, a model developed by the Tax Foundation forecast GDP rising by an average of 0.29 percent per year over the next decade, enough that the plan would decrease federal revenues by $1.47 trillion without considering any induced growth but by just $448 billion with estimated growth effects, so that the growth effect offset about 70 percent of the static revenue loss. (The much bigger estimated effect is apparently largely due to its assumption of an infinitely

[10] The descriptor "static" does not mean that these estimates ignore likely behavioral responses—they do not—just that they do not account for revenue feedback due to any induced change in aggregate growth.

elastic supply of capital to the US economy.) A letter to the Treasury Secretary Steven Mnuchin signed by nine prominent economists suggested a similarly large response, a gain in the long-run level of GDP of just over 3 percent, or 0.3 percent per year for a decade.[11]

It is important to keep in mind some limitations of these models. These estimates are not about welfare effects, because they do not allow for the value of the decreased amount of leisure or the value of the reduced consumption to finance saving. Moreover, they do not account for that part of the income that is owed to foreigners who lend funds to finance the excess of induced domestic investment over induced domestic saving. The significance of this factor is reflected in the Congressional Budget Office (2018) projection that real GDP (income generated in the United States) would be 0.7 percent higher over 2018–2027, but that real GNP (income of Americans) would rise by only 0.4 percent, implying that nearly half of the increased income generated within the United States would accrue to foreigners.

Finally, the growth effects of these models do not always take into account the effects of larger budget deficits. In their careful review of the empirical literature, Gale and Samwick (2017) conclude that the weight of the evidence suggests that deficit-financed tax cuts *reduce* national income in the long run, although if coupled with appropriate reforms, the overall mixture could expand the size of the economy.

**Implications for Simplicity**

On November 2, 2017, chairman of the House Ways and Means Committee Kevin Brady (R-TX) and Speaker of the House Paul Ryan (R-WI) said that, under the new tax system being contemplated, nine of ten taxpayers could file on a postcard. That same day, stock in the tax preparation company H&R Block dropped 2.8 percent to its lowest point in over half a year. The share price soon bounced back, but in June 2018, the company announced it was closing 400 locations. The rhetorical emphasis on a postcard-sized tax form has become quaint, now that more than 90 percent of tax returns are filed using software. In any event, a draft version of a new postcard-sized Form 1040 was unveiled in June 2018. The two-sided form is too large to qualify for the standard postcard rate and would reveal the filer's Social Security number unless enclosed in an envelope. It does eliminate more than half of the 78 line items on the previous form, but may require as many as six new worksheets not on the postcard (as reported in Tankersley 2018).

As with the Tax Reform Act of 1986,[12] a substantial simplification was achieved in the 2017 legislation by the near-doubling of the standard deduction, which is

[11] The letter is available at https://www.treasury.gov/press-center/press-releases/Documents/Economist_Letter_STM_11252017.pdf. A subsequent article coauthored by one of the nine signatories of this letter, Robert J. Barro, and a critic of the Tax Cuts and Jobs Act, Jason Furman, suggested that the growth rate stimulus would be between 0 and 0.2 percent per year (Barro and Furman 2018).

[12] The net simplification in the Tax Reform Act of 1986 was unclear. In this journal, in Slemrod (1992), I concluded that, despite a few scattered signs that tax-related financial planning declined, the compliance cost of the income tax system was probably higher circa 1992 than it was in the early 1980s, which suggests that the 1986 reform achieved little, if any, simplification in the income tax system.

predicted to reduce the fraction of itemizing returns from about 30 percent to about 12 percent (Tax Foundation 2017).[13] Undoubtedly, this change will reduce both administrative and compliance costs. To the extent that the itemized deductions should be subtracted from income to obtain a better measure of well-being, as is probably true for most medical expenses (as a counterexample, elective cosmetic surgery expenses are not deductible), this provision will also erode horizontal equity. Expanding the standard deduction also means that less than half as many people will receive the implicit subsidy to charitable giving, state and local government spending, mortgage borrowing, and the other expenses that deductibility provides. However, the big spenders among the itemizers will still itemize, so the weighted average increase in the cost of these activities will not be nearly as large as the drop in the fraction of itemizing households suggests. Eliminating the corporate alternative minimum tax is certainly a simplification, as is substantially limiting the scope of the individual alternative minimum tax so that it affects many fewer taxpayers. The Joint Committee on Taxation (2018, p. 11) estimated that the number of taxpayers subject to the individual alternative minimum tax will fall from over four million to approximately 600,000.

The Tax Cuts and Jobs Act also adds some nontrivial complications, especially in certain provisions affecting choice of corporate form and the attempts to block income shifting of corporate taxes. Enforcing the new rules that allow a deduction of up to 20 percent of the income of pass-through entities subject to myriad limits and restrictions will be very tough.[14] One newspaper article described the "crack and pack" strategies that businesses are exploring to get around the provision excluding high-income lawyers, doctors, and other professionals from this deduction (reported in Simon and Rubin 2018). "Crack and pack" refers to splitting operations apart, reclassifying, and re-categorizing their operations. One lawyer quoted in the article spoke of splitting his law firm into one entity holding four lawyers and the other holding the 26-person administrative staff, presumably in the hope that the latter could be exempt from the statutory limitation. Some business owners are considering giving shares of the business to family members, each of whom files a tax return that falls below the income threshold. The cat-and-mouse game between the IRS issuing regulations and new ways around them has just begun (for further discussion of the gaming opportunities opened up by the Tax Cuts and Jobs Act, see Kamin et al. forthcoming). An even bigger problem might be monitoring the larger incentives now built in to reclassify labor income as

[13] As with the Camp proposal mentioned above, the doubling of the standard deduction was coupled with the repeal of personal and dependent exemptions, which are contingent on family size. The pro-family nature of the income tax system was largely retained by an expansion of the existing child credit and the introduction of a new credit for any dependent, including children who are too old to be eligible for the child tax credit.

[14] The effectiveness of limiting the pass-through deduction has implications for distribution, as well; the Joint Committee on Taxation (2018, p. 4) estimates that more than half of the benefit of this provision will go to taxpayers with income in excess of $500,000 per year.

business income subject to lower tax rates.[15] In addition, the new rules designed to limit income shifting in the new world of a modified territorial system, like the "base erosion and anti-abuse tax" (BEAT), a sort of limited scope alternative minimum tax, and the GILTI mentioned earlier, are reputed to be labyrinthine.

**What the Tax Cuts and Jobs Act Didn't Do**

Given that comprehensive tax reform comes along rarely in the United States, it behooves us to consider what the Tax Cuts and Jobs Act might have accomplished but did not.

I begin with base broadening. Although the mantra of tax reformers has long been "broaden the base, lower the rates" on the grounds that deviations will cause inefficient distortions of resources, this argument is not airtight: the broadest base is not necessarily the best. Optimal commodity tax models show that a uniform tax on all goods and services (equivalent in a one-period model to an income tax with no deductions for particular expenditures) minimizes efficiency cost only in very special circumstances—implicit separability between leisure and goods.[16] Otherwise, the optimal tax pattern should take advantage of commodities' relative substitutability or complementarity with leisure: a complement to leisure such as skis should be taxed relatively heavily and a substitute for leisure such as work uniforms should be taxed relatively lightly or even subsidized. It might also be optimal to single out particular activities for special tax treatment if they generate externalities. However, allowing an income tax deduction for an expenditure provides a subsidy equal to the taxpayer's marginal income tax rate, which is only by chance equal to the marginal social benefit, especially given that the implicit rate of subsidy is higher for those with higher marginal tax rates. Finally, income net of, for example, medical expenses almost certainly provides a better measure of well-being than income not subtracting them, and so *not* allowing this deduction would violate horizontal equity.

These arguments imply that base broadening proposals must be considered on a case-by-case basis. For example, a deduction for state and local taxes, which was scaled back by the Tax Cuts and Jobs Act, could arguably be justified on the grounds that some state and local spending provides externalities beyond the taxing jurisdiction's borders. However, because much state and local spending provides no such externality, an untargeted deduction seems like a blunt instrument for this purpose.

The 2017 tax act did not make a big dent in some of the biggest tax expenditures—spending programs that operate through the tax system—with the tax preference to employer-provided health insurance being the biggest. In fact, the Tax Cuts and Jobs Act postponed from 2020 to 2022 the "Cadillac Tax" of the Affordable Care Act that would levy a 40 percent tax on the costs of health plans

[15] This issue is faced by the dual income tax systems used in Nordic countries, under which labor income is generally taxed at a higher rate than capital and business income (Sørensen 2010).

[16] An early and elegant introduction to this argument is provided in Sandmo (1976), and Kaplow (2008) provides a comprehensive treatment.

that exceed $10,200 per individual or $27,500 for family coverage. Similarly, it did not directly address, other than via the less-widespread itemization that limits both property tax and mortgage interest deductions, the substantial income tax preference for owner-occupied housing arising from the complete exemption of the return (implicit rent) the asset provides. The cap on the home value eligible for the mortgage interest deduction was tightened up a bit, from $1 million to $750,000 (for first and second homes), which reduces the subsidy for some high-value debt-financed properties, especially those more likely to be located in blue states. In addition, the Tax Cuts and Jobs Act suspends until 2026 the deduction for interest paid on home equity loans and lines of credit unless they are used to buy, build, or substantially improve the taxpayer's home that secures the loan.

It did not increase the federal gasoline tax, stuck at 18.4 cents per gallon since 1993, or implement some form of carbon tax. Such changes have been favored even by many generally tax-phobic conservative economists and groups, such as the Chamber of Commerce, as a way to fund road repair and to address climate damage (as discussed in Mankiw 2006). A gas tax hike could be a potent source of revenue; every one-cent per gallon increase would raise more than $1.3 billion of revenue per year (Congressional Budget Office 2016, p. 198).

It did not abolish the estate tax, as many Republicans favored and most of their precursor proposals featured. However, it did (through 2025) approximately double the exemption to $11.2 million per individual and $22.4 million per married couple, reducing the number of estates that will be subject to the tax by about two-thirds, from about 5,500 to 1,700 (Tax Policy Center 2017a).

It did not eliminate the tax treatment of carried interest, so that investment fund managers will continue to be subject to the lower capital gains tax rate on their share of the fund's qualified dividends and long-term capital gains. The Obama administration favored this change, and Trump expressed support for it during the presidential campaign, but it did not make it into the Tax Cuts and Jobs Act. However, the Tax Cuts and Jobs Act did increase from one to three years the minimum holding period for carried interest profits to qualify for the capital gains tax rate.

Overall, the Tax Cuts and Jobs Act did not move to a radically different tax system. Some of those economists who gave the Tax Reform Act of 1986 a low grade did so because they would have preferred moving to a consumption tax (for example, Shoven 1990). The US Department of the Treasury (1984) report had included a separate volume on the value-added tax, and concluded that its advantages did not justify its incremental administrative cost. In 2017, this idea didn't get very far, although two Republican presidential candidates (Ted Cruz and Rand Paul) supported a value-added tax, albeit by other names; Cruz called it a "business flat tax," and Paul called it a "business activity tax." Of special interest to tax cognoscenti was the proposal made by Republican presidential candidate Marco Rubio, who favored a version of David Bradford's X Tax, a modification of the Hall–Rabushka flat tax that levies graduated rates of 15 percent, 25 percent, and 35 percent on labor income, while maintaining a flat (25 percent) rate on business receipts minus expensed costs of doing business (as reported in Pomerleau 2016).

For a few months, there was discussion of replacing the corporate income tax with a destination-based cash flow tax, the base of which is revenues minus costs and payments to labor. In essence, it is a value-added tax plus a subsidy to labor income at the same rate; as with most value-added taxes, imports are subject to tax while exports are not, providing the destination basis. (Alas, the DBCFT has neither a mellifluous acronym like VAT nor a super-heroic moniker like the X Tax.) For more detail and arguments in favor of the destination-based cash-flow tax, see Auerbach, Devereux, Keen, and Vella (2017). It did not become part of the Tax Cuts and Jobs Act.

## What the Tax Cuts and Jobs Act Got Wrong

I've argued that the Tax Cuts and Jobs Act featured a few of the same kinds of base broadeners and simplification features that tax reform advocates supported in the Tax Reform Act of 1986 and moved the corporate tax system in some directions that are widely supported, and were in fact advocated by the Obama administration. But in two fundamental aspects it pushed the tax system in what I believe is the wrong direction.

### Deficit

Although the Tax Reform Act of 1986 was designed to be revenue neutral,[17] the Tax Cuts and Jobs act was not—not even close. The Joint Committee on Taxation (2017) estimated that it will reduce revenue by $1.456 trillion over the 10 years 2018 through 2027, not including debt service costs or macroeconomic feedback. The Congressional Budget Office (2018, pp. 128–129) estimated the reduction in federal revenue over 2018–2028 at $1.854 trillion, including both macroeconomic effects and the effect on debt-service costs. This is clearly the wrong direction. Deficit spending is not needed at present for short-term stimulus purposes. The increased debt reduces budget flexibility for the next recession and worsens the massive long-term fiscal imbalance already faced by the United States.

There were political concerns over the deficit in 1984–86, too: a big tax cut in 1981 followed by smaller tax increases in 1982 and 1984 had deficit hawks concerned. After all, the ratio of publicly held debt-to-GDP ratio had grown from 25.2 percent in 1981 to what seemed at the time to be a scary 38.4 percent in 1986. This ratio is now twice as high. In the mid-1980s, the leading alternative to maintaining revenue neutrality in the gestation of the Tax Reform Act of 1986 was to *increase* revenue.

Of course, the reductions in tax revenue from the Tax Cuts and Jobs Act adds to people's paychecks and businesses' cash flows. However, borrowing is not a way to raise resources for the government; instead, it puts off assigning the tax burden and shifts it to future generations. The political benefit of cutting taxes is obvious, as the

[17] Whether it succeeded is another question; for discussion, see Wallace, Wasylenko, and Weiner (1991).

burden passed on to future generations is abstract and uncertain. The Tax Reform Act of 1986 leveraged a similar feature to a much lesser degree, because it reduced individual income taxes overall, and most popular accounts of winners and losers ignored that the increases in corporation tax would be borne by *some* people, even if it was difficult to say exactly which ones.

Higher deficits will tend to increase interest rates and crowd out private investment to an extent that depends on how much the higher interest rates induce greater private saving and inflows of foreign capital. Reviewing the empirical evidence, the Congressional Budget Office's central estimate is that each dollar of increase in the budget deficit reduces domestic investment by 33 cents (Huntley 2014).

This issue was noted, after the passage of the Tax Cuts and Jobs Act, by some prominent economists, including some who lauded the growth-enhancing properties of the Tax Cuts and Jobs Act. Martin Feldstein (2018) took to the op-ed pages of the *Wall Street Journal* to call the exploding debt and deficit the federal government's most urgent domestic challenge, and recommended slowing the growth of Social Security and Medicare. Feldstein's argument raises the "starve-the-beast" theory of deficit-financed tax cuts: that they are a tactic to force down government spending. However, the historical record suggests that this tactic does not work (Romer and Romer 2009). Laurence Kotlikoff (2017), who had argued that the Tax Cuts and Jobs Act would ultimately raise GDP by 5 percent and be approximately revenue neutral, also argued that the United States didn't need a tax bill that was revenue neutral at best, but rather needed a tax bill that raised federal revenues dramatically—by 60 percent.

**Distribution**

While the Tax Reform Act of 1986 was designed from the start to be distributionally neutral, no such constraint was put on the Tax Cuts and Jobs Act of 2017. Given that the 2017 tax act was not revenue neutral, what distributional neutrality even means is not clear. Here is some relevant information. Using conventional incidence assumptions, the Tax Policy Center (2017b) calculated that the top 1 percent of income earners on average gets a tax cut in 2018 of $51,140, the average member of the middle quintile receives $930, and the lowest quintile gets just $60 on average; these dollar figures amount to a 3.4, 1.6, and 0.4 percent increase in after-tax income, respectively. Overall, more than half of the benefits go to the top 10 percent of earners, so the tax cuts accrue disproportionately to high-income households. But high-income households are liable for a disproportionate fraction of tax liability, so how the current federal tax burden is allocated across income groups does not change much, at least in the short run before various provisions begin to sunset.[18] Nevertheless, put in the context of the sustained increase in inequality in the last few decades, this massive tax cut for the richest members of society strikes me as the wrong direction for policy. Moreover, the distributional consequences over the long

[18] This can be inferred from the calendar year 2019 table in Joint Committee on Taxation (2017, p. 1); note that this starts to change as early as calendar year 2021 (p. 2).

run depend on how the deficits are financed; if, for example, they are financed by cuts in entitlement programs such as Medicare and Medicaid, the balanced-budget consequences are massively regressive (Gale, Gelfond, Krupkin, Mazur, and Toder 2018, pp. 23–26).

The distributional impact of the new tax legislation depends crucially on what one believes about the incidence of the corporate tax rate cut. The Tax Policy Center (2017b) incidence calculations assume that 60 percent of the benefit of a corporation income tax rate cut ends up benefiting owners of corporate shares, and only 20 percent goes to workers (with the other 20 percent going to capital owners generally). Both the Congressional Budget Office and Treasury's Office of Tax Analysis come to similar conclusions. Of course, just because these highly respected government offices assert something doesn't make it correct; the more the corporate tax rate cut would ultimately benefit workers rather than share owners or capital owners, the more progressive will be the ultimate distributional impact of the Tax Cuts and Jobs Act. Among academic tax economists, this issue is *very* controversial.

More than a half century ago, Harberger (1962) developed and parameterized a general equilibrium, closed-economy model of this question and concluded that over a horizon he called the long run (although the capital stock was assumed to be fixed), the burden of the tax would be borne by capital owners in general, and not passed along to people in their role as workers or consumers. The more open an economy, the greater the share of burden from a source-based tax on capital that will fall on immobile domestic factors such as land and labor; in the extreme case of a small, open economy facing an elastic supply of capital, none of the burden of such a tax will be borne by capital owners. If profits are partly due to location-specific rents, the tax burden may fall in part on the claimants to these rents, although if the rents are location-specific and firms are mobile, the burden can be shifted to the immobile factors (early models are described in Kotlikoff and Summers 1987; see also Auerbach 2006). Over the very long run, to the extent that a corporate tax cut induces capital accumulation, which in turn raises the productivity of workers, a competitive labor market will generate real wage increases. The distributional impact will depend on which groups of workers, skilled versus unskilled, see their wages rise, which in turn depends on their relative complementarity with the accumulated capital.

A more recent literature revisits the question of who bears, or benefits from, changes in the corporate tax rate in static models without capital accumulation where wage-setting institutions and labor market frictions matter (for a skeptical review of this literature, see Gravelle 2017). For example, in wage bargaining models, firm owners and workers share any surplus generated by the firm and, if corporate tax cuts increase that surplus, the models predict that some of that increase will be shared with employees through higher compensation. A similar result obtains in some efficiency wage models. Recent empirical analyses with this kind of model in mind have found substantial pass-through to workers, as much as one-half in the Arulampalam, Devereux, and Maffini (2012) study of federal corporate taxation in nine European countries and in the Fuest, Peichl, and Siegloch (2018) study of

German municipal business taxes. These effects are certainly context-dependent, however, and their application to US federal corporate taxation has not been demonstrated. For example, Azémar and Hubbard (2015), who conclude based on an analysis of 13 OECD countries that on average 60 percent of corporate taxes are passed to labor, find that the pass-through is more than ten times as high in countries with the highest union density compared to the average union density; the pass-through would presumably be even lower, and perhaps negligible, for the largely nonunion United States.

The unsettled question of the incidence of the corporation income tax means that the overall distributional impact of the Tax Cuts and Jobs Act is also uncertain. But to assert that it will largely benefit workers is, in my opinion, a stretch that the empirical literature does not substantiate.

How the tax burden *should* be assigned to people of different levels of well-being is a delicate problem because it inescapably involves interpersonal comparisons of utility. In the seminal modern paper of optimal taxation theory in this area, Mirrlees (1971) investigated how a government seeking to maximize a utilitarian social welfare function in a society of people with heterogeneous ability to earn income should choose the rate structure, and therefore the progressivity, of an income tax that must raise some given amount of total revenue. In this class of models (an early example is Helpman and Sadka 1978), increased dispersion in the distribution of earning power increases optimal progressivity; thus, the large increase in pre-tax income inequality in the past three or four decades would have, ceteris paribus, increased optimal progressivity. However, these models cannot prescribe the level of optimal progressivity without additional value-laden assumptions about society's willingness to trade off equity and efficiency.

## A Teachable Moment

Whether or not the Tax Cuts and Jobs Act is good for the US economy and its population, it is clearly good for those of us who study taxation. Indeed, it is a wonderfully generous gift because it provides scores of natural experiments that could help provide credible estimates of the causal effects of tax policy. Indeed, the recent move away from studying natural experiments in taxation, toward other research designs such as bunching analysis and randomized controlled trials may have been partly caused by the dearth of recent comprehensive tax reform episodes.

It is, of course, far too early to assess the full economic impact of the Tax Cuts and Jobs Act, but even within the first few years one can learn quite a bit (for a set of studies a few years after the passage of the Tax Reform Act of 1986, see Slemrod 1990). Two aspects of the immediate aftermath of its enactment are of particular interest, however. One is its degree of popularity. Between its enactment in December 2017 and mid-April 2018, over 40 polls were taken by major polling organizations. In that period, the fraction approving of the Tax Cuts and Jobs Act never exceeded the fraction disapproving by more than a sliver. The excess of the

percent disapproving over approving peaked at the time of passage at just over 21 points, fell to be about even in mid-February as paychecks swelled due to reduced income tax withholding by employers, but then started to diverge, with disapprovers exceeding approvers by almost 8 percentage points by the end of April (as summarized at Real Clear Politics, "Trump, Republicans' Tax Reform Law," https://www.realclearpolitics.com/epolls/other/trump_republicans_tax_reform_law-6446.html). However, the Tax Reform Act of 1986 was also viewed less favorably over time: in 1990 compared to 1986, on average people were less inclined to think it had a positive effect on the economy, was fairer, less complicated, and provided tax reduction (Auerbach and Slemrod 1997, table 6, p. 618).

Second, because the Tax Cuts and Jobs Act was much more partisan than the Tax Reform Act of 1986—no Democrats in either House of Congress voted for the 2017 legislation—its perceived success is much more tied to partisan political bragging rights. This might partly explain why, in the immediate aftermath of its passage, and in some cases even before it had become law, hundreds of companies announced actions they say were induced by the new law—raising wages, providing bonuses, and announcing hiring and investment. Some have viewed these disclosures as attempts to curry favor with the Trump administration and Republican Party or as trying to boost the public image of tax reform so that, in the event Democrats regained control of the House or Senate, they would not try to dismantle key pieces of the tax legislation (as Republicans did with the Patient Protection and Affordable Care Act of 2010). Analyzing these announcements, in Hanlon, Hoopes, and Slemrod (forthcoming), we conclude they are systematically related both to economic factors such as the size of the tax cut a company received and also to political factors such as whether the company is in a politically sensitive sector.

Looking ahead to the future research agenda, several aspects of the process leading to the new law render it an especially promising laboratory. The potential endogeneity of timing that has so concerned macroeconomists studying fiscal policy (for example, see the discussion in Romer and Romer 2010) is not an obvious problem in this case, as there was nothing about the state of the business cycle that precipitated the passage of the Tax Cuts and Jobs Act—instead, it passed because a close election tipped in one direction. The macro economy being well-behaved in the few years leading up to the Tax Cuts and Jobs Act also helps, because plausibly the post-tax-reform counterfactual is also fairly placid (setting aside the disruptions caused by, for example, a trade war).

On the other side, the extensiveness of the Tax Cuts and Jobs Act is a mixed bag for research purpose. It changes scores of relative prices and rewards, as did the Tax Reform Act of 1986, making identification of the impact of any one aspect more difficult. For instance, how can the effect of moving to a territorial corporate tax system be separated from the corporate tax rate cut or the anti-income-shifting measures? The Tax Cuts and Jobs Act also does not provide a natural difference-in-differences design of the impact of individual marginal tax rates as did the Tax Reform Act of 1986, which cut the top rate substantially, from 50 percent to 28 percent, and other rates much less. Feldstein (1995), the seminal difference-in-differences estimate of

the elasticity of taxable income, was based on the much larger tax rate cut for the highest-income group compared to everyone else.

Reference to the elasticity of taxable income literature reminds us that it is not only the world and the tax system that have changed in the past 30 years—the economics of taxation has, too. The tax research agenda beginning now should absorb what we have learned in the past 30-plus years, both in theoretical perspective and empirical methods. In the past two decades, tax economists have focused on the concept of the elasticity of taxable income (ETI), the idea that, under some assumptions, all the behavioral responses to taxation—real decisions such as labor supply as well as avoidance and evasion—are symptoms of inefficiency, and the anatomy of the total response does not matter for that issue (although it might matter for evaluating further policy changes). This is not just a matter of estimating compensated ETIs instead of labor supply elasticities. Since Slemrod and Kopczuk (2002), we have also realized that the ETI is not a structural parameter, but rather depends on aspects of the tax system such as how broad the tax base is and how it is enforced. For example, the response of wealthy Americans to an increase in the tax rate on capital income arguably depends on the effectiveness of the IRS in monitoring evasive foreign accounts.

## Is It Tax Reform, or Just Confusion?

Neither choice is accurate, in my opinion.

The Tax Cuts and Jobs Act is not tax reform, at least not in the traditional sense of broadening the tax base and using the revenue so obtained to lower the rates applied to the new base. Nor, based on its unofficial title, did it aspire to this approach as a main objective. It does, though, contain several base-broadening features long favored by tax reform advocates.

Nor is the Tax Cuts and Jobs Act just confusion. There are coherent arguments buttressing the centerpiece cut in the corporation tax rate. To the extent that the new legislation reduces the cost of capital (which is not obvious), business investment will be higher than otherwise.

Its serious downsides are the contribution to deficits and to inequality. The former is less of a concern to the extent that the Tax Cuts and Jobs Act turns out to stimulate growth; the latter is less of an issue the more its centerpiece cuts in business taxation will be shifted to the benefit of workers, especially low-income workers. In both cases, the Tax Cuts and Jobs Act represents a huge gamble on the magnitude of these effects, about which the evidence is not at all clear. My own view is that the stimulus to growth will be modest, far short of many supporters' claims, and so the Tax Cuts and Jobs Act will increase federal deficits by nearly $2 trillion over the next decade, a nontrivial stride in the wrong direction that promises to shift the tax burden to future generations. How it will affect the within-generation distribution of welfare is the most controversial question of all. Although according to conventional wisdom, the Tax Cuts and Jobs Act delivers the bulk of the tax cuts

to the richest Americans, whose relative well-being has been rising continuously in recent decades, other plausible models of the economy, supported by some new empirical evidence, raise the possibility that the gains will be more widely shared. This is the most important question about which we know too little.

■ *I have benefited from conversations with Bill Gale, Itai Grinberg, and Jim Hines. Able research assistance was provided by Kendra Robbins. Helpful comments on an earlier draft were received from the JEP editorial team, Alan Auerbach, Tom Barthold, Lily Batchelder, Len Burman, Kim Clausing, Naomi Feldman, Bill Gale, Jane Gravelle, Michelle Hanlon, Jim Hines, Jeff Hoopes, Jay Mackie, Peter Merrill, Tom Neubig, Richard Rubin, Matthew Shapiro, Gene Steuerle, Eric Toder, Tejaswi Velayudhan, and Alan Viard.*

## References

**Alvaredo, Facundo, Anthony B. Atkinson, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman.** No date. "The World Wealth and Income Database." http://www.wid.world.

**Arulampalam, Wiji, Michael P. Devereux, and Giorgia Maffini.** 2012. "The Direct Incidence of Corporate Income Tax on Wages." *European Economic Review* 56(6): 1038–54.

**Auerbach, Alan J.** 2006. "Who Bears the Corporate Tax? A Review of What We Know." In *Tax Policy and the Economy*, vol. 20, edited by James M. Poterba, pp. 1–40. MIT Press.

**Auerbach, Alan J., Michael P. Devereux, Michael Keen, and John Vella.** 2017. "Destination-Based Cash Flow Taxation." Available at SSRN: https://ssrn.com/abstract=2908158.

**Auerbach, Alan J., and Joel Slemrod.** 1997. "The Economic Effects of the Tax Reform Act of 1986." *Journal of Economic Literature* 35(2): 589–632.

**Auten, Gerald, and David Splinter.** 2017. "Income Inequality in the United States: Using Tax Data to Measure Long-Term Trends." Draft subject to change. http://davidsplinter.com/AutenSplinter-Tax_Data_and_Inequality.pdf.

**Azémar, Céline, and R. Glenn Hubbard.** 2015. "Country Characteristics and the Incidence of Capital Income Taxation on Wages: An Empirical Assessment." *Canadian Journal of Economics/Revue canadienne d'économique* 48(5): 1762–1802.

**Bach, Natasha.** 2017. "Trump Thinks the U.S. Could See 6% Economic Growth. The Data Says Otherwise." *Fortune*, December 7.

**Barro, Robert J., and Jason Furman.** 2018. "The Macroeconomic Effects of the 2017 Tax Reform." Conference draft presented at the Brookings Papers on Economic Activity meeting, Washington, DC, March 8–9. https://www.brookings.edu/bpea-articles/macroeconomic-effects-of-the-2017-tax-reform/.

**Birnbaum, Jeffrey H., and Alan S. Murray.** 1987. *Showdown at Gucci Gulch: Lawmakers, Lobbyists, and the Unlikely Triumph of Tax Reform*. Random House.

**Blinder, Alan S.** 2017. "Almost Everything Is Wrong with the New Tax Law." *Wall Street Journal*, December 28.

**Blouin, Jennifer, and Linda Krull.** 2009. "Bringing It Home: A Study of the Incentives Surrounding the Repatriation of Foreign Earnings under the American Jobs Creation Act of 2004." *Journal of Accounting Research* 47(4): 1027–59.

**Branstetter, Lee G., and Daniel Sichel.** 2017. "The Case for an American Productivity Revival." Peterson Institute for International Economics Policy Brief 17-26.

**Burman, Leonard E., Kimberly A. Clausing, and John F. O'Hare.** 1994. "Tax Reform and Realizations of Capital Gains in 1986." *National Tax Journal* 47(1): 1–18.

**Carter, Jimmy.** 1976. "Our Nation's Past and Future: Address Accepting the Presidential Nomination at the Democratic National Convention in

New York City." The American Presidency Project, July 15. http://www.presidency.ucsb.edu/ws/index.php?pid=25953.

**Chen, Huafeng, and Shaojun Chen.** 2012. "Investment–Cash Flow Sensitivity Cannot Be a Good Measure of Financial Constraints: Evidence from the Time Series." *Journal of Financial Economics* 103(2): 393–410.

**CNBC.** 2017. "An Open Letter to Congress Signed by 137 Economists Supporting GOP Tax Reform Bill." November 29. https://www.cnbc.com/2017/11/29/137-economists-support-gop-tax-reform-bill-in-open-letter-commentary.html.

**Congressional Budget Office.** 2016. *Options for Reducing the Deficit: 2017 to 2026*. December.

**Congressional Budget Office.** 2018. *The Budget and Economic Outlook: 2018 to 2028*. April 9.

**Council of Economic Advisers.** 2017. *The Growth Effects of Corporate Tax Reform and Implications for Wages*. October.

**DeBacker, Jason, and Roy Kasher.** 2018. "Effective Tax Rates on Business Investment Under the Tax Cuts and Jobs Act." *AEI Economic Perspectives*, American Enterprise Institute, May.

**Faulkender, Michael, and Mitchell Petersen.** 2012. "Investment and Capital Constraints: Repatriations under the American Jobs Creation Act." *Review of Financial Studies* 25(11): 3351–88.

**Fazzari, Steven M., R. Glenn Hubbard, and Bruce C. Petersen.** 1988. "Financing Constraints and Corporate Investment." *Brookings Papers on Economic Activity,* no. 1, pp. 141–95.

**Feldstein, Martin.** 1995. "The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act." *Journal of Political Economy* 103(3): 551–72.

**Feldstein, Martin.** 2018. "Reagan's Cure for America's Debt Disease." *Wall Street Journal*, March 5.

**Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch.** 2018. "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany." *American Economic Review* 108(2): 393–418.

**Gale, William G., Hilary Gelfond, Aaron Krupkin, Mark J. Mazur, and Eric Toder.** 2018. "Effects of the Tax Cuts and Jobs Act: A Preliminary Analysis." June 13. Tax Policy Center, Urban Institute & Brookings Institution. https://www.brookings.edu/wp-content/uploads/2018/06/ES_20180608_tcja_summary_paper_final.pdf.

**Gale, William G., and Andrew A. Samwick.** 2017. "Effects of Income Tax Changes on Economic Growth." Chap. 2 in *The Economics of Tax Policy*, edited by Alan J. Auerbach and Kent Smetters. Oxford University Press.

**Gravelle, Jane G.** 2017. "Corporate Tax Reform: Issues for Congress." Congressional Research Service, September 22.

**Hanlon, Michelle, Jeffrey Hoopes, and Joel Slemrod.** Forthcoming. "Tax Reform Made Me Do It!" In *Tax Policy and the Economy*, vol. 33, edited by Robert A. Moffitt. University of Chicago Press.

**Harberger, Arnold C.** 1962. "The Incidence of the Corporation Income Tax." *Journal of Political Economy* 70(3): 215–40.

**Helpman, Elhanan, and Efraim Sadka.** 1978. "The Optimal Income Tax: Some Comparative Statics Results." *Journal of Public Economics* 9(3): 383–93.

**Huntley, Jonathan.** 2014. "The Long-Run Effects of Federal Budget Deficits on National Saving and Private Domestic Investment." Congressional Budget Office Working Paper 2014-02.

**ITEP.** 2017. "Fortune 500 Companies Hold a Record $2.6 Trillion Offshore." March. https://itep.org/wp-content/uploads/pre0327.pdf.

**Joint Committee on Taxation.** 2017. *Distributional Effects of the Conference Agreement for H.R. 1, The "Tax Cuts and Jobs Act."* JCX-68-17, December 18.

**Joint Committee on Taxation.** 2018. *Overview of the Federal Tax System as in Effect for 2018*. JCX-3-18, February 7.

**Jorgenson, Dale W., and Kun-Young Yun.** 1990. "Tax Reform and US Economic Growth." *Journal of Political Economy* 98(5, Part 2): S151–S193.

**Kamin, David, David Gamage, Ari Glogower, Rebecca Kysar, Darien Shanske, Reuven Avi-Yonah, Lily Batchelder, J. Clifton Fleming, Daniel Hemel, Mitchell Kane, David Miller, Daniel Shaviro, and Manoj Viswanathan.** Forthcoming. "The Games They Will Play: Tax Games, Roadblocks, and Glitches Under the 2017 Tax Legislation." *Minnesota Law Review*.

**Kaplow, Louis.** 2008. *The Theory of Taxation and Public Economics*. Princeton University Press.

**Kotlikoff, Laurence.** 2017. "The US Is Dead Broke and the Tax Bill Does Nothing to Help Matters." *The Hill*, December 21.

**Kotlikoff, Laurence J., and Lawrence H. Summers.** 1987. "Tax Incidence." Chap. 16 in *Handbook of Public Economics*, vol. 2, edited by Alan J. Auerbach and Martin Feldstein. Elsevier.

**Mankiw, N. Gregory.** 2006. "Raise the Gas Tax." *Wall Street Journal*, October 12, p. A12.

**McLure, Charles E., Jr., and George R. Zodrow.** 1987. "Treasury I and the Tax Reform Act of 1986: The Economics and Politics of Tax Reform." *Journal of Economic Perspectives* 1(1): 37–58.

**Mirrlees, James A.** 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38(2): 175–208.

**Moody's Investor Service.** 2017.

"Announcement: Moody's: US Corporate Cash Pile Grows to $1.84 Trillion, Led by Tech Sector." July 19. https://www.moodys.com/research/Moodys-US-corporate-cash-pile-grows-to-184-trillion-led–PR_369922.

**Pomerleau, Kyle.** 2016. "Why Marco Rubio's Tax Plan Might Stick Around." March 16. Tax Foundation. https://taxfoundation.org/why-marco-rubios-tax-plan-might-stick-around/.

**Romer, Christina D., and David H. Romer.** 2009. "Do Tax Cuts Starve the Beast? The Effect of Tax Changes on Government Spending." *Brookings Papers on Economic Activity,* no. 1, pp. 139–200.

**Romer, Christina D., and David H. Romer.** 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review* 100(3): 763–801.

**Rubin, Richard, and Mike Vilensky.** 2018. "New York Found a Fix for Some Hit by Tax Law. Employers Are Skeptical." *Wall Street Journal*, May 2.

**Sandmo, Agnar.** 1976. "Optimal Taxation: An Introduction to the Literature." *Journal of Public Economics* 6(1–2): 37–54.

**Shoven, John B.** 1990. "The U.S. Tax Reform of 1986: Is It Worth Copying?" Chap. 10 in *World Tax Reform: Case Studies of Developed and Developing Countries*, edited by Michael J. Boskin and Charles E. McLure, Jr. ICS Press.

**Simon, Ruth, and Richard Rubin.** 2018. "Crack and Pack: How Companies Are Mastering the New Tax Code." *Wall Street Journal*, April 3.

**Slemrod, Joel, ed.** 1990. *Do Taxes Matter? The Economic Consequences of the Tax Reform Act of 1986.* Cambridge and London: MIT Press.

**Slemrod, Joel.** 1992. "Did the Tax Reform Act of 1986 Simplify Tax Matters?" *Journal of Economic Perspectives* 6(1): 45-57.

**Slemrod, Joel.** 1996. "High-Income Families and the Tax Changes of the 1980s: The Anatomy of Behavioral Response." Chap. 6 in *Empirical Foundations of Household Taxation*, edited by Martin Feldstein and James M. Poterba. University of Chicago Press and NBER.

**Slemrod, Joel, and Wojciech Kopczuk.** 2002. "The Optimal Elasticity of Taxable Income." *Journal of Public Economics* 84(1): 91–112.

**Sørensen, Peter Birch.** 2010. "Dual Income Taxes: A Nordic Tax System." Chap. 5 in *Tax Reform in Open Economies: International and Country Perspectives,* edited by Iris Claus, Norman Gemmell, Michelle Harding, and David White. Edward Elgar.

**Steuerle, C. Eugene.** 2008. *Contemporary U.S. Tax Policy.* 2nd edition. Washington, DC: *The Urban Institute Press.*

**Tankersley, Jim.** 2018. "The New Tax Form Is Postcard-Size, But More Complicated Than Ever." *New York Times*, June 25.

**Tax Foundation.** 2017. "The State and Local Tax Deduction: A Primer." March 15.

**Tax Policy Center.** 2017a. "T17-0308 - Estate Tax Returns and Liability Under Current Law and the House and Senate Versions of the Tax Cuts and Jobs Act, 2018-2027." December 6. http://www.taxpolicycenter.org/model-estimates/estate-tax-and-tax-cuts-and-jobs-act-dec-2017/t17-0308-estate-tax-returns-and.

**Tax Policy Center.** 2017b. "Distributional Analysis of the Conference Agreement for the Tax Cuts and Jobs Act." December 18.

**Tax Reform Task Force.** 2016. *A Better Way—Our Vision for a Confident America: Tax.* June 24. https://abetterway.speaker.gov/_assets/pdf/ABetterWay-Tax-PolicyPaper.pdf.

**US Department of the Treasury.** 1984. *Tax Reform for Fairness, Simplicity and Economic Growth.* The Treasury Department Report to the President, November.

**Wallace, Sally, Michael Wasylenko, and David Weiner.** 1991. "The Distributional Implications of the 1986 Tax Reform." *National Tax Journal* 44(2): 181–98.

**White House, and US Department of the Treasury.** 2012. *The President's Framework for Business Tax Reform.* Report, February.

**White House, and US Department of the Treasury.** 2016. *The President's Framework for Business Tax Reform: An Update.* Report, April.

**York, Erica.** 2018. "Lawmakers May Vote on Making Key Provisions of the TCJA Permanent." Tax Foundation. April 5. https://taxfoundation.org/lawmakers-may-vote-making-key-provisions-tcja-permanent/.

**Zakaria, Fareed.** 2017. "The GOP Tax Bill May Be the Worst Piece of Legislation in Modern History." *Washington Post*, December 14.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 97–120*

# Measuring the Effects of Corporate Tax Cuts

Alan J. Auerbach

On December 22, 2017, President Donald Trump signed the Tax Cuts and Jobs Act (TCJA), the most sweeping revision of US tax law since the Tax Reform Act of 1986. The law introduced many significant changes. However, perhaps none was as important as the changes in the treatment of traditional "C" corporations—those corporations subject to a separate corporate income tax. Beginning in 2018, the federal corporate tax rate fell from 35 percent to 21 percent, some investment qualified for immediate deduction as an expense, and multinational corporations faced a substantially modified treatment of their activities.

In the views of its critics, the previous US corporate tax system discouraged companies from being US corporations, discouraged US corporations from repatriating the earnings from their overseas operations, and discouraged both US and foreign companies from operating in the United States—or at least from reporting the profits from their US operations in the United States. Additional to these concerns was a more traditional focus from the standpoint of economic research on the possible effects of the tax system on the composition of investment within the United States and the incentive for borrowing due to its favorable tax treatment.

The debate leading up to the bill's passage included some heated discussion among economists regarding the benefits of the corporate tax cut and who would receive them. Notably, the White House Council of Economic Advisers (2017) forecast that reducing the corporate tax rate to 20 percent (as the original version of

■ *Alan J. Auerbach is the Robert D. Burch Professor of Economics and Law, University of California, Berkeley, California. His email address is auerbach@econ.berkeley.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.97 doi=10.1257/jep.32.4.97

the legislation proposed) would lead to a substantial rise in wages. The result would be to "increase average household income in the United States by, very conservatively, $4,000 annually. … Moreover, the broad range of results in the literature suggest that over a decade, this effect could be much larger." Skeptics argued that the implied aggregate increase in income was implausibly large relative to the size of the tax cut. For example, Furman (2017) estimated that an increase in average income of between $4,000 and $9,000 (the upper bound for the CEA's estimated income increase) would be "between 275% and 550% of the total cost of the $200 billion corporate tax cut—implying a supply-side effect that's more than a little far-fetched." As discussed below, most other economic forecasts were for increases in income only a small fraction as large. Yet, in the weeks after the bill's passage, many leading corporations announced plans to provide $1,000 bonuses to their employees, commonly citing the tax cut as the reason for their actions (as reported in Shell 2018).

This paper seeks to evaluate these and other claims about the impact of the Tax Cuts and Jobs Act[1] to understand its effects on resource allocation and distribution. It begins by setting the stage with discussions of how corporate tax incidence has been studied in the past, comparisons of US corporate tax rates to other countries before the 2017 tax law, and some ways in which the US corporate sector has evolved that are especially relevant to tax policy—specifically, the decline in the share of business income accounted for by C corporations and the rising share of business income from international operations.

The discussion then turns to an explanation of the main changes of the Tax Cuts and Jobs Act of 2017 for the corporate income tax. A range of estimates suggests that the law is likely to contribute to increased US capital investment and, through that, an increase in US wages. The magnitude of these increases is extremely difficult to predict, because of the many channels through which investment may respond, the mechanisms connecting wage increases to increased investment and profitability, the instability of the law itself (because some of its provisions are explicitly temporary), how the law exacerbates the underlying US fiscal imbalance, and the possible international tax and trade policy reactions. Indeed, the public debate about the benefits of the new corporate tax provisions enacted (and the alternatives not adopted) has highlighted the limitations of standard approaches in distributional analysis to assigning corporate tax burdens. In particular, while such analyses have traditionally been framed in terms of the impact of the tax rate facing corporate fixed investment in a closed economy, such analysis must now be applied to the responses of multinational corporations, with worldwide operations and shareholder bases and a growing dependence on ideas rather than tangible assets. Such

[1] Although the law is commonly known as the Tax Cuts and Jobs Act, this name was stricken from the bill shortly before passage, its adoption deemed as not germane, according to Senate rules, to the budget reconciliation process used to pass the bill using a simple majority. The law's official name is considerably longer and not particularly memorable.

analysis requires a more sophisticated approach and more empirical evidence on the many potential margins of taxpayer response.

## Measuring Corporate Tax Incidence: Some Preliminaries

Tax incidence analysis involves estimating the effects of tax policy changes on different groups of individuals via the effects on prices and returns to labor and capital. The starting point for discussions of corporate tax incidence—that is, who bears the corporate tax?—dates back to Harberger's (1962) classic analysis of a two-sector general equilibrium model, which found that, in a closed economy with fixed factor supplies, the corporate tax fell approximately 100 percent on owners of all capital. The underlying intuition was that the corporate tax causes capital to shift from the corporate sector to the noncorporate sector (consisting of all businesses not subject to the corporate tax), depressing after-tax returns equally in both sectors but, for reasonable parameter assumptions, not shifting any of the tax burden to labor. Put equivalently, rebating all corporate tax revenues to owners of capital would leave them no better off than in the absence of the tax, and wage-earners no worse off.

With some modifications, the influence of Harberger's (1962) basic approach continues. For example, until relatively recently the distributional analysis of the Congressional Budget Office adopted the Harberger result in assigning 100 percent of the burden of corporate taxes collected to individuals in proportion to their receipt of capital income. Congressional Budget Office (2012) modified this assumption, citing its review of the empirical literature, and now follows the practice of assigning 25 percent of the burden of corporate taxes to individuals in proportion to their receipt of wage and salary income and 75 percent in proportion to their receipt of capital income.

A main source of the assumed shift in some of the burden toward labor is the consideration of international capital flows. Standard incidence analysis indicates that immobile factors such as labor bear some of the capital income taxes imposed within a country as the result of capital flight, with this burden approaching full shifting to labor as a country's size diminishes (for example, Kotlikoff and Summers 1987). More sophisticated open-economy general equilibrium models indicate a range of possible sharing of the burden of corporate taxes between labor and capital, depending on the degree of international capital mobility and the substitutability of foreign and domestic products, suggesting a share borne by labor of perhaps 40 percent based on calibration assumptions applicable to the United States (Gravelle 2013).

To pin down the effects of corporate tax changes more precisely, one would ideally look directly at empirical evidence on the effects of corporate tax changes on factor incomes in an international context. For an exogenous change in a country's corporate tax system, one would compare changes in after-tax incomes of different factors (for example, capital and labor) in different locations, or preferably changes

in the purchasing power of such after-tax incomes, to determine the distribution of burdens of the tax change. A handful of studies have tried to approximate this type of experiment to determine the share of the burden falling on labor, using panel data on countries, labor compensation, and tax rates. Unfortunately, the results of such analyses fall within a very wide range, from finding virtually no effect (Clausing 2013) to finding that "a 1% increase in corporate tax rates leads to a 0.5% decrease in wage rates" (Hassett and Mathur 2015). Because corporate profits are small relative to wages in the average economy in their sample, the Hassett–Mathur results imply, for a given level of corporate profits, that the reduction in wages resulting from a corporate tax rate increase would far exceed the revenue raised. That is, the incidence of the corporate tax change on labor, in this framework, would be considerably higher than 100 percent.

This range of findings for national corporate tax changes in a global economy and the small number of recent published studies in this literature hint at the empirical challenges involved. It has proved difficult to identify credible natural experiments for corporate tax reforms or to control for the many developments occurring within countries at the same time as corporate tax changes. A larger recent literature on corporate tax incidence looks within countries, considering differences across industries and across states or regions, for the United States as well as other countries, and using a range of models and assumptions (or examples, see Arulampalam et al. 2012; Liu and Altshuler 2013; Suárez Serrato and Zidar 2016; Fuest, Peichl, and Siegloch 2018). The findings are typically that a large share of the corporate tax falls on labor—quite plausible for changes adopted in a small part of a country in which there is considerable capital mobility, but not directly applicable to the issue in the recent debate of how the Tax Cuts and Jobs Act would affect US wages.

In attempting to translate results from the incidence literature into predictions about the effects of the Tax Cuts and Jobs Act on wages, it is useful to keep several other points in mind. First, while the literature has typically focused on changes in some measure of tax rates, with the overall tax structure fixed, the Tax Cuts and Jobs Act contained important changes in the structure of taxation itself. Standard distributional analysis such as those from the Congressional Budget Office commonly assume the same relative impact on labor and capital of changes in corporate tax revenues regardless of the way in which corporate taxes change, but economic theory and evidence suggests otherwise.

Second, in using incidence assumptions to break down projected changes in tax revenue into the shares borne by different groups, one is effectively equating the burden of tax changes to changes in tax revenue. However, these two measures differ conceptually and in practice because of behavioral responses to taxation. The change in tax revenue is calculated as the difference between the original tax base multiplied by the original tax rate, and the new tax base multiplied by the new tax rate. However, the starting point for thinking about the burden of a tax change would look at the change in tax rate multiplied by the initial tax base (as discussed in Joint Committee on Taxation 1993, p. 26). Put another way, the difference

between the change in revenue and the change in burden is equal to the change in the tax system's deadweight loss—the change in tax burden over and above revenue raised.[2] If a tax cut resulted in a substantial increase in the tax base, as was argued by many of those supporting the 2017 corporate tax changes, then the measured impact on wages would be quite a bit larger if based on the net change in revenue rather than on the change in revenue holding the tax base fixed.[3]

A third point to keep in mind is that constructing distributional estimates implicitly requires filling in important details not provided in the legislation. As an illustration, estimates of the effects of the 2017 law on tax revenue, even those taking both firm-level and economy-wide behavioral responses into account, projected substantial increases in the federal budget deficit over the next decade (Joint Committee on Taxation 2017b). These deficits, as well as subsequent fiscal and monetary responses to them, have economic effects as well, which the Congressional Budget Office and other forecasters must confront in forming their economic forecasts about the effects of the new law. If deficits crowd out capital accumulation, for example, this will likely depress wage growth.

Finally, estimates of the allocation of the burdens of taxation generally reflect a long-run, equilibrium analysis without necessarily taking account of the adjustment process. For example, a forecast may find that, in the end, all capital equally bears a proportional share of a corporate tax change. But in the very short run with a given capital allocation, the change may fall largely on corporate shareholders (for further discussion of such dynamic incidence issues, see Auerbach 2006). Likewise, to the extent that wage growth results from capital accumulation, any effect of a tax change on wages ought to occur over time, rather than immediately. While it is possible that a tax cut may generate economic stimulus that pushes up real wages, this is typically ignored in distributional analysis due to the focus on longer-run estimates. In that sense, most analyses of tax incidence adopt a "supply side" approach.

Even in light of all of these qualifications, a primary mechanism through which a corporate tax cut may influence wages is still likely to be capital deepening in the corporate sector that leads to increased labor productivity. One may trace the controversy over distributional effects of the 2017 tax cut (or other potential tax corporate cuts) to differences over the effectiveness of such tax cuts at promoting capital deepening, differences over the extent to which any such capital deepening would generate increases in wages, and differences over whether a corporate tax cut might increase wages through other significant channels.

[2]Harberger's (1962) original analysis considered a small corporate tax introduced into an economy without such a tax initially; for this case, there would be no difference between the two measures (the burden of tax changes versus the changes in tax revenue)—that is, no first-order deadweight loss.

[3]Traditional revenue estimates of tax legislation, like those produced by the Joint Committee on Taxation, typically incorporate some behavioral responses, although they exclude macroeconomic feedback effects that characterize "dynamic" scoring exercises. Thus, they differ from the fully "static" estimates one might wish to use for incidence and welfare analysis but do not involve the full impact incorporating dynamic scoring.

*Figure 1*
**G-7 Corporate Tax Rates**
(*percent*)

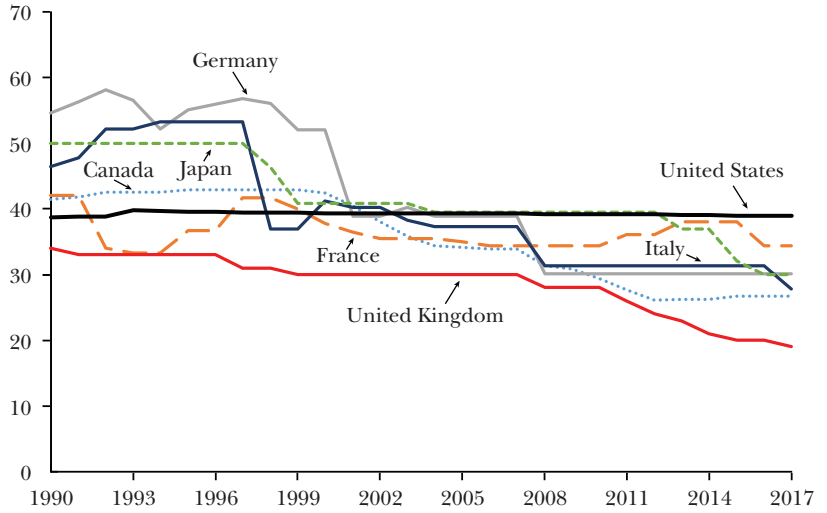


*Source:* Author using data from OECD Tax Database.
*Note:* Tax rates are for combined central and subcentral governments.

## US Corporate Tax Rate(s) as of 2017: High or Low?

Perhaps the simplest and most familiar argument for cutting the US corporate tax rate during the years leading up to 2017 was the changing landscape of corporate tax systems in other countries. Over the past few decades, developed countries have generally reduced their statutory corporate tax rates. Figure 1 shows the evolution of statutory tax rates (including subnational corporate taxes) for the G-7 countries between 1990 and 2017. Over this period, the United States made the transition from low-tax-rate country to high-tax-rate country without undertaking any significant policy changes, as alone among the G-7 countries it did not reduce its federal corporate tax rate (which actually rose from 34 percent to 35 percent in 1993).

While this comparison of statutory tax rates is striking, it ignores important differences among tax systems. Prior to the 2017 legislation, a common criticism of the argument for cutting the US corporate tax rate was that the effective US tax rate was actually not all that high if one took account of various provisions that narrowed the corporate tax base and lowered actual tax payments. As a simple illustration of the concept, if only half of US corporate income were subject to tax, then the tax system's effects would be the same as one that taxed all corporate income at half the statutory rate. However, other than the statutory rate itself, the provisions that affect corporate taxes are complicated and the implications for calculating the relevant effective tax rate are not clear. In the end, different effective tax rate measures will be useful, depending on the question that one seeks to address.

To illustrate the difficulty of constructing an aggregate effective tax rate measure, consider the common and apparently simple calculation that relates corporate taxes paid to corporate income, say $T/Y$. In 2013, the last year for which data are publicly available, C corporations in the United States had $1.258 trillion in taxable income and paid $293 billion in federal taxes, representing an average tax rate of 23 percent (Statistics of Income, 2013, Table 21), well below the statutory rate of 35 percent. This low rate, moreover, fails to account for deductions that had already reduced the denominator, taxable income, $Y$, relative to what some would argue is appropriate for measuring income; for example, accelerated depreciation that provides more generous deductions than economic depreciation. In this calculation, the only reason for the gap between the statutory and effective tax rate is, mechanically, the use of tax credits that reduce tax liability.

However, a closer look will discover that the bulk of these tax credits were foreign tax credits, intended to offset taxes already paid on foreign-source income to other countries. If one views the foreign tax credit as a loophole and considers as an ideal norm the full US taxation of the worldwide income of corporations without any credits for foreign taxes, then it makes sense to view the 23 percent tax rate as reflecting a low rate of tax. On the other hand, as of 2017 all other G-7 countries had largely dispensed with taxing foreign-source corporate income at all. The other six had all adopted a so-called "territorial" approach of exempting corporate foreign-source income from tax, which in this calculation would be equivalent to allowing foreign tax credits to eliminate all US taxes on foreign source income. Taking that territorial approach as the norm would suggest that US taxation of foreign-source income was high relative to other countries. In short, corporations faced a lower US tax rate on foreign-source income than on domestic-source income, but how one should interpret this fact is not obvious.

This difficulty increases once one recognizes that the location of profits reported by companies (domestic versus foreign) may differ from where profits are earned. One of the arguments for tax reform as of 2017 was that the US tax system encouraged companies to shift the location of reported profits to low-tax foreign countries, through cross-border transactions with related parties in these countries and shifts in the location of deductible expenses such as interest. To the extent that such profits then faced a lower tax rate, this would effectively represent a lower tax rate on US domestic profits.

The issue of how to view a lower tax rate on foreign-source income also arises in effective tax rate calculations presented in corporations' public financial statements. For example, Apple's 2017 Form 10-K reports (on p. 56) an effective tax rate (defined here as total worldwide taxes divided by total worldwide earnings) of 24.6 percent. Most of the reduction from the 35 percent US statutory rate is attributable to "indefinitely reinvested earnings of foreign subsidiaries"—the foreign-source earnings that financial accounting treats as having no deferred US tax liability associated with future taxes on earnings repatriation. In this instance, foreign taxes are included in the calculation, and so the lower tax rate reflects a lower overall tax on foreign-source income, rather than just a lower US tax.

However, even after adjusting corporate taxable income for provisions that reduce it relative to economic income, the tax rate on measured domestic corporate income does not appear to be low in relation to the statutory rate. In Auerbach (2007), I estimated average annual tax rates for US nonfinancial corporations, comparing taxes paid on domestic earnings to income as measured based on the National Income and Product Accounts, rather than the taxable income reported on tax returns. For the period from 1993 to 2003, when the corporate tax rate was 35 percent, nonfinancial corporations faced annual average tax rates ranging from 29.2 percent to 49.2 percent, as provisions that reduced tax rates, such as accelerated depreciation, were in many years more than offset by provisions that raised tax rates (notably, the limited deductibility of net operating losses). These limits on the deduction of losses raise average tax rates because the denominator (income) falls by a greater proportion than the numerator (taxes): in the extreme case where losses are completely nondeductible, losses affect only the denominator. Hines (2017), using a related approach based on the reported magnitudes of domestic tax expenditures—tax provisions that reduce the tax base—finds in more recent calculations only a small reduction in the average corporate tax rate relative to its statutory value, even without taking account of the increases associated with the limits on deduction of losses.

Thus, leading up to the tax debate in 2017, US corporations faced a very high statutory tax rate relative to other countries, a much lower US tax rate on foreign-source income relative to domestic-source income, and a reasonably high average tax rate on reported domestic-source equity income, even taking into account deductions and credits that lowered tax liabilities. One should note that none of these alternative tax rate measures accounts for the additional taxes paid by shareholders of US corporations, or the treatment of corporate borrowing and interest deductibility. Nor do they distinguish among industries or the types of assets in which companies might be investing.

In summary, it is difficult to know which tax rate calculation is appropriate without first identifying the question one wishes to answer, which in turn relates to the behavioral response, or responses, of interest. For example, one might wish to evaluate the responses of companies deciding how much to invest in the United States, or in which types of assets to invest, or whether to invest in the United States or another country, or of individuals choosing whether to invest in US corporate stock, corporate bonds, or noncorporate businesses. Each of these decisions involves a different tax rate calculation, and the decisions will naturally differ in their implications for corporate tax incidence. In addition, the importance of decisions on different margins has changed over time, not just because of changes in US and foreign tax provisions, but also because of the changing nature of corporate activity.

## Modeling Behavior of the Evolving US Corporate Sector

Economic analyses of the behavioral responses to corporate taxation commonly begin by considering the tax wedge imposed on investment in the corporate sector.

This wedge depends, of course, on the corporate tax rate itself, but on other tax provisions as well, although empirical studies of corporate tax incidence have not always taken account of these other provisions. Doing so leads one to compute a forward-looking *marginal effective tax rate* on new investment—in present value, the share of the before-tax rate of return on an incremental new investment going to federal (or federal plus state) taxes.

The marginal effective tax rate generally takes into account a number of factors: corporate-level taxes; incentive provisions applying to particular types of investment; interest deductibility at the corporate level; and the taxes of shareholders and bondholders on dividends, capital gains, and interest income. Estimates of the incentives to shift between corporate and noncorporate activity involve calculating a similar marginal effective tax rate for noncorporate investment. Because the marginal effective tax rate is a prospective tax rate applying to a particular investment decision, it relates more directly to specific behavioral responses than the average effective tax rates discussed in the previous section. The approach to calculating marginal effective tax rates has been refined for decades; Congressional Budget Office (2014) is a good recent example. The marginal effective tax rate is one component of the user cost of capital facing investment, along with actual depreciation and the required return to investors.

Having computed various marginal effective tax rates for different kinds of investments, one can estimate the effects of taxation on investment incentives in different sectors and the impact on returns to labor and capital. Among the insights one gets from such an analysis are that the corporate sector marginal effective tax rate has historically been substantially reduced by interest deductibility, given that interest income is typically received by individuals or entities (such as pension funds or retirement accounts) in much lower tax brackets than corporations. As a result, the net tax burden on interest payments, taking account of corporate tax deductions and taxes paid by recipients, is negative. For example, a Congressional Budget Office (2014) study found an overall marginal effective tax rate on corporate investment in tangible capital of 31 percent, equal to a weighted average of a 38 percent tax rate on equity-financed investment and -6 percent on debt-financed investment, confirming the strong tax incentive to use debt finance. In addition, effective tax rates vary considerably across assets, because of provisions for depreciation and other incentives that apply differentially. The Congressional Budget Office analysis found a 30 percentage-point range in effective tax rates for C corporations by asset type, varying from 12 percent for replacement railroad track to 42 percent for nuclear fuel.

The exact methodology varies from study to study; for example, Gravelle (2016) includes intangible assets (which could be expensed and in some cases qualified for the research and experimentation credit) in her calculations and accordingly finds lower marginal effective tax rates. But such calculations typically share a number of important common elements in addition to the focus on marginal investment decisions. These common elements include a closed-economy perspective that looks at the overall tax wedge faced when domestic savers provide funds for domestic

*Figure 2*
**Shares of Total Business Net Income (Less Deficit) 1980–2013**
*(in percent)*

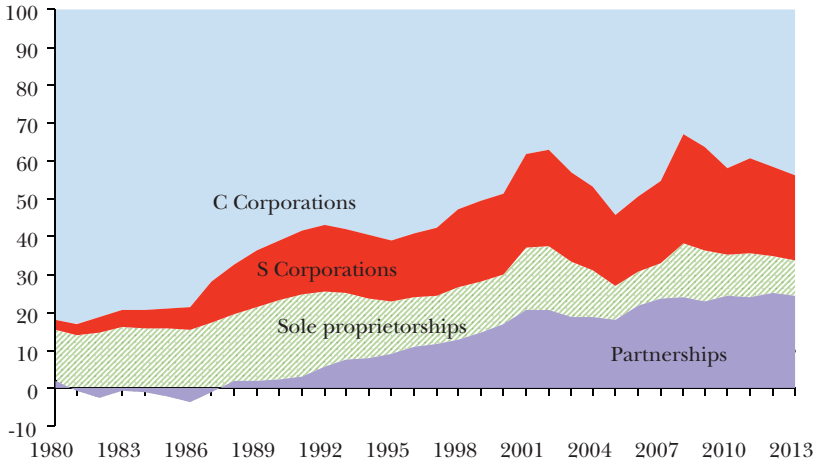


*Source:* The data come from Looney and Krupkin (2017).

investment. Such simplifying assumptions are increasingly restrictive because of how the corporate sector has evolved, and are limited in their usefulness in evaluating some tax reform proposals, particularly those affecting international corporate activities. To understand the nature of these restrictions and limitations, it is useful to highlight some features that now characterize the US corporate sector.

Figure 2 shows the components of overall US business sector income. The top area shows the C corporate sector, which fluctuates procyclically more than other business income, and hence has accounted for a lower share of business income during recessions. Also, the share of business income going to C corporations has fallen from around 80 percent in 1980 to about half of all business income. This trend is one factor underlying the relatively low share of federal tax revenue for which the corporate sector currently accounts.

It also highlights the importance of tax provisions affecting the remaining, "pass-through" entities: sole proprietorships, partnerships, and S corporations. The S corporations have corporate legal status, but pass through their earnings and tax liability to owners and face no entity-level federal tax; the sole proprietorships and partnerships pass through their earnings and tax liability to owners but do not have corporate legal status. (It is customary to lump S corporations with other pass-through entities when referring to the noncorporate sector, because they are taxed in the same way that noncorporate entitites are.) The negative income of partnerships in the early 1980s is attributable to the tax shelters largely eliminated by the Tax Reform Act of 1986 (Auerbach and Slemrod 1997).

As of 2017, the differential tax wedge associated with the C corporation sector as compared to the pass-through sector was not especially large: Congressional Budget Office (2014) estimated the pass-through sector to have a marginal effective tax rate of 27 percent, relative to 31 percent for the corporate sector. While some of the growth of pass-through entities has been attributed in the past to attempts to avoid corporate double-taxation (for example, Auerbach and Slemrod 1997), net increases in top individual marginal tax rates between 1993 and 2017 along with lower tax rates on capital gains and dividends adopted over the same period reduced the tax gap between the two sectors. Thus, the tax wedge between corporate and pass-through sectors may be less important than in years past as a factor influencing the allocation of capital.

On the other hand, international capital flows have become more important over time, not only overall, but especially within the corporate sector, through the activities of multinational corporations. In the five decades between 1966 and 2016, the share of the income of US resident corporations that was accounted for by foreign operations rose from 6.3 to 31.1 percent (Auerbach 2017). The increasing importance of international capital flows is one factor underlying the shift in consensus that a lower share of the US corporate tax is now being borne by capital (typified by the 2012 change in assumptions by the Congressional Budget Office mentioned earlier). The intuition is that a higher US corporate tax rate may now more easily lead to a shift of capital to other countries. However, the magnitude of any such response depends on the tax rules that apply to international investment, and the complexity of such rules gives rise to a range of behavioral responses among firms potentially operating in more than one jurisdiction. These responses include discrete location decisions, rather than just investment levels, along with the financial and accounting strategies firms use to shift reported profits among countries. For both discrete location decisions and profit-shifting, the statutory tax rate may be more relevant than a computed marginal effective tax rate. The reason is that discrete location decisions may involve choosing where to locate profitable existing activities, which are subject to the statutory rate, in addition to new capital investment. Moreover, shifting reported profits need not coincide with changes in the actual location of investments. Indeed, differences in statutory tax rates among countries seem to affect both types of decisions (Devereux and Griffith 1998; Dharmapala 2014).

Another relevant aspect of the increase in international capital flows is the growth in cross-border ownership of corporations. The traditional view that equates the nationality of corporations and their owners is now far from accurate; foreign shareholders owned roughly a quarter of US corporate equity in 2015 (according to Rosenthal and Austin 2016). This pattern calls into question the suitability of marginal effective tax rate or related calculations that combine US corporate and shareholder taxes in constructing an overall tax wedge. To the extent that companies draw from a worldwide shareholder base, incentives to invest in the United States may depend more on corporate-level taxes than those at the US shareholder level.

The internationalization of shareholding also suggests that companies may change residence even if their shareholders do not—an issue that has arisen as

corporations engaged in so-called "inversions" undertaken to shift their corporate residence away from the United States. Tax incentives to invert relate not only to the US corporate tax rate, but also to the traditional US approach to taxing foreign source income. As already discussed, the United States has attempted to impose some taxes on the foreign source income of US corporations—but only US corporations—while other countries have increasingly adopted a territorial tax system in which they do not seek to tax the foreign source income of their resident corporations. This difference meant that a US company would face a higher overall tax rate on its investments in low-tax countries than would a non-US company, even if the non-US company resided in a country with a tax rate as high as or higher than the US tax rate.

Finally, the composition of investments by the US corporate sector has changed over time, with an increasing share devoted to intangible assets. Measuring the total value of such assets is difficult because one can view many business expenditures (such as advertising) as creating value. However, based on relatively narrow definitions of purchased intellectual property assets, both the Bureau of Economic Analysis Fixed Assets Accounts and the Federal Reserve Board's Flow of Funds Accounts show a doubling of the share of intangible assets as a share of business capital over the 50 years from 1966 to 2016 (Auerbach 2017).

The growing dependence of production on intangible assets presents an increasing challenge to the enforcement of international tax rules for multinational corporations. Such rules rely on determining the location and profitability of a firm's assets, but this determination is especially difficult for intangible assets with no physical presence and with firm-specific characteristics that determine profitability. These characteristics facilitate the responsiveness of profit-shifting to differences in statutory tax rates, and companies that rely heavily on intellectual property have been among the highest-profile firms criticized for international tax avoidance. This enhanced opportunity for profit shifting, as well as perceived spillover benefits from the development and use of intellectual property, has led many countries, including Ireland, the United Kingdom, Belgium, the Netherlands, and Luxembourg, to adopt favorable tax rates for income attributed to intellectual property tax regimes, sometimes called "patent boxes."

In summary, the rise of the multinational corporation, with cross-border ownership and operations, and the growing importance of intellectual property in production have broadened the set of relevant behavioral responses to corporate taxation and led governments to participate in a multidimensional tax competition game. In this game, each country chooses not only its statutory corporate tax rate, but also asset-specific provisions applying to domestic investment and rules applying to cross-border investments. Changes in any one instrument may affect firms on several decision margins, and policy changes might influence US investment through several direct and indirect channels. While one may expect a reduction in the US corporate tax rate to encourage US-based investment and production, the effects of other policy changes may be more complex.

For example, consider an increase in the US tax rate applicable to the foreign source income of US companies. This could affect US domestic investment and

production in at least three ways. First, it might discourage those companies from producing abroad, because of a lower after-tax rate of return on foreign production. This may cause an increase or decrease in their investment and production in the United States, depending on whether the foreign and domestic activities of multinational companies are gross substitutes or gross complements in production—although available evidence suggests overall complementarity of foreign and domestic operations (for example, Desai, Foley, and Hines 2009; Becker and Riedel 2012). Second, it could reduce the tax benefits US companies get by shifting their US source profits to low-tax foreign countries, which raises the effective tax rate on US profits and therefore discourages US production. Evidence from other countries confirms this effect, finding that strengthening provisions aimed at limiting profit shifting reduces domestic investment (Overesch 2009; de Mooij and Liu 2018). Third, an increase in the US corporate tax rate applicable to foreign-source income might encourage companies to relinquish US residence through corporate inversions, because only US companies are subject to this higher tax on foreign operations. Although there is little empirical evidence on this question, some have argued that shifting residence may also lead to shifts in certain activities away from the United States.

More generally, it is useful to think of international tax provisions and their effects in terms of the extent to which they conform to three different ways of taxing companies: residence-based corporate taxation, which countries define in various ways but usually involves factors like location of key management and headquarters activities and place of incorporation; source-based taxation, which looks at where the companies' production takes place; and destination-based taxation, which looks at where the companies sell their products. The US corporate tax system before the 2017 legislation was a hybrid of residence- and source-based taxation. It imposed corporate income tax on all profits resulting from production occurring in the United States, whether by US or foreign companies, but also taxed the foreign-source earnings of US companies when repatriated (with a credit for foreign tax paid). The increase in multinational activity and the greater reliance on intellectual property in production had made it easier for companies to shift the location of their production or at least the reported location of their profits, challenging a reliance on source-based taxation. The internationalization of companies and their ownership had made corporate residence less of a fixed characteristic, challenging a reliance on residence-based taxation.

The challenges to relying on source or residence as a basis of taxation has led some to consider a move in the direction of destination-based taxation, with the rationale that the location of consumers is more determinate. Indeed, for a time during the recent tax reform debate, the US Congress considered a proposal for doing so in the form of a destination-based cash-flow tax (Ways and Means Committee 2016).[4] (Though the proposal ultimately was not adopted, the Tax Cuts and Jobs

[4] A related approach would be to apportion profits to particular jurisdictions in proportion to the location of the company's sales. Among US states, the use of sales to apportion corporate taxes has been growing in importance over time, for related reasons (Suárez Serrato and Zidar 2016, figure 3).

Act includes provisions, discussed below, that are related in form and motivation.) Such a tax is still formally on corporations and other businesses, but its incidence should be quite different from traditional corporate taxes, precisely because of its use of a destination basis. To understand why, consider the destination-based cash-flow tax in the context of the national income identity.

Start with the national income identity that GDP equals the sum of consumption ($C$), domestic investment ($I$), government purchases ($G$) and exports ($X$) less imports ($M$). It follows that $C = GDP - I - G - X + M$: taxing consumption can be achieved by taxing income net of exports, also taxing imports, allowing expensing of investment, and not taxing government purchases. (This explanation follows Auerbach 2017).[5] If one divides private GDP ($GDP - G$) into returns to labor, $W$, and returns to capital, $R$, then the consumption tax can be implemented in two pieces, as a tax on returns to labor, $W$ (already effectively covered by the existing personal income tax), plus a border-adjusted tax (that is, allowing a deduction for exports and imposing a tax on imports) on business cash flows, $R - I - X + M = C - W$. This latter component is the destination-based cash-flow tax base. In its operation, companies would pay tax on their domestic cash flows: receipts from domestic purchasers less domestic labor costs and purchases of intermediate and investment goods from domestic sellers.

Because the destination-based cash-flow tax is a tax on domestic consumption net of returns to labor, one would expect it to have incidence similar to a tax on domestic consumption that exempts consumption financed by wage and salary income—roughly approximating a tax on domestic wealth. Its incidence would therefore be quite different from what is usually assumed for the corporate tax, as this form of tax would omit the main channel through which the corporate tax is modeled as being shifted to labor—capital flight—precisely because there would be no tax discouraging domestic production or lowering the rate of return to domestic investment. Companies would face no additional tax because of producing in the United States. They would also have no capacity to shift profits to other countries, because the destination-based cash-flow tax would ignore all of the transactions typically used to do so, including interest deductions and cross-border transactions with related parties. However, as a tax on domestic US consumption, the destination-based cash-flow tax would also not fall on foreign-owned capital, whereas taxes based on US production or the ownership of US companies could do so to some extent. This is the case, even though distributional analyses dividing the corporate tax burden between labor and capital often ignore the distinction between foreign and domestic owners of capital.

One further complication in thinking about the effects of any US policy change, especially those relating to international activity, is the need to account for the responses of other countries. Although a number of multilateral initiatives

[5] Note that expensing of investment, rather than the schedule of depreciation deductions traditionally allowed under an income tax, is required to achieve a tax on consumption but that a destination-based income tax is also possible.

have sought to limit tax competition by tightening the rules applicable to multinational companies, most recently through the OECD project on Base Erosion and Profit Shifting, such restrictions may change the nature of tax competition without necessarily restraining it. For example, some forms of tax competition, such as the "patent box" regimes mentioned above, may still be allowed. Also, more restrictions on provisions that allow companies to avoid taxes through profit-shifting might lead to more intense tax competition between countries with respect to statutory tax rates, and less efficient overall outcomes, by focusing tax reductions more on activities with relatively low responsiveness to taxation (Keen 2001). In the present setting, one may expect the significant changes in the US tax system to prompt responses in other countries, although these effects are generally not considered in the various projections that have been done for the effects of the Tax Cuts and Jobs Act of 2017.

## The Tax Cuts and Jobs Act and its Potential Effects

Several aspects of the US corporate tax and its effects were under scrutiny during the tax reform process, including the high statutory corporate tax rate and continued worldwide taxation of the profits of US corporations. Using a worldwide base for corporate taxation was cited not only as a driving factor behind corporate inversions, but also—because the additional US tax would be due only upon the repatriation of foreign-source income—as a reason behind the large accumulation of retained earnings offshore—estimated as of 2015 at $2.6 trillion (Joint Committee on Taxation 2016). The worldwide basis for US corporate taxation has also been linked to the large concentration of US companies' offshore earnings in tax havens—generally defined as countries with very low corporate tax rates and flexible rules regarding the transactions of multinationals (in this journal, Zucman 2014).

In the end, the new tax law introduced a substantial cut in the corporate tax rate and a temporary allowance for investment expensing of certain classes of capital, both of which one would expect to encourage domestic investment by lowering the marginal effective tax rate on new investment.[6] Working in the same direction was the elimination of the corporate Alternative Minimum Tax (AMT). This investment stimulus is somewhat offset by a new limit on interest deductibility and new restrictions on the ability of companies to use net operating losses to offset past or future income (through tax loss carrybacks and carryforwards).

[6] It is worth noting that the effects of these two provisions on the effective tax rate on new investment interact negatively. That is, a reduction in the statutory corporate tax rate makes the accelerated deductions from expensing less valuable; for assets qualifying for full expensing, a well-known result is that the marginal effective tax rate on equity-financed investments is zero, and therefore is not reduced at all by cuts in the statutory tax rate. Indeed, to the extent that investment assets are debt-financed, assets that are expensed face a negative marginal effective tax rate at the corporate level, so that a reduction in the statutory tax rate actually raises the marginal effective tax rate.

The new law also adopted a 20 percent deduction for the qualifying income of pass-through entities, with the determination of qualifying income being subject to a very complicated set of provisions aimed primarily at preventing wage and salary income from being reclassified as business income. Even with this tax benefit, the sharp drop in the corporate rate likely reduces the tax incentive for a business to operate as a pass-through business rather than as a C corporation. For example, Congressional Budget Office (2018) estimates that the marginal effective tax rate for corporate business capital will fall by around 8 percentage points during the early years under the new tax law, before investment incentives decline, while the marginal effective tax rate for noncorporate business capital will fall by about half as much.

In addition, the law includes three key provisions aimed specifically at influencing the behavior of multinationals, each with its own new acronym.

First, the Global Intangible Low Tax Income (GILTI) provision replaced the tax on *repatriated* foreign-source income with an annual tax at half the domestic rate (10.5 percent) levied on *accrued* foreign-source income above a 10 percent rate of return on foreign plant and equipment and subject to only a partial (80 percent) foreign tax credit. As a consequence, the GILTI left earnings taxed abroad at less than 13.125 percent still subject to some US tax.[7] While the effect of a change in the present-value tax rate on foreign-source income is ambiguous (an immediate tax on accrual of income replacing a higher-rate tax on deferred realization of income), eliminating the tax consequences of repatriation reduces the incentive to keep earnings offshore, the so-called "lock-out effect." This change could increase domestic investment, to the extent that companies facing a lower tax barrier to repatriating earnings are liquidity-constrained. However, evidence from a previous episode during which the tax incentive to repatriate earnings was temporarily enhanced found that the induced repatriations led to little additional domestic investment (Dharmapala, Foley, and Forbes 2011).

Second, the Base Erosion Anti-Abuse Tax (BEAT) is a new minimum tax at a rate of 10 percent on the income of companies operating in the United States. The minimum tax base calculation disallows deductions for some imports from related parties.[8] The BEAT aims to limit the extent to which companies can shift reported

[7] Accompanying the elimination of the tax on any earnings repatriated in the future was a one-time tax on previously accumulated offshore earnings, which is a lump-sum tax (if one ignores the possibility of any behavioral impact coming through induced changes in expectations about future tax policy). The common rationale for this transition tax is that it substitutes for the taxes that companies would have been expected to pay on repatriations of previously accumulated offshore earnings under the old system. However, the tax rates associated with this measure rose throughout the brief legislative process, suggesting that the rates were determined by a need to hit a tax revenue target as much as to satisfy a specific policy aim. Moreover, there was no attempt to offset other windfalls associated with changes in business taxation, in particular the reduced corporate tax rate applied to the income from past domestic investments.

[8] In particular, import costs falling in the accounting category "cost of goods sold"—referring to intermediate goods—still are deductible, whereas other imports from related parties, such as accounting and financial services, are not.

profits out of the United States using internal transfer pricing manipulation by inflating the cost of their imports from related foreign parties. However, the tax applies to imports even from high-tax countries. Companies subject to the BEAT would face higher costs of operating in the United States, as well as an incentive to spin off foreign-related parties so that imports would no longer be covered by the tax.

Third, the Foreign Derived Intangible Income (FDII) provision introduces a lower tax rate (13.125 percent) on the share of domestic earnings from foreign-derived intangible income in excess of 10 percent of assets and attributable to exports (based on the share of export sales in all sales). This provision resembles the "patent box" tax rules adopted in other countries, which as discussed above aim to reduce the tax rate on one kind of especially mobile activity. However, it does not refer specifically to income generated by intellectual property, and in applying only to export income it is more limited in scope than existing patent boxes. It would encourage companies to locate export-related activities in the United States, particularly those involving intangible assets (which typically will not show up in the asset base calculation and so will not raise the threshold above which earnings are tax favored).[9]

Heightening the usual uncertainty about the effects of a major new tax law was the fact that it calls for many additional changes during the next 10 years. These include phasing out investment expensing, tightening the interest deduction limit, introducing a requirement that companies amortize rather than immediately deduct expenditures on research and development, and raising the tax rates associated with all three of the key international provisions just discussed.[10] These changes are central to the budget chicanery that has become a central part of the US legislative process, enabling Congress to hit a predetermined 10-year revenue-loss target and to avoid increasing deficits after 10 years; otherwise, the new tax law would have raised deficits in a way that required an unattainable supermajority vote in the Senate. Further, the large increase in budget deficits even under the bill as enacted could lead to other modifications, such as an eventual increase in the corporate tax rate, which would lessen any immediate positive impact on domestic investment.

Table 1 summarizes the new tax law provisions, indicating the anticipated impact on domestic investment. The last column of the table provides the associated 10-year revenue estimates from Joint Committee on Taxation (2017a) as well as those for fiscal year 2020, a year after initial phase-ins and before most changes

[9] In terms of US dependence on residence, source, or destination as a basis of taxation, one may view the first of the three changes listed in this section as maintaining but reforming the pre-existing approach that combines residence- and source-based taxation. The second and third provisions involve rather limited steps in the direction of destination-based taxation, with a similar motivation of curbing tax avoidance as the border adjustment that would have been part of the destination-based cash flow tax, as they reduce both the tax deduction for imports and the tax on exports.

[10] The requirement for research and development amortization is particularly puzzling in light of the apparent intent of the FDII provision to encourage the location of intellectual property in the United States.

*Table 1*
**Key Provisions of the Tax Cut and Jobs Act Affecting Multinational Corporations**

| *Provision* | *Policy change* | *Predicted economic impact* | *Revenue impact, JCT ($billions), 2020 (above) & 10-year (below)* |
|---|---|---|---|
| Corporate tax rate cut | Reduction from 35% to 21% | Increased domestic investment, from lower marginal effective tax rate (intensive margin) and lower average tax rate (extensive margin) | – 130.5<br>–1,348.5 |
| Investment expensing | Full through 2022; gradually phased out by 2027 | Increased domestic investment; possibly larger if temporary | – 24.6<br>–86.3 |
| Limitation on interest deductions | 30% of EBITDA through 2021; 30% of EBIT thereafter | Reduced domestic investment; increased borrowing abroad | +19.7<br>+253.4 |
| Net operating loss deductions | Elimination of 2-year loss carrybacks; limit of use of loss carryforwards to 80% of taxable income; elimination of 20-year expiration of loss carryforwards | Reduced domestic investment, especially in more cyclical industries | +11.1<br>+201.1 |
| Tax on offshore earnings | Elimination of tax on earnings repatriation; one-time tax on previously accumulated offshore earnings (15.5% for cash; 8% for assets) subject to scaled foreign tax credit; new tax on earnings in excess of 10% of offshore assets (GILTI) at 10.5% through 2025 and 13.125% thereafter | Increased earnings repatriation; uncertain impact on foreign and domestic investment | +6.0<br>+227.6 |
| Minimum tax on domestic earnings | Tax on expanded base (BEAT) that eliminates deduction of cost of imports (except for "cost of goods sold") from related foreign parties, at a rate of 5% in 2018, 10% from 2019–2025, and 12.5% starting in 2026 | Reduced domestic investment; spinoff of foreign operations | +13.3<br>+149.6 |
| Tax benefit for exports | Reduced tax rate, at 13.125% through 2025 and 16.406% starting in 2026, on foreign-derived intangible income (FDII), which is earnings above 10% of assets, multiplied by the fraction of domestic earnings apportioned to export sales | Increased location of intellectual property in the United States, to the extent that provision is expected to survive | +6.9<br>– 63.8 |

*Note:* JCT is the Joint Committee on Taxation. EBITDA is Earnings Before Interest, Taxes, Depreciation, and Amortization. EBIT is Earnings Before Interest and Tax. GILTI stands for Global Intangible Low Tax Income. BEAT stands for Base Erosion and Anti-Abuse Tax.

and expirations.[11] For some provisions, such as the permanent corporate tax rate cut, the 10-year and 2020 revenue effects are in close alignment. For others, the relationship between the two revenue effects is affected by phase-out (as in the

[11] For the change in the treatment of offshore earnings, the numbers in the table are the sum of those from three changes: elimination of the tax on repatriations, introduction of the tax on accrued offshore income, and the transition tax on previous accumulations by foreign subsidiaries.

case of expensing) or phase-in (as in the case of interest deduction limits). In the case of the Foreign-Derived Intangible Income (FDII) provisions, the one-year and 10-year revenue effects are of opposite sign because of a short-run timing response, presumably due to an immediate reduction in profit-shifting and a relocation of some intellectual property to the United States.

Assessing the net impact of these provisions on investment is very challenging, as it requires one to account for the interaction of a broad range of provisions, with little evidence regarding many behavioral responses. The task is much more difficult than for the case of a simple cut in the corporate tax rate, the conceptual experiment that many have in mind when predicting the effects of corporate taxation on production, investment, and wages. Further, overall assessments, especially with regard to the short run, must also account for the pace of adjustment, demand stimulus, the effects of increased deficits on national saving and capital inflows, and the potential response of monetary policy.

With these concerns duly noted, there have been attempts to quantify the legislation's impact on domestic investment. The Joint Committee on Taxation (2017b) "projects an increase in investment in the United States, both as a result of the proposals directly affecting taxation of foreign source income of US multinational corporations, and from the reduction in the after-tax cost of capital in the United States." The average increase in the capital stock over the 10-year budget window is 0.9 percent and the average increase in GDP is 0.7 percent, although the increases are smaller at the end of the period because of the changes in provisions noted above. Congressional Budget Office (2018) projects an average increase in GDP of 0.7 percent over the 10-year budget period. A relatively similar private-sector assessment by Macroeconomic Advisers (2018) finds that potential GDP rises by 0.6 percent by the end of the budget period, "mainly by encouraging an expansion of the domestic capital stock." The Penn Wharton Budget Model (2017) estimates a 10-year growth in GDP of between 0.6 and 1.1 percent, depending on assumptions about the composition of returns to capital. Barro and Furman (forthcoming, Table 11) estimate that GDP would be higher as a result of an increased capital-labor ratio, by 0.4 percent after 10 years under the law as written, and 1.2 percent if initial provisions were made permanent, with the effects being smaller if deficit-induced crowding out is taken into account.[12]

Based on a production-function approach, using the standard Cobb–Douglas constant-income-shares assumption, a GDP growth estimate in this range, say 0.6 percent, also suggests an increase in annual labor income of 0.6 percent, or approximately $500 per household at current income levels.[13] An increase in compensation of $500 per family for 125 million families equals $62.5 billion, which

[12] The estimates by Barro and Furman explicitly do not take account of the effects of the international tax provisions. It is unclear how important a role these provisions play in the estimates by Macroeconomic Advisers and the Penn Wharton Budget Model.

[13] This uses Furman's (2017) estimate of 125 million households and 2017 compensation of employees of $10.3 trillion, and follows footnote 3 in using the revenue estimate exclusive of dynamic scoring.

compares to the fiscal year 2020 revenue loss from the corporate tax rate reduction alone, listed in Table 1, of $130.5 billion. This increase amounts to about half of the reduction in corporate tax revenues, or roughly double what one would get by applying the common assumption that 25 percent of the corporate tax cut goes to labor.

An effect of this size is certainly plausible, given the many other provisions that may stimulate investment and the initial deficit-induced demand stimulus, but it is a far smaller number than some of those discussed in the introduction. Might these analyses fail to account for important potential channels or macroeconomic responses through which the tax reform might affect output and wages? The impact of the international provisions is especially subject to uncertainty, but many other modeling assumptions are involved in the overall estimates. Moreover, how does one square these predictions of a gain in the range of $500 per household with numerous companies having announced immediate $1,000 bonuses to their workers?

From the basic perspective of a competitive economy without frictions, treating labor markets as spot markets, one would predict that firms would raise wages only when labor demand increases, which in turn would require an increase in labor productivity or a spur in demand for the firm's products. Even accounting for the impact of fiscal stimulus on product demand, one would not expect this increase in labor demand to occur immediately. Moving to a slightly more complex view of labor markets, involving costs of training and adjustment and other frictions, could help explain why firms might raise wages in anticipation of stronger future labor demand, as firms would wish to smooth fluctuations in their employment by building up their workforce. (These actions could also have been encouraged to the extent that companies could deduct payments at the higher, 35 percent tax rate.) Whether such an explanation suffices is unclear, given the magnitude of immediate bonus payments.

However, in a less-competitive setting, in which firms earn rents, the possibility arises that firms may share some of those rents with workers. For a recent survey of the extent of rent sharing by firms, see Card, Cardoso, Heining, and Kline (2018, especially Table A1). To the extent that such rents have increased relative to normal returns to capital (as suggested, for example, by Furman and Orszag 2015), rent sharing could play an important role in determining the incidence of corporate tax changes.[14] In this setting, a corporate tax rate reduction could potentially lead to an increase in labor compensation by immediately increasing after-tax corporate

[14]The rent-sharing hypothesis with respect to corporate tax changes stands in contrast to other assumptions about corporate tax incidence based on the existence of rents. Notably, the US Treasury assumes that the corporate tax collected on supernormal returns (which it estimates to account for 63 percent of corporate taxable income) is borne by owners of corporate capital (Cronin, Lin, Power, and Cooper 2013). Even without direct rent sharing with workers, the incidence of a tax on rents could still fall partially on workers to the extent that the rents are specific to a company rather than to a location, for then the company could move the rent-producing activities to other jurisdictions, thereby lessening worker productivity (Auerbach and Devereux 2018).

profitability, prior to the occurrence of any increase in labor productivity. Indeed, some of the recent empirical literature on corporate tax incidence within countries adopts the rent-sharing approach, including Arulampalam, Devereaux, and Maffini (2012) and Liu and Altshuler (2013). Why such a sharing of rents with workers should take the form of bonuses rather than wage increases is unclear, without knowing the process underlying the determination of rent sharing, although the uncertain fate of various provisions of the Tax Cuts and Jobs Act could potentially be playing a role.[15]

## Discussion

While there is no simple consensus framework in which to evaluate the Tax Cuts and Jobs Act of 2017, one can reach some plausible conclusions about the rough magnitudes of the effects of the tax reform on US labor and capital income. But the potential for disagreement with these estimates is large. One source of disagreement is over whether one assumes that the changes in the Tax Cuts and Jobs Act supposed to occur within the next ten years (such as expiring incentives or changes in tax rates) will be sustained or modified. Another set of disagreements can arise because of differences in behavioral models of corporate responses or assumed parameter values. At a more subtle level, differences can also occur in the hypothetical experiments that individuals have in mind. For example, what changes other than a reduction in the corporate tax rate are included in a given study? Measuring the potential effects of the legislation requires accounting for myriad other provisions affecting investment decisions and international activity, which the law substantially altered. To do this, one must calculate tax wedges and trace out potential behavioral effects on several margins, for which there may be relatively little or no direct empirical evidence, or for which historical evidence may be of limited use given the changing characteristics of the US corporate sector. In addition, one must take account of interactions among different provisions, some of which may be subtle and not even intended. Finally, one must decide how to address the possibility that monetary and fiscal policy will be altered in the future to deal with projected deficit increases.

There are other important questions not even addressed in the recent debate and analysis. For example, even if workers gain as a group from the legislation, the recent growth in earnings inequality highlights that one should not think about wage and salary earners as a monolithic group. Whether through differences in rent sharing across the income distribution, or differences in capital-labor

[15] The Council of Economic Advisors (2018) argues that the reduction in profit shifting by US corporations induced by a lower corporate tax rate would lead to additional rent-sharing by US workers. This analysis suggests that the workers' share of rents depends on the magnitude of those rents reported as domestic US profits, rather than a firm's overall profits: for example, the argument is that if a US multinational shifted its reported profits from tax havens to the United States, or repatriated earnings recorded abroad, the higher measured US profitability would directly benefit workers.

complementarity that lead to differences in gains (and losses) from capital deepening, the effects of corporate taxation on different groups of wage earners is another direction in which distributional analysis needs to develop.

■ *The author thanks Rosanne Altshuler, Tom Barthold, Dhammika Dharmapala, Gordon Hanson, Jim Hines, Enrico Moretti, Joel Slemrod, Timothy Taylor, and Owen Zidar for comments on earlier drafts.*

## References

**Apple Inc.** 2017. Form 10-K (Annual Report).

**Arulampalam, Wiji, Michael P. Devereux, and Giorgia Maffini.** 2012. "The Direct Incidence of Corporate Income Tax on Wages." *European Economic Review* 56(6): 1038–54.

**Auerbach, Alan J.** 2006. "Who Bears the Corporate Tax? A Review of What We Know." In *Tax Policy and the Economy,* vol. 20, edited by James M. Poterba, 1–40. MIT Press.

**Auerbach, Alan J.** 2007. "Why Have Corporate Tax Revenues Declined? Another Look." *CESifo Economic Studies* 53(2): 153–71.

**Auerbach, Alan J.** 2017. "Demystifying the Destination-Based Cash-Flow Tax." *Brookings Papers on Economic Activity* no. 2, pp. 409–432.

**Auerbach, Alan J., and Michael P. Devereux.** 2018. "Cash-Flow Taxes in an International Setting." *American Economic Journal: Economic Policy* 10(3): 69–94.

**Auerbach, Alan J., and Joel Slemrod.** 1997. "The Economic Effects of the Tax Reform Act of 1986." *Journal of Economic Literature* 35(2): 589–632.

**Barro, Robert J., and Jason Furman.** Forthcoming. "The Macroeconomic Effects of the 2017 Tax Reform." *Brookings Papers on Economics Activity.*

**Becker, Johannes, and Nadine Riedel.** 2012. "Cross-Border Tax Effects on Affiliate Investment—Evidence from European Multinationals." *European Economic Review* 56(3): 436–50.

**Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline.** 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 36(S1): S13–70.

**Clausing, Kimberly A.** 2013. "Who Pays the Corporate Tax in a Global Economy?" *National Tax Journal* 66(1): 151–84.

**Congressional Budget Office.** 2012. *The Distribution of Household Income and Federal Taxes, 2008 and 2009.* July.

**Congressional Budget Office.** 2014. *Taxing Capital Income: Effective Marginal Tax Rates Under 2014 Law and Selected Policy Options.* December.

**Congressional Budget Office.** 2018. *The Budget and Economic Outlook: 2018 to 2028.* April.

**Council of Economic Advisers.** 2017. "Corporate Tax Reform and Wages: Theory and Evidence." October.

**Cronin, Julie Anne, Emily Y. Lin, Laura Power, and Michael Cooper.** 2013. "Distributing the Corporate Income Tax: Revised U.S. Treasury Methodology." *National Tax Journal* 66(1): 239–62.

**de Mooij, Ruud, and Li Liu.** 2018. "At a Cost: the Real Effects of Transfer Pricing Regulations." Working Paper No. 18/69, International Monetary Fund.

**Desai, Mihir A., C. Fritz Foley, and James R. Hines Jr.** 2009. "Domestic Effects of the Foreign Activities of US Multinationals." *American Economic Journal: Economic Policy* 1(1): 181–203.

**Devereux, Michael P., and Rachel Griffith.** 1998. "Taxes and the Location of Production: Evidence from a Panel of U.S. Multinationals." *Journal of Public Economics* 68(3): 335–67.

**Dharmapala, Dhammika.** 2014. "What Do We Know about Base Erosion and Profit Shifting? A Review of the Empirical Literature." *Fiscal Studies* 35(4): 421–48.

**Dharmapala, Dhammika, C. Fritz Foley, and Kristin J. Forbes.** 2011. "Watch What I Do, Not What I Say: The Unintended Consequences of the Homeland Investment Act." *Journal of Finance* 66(3): 753–87.

**Fuest, Clemens, Andreas Peichl, and Sebastian Siegloch.** 2018. "Do Higher Corporate Taxes Reduce Wages? Micro Evidence from Germany." *American Economic Review* 108(2): 393–418.

**Furman, Jason.** 2017. "No, the GOP Tax Plan Won't Give You a $9,000 Raise." *Wall Street Journal*, October 22.

**Furman, Jason, and Peter Orszag.** 2015. "A Firm-Level Perspective on the Role of Rents in the Rise in Inequality." Unpublished paper, October 16.

**Gravelle, Jane G.** 2016. *Corporate Tax Integration and Tax Reform*. Report R44638, September 16. Congressional Research Service.

**Gravelle, Jennifer.** 2013. "Corporate Tax Incidence: Review of General Equilibrium Estimates and Analysis." *National Tax Journal* 66(1): 185–214.

**Harberger, Arnold C.** 1962. "The Incidence of the Corporation Income Tax." *Journal of Political Economy* 70(3): 215–40.

**Hassett, Kevin A., and Aparna Mathur.** 2015. "A Spatial Model of Corporate Tax Incidence." *Applied Economics* 47(13): 1350–65.

**Hines, James R., Jr.** 2017. "Business Tax Burdens and Tax Reform." *Brookings Papers on Economic Activity*, no. 2, pp. 449–71.

**Joint Committee on Taxation.** 1993. *Methodology and Issues in Measuring Changes in the Distribution of Tax Burdens*, June 14.

**Joint Committee on Taxation.** 2016. Letter from Thomas Barthold, Chief of Staff, to the Honorable Kevin Brady and Richard Neal, August 31.

**Joint Committee on Taxation.** 2017a. "Estimated Budget Effects of the Conference Agreement for H.R. 1, the 'Tax Cuts and Jobs Act.'" JCX-67-17, December 18.

**Joint Committee on Taxation.** 2017b. "Macroeconomic Analysis of the Conference Agreement for H.R. 1, the 'Tax Cuts and Jobs Act.'" JCX-69-17, December 22.

**Keen, Michael.** 2001. "Preferential Regimes Can Make Tax Competition Less Harmful." *National Tax Journal* 54(4): 757–62.

**Kotlikoff, Laurence, and Lawrence H. Summers.** 1987. "Tax Incidence." In *Handbook of Public Economics,* vol. 2, edited by Alan Auerbach and Martin Feldstein, 1043–1092. Elsevier.

**Liu, Li, and Rosanne Altshuler.** 2013. "Measuring the Burden of the Corporate Income Tax under Imperfect Competition." *National Tax Journal* 66(1): 215–38.

**Looney, Adam, and Aaron Krupkin.** 2017. "9 Facts about Pass-through Businesses." Tax Policy Center, May 15.

**Macroeconomic Advisers.** 2018. *The Macroeconomic Impact of the Tax Cuts and Jobs Act*, February 6. https://bea.gov/about/pdf/acm/2018/The-Macroeconomic-Impact-of-the-Tax-Cuts-and-Jobs-Act.pdf.

**Overesch, Michael.** 2009. "The Effects of Multinationals' Profit Shifting Activities on Real Investments." *National Tax Journal* 62(1): 5–23.

**Penn Wharton Budget Model.** 2017. "The Tax Cuts and Jobs Act, As Reported by Conference Committee (12/15/17): Static and Dynamic Effects on the Budget and the Economy." December 18. http://budgetmodel.wharton.upenn.edu/issues/2017/12/18/the-tax-cuts-and-jobs-act-reported-by-conference-committee-121517-preliminary-static-and-dynamic-effects-on-the-budget-and-the-economy?rq=the%20tax%20cuts%20and%20jobs%20act.

**Rosenthal, Steven M., and Lydia S. Austin.** 2016. "The Dwindling Taxable Share of U.S. Corporate Stock." *Tax Notes* 151(6): 923–34.

**Shell, Adam.** 2018. "List of Companies that Paid Bonuses or Boosted Pay since Tax Bill Passed." *USA Today*, January 11, 2018; updated January 25, 2018.

**Statistics of Income.** 2013. *Corporation Income Tax Returns*. Washington, DC: Internal Revenue Service.

**Suárez Serrato, Juan Carlos, and Owen Zidar.** 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review* 106(9): 2582–2624.

**Ways and Means Committee, US House of Representatives.** 2016. *A Better Way: Our Vision for a Confident America*. June, 24. http://abetterway.speaker.gov/_assets/pdf/ABetterWay-Tax-Policy-Paper.pdf.

**Zucman, Gabriel.** 2014. "Taxing across Borders: Tracking Personal Wealth and Corporate Profits." *Journal of Economic Perspectives* 28(4): 12–48.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 121–146*

# Outside the Box: Unconventional Monetary Policy in the Great Recession and Beyond

Kenneth N. Kuttner

In November 2008, the Federal Reserve faced a deteriorating economy and a financial crisis. The federal funds rate had already been reduced to virtually zero. Thus, the Federal Reserve turned to unconventional monetary policies. Through "quantitative easing," the Fed announced plans to buy mortgage-backed securities and debt issued by government-sponsored enterprises. Subsequent purchases would eventually lead to a five-fold expansion in the Fed's balance sheet, from $900 billion to $4.5 trillion, and leave the Fed holding over 20 percent of all mortgage-backed securities and marketable Treasury debt (as reported in the Fed's Z.1 release, table L.211, and Treasury Bulletin, table OFS-1). In addition, Fed policy statements in December 2008 began to include explicit references to the likely path of the federal funds interest rate, a policy that came to be known as "forward guidance."

The Fed ceased its direct asset purchases in late 2014. Starting in October 2017, it has allowed the balance sheet to shrink gradually as existing assets mature. From December 2015 through June 2018, the Fed has raised the federal funds interest rate seven times.

Thus, the time is ripe to step back and ask whether the Fed's unconventional policies had the intended expansionary effects—and by extension, whether the Fed should use them in the future.

■ *Kenneth N. Kuttner is Professor of Economics, Williams College, Williamstown, Massachusetts, and Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. His email address is kenneth.n.kuttner@williams.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.121

doi=10.1257/jep.32.4.121

The aim of this paper is to take stock of what we have learned about unconventional monetary policy in the nine years since its inception, and to highlight some open questions. It begins with a review of the key features of unconventional policy. Next, it discusses the transmission of unconventional policy to financial markets, institutions, and the economy more broadly. Then it addresses the question of effectiveness with a selective survey of empirical work on the financial and economic impact of these policies, and it takes up the issue of the policies' unintended side effects. The paper concludes with some thoughts on the shape unconventional monetary policy might take in the future.

## What Were the Unconventional Federal Reserve Policies?

### Quantitative Easing

Quantitative easing refers to a set of four asset purchase programs: the three Large-Scale Asset Purchases (LSAPs), commonly known as QE1, QE2, and QE3; and the Maturity Extension Program (MEP), also known as the second "Operation Twist."[1] Table 1 summarizes the key features of these programs.

QE1 was announced in November 2008.[2] Initially, it was limited to purchasing $100 billion of debt issued by the government-sponsored enterprises Fannie Mae, Freddie Mac, and Ginnie Mae, plus $500 billion in agency-backed mortgage-backed securities.[3] Its stated purpose was to "reduce the cost and increase the availability of credit for the purchase of houses . . ."[4] On March 18, 2009, the Federal Open Market Committee announced that it would expand its purchases of agency debt and mortgage-backed securities, and would also purchase $300 billion of longer-term Treasury securities "to help improve conditions in private credit markets" more generally.[5]

QE2 was announced on November 3, 2010. The program entailed the purchase of $600 billion in longer-term Treasuries, but no agency debt or mortgage-backed securities.

The Maturity Extension Program was announced on September 21, 2011. The program initially involved the purchase of $400 billion of 6- to 30-year Treasuries, accompanied by the sale of the same quantity of 1- to 3-year securities, with the intention "to put downward pressure on longer-term interest rates and help

[1] The first "Operation Twist" was a short-lived episode in 1961.

[2] Excluded from the list of quantitative easing episodes that follow are the assets acquired by the Federal Reserve in its capacity as lender of last resort, such as the asset-backed commercial paper purchased as part of the Commercial Paper Funding Facility, which was operated from October 2008 to February 2010 in an effort to avert a liquidity crisis.

[3] To put this into perspective, in the five years prior to the crisis, the Fed would purchase $2.75 billion of Treasury securities in a typical month.

[4] Press Release, November 25, 2008, at https://www.federalreserve.gov/newsevents/pressreleases/monetary20081125b.htm.

[5] Press Release, March 18, 2009, at https://www.federalreserve.gov/newsevents/pressreleases/monetary20090318a.htm.

*Table 1*
**Characteristics of the Four Asset Purchase Programs**

| *Program* | *Dates* | *Assets purchased* | *Size (billions)* | *Sterilized?* |
|---|---|---|---|---|
| First LSAP (QE1) | 11/2008 to 3/2009 | Agency debt | $200 | No |
| | | Agency MBSs | $1,250 | |
| | | Treasuries | $300 | |
| Second LSAP (QE2) | 11/2010 to 6/2011 | Longer-dated Treasuries | $600 | No |
| MEP ("Twist") | 9/2011 to 12/2012 | 6- to 30-year Treasuries | $667 | Yes |
| Third LSAP (QE3) | 9/2012 to 10/2014 | MBSs | $40/month | No |
| | 12/2012 to 10/2014 | Longer-dated Treasuries | $45/month | |

*Note:* Quantitative easing refers to a set of four asset purchase programs: the three Large-Scale Asset Purchases (LSAPs), commonly known as QE1, QE2, and QE3; and the Maturity Extension Program (MEP), also known as the second "Operation Twist." The table summarizes the key features of these programs. MBSs are mortgage-backed securities.

make broader financial conditions more accommodative."[6] The Fed announced an extension of the program June 20, 2012, which ultimately amounted to $667 billion. In contrast to the three large-scale asset purchases, all of which entailed balance sheet expansions, this program "sterilized" the asset purchases with offsetting asset sales, leaving unchanged the overall size of the balance sheet.

QE3, which commenced in September 2012, initially involved the purchase of $40 billion per month of mortgage-backed securities in a renewed effort to "support mortgage markets." In December 2012, the program was expanded to include $45 billion per month of Treasury securities. Unlike the other three quantitative easing policies, QE3 was open-ended and did not set a dollar limit at the time of the program's launch.

These quantitative easing policies differ in clear ways from conventional monetary policy. For example, Figure 1 shows that quantitative easing drastically enlarged and altered the composition of the Fed's System Open Market Account portfolio. In contrast, the quantitative aspects of conventional policy, in terms of the Fed's balance sheet or the money supply, had always been negligible. The magnitude of the open market operations (essentially, temporary asset purchases) required to move the federal funds rate was vanishingly small—virtually undetectable in the Fed balance sheet (Friedman and Kuttner 2010).

[6]Press Release, September 21, 2011, at https://www.federalreserve.gov/newsevents/pressreleases/monetary20110921a.htm.

*Figure 1*
**The Composition of the Federal Reserve System Open Market Account Portfolio**
*(in trillions of dollars)*

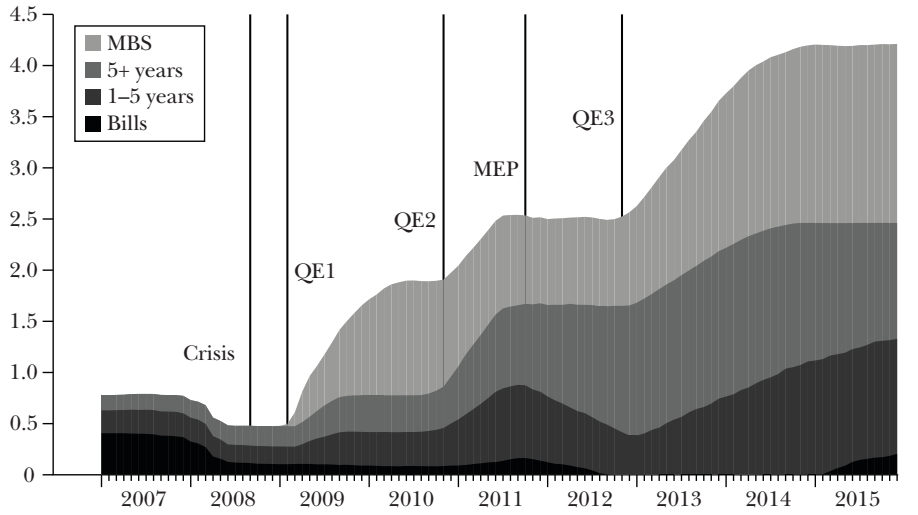


*Note:* Excludes assets associated with temporary liquidity facilities and US Treasury floating rate notes. "MBS" stands for mortgage-backed securities; "5+ years" stands for Treasuries with maturities of 5 or more years; "1–5 years" stands for Treasuries with maturities of 1–5 years. QE1, QE2, and QE3 are three quantitative easing programs. MEP is the Maturity Extension Program.

Another difference is that the goal of quantitative easing was not stated in terms of an explicit interest rate target.[7] And because a $100 billion purchase of mortgage-backed securities is not necessarily equivalent to a $100 billion sterilized purchase of 10-year Treasuries, it is not straightforward to distill the effects of the various quantitative easing programs into an interest rate equivalent.

A common misconception is that the purpose of quantitative easing was to increase bank reserves and the money supply. The Fed's pronouncements clearly contradict this view. For example, in the December 16, 2008, meeting of the Federal Open Market Committee, then-Fed Chair Ben Bernanke characterized the approach of the Bank of Japan as based on the theory "that providing enormous amounts of very cheap liquidity to banks ... would encourage them to lend and that lending, in turn, would increase the broader measures of the money supply ..." Contrasting this with the Fed's approach, Bernanke stated, "[W]hat we are doing is different from quantitative easing because, unlike the Japanese focus

[7] In this respect, the Fed's version of quantitative easing differs from the Bank of Japan's current "QQE with Yield Curve Targeting" policy, and from a proposal originally floated by Ben Bernanke (2002).

on the liability side of the balance sheet, we are focused on the asset side of the balance sheet."[8]

**Forward Guidance**

The Fed's conventional modes of communication were already providing markets with a great deal of information relevant to forming expectations about future policy expectations. Statements and minutes of the Fed Open Market Committee included assessments of economic conditions, for example, along with the economic projections of board members and regional bank presidents. What distinguished forward guidance was its explicit reference to the likely path of the target interest rate. The tactic sought to communicate a lengthening of the anticipated period of time over which interest rates were likely to remain low.

The early forward guidance statements were qualitative and vague. The December 16, 2008, statement said that rates were likely to remain low for "some time." The March 18, 2009, statement referred to an "extended period." The statements used the word "anticipate" and were conditioned on unspecified "economic conditions." In 2011, forward guidance began to involve calendar-based statements and explicit time horizons. But the horizons were repeatedly extended as the economy languished, and continued to be framed in terms like "are likely" and conditioned on economic developments.

In the Federal Open Market Committee statement of December 12, 2012, forward guidance became more explicit. It said that the low interest rate policy would remain in place so long as unemployment remained above 6.5 percent and the inflation forecast was below 2.5 percent.

With the unemployment rate at 6.7 percent in December 2013, the Federal Open Market Committee began to include, in its policy statement, language indicating its intention to keep the federal funds rate low "well past the time that the unemployment rate declines below 6-1/2 percent."[9] As time progressed, the reversion to qualitative, open-ended forward guidance led to considerable speculation regarding the date of the first rate increase. "Lift-off" eventually occurred 18 months after the unemployment rate crossed the 6.5 percent threshold, by which time the rate had declined to 5 percent.

## Monetary Policy Transmission

Actions by the Federal Reserve affect a constellation of interest rates and asset prices, which in turn influence spending decisions by households and firms, and

[8]Bernanke's distinction notwithstanding, I will follow common usage in this paper in referring to the Fed's policies as "quantitative easing." The transcript of the meeting is at https://www.federalreserve.gov/monetarypolicy/files/FOMC20081216meeting.pdf.

[9]Press Release, December 18, 2013, https://www.federalreserve.gov/newsevents/pressreleases/monetary20131218a.htm.

lending decisions by financial institutions. Many of these mechanisms, although not all, operate in the same way under conventional and unconventional monetary policies. But the arrival of unconventional policies has prompted a reexamination of the linkages between monetary policy and financial markets and led to renewed interest in models characterized by imperfect substitutability between assets.

**The Transmission of Conventional Monetary Policy**

Before the federal funds rate was reduced to virtually zero in late 2008, it was the sole tool of US monetary policy. However, little or no economic activity depends directly on the funds rate, as it applies only to overnight borrowing and lending between banks. Instead, the funds rate affects spending indirectly, through a number of distinct channels.

One is through the interest rates on longer-maturity obligations, such as mortgages and corporate bonds, which are more relevant to spending decisions than the overnight funds rate. Interest rates also affect the prices of assets, such as equities and houses, creating wealth effects that influence households' spending decisions. Similarly, interest rate changes affect imports and exports through their impact on the exchange rate.

It is important to note that long-term rates, asset prices, and the exchange rate depend on the market's forecast of *future* short-term rates, not just the current funds rate target. Therefore, Fed communication—announcements, speeches, press conferences, and the like—will affect spending to the extent that they provide information about the likely path of future policy.

Conventional policy can also affect spending through the banking system. In the traditional bank lending channel advanced by Kashyap and Stein (1994), the increase in bank reserves associated with expansionary policy increases loan supply. Moreover, for a bank that finances long-term assets with short-term liabilities, a rate reduction will increase the market value of its equity, promoting lending. (Working in the opposite direction, lower rates crimp banks' net interest margin, which tends to reduce loan supply.)

Finally, in the credit channel described by Bernanke and Gertler (1995), expansionary policy ameliorates informational frictions and reduces firms' external finance premiums, thus enhancing the real effects of rate cuts.

**The Transmission of Forward Guidance**

Forward guidance affects interest rates and asset prices by conveying information about the likely trajectory of future interest rates. In that respect, it does not differ qualitatively from other forms of Fed communication that hint at future policy. The main difference is that the interest rate path communicated as part of forward guidance was more explicit than under the conventional policy regime.

There are two reasons why forward guidance may affect interest rate expectations. One interpretation, dubbed "Odyssean" by Campbell, Evans, Fisher, and Justiniano (2012), is that forward guidance would commit the Fed to pursuing the time-inconsistent policy of allowing the inflation rate to exceed the Fed's objective

for some period of time. A credible commitment to higher inflation in the future would reduce future short-term real interest rates (Eggertsson and Woodford 2003). Odyssean forward guidance is therefore unambiguously expansionary.

Alternatively, forward guidance may convey information without implying a commitment, the case Campbell et al. (2012) referred to as "Delphic." There are two possibilities as to the type of information that could be transmitted. One possibility is that an expansionary forward guidance announcement reveals to the private sector proprietary Fed information that the economy is weaker than previously thought, which in turn implies that interest rates are likely to remain low for a longer time. However, as noted by Woodford (2012), if current real expenditures depended on expected future income, then an announcement that led to a more downbeat view of the economy could be contractionary.

A second way in which forward guidance could affect expectations is by communicating information about the Fed's policy rule. This channel may be especially important when markets had no clear sense of how economic conditions would affect how long interest rates would remain near zero. Consistent with this view, using information gleaned from the New York Fed's surveys of primary dealers, Femina, Friedman, and Sack (2013) showed that successive forward guidance statements pushed back the date of the expected first interest rate increase. Also consistent with this view is the finding by Swanson and Williams (2014) of a decreased sensitivity, beginning in late 2011, of medium-term interest rates to macroeconomic news.

**The Transmission of Quantitative Easing**

Quantitative easing entails the use of the Fed's balance sheet to influence long-term and private sector interest rates. This could occur through three mechanisms: imperfect substitutability, signaling about future policy, and improvements in financial balance sheets.

If assets are perfect substitutes, then arbitrage will mean that all assets have equal expected returns. But with imperfect substitutability, each asset class has its own downward-sloping demand curve, allowing changes in the relative supplies of assets to affect prices and yields. This supply-and-demand mechanism is what accounts for portfolio balance effects that were integral to macro models from the 1960s and 1970s, such as those developed by Tobin (1963).

Imperfect asset substitutability may arise from two sources. One comes from the fact that the prices of long-maturity bonds are more sensitive to interest rate fluctuations than those with shorter maturities. Investors with an aversion to interest rate risk will require a higher expected return on long-term bonds, relative to what they would have earned from investing in short-term debt (a "term premium"). Using asset purchases to reduce the supply of long-term bonds should therefore lower their yields by narrowing the term premium.

Market segmentation can also underpin imperfect substitutability. This may arise from investors' preferences for specific types of assets or "preferred habitats" (as hypothesized by Modigliani and Sutch 1966), or by incentives that investors have

to hold a minimum share of portfolios in a certain form like securities free from default risk. Vayanos and Vila (2009), for example, developed a model incorporating features of both preferred habitat and portfolio balance models.

Quantitative easing could also affect interest rates by sending a signal about future policy. The idea is that significant purchases of long-maturity bonds signal the Fed's intention to keep the policy interest rate near zero for a longer period of time. As with forward guidance, there are both Delphic and Odyssean interpretations of how the signaling channel could operate. One Delphic view is that asset purchases reveal a downgrading of the Fed's view of economic conditions, and thus should lead to expectations of lower future rates. Another is that signaling conveys information about a change in the Fed's policy rule—for example, that it is placing a higher weight on unemployment or lower-than-intended inflation. The Odyssean interpretation is that a large balance sheet would provide a strong incentive for the Fed to maintain a highly expansionary policy for a longer period of time than it might otherwise have desired, perhaps because the Fed would want to sell off many of the assets it owns before raising rates.

In addition to putting downward pressure on interest rates, asset purchases also may have stimulated spending by increasing loan supply. The purchases effectively raised banks' capital ratios by increasing the value of the existing assets on their balance sheets. In addition, the purchases of mortgage-backed securities (especially under QE1, when many investors were anxious to reduce their exposure to housing-related risk) increased the liquidity of the market for those securities. Both mechanisms would have made banks more willing to lend.

## Unconventional Monetary Policy and Interest Rate Effects

The main challenge in assessing the impact of monetary policy is isolating exogenous policy changes that can be used to identify causal effects.[10] In the study of conventional monetary policy, the monetary policy "shocks" used to identify the causal effects of changes in the federal funds rate are typically modeled as deviations from the Fed's normal response to economic conditions, most commonly derived from a structural vector autoregression econometric model.

Assessing the impact of unconventional policy is more difficult than it is for conventional policy, for at least two reasons. First, it is not clear what variable to use as a summary measure of monetary policy, given the heterogeneity of the asset purchases and differences in the framing of the forward guidance announcements. Second, defining "shocks" is problematic. Because the financial crisis was such a singular event, it is hard to know what the Fed's "normal" response to it would have been. And in any case, in gauging the macroeconomic effects of unconventional

[10] See Nakamura and Steinsson (in the Summer 2018 issue of this journal) for an in-depth discussion of the identification issues bedeviling efforts to measure the effects of monetary policies.

policy, the comparison to a "no policy" counterfactual will be more relevant than one that looks at deviations from the usual policy rule.

Given these obstacles, it is not surprising that research on quantitative easing and forward guidance has tended to focus narrowly on how such policies affect interest rates on Treasury bonds and mortgage-backed securities, rather than on their ultimate macroeconomic impact.[11] The two most common approaches to assessing the interest rate effects are event studies using high-frequency data and time series models of term premiums, both of which have their limitations.

**Event studies**

A typical event study for estimating the effects of unconventional monetary policies on interest rates examines changes in bond yields over a one- or two-day window around which the policies are announced. This approach relies on two identifying assumptions. The first is that the announcement was unanticipated. This seems plausible for the early stages of the first large-scale asset purchases. However, lacking a market-based measure of financial markets' expectations, such as the prices of federal funds futures I used in Kuttner (2001), there is no satisfactory way to confirm this. Subsequent large-scale asset purchases and the Maturity Extension Program may have been anticipated to some extent, in which case, the measured financial market reactions in the few days around the announcement of a policy may understate their true effects.

The second key assumption is that the announcement was not interpreted as revealing the Fed's proprietary information about the state of the economy, which in turn would have affected bond yields. This could be problematic, in light of the Campbell et al. (2012) finding that expansionary policy surprises have historically been associated with upward revisions in private-sector unemployment rate forecasts.

Table 2 summarizes some estimates of cumulative effects from a selection of event studies. The results vary somewhat across studies, due to differences in the length of the event window, the choice of interest rate data, and the selection of events, but all tell roughly the same story.

The most salient result is that the QE1 announcements had very large, negative effects on long-term interest rates: approximately 100 basis points for Treasuries and mortgage-backed securities and upwards of 150 basis points (depending on the horizon) for agency issues. The reactions represent extreme tail events, the largest one-day changes observed in the entire post-crisis period. The effects of subsequent programs on yields were materially smaller. The estimated two-day effects of the second large-scale asset purchase announcements are in the –30 to –40 basis point range with comparable figures for the Maturity Extension Program. The QE3 announcements appear to have had only a small impact on yields.

[11] The literature on the interest-rate and economic effects of unconventional monetary policy is vast, and the studies mentioned here are intended to illustrate main themes, not to offer a literature review. Bhattarai and Neely (2016) provide a more comprehensive survey.

*Table 2*
**Estimated Event-Study Interest Rate Effects**

| *Study* | *Window (days)* | *Yield on:* | *QE1 (basis points)* | *QE2 (basis points)* | *MEP (basis points)* | *QE3 (basis points)* |
|---|---|---|---|---|---|---|
| Gagnon, Raskin, Remache, and Sack (2011) | 1 | T10 | –91*** | | | |
| | | Agency | –156*** | | | |
| | | MBS | –113*** | | | |
| Krishnamurthy and Vissing-Jorgenson (2011) | 2 | T10 | –107* | –30*** | | |
| | | Agency | –200*** | –29*** | | |
| | | MBS | –88 | –13** | | |
| Ehlers (2012) | 1 | T10 | | –14 | –27*** | |
| | 2 | T10 | | –40*** | –46*** | |
| Bauer and Neely (2014) | 1 | T10 | –123** | –23 | | –14 |

*Notes:* "T10" refers to the 10-year Treasury, MBS to the 15-year Agency mortgage-backed securities, and "Agency" to the debt issued by Ginnie Mae, Fannie Mae, and/or Freddie Mac. QE1, QE2, and QE3 are three quantitative easing programs. MEP is the Maturity Extension Program. Asterisks indicate the magnitude of the ratio of the observed event-day relative to the standard deviation of the yield changes at the indicated horizon, as reported by the authors:
***denotes ratios greater than 2.58 in absolute value (1 percent tail),
**ratios greater than 1.96 (5 percent tail), and
*greater than 1.69 (10 percent tail).

Taken together, the event studies suggest that the four policies' cumulative effects on the 10-year Treasury yield totaled at least –150 basis points. The evidence should be interpreted with caution, however. There are five reasons why the results could be inaccurate or not fully generalizable to other situations.

First, QE1 was launched at a time of high stress levels in financial markets. The initial November 25, 2008, announcement cited widening spreads on the debt of government-sponsored entities and on the mortgages they guaranteed. It stated that the action was being "taken to reduce the cost and increase the availability of credit for the purchase of houses," saying nothing about long-term interest rates more broadly. Similarly, the December 16, 2008, minutes of the Federal Open Market Committee called attention to soaring risk spreads on corporate bonds and rising premiums for on-the-run (most recently issued) Treasuries, and described the functioning of Treasury markets as "impaired." Therefore, much of the impact of the first large-scale asset purchases probably came from a restoration of market functioning, rather than a reduction in either expected future interest rates or the term premium.

Second, several announcements of quantitative easing also contained forward guidance. Most conspicuously, the December 16, 2008, and March 18, 2009, announcements both stated an intention to keep the federal funds rate at "exceptionally low levels." Some efforts to disentangle these effects are discussed below.

Third, the paucity of announcements means that the results are sensitive to individual observations. For example, the 51 basis point drop in the 10-year Treasury yield on March 18, 2009, is, by a wide margin, the largest in the past 20 years (the runner-up is only –28 basis points). Excluding this observation reduces the estimated impact of QE1 by more than half. Moreover, the small number of observations is an invitation to "cherry pick" dates, and studies that that find a reason to exclude observations with small or perverse reactions are likely to be biased towards finding larger effects.

Fourth, the statistical precision of the event study approach is unclear. If one makes the dubious assumption of equal variance on event and non-event days, then it would be legitimate to use the variance of non-event-day changes in assessing the precision of the estimated effects. Dropping this assumption requires using only event days to calculate the variance, which is problematic given the small number of observations.[12]

Fifth, and perhaps most important, it can be hard for an event study to measure persistence. It may take some time before changes in asset supplies are fully reflected in prices and yields (Greenwood, Hanson, and Liao 2018). The dilemma is that an event window of sufficient length to account for a gradual response will include "noise" resulting from the arrival of additional information and events, making it less likely to discern a statistically significant impact of the policy. The findings for QE1 and QE2 summarized in Table 2 are *so* large, however, that they remain clearly discernible (in the sense that the cumulative responses exceed two standard deviations) for at least one or two weeks.

Assessing the policies' persistence at longer horizons requires imposing a parametric structure on the responses. In an effort to get at the persistence issue, Wright (2012) estimated a vector autoregression on daily data encompassing all four of the quantitative easing programs (but not distinguishing observations according to whether they were associated with forward guidance statements). He detected measurable responses over several weeks, but found that the effects wore off after two to three months.

In another effort, Swanson (2017) addressed the issue of persistence by using a two-factor model to differentiate between the effects of forward guidance and quantitative easing, and also fitted an exponential function to the responses as a way to parameterize the rate of decay. Like Wright (2012), he found that the effects of both policies were relatively short-lived. He also found that dropping the outsize reaction of March 18, 2009, significantly decreased the magnitude but increased the persistence of the effects of the large-scale asset purchases (again illustrating the fragility of results based on a small number of announcements).

[12] For example, the standard deviation of the cumulative effect of the eight QE1 announcements on the 10-year Treasury yield is 58 basis points. Using the *t*-distribution with 7 degrees of freedom, this gives a 95 percent confidence interval ranging from –20 to 208 basis points.

*Figure 2*
**Kim–Wright Estimated 10-year Term Premium and 10-year Treasury Yield**
*(percent)*

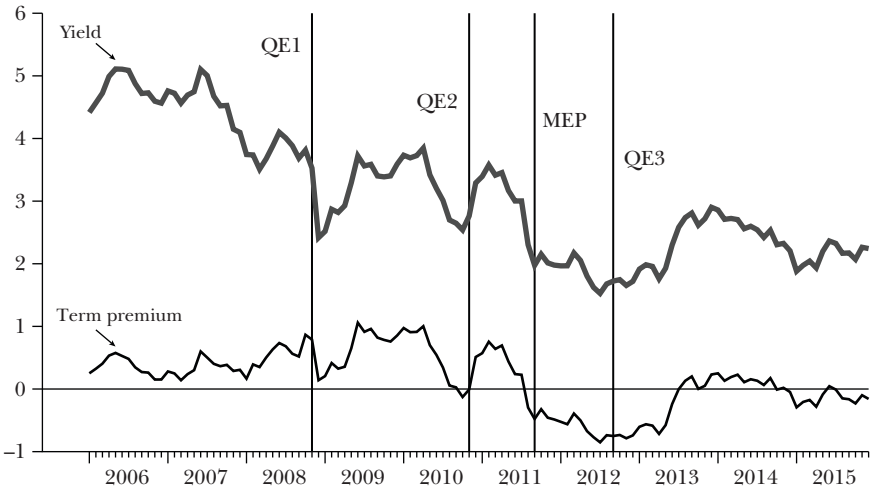


*Note:* QE1, QE2, and QE3 are three quantitative easing programs. MEP is the Maturity Extension Program.

**Time Series Analysis of Term Premiums**

Time series econometric methods can be also used to assess the effects of the large-scale asset purchases on bond yields—and in particular on term premiums. Term premiums cannot be observed directly, however, so estimating the policies' effects requires the additional step of fitting a term structure model to the data.

The "affine term structure models" used for this purpose involve specifying the vector of bond yields over different term structures as a function of a small number of factors, which are assumed to follow a first-order vector autoregressive process. The one-period risk-free interest rate is assumed to be a function of the same factors. The structure means that all co-movements between bond returns of different terms are attributed to the factors, and further implies that only the risk associated with those factors is priced.

Figure 2 plots the fitted 10-year term premium, interpretable as an estimate of the difference between the 10-year Treasury yield and the average of forecast short-term interest rates over the life of the bond, derived from the Kim–Wright (2005) method. Already quite low by historical standards prior to the financial crisis, the term premium declined by approximately 200 basis points from mid-2009 to mid-2012. The premium actually fell into *negative* territory, implying that investors were willing to sacrifice some return for the hedge provided by 10-year Treasuries. The yield and term premium fell more or less in lockstep over the quantitative easing period, and the correlation between monthly changes is 0.97. It seems that that a shrinking term premium accounts for almost the entire

*Table 3*
**Estimated Effects of Quantitative Easing on 10-year Term Premiums**
*(basis points)*

| *Study* | *QE1* | *QE2* | *MEP* | *QE3* |
|---|---|---|---|---|
| Gagnon, Raskin, Remache & Sack (2011) | –38[a] | | | |
| D'Amico, English, López-Salido & Nelson (2012) | –35 | –45 | | |
| Ihrig, Klee, Li, Schulte & Wei (2012) | –40 | –40 | –17 | –50[b] |
| Hamilton & Wu (2012) | | | –27[c] | |

*Notes:* QE1, QE2, and QE3 are three quantitative easing programs. MEP is the Maturity Extension Program.
[a] The smallest of the range of estimates reported.
[b] Estimated by Engen, Laubach, and Reifschneider (2015) using the Ihrig, Klee, Li, Schulte, and Wei (2012) model.
[c] The reported impact of a $400 billion maturity swap, scaled up to the $667 billion size of the Maturity Extension Program.

decline in the yield of a 10-year bond, with very little attributable to falling interest rate expectations.

The time series method has several advantages over the event study approach. First, it makes use of more information. Rather than relying on a handful of announcements, it uses the entire time path of interest rates and asset quantities. The underling analytical structure makes possible a quantitative assessment—that is, the yield change, in basis points, for a given $100 billion in asset purchases—which is hard to do in an event study framework. Also, the effects of the policy can be estimated regardless of whether asset purchase programs were anticipated.

The identifying assumption underlying this approach is that changes in supplies of assets of a specific maturity result from factors such as the Treasury's debt management or Fed portfolio allocation decisions and are otherwise unrelated to expected interest rates or term premiums. As an example, Greenwood and Vayanos (2014) cite the drop in the average maturity of outstanding Treasury debt in the 1960s and 1970s, which resulted from a 4.5 percent regulatory ceiling on bonds' coupon rates at that time. There is no evidence that either the Treasury or the Fed (at least pre-quantitative easing) adjusted asset supplies in response to term premiums, so it is probably legitimate to treat the supply variables as exogenous.

The identifying assumption would also be violated if asset supplies and term premiums were both a function of an omitted variable, such as macroeconomic conditions and/or the state of the financial system. This is a concern for the quantitative easing period, when the Fed's asset purchases were clearly an endogenous response to the deteriorating state of the economy (just as the federal funds interest rate was endogenous before quantitative easing). For this reason, studies taking this approach generally fit the models to data before quantitative easing occurred.

Table 3 summarizes the findings from four well-known studies looking at the effects of quantitative easing policies on term premiums. Taken together, the studies suggest that the policies collectively reduced the 10-year term premium by

as much as 150 basis points—remarkably similar to event-study results surveyed previously.

Gagnon, Raskin, Remache, and Sack (2011) and D'Amico, English, López-Salido, and Nelson (2012) both used reduced-form regressions of the Kim–Wright (2005) term premium on measures of relative asset supplies. The two studies' regressions differ in several respects, such the construction of the supply measures and the inclusion of control variables. Despite these differences, both studies have QE1 subtracting at least 35 basis points from the 10-year term premium. D'Amico, English, López-Salido, and Nelson (2012) put the impact of QE2 at –45 basis points.

Ihrig, Klee, Li, Schulte, and Wei (2012) extended an otherwise standard affine term structure model to include asset supplies as additional factors. Their estimates for QE1 and QE2 are quite similar to those just mentioned. They also report a sizable –50 basis point effect of QE3, reflecting the very large magnitude of the asset purchases at that time. The estimated Maturity Extension Program effects are roughly half the size of the other programs. Also employing a modified affine term structure model, Hamilton and Wu (2012) used measures of asset supplies to forecast the three factors on which the term premiums depend. They put the impact of the Maturity Extension Program at –27 basis points—somewhat larger than the Ihrig et al. (2012) estimate, but still smaller than the effect of the large-scale asset purchases.

There are several reasons to use caution in interpreting the time series results. First, estimates of the term premium can differ a great deal across models, as illustrated in Rudebusch, Sack, and Swanson (2007). Second, the confidence intervals associated with the term premium estimates are wide. As Li, Meldrum, and Rodriguez (2017) note, it is hard to estimate the long-run average yields and the parameters characterizing the speed of mean reversion.[13] Third, the term structure models assume stable parameters, which may be unwarranted during a financial crisis with unprecedented policy tools being introduced.

**What Explains the Interest Rate Declines?**

There are competing explanations for what channels were most important in connecting unconventional monetary policy and falling interest rates. In late 2008 and early 2009, improvement in market functioning probably accounted for much of the sharp initial drop in yields under QE1. Gagnon, Raskin, Remache, and Sack (2011) argue this case by citing the large spreads between mortgage-backed security and Treasury yields as symptomatic of market dysfunction prevailing at the time.

But remaining somewhat unsettled is the question of the importance of the signaling channel, working through expectations of future short-term rates, and the effects of large-scale asset purchases in leading to a rebalancing of portfolios, which would have affected term premiums. Disentangling these two is inherently difficult. Further complicating matters is the fact that several early announcements

[13] Li, Meldrum, and Rodriguez (2017) also showed that the use of professional forecasts in the Kim–Wright (2005) model ameliorates these problems.

of large-scale asset purchases, the ones associated with the most extreme market reactions, coincided with forward guidance statements.

Some inferences can be drawn using direct market-based measures of interest rate expectations. Gagnon, Raskin, Remache, and Sack (2011) found that there was no change in the one-year-ahead forward rate on December 16, 2008; and that the 28 basis point drop on March 18, 2009, was reversed shortly thereafter. Thus, they attributed the change in the yields to the large-scale asset purchases, rather than forward guidance. Similarly, Swanson's (2017) model attributed most of the March 18 yield decline to the large-scale asset purchase factor.

However, looking at the overall impact of QE1, Krishnamurthy and Vissing-Jorgenson (2011) ascribed a larger share of the market reaction to the signaling channel. Observing that the announcements were collectively associated with a 40-basis-point reduction in the two-year federal funds futures rate, they concluded that the signaling effect accounted for a nonnegligible 20–40 basis points of the 107-basis point drop in the 10-year Treasury yield. Bauer and Rudebusch (2014) reached a similar conclusion using Eurodollar futures.

Another way to address the relative importance of signaling and the expected future short-term rate, versus portfolio balance effects from large-scale asset purchases and term structure effects, is to look at the results implied by an affine term structure model. Using the Kim–Wright estimates of the term premium, Bauer and Rudebusch (2014) calculated that 22 percent of the QE1-induced reduction in the 10-year yield was attributable to signaling, with 78 percent coming from the term premium. However, the estimated impact of QE1 on conventionally estimated term premiums was very imprecise, and much larger signaling effects could not be ruled out. Their favored model (with restricted risk prices) put the contribution of the signaling effect at 36 percent (and in the 30–56 percent range), which suggests that the majority of the yield decline can be attributed to a reduction in the term premium.

An additional question relating to the transmission mechanism has to do with whether it is the *stock* of outstanding assets that affects yields, stemming from market segmentation; or the *flow* of asset purchases, which could result from transitory liquidity or market functioning effects. In an effort to address this issue, D'Amico and King (2013) study how the purchase of a specific bond affected its price, as well as those of close substitutes. Comparing yields pre- and post-QE1 and aggregating over the relevant set of bonds, they estimated a "stock effect" yield reduction of 30 basis points. Transitory "flow effects" of bond purchases were also detectable in daily data, but of a much smaller magnitude. Significantly, this micro-level evidence does not speak to the aggregate effect of reducing the supply of long-term interest-sensitive bonds ("removing duration"), implying that the overall impact of QE1 is likely to have been larger. On the other hand, the authors note that market segmentation was likely to have been stronger during the period of QE1, when financial markets were under a great deal of stress, and consequently that supply effects are likely to have been smaller during subsequent large-scale asset purchases. Using methods similar to those employed by D'Amico and King (2013), Meaning and

Zhu (2011) found that QE2 shifted the Treasury yield curve down by roughly 20 basis points—a smaller "bang for the buck," given that the volume of Treasuries purchased was twice that of QE1.

## Unconventional Monetary Policy and Effects on Economic Outcomes

The evidence discussed so far points to a meaningful impact of unconventional monetary policy. But lowering interest rates is not an end unto itself; it matters only to the extent that it affects the decisions of financial institutions, firms, and households.

In the context of unconventional monetary policy, it is especially important to be cautious about treating interest rate reductions as an end in themselves. First, in an environment of financial stress, uncertainty, and scarce investment opportunities, it is not a foregone conclusion that interest rate reductions will have the same effects on spending as at other times. Perhaps in a time of economic stress, the cost of funds is of second-order importance for potential borrowers. Second, a change in term premiums may have a smaller effect than a lowering of the expected path of future short-term interest rates. Stein (2012) argues that a risk-neutral firm might adjust its capital structure to take advantage of the lower term premium without altering its real economic decisions. Indeed, Kiley (2014) finds, using a quantitative macro model, that term premium reductions had substantively smaller expansionary effects than reductions of expected future interest rates.

Thus, in this section we discuss evidence about the effects of unconventional monetary policies on bank lending and firm behavior, and also consider some studies that try to model the overall macroeconomic effects.

### Bank Lending

Two recent papers have uncovered micro-level evidence that quantitative easing increased bank lending. Rodnyansky and Darmouni (2017) used a difference-in-difference model to study the effects of large-scale asset purchases on bank lending. They regressed loan growth on indicator variables for large-scale asset purchases, which do not vary across banks, interacted with a measure of exposure of each bank to mortgage-backed securities. They found that banks with higher initial holdings of mortgage-backed securities were more likely to increase lending following QE1 and QE3, both of which (and unlike QE2 and the Maturity Extension Program) entailed significant purchases of mortgage-backed securities.

Luck and Zimmerman (2017) provide parallel findings for total loan growth. Using data on mortgage originations and small business lending data reported by banks to comply with the Home Mortgage Disclosure Act and the Community Reinvestment Act, they were able to distinguish the policies' effects on mortgage refinancing versus commercial and industrial lending. While QE1 and QE3 both

encouraged banks to extend credit, only QE3 increased commercial and industrial lending. They also exploited spatial variation in banks' holdings of mortgage-backed securities to assess the effects of the large-scale asset purchase on county-level employment growth. The main finding is that counties whose banks had relatively large holdings of mortgage-backed securities tended to experience more rapid employment growth following QE3, relative to those with smaller exposures. The same was not true for QE1, however, whose effects were limited to mortgage refinancing.

**Firm behavior**

Using firm-level micro data, Foley-Fischer, Ramcharan, and Yu (2016) found empirical support for the hypothesis that the reduction in bond yields resulting from the Maturity Extension Program materially affected firms' financing and investment decisions. They used a difference-in-difference approach, with firms' long-term debt levels before the Maturity Extension Program as the treatment variable—the idea being that those relying more on long-term debt would have benefitted more from reductions in long-term interest rates. The identifying assumption is that firms' preference for long-term debt is exogenous, and unrelated to any factors that might have affected their response to interest rates generally, or the Maturity Extension Program specifically.

Additionally, they found that firms with a relatively heavy reliance on long-term debt experienced positive excess stock returns on September 22, 2011, the day of the announcement of the Maturity Extension Program. The program also seems to have affected firm's financing decisions. In the year following the commencement of the Maturity Extension Program, firms with high levels of long-term debt tended to issue even more of it. More importantly, a greater reliance on long-term debt was associated with larger increases in capital spending and employment following the Maturity Extension Program. The asset purchases therefore appear to have affected firms' real economic decisions, not just their capital structure.

**Macroeconomic Impact**

Ultimately, we care about the effect of quantitative easing on macroeconomic variables like GDP and the unemployment rate. A first step towards gauging its macroeconomic implications is to translate the decline in bond yields into an equivalent reduction in the federal funds rate. Previous studies, such as Kuttner (2001), have found that a 100 basis point surprise cut in the funds rate target results in a reduction in the 10-year yield of approximately 33 basis points. Using this as a rule of thumb, it would have taken 450 basis points of funds rate cuts to produce the 150 basis point reduction in the Treasury yield that seems to have resulted from quantitative easing.

A more rigorous approach is to use a term structure model to back out the value of the (negative) latent federal funds rate that is consistent with the observed behavior of the term structure of interest rates. Wu and Xia (2016) propose a model of the "shadow federal funds rate" by truncating from below the distribution of forward interest rates, thus introducing a nonlinearity into what would otherwise

have been a linear relationship between forward rates and the underlying factors. According to their calculations (reported at https://www.frbatlanta.org/cqer/research/shadow_rate.aspx), the shadow federal funds rate reached a nadir of –3 percent in May 2014.

Wu and Xia (2016) then used a factor-augmented vector autoregression to assess the impact of shocks to the shadow funds rate on various measures of real activity. According to their calculations, the reduction in the shadow rate reduced the unemployment rate by a full percentage point from July 2009 to December 2013, relative to a counterfactual with no quantitative easing.

Using a very different econometric model, Engen, Laubach, and Reifschneider (2015) obtained results similar to those of Wu and Xia (2016). Feeding the 120-basis-point reduction in term premium from Ihrig, Klee, Li, Schulte, and Wei (2012) into the Federal Reserve Board's FRB/US model, they concluded that the four quantitative easing policies combined reduced the unemployment rate by 1.2 percentage points relative to what it would have been in the absence of quantitative easing.

Yet another approach to gauging the policies' aggregate effects is to use dynamic stochastic general equilibrium models that incorporate some sort of financial friction. In Gertler and Karadi (2013), the friction takes the form of limited arbitrage, either between risk-free government and privately issued risky assets, or across different maturities of risk-free assets. Quantitative easing is modeled as a policy in which the central bank steps in and performs intermediation between different assets that private financial institutions are unwilling to do. Under the assumption of a zero short-term interest rate, their calibration indicates that QE1 reduced the magnitude of the GDP contraction by 3.5 percentage points (quite substantial, relative to the actual peak-to-trough contraction of 4.3 percent), with QE2 increasing GDP by 1 percent within the span of a year. Quantitative DSGE results can be sensitive to model specification, however. For example, the simulations in Chen, Cúrdia, Vasco, and Ferraro (2011) put the impact on GDP of QE3 at only 0.4 percent, with considerably more market segmentation required to obtain larger effects.

## Side Effects of Unconventional Monetary Policy

The evidence summarized to this point supports the view that the Fed's unconventional policies largely achieved their purpose of reducing long-term interest rates and stimulating economic activity. Concerns have been raised about the possibility of adverse unintended consequences, such as inflation, financial instability, and international spillovers, but such outcomes seem to have been modest.

### Two Nonissues

One concern was that the vast expansion in bank reserves and the monetary base would be inflationary. A number of prominent economists went so far as to write in 2010 an open letter to Ben Bernanke predicting that QE2 would risk "currency debasement and inflation" (e21 Staff 2010). This outcome did not occur, of course.

Another concern was that the large balance sheet might complicate the process of "normalizing" monetary policy—that is, switching back to the use of the federal funds interest rate as the short-term interest rate. This fear also turns out to have been misplaced. As discussed by Ihrig, Meade, and Weinbach (2015), paying interest on reserves has allowed the Fed to raise short-term interest rates, even with banks holding $2.5 trillion of excess reserves.

**Risk-taking**

Less easily dismissed is the concern that unconventional monetary policy encouraged excessive risk-taking by firms and financial intermediaries. For example, while acknowledging that low interest rates are intended to encourage some risk-taking, Fed Chair (then Governor) Jerome Powell (2017) raised the question of whether or not "low rates have encouraged excessive risk-taking through the buildup of leverage or unsustainably high asset prices."

Excessive risk-taking is especially relevant to institutions, such as insurance companies, with commitments to streams of fixed future payments (Rajan 2005). It also applies to money market mutual funds, which require an interest margin of sufficient size to cover management fees. Such institutions may feel compelled to "reach for yield," investing in riskier assets in order to hit targets for investment income.[14]

Several recent studies examining the effects of quantitative easing on financial institutions find little reason for concern over additional risk-taking. Foley-Fischer, Ramcharan, and Yu (2016) found that spreads narrowed between A– rated corporate bonds and Treasury yields after the Maturity Extension Program, suggesting that insurance companies were shifting towards somewhat riskier (but still high-quality) assets. (It may also have been the case that the A– securities were perceived to have become less risky as a result of the expansionary policy.) Importantly, the effect did not extend to lower-rated bonds, which typically imposed on institutional investors a more stringent capital requirement. Thus, while some reaching-for-yield may have occurred, it certainly didn't qualify as reckless.

Focusing on banks, Kurzman, Luck, and Zimmerman (2017) found that those with higher initial holdings of mortgage-backed securities were more likely to relax lending standards following QE1 and QE3. On the face of it, this suggests riskier behavior by banks. However, observing that QE1 resulted in relatively larger gains in the value of banks laden with mortgage-backed securities, they attributed the increased lending to the improvement in the banks' capital positions. Increased liquidity of mortgage-backed securities resulting from QE3 also seems to have played a role. There is nothing to indicate that the risk-taking was excessive.

Looking at several different types of financial institutions, Chodorow-Reich (2014) examined how large-scale asset purchases might affect risks. In an event study framework, he found that for insurance companies and bank holding companies,

[14] An extensive literature, too rich to do justice to here, has examined how low interest rates affect the risk-taking of financial institutions in contexts that do not involve quantitative easing; for a survey, see De Nicolò, Dell'Ariccia, Laeven, and Vaencia (2010).

stock prices rose and spreads on credit default swaps (a proxy for market-perceived credit risk) fell immediately following the announcements of large-scale asset purchases. He attributed this to an improvement in the value of the assets already on the institutions' books, which lessened solvency concerns. He also examined, for the money market mutual funds, the relationship between fixed "structural" expenses and gross yield, which is inversely related to asset quality. He detected a statistically significant tendency for high-cost funds to reach for yield, but the effect was economically small, and dissipated by 2013.

Indeed, in a number of settings, a moderate increase in risk tolerance may be beneficial. For example, in an economy recovering from a financial crisis, some additional reaching for yield could be welfare-improving if other distortions have resulted in too *little* risk-taking (Chodorow-Reich 2014). In this case, increasing risk tolerance should promote lending and economic recovery.

Moreover, a shift of financial institutions towards riskier investments is not necessarily accompanied by a reduction in financial stability. Very low interest rates may have reduced institutions' risk in a manner that increased the value of legacy assets and net worth. In addition, the improvement in macroeconomic conditions brought about by the expansionary policy may have decreased the credit risk associated with many of those assets. In this environment, the additional risk-taking would be beneficial.

**International Spillovers**

Seven years of quantitative easing and near-zero interest rates had had far-reaching effects on other economies. The evidence in Neely (2015) and Bauer and Neely (2014) shows that the Fed's QE1 announcements significantly reduced bond yields in other developed countries by amounts roughly half that in the United States. However, the expansionary impact of the yield reductions on other countries was offset by a depreciation of the US dollar, which fell by amounts ranging from 3.5 percent for the British pound to 7.8 percent for the euro.

But the main concerns arose because with the near-zero or even negative interest rates in Japan, the United Kingdom, and the euro area, the comparatively high rates of interest in emerging market economies attracted very large capital inflows, much in the form of portfolio investment, which put pressure on their exchange rates to appreciate.

The influx of funds presented central banks of emerging market economies with a dilemma. They were reluctant to let their exchange rates appreciate, for fear that it would lead to excessive current account deficits. They were reluctant to limit their currencies' appreciation by allowing domestic interest rates to fall along with those of developed economies, because it would have led to monetary policy that was excessively expansionary. They were reluctant to try to limit appreciation by purchasing US dollars and holding a large amount of foreign exchange reserves, which would have had a high opportunity cost in a low-interest-rate environment.

A number of studies have documented how the Fed's unconventional policies affected emerging market economies. For example, using a Bayesian vector

autoregression with monthly data, Bhattarai, Chatterjee, and Park (2015) confirmed that a portion of the capital flows into emerging market economies were attributable to quantitative easing, and that the policy led to exchange rate appreciation, reduction in bond yields, and stock market booms.

Studies using high-frequency data to assess the effects of specific large-scale asset purchases also confirm that they led to capital flows into emerging market economies, although the picture arising from these studies is more nuanced. Using an event-study approach similar to those looking at bond yields, Bowman, Londono, and Sapriza (2015) found that QE1 reduced sovereign bond yields in emerging market economies, just as it did in the United States. Subsequent large-scale asset purchases had no distinguishable effect on yields in emerging market economies. The exchange rate index for emerging market economies showed no statistically significant response to any of the large-scale asset purchases, although some individual countries experienced large movements.

To assess the magnitude of the capital flows caused by the quantitative easing policies, Fratzscher, Lo Duca, and Straub (2018) used high-frequency data on flows into more than 16,000 equity and 8,000 bond funds. They regressed daily flows on a set of three variables capturing the Fed's policies: purchases of Treasury securities, liquidity operations, and indicators for the various announcements of large-scale asset purchases. They found that the effects varied a great deal across the different asset purchases and types of assets. For example, QE1 seems to have led to an *outflow* from emerging market economies bonds and into US equities, with roughly half of those funds returning after QE2. All three of the large-scale asset purchases were also associated with some inflows into emerging market economy equity funds. In the end, the composition of mutual fund flows into emerging market seems to have been affected more than the total volume.

The large volume of portfolio investment naturally raised concerns that the Fed's inevitable normalization of policy would lead to an abrupt outflow of capital. There is some evidence suggesting that this was the case, although the overall impact was less than feared.

Aizenman, Binici, and Hutchison (2016) provided some event-study evidence that news items about the Fed's intentions to unwind its large-scale asset purchases were somewhat disruptive to financial markets in emerging market economies. Specifically, they found that remarks by Ben Bernanke hinting at tapering led to exchange rate depreciation and a widening of spreads on credit default swaps among emerging market economies.[15] The response was neither uniform nor long-lived, however. Paradoxically, emerging market economies with strong fundamentals (small current account deficits, low external debt, and larger foreign exchange rate

[15] Bernanke's May 22, 2013, congressional testimony was widely blamed for precipitating the infamous "taper tantrum," in which markets reacted strongly to the news that the pace of quantitative easing might slow. Bernanke said: "As the economic outlook, and particularly the outlook for the labor market, improves in a real and sustainable way the Committee will gradually reduce the flow of purchases." At https://www.gpo.gov/fdsys/pkg/CHRG-113shrg81472/pdf/CHRG-113shrg81472.pdf.

reserves) reacted more strongly to statements hinting at future tapering than those with weak fundamentals.

## Conclusions

No study of the effects of unconventional policy is definitive, and all of those surveyed in this article have their limitations. A preponderance of evidence nonetheless suggests that forward guidance and quantitative easing succeeded in lowering long-term interest rates. Studies using micro data have documented tangible effects of quantitative easing on firms and financial intermediaries. Macro models suggest that the interest rate reductions are likely to have had a meaningful impact. The adverse side effects appear to have been mild, and are dwarfed by the costs of the more protracted recession in the United States that likely would have occurred in the absence of the unconventional policies. The benefits of unconventional policy therefore probably outweighed the costs.

Some questions are not entirely settled. First, the persistence of the effects on interest rates remains unclear. Second, disentangling the effects of quantitative easing from those of forward guidance is difficult. Third, the effects of these policies may have been in part a function of turbulent financial conditions, or may have diminished over time as the novelty wore off.

Given the uncertainties and weaknesses of the evidence, what have the past nine years taught us about the appropriate design of unconventional policies, should they be needed in the future? Six tentative lessons can be drawn from the US experience.

First, unconventional monetary policy should be conducted in a rule-like manner to the extent possible. In practice, this means clearly relating asset purchases and/or forward guidance to the Fed's objectives and forecasts. A policy articulated on a flow basis conditioned on ongoing economic developments, like QE3, is likely to be more amenable to expression in terms of a rule than one involving large, infrequent discrete adjustments to the balance sheet targets.

Second, if the research is correct in indicating that quantitative easing functions primarily through the removal of duration risk from the market, policy objectives could be accomplished either by reallocating a central bank portfolio of a fixed size or by expanding the balance sheet. Given that the purpose of quantitative easing was not to increase bank reserves, it would make sense to use portfolio reallocation as the first step in implementing quantitative easing. However, given that there have been no discernible ill effects from expanding the balance sheet (independent of any that may have resulted from very low interest rates), the unsterilized purchase of long-term bonds is a perfectly viable policy option, too.

Third, forward guidance and quantitative easing are not substitutes, as they operate through different transmission mechanisms: expectations of future interest rates for the former, the portfolio balance effect (primarily) for the latter. Thus, the two policies could be implemented independently. There could also be

complementarities between them. For example, to the extent that market participants interpreted the large-scale asset purchases as communicating the Fed's interest rate intentions, they may have reinforced the impact of forward guidance on interest rate expectations.

Fourth, a central bank that engages in large-scale asset purchases faces a large exposure if interest rates rise. This can be viewed as a positive, in the sense that the interest-rate risk could commit the Fed to a larger or more sustained monetary expansion. However, wagering central bank independence is probably best reserved for truly dire circumstances. Although it would undermine the commitment value of asset purchases, an agreement with the US Treasury that would indemnify the Fed against any losses might increase its willingness to pursue quantitative easing on a large scale.

Fifth, the appropriate choice of assets to purchase will depend on the circumstances. If asset purchases operated solely via the removal of duration risk, and if changes in Treasury yields were fully passed through to those on other debt securities, then there would be no reason to purchase any assets other than Treasuries. However, the purchase of mortgage-backed securities in QE1 was appropriate as a means to improve functioning of that market. Similarly, one can conceive of circumstances that might call for central bank purchases of other types of securities, such as corporate bonds and equities. But venturing into this territory would require an amendment to the Federal Reserve Act, and would raise a number of thorny issues (far beyond the scope of this paper), such as the appropriate role of a central bank in allocating credit.

Sixth, the Fed could have two distinct policy tools: setting short-term interest rates by paying interest on excess reserves, while managing the size and composition of its balance sheet. Having two instruments at its disposal would give the Fed greater flexibility to pursue multiple policy objectives. For example, Greenwood, Hanson, and Stein (2016) argued that the ongoing provision of a large volume of short-term risk-free assets would reduce the potentially destabilizing overreliance of the private sector on short-term funding, and thus enhance financial stability. Understanding the operation and appropriate use of balance sheet policies is an important topic for future research.

■ *The paper has benefitted immeasurably from comments from numerous participants at a seminar convened by the Hutchins Center on Fiscal and Monetary Policy at Brookings, especially Ben Bernanke, Olivier Blanchard, Steve Cecchetti, Bill English, Joe Gagnon, Michael Kiley, Donald Kohn, Arvind Krishnamurthy, Brian Sack, Fergal Shortall, Min Wei, and David Wessel; and from the editors of the* Journal of Economic Perspectives, *Mark Gertler, Gordon Hanson, and Timothy Taylor.*

## References

**Aizenman, Joshua, Mahir Binici, and Michael M. Hutchison.** 2016. "The Transmission of Federal Reserve Tapering News to Emerging Financial Markets." *International Journal of Central Banking* 12(2): 317–56.

**Bauer, Michael D., and Christopher J. Neely.** 2014. "International Channels of the Fed's Unconventional Monetary Policy." *Journal of International Money and Finance* 44: 24–46.

**Bauer, Michael D., and Glenn D. Rudebusch.** 2014. "The Signaling Channel for Federal Reserve Bond Purchases." *International Journal of Central Banking,* September, 10(3): 233–89.

**Bernanke, Ben S.** 2002. "Deflation: Making Sure 'It' Doesn't Happen Here." Speech before the National Economists Club, Washington, DC, November 21. https://www.federalreserve.gov/boarddocs/speeches/2002/20021121/.

**Bernanke, Ben S., and Mark Gertler.** 1995. "Inside the Black Box: The Credit Channel of Monetary Policy Transmission." *Journal of Economic Perspectives* 9(4): 27–48.

**Bhattarai, Saroj, Arpita Chatterjee, and Woong Yong Park.** 2015. "Effects of US Quantitative Easing on Emerging Market Economies." Globalization and Monetary Policy Institute Working Paper 255, Federal Reserve Bank of Dallas.

**Bhattarai, Saroj, and Christopher J. Neely.** 2016. "A Survey of the Empirical Literature on U.S. Unconventional Monetary Policy." Federal Reserve Bank of St. Louis Working Paper 2016-021A.

**Bowman, David, Juan M. Londono, and Horacio Sapriza.** 2015. "U.S. Unconventional Monetary Policy and Transmission to Emerging Market Economies. *Journal of International Money and Finance* 55: 27–59.

**Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano.** 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity,* no. 1, pp. 1–80.

**Chen, Han, Vasco Cúrdia, and Andrea Ferraro.** 2011. "The Macroeconomic Effects of Large-Scale Asset Purchase Programs." Staff Report 527, Federal Reserve Bank of New York.

**Chodorow-Reich, Gabriel.** 2014. "Effects of Unconventional Monetary Policy on Financial Institutions." *Brookings Papers on Economic Activity*, no. 1, pp. 155–227.

**D'Amico, Stefania, William English, David López-Salido, and Edward Nelson.** 2012. "The Federal Reserve's Large-Scale Asset Purchase Programmes: Rationale and Effects." *Economic Journal* 122(564): F415–F446.

**D'Amico, Stefania, and Thomas B. King.** 2013. "Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply." *Journal of Financial Economics* 108(2): 425–48.

**De Nicolò, Gianni, Giovanni Dell'Ariccia, Luc Laeven, and Fabian Vaencia.** 2010. "Monetary Policy and Bank Risk Taking." IMF Staff Position Note SPN/10/09, International Monetary Fund.

**e21 Staff**. 2010. "An Open Letter to Ben Bernanke." Economics21.org, November 15 https://economics21.org/html/open-letter-ben-bernanke-287.html.

**Eggertsson, Gauti B., and Michael Woodford.** 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* no. 1, pp. 139–211.

**Ehlers, Torsten.** 2012. "The Effectiveness of the Federal Reserve's Maturity Extension Program—Operation Twist 2: The Portfolio Rebalancing Channel and Public Debt Management." In *BIS Papers*, no. 65, *Threat of Fiscal Dominance?* pp. 245–256. Bank for International Settlements. https://www.bis.org/publ/bppdf/bispap65.pdf.

**Engen, Eric M., Thomas Laubach, and David Reifschneider.** 2015. "The Macroeconomic Effects of the Federal Reserve's Unconventional Monetary Policies." Finance and Economics Discussion Series 2015-005, Board of Governors of the Federal Reserve Board, Washington, D.C.

**Femina, Katherine, Steven Friedman, and Brian Sack.** 2013. "The Effects of Policy Guidance on Perceptions of the Fed's Reaction Function." Federal Reserve Bank of New York Staff Report 652.

**Foley-Fisher, Nathan, Rodney Ramcharan, and Edison Yu.** 2016. "The Impact of Unconventional Monetary Policy on Firm Financing Constraints: Evidence from the Maturity Extension Program." *Journal of Financial Economics* 122(2): 409–429.

**Fratzscher, Marcel, Marco Lo Duca, and Roland Straub.** 2018. "On the International Spillovers of US Quantitative Easing." *Economic Journal* 128(608): 330–77.

**Friedman, Benjamin M., and Kenneth N. Kuttner.** 2010. "Implementation of Monetary Policy: How Do Central Banks Set Interest Rates?" Chap. 24 in *Handbook of Monetary Economics,* vol. 3B, edited by Benjamin M. Friedman and Michael Woodford. Amsterdam: North Holland.

**Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack.** 2011. "The Financial Market Effects of the Federal Reserve's Large-Scale

Asset Purchases." *International Journal of Central Banking* 7(1): 3–43.

**Gertler, Mark, and Peter Karadi.** 2013. "QE 1 vs. 2 vs. 3 … : A Framework for Analyzing Large-Scale Asset Purchases as a Monetary Policy Tool." *International Journal of Central Banking* 9(S1): 5–53.

**Greenwood, Robin, Samuel G. Hanson, and Gordon Y. Liao.** 2018. "Asset Price Dynamics in Partially Segmented Markets." *Review of Financial Studies*, corrected proof: https://doi.org/10.1093/rfs/hhy048.

**Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein.** 2016. "The Federal Reserve's Balance Sheet as a Financial-Stability Tool." In *Designing Resilient Monetary Policy Frameworks for the Future*, Proceedings of the Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City, August 25–27, 2016, pp. 335–97.

**Greenwood, Robin, and Dimitri Vayanos.** 2014. "Bond Supply and Excess Bond Returns." *Review of Financial Studies* 27(3): 663–713.

**Hamilton, James D., and Jing Cynthia Wu.** 2012. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." *Journal of Money, Credit and Banking* 44(s1): 3–46.

**Ihrig, Jane, Elizabeth Klee, Canlin Li, Brett Schulte, and Min Wei.** 2012. "Expectations about the Federal Reserve's Balance Sheet and the Term Structure of Interest Rates." Finance and Economics Discussion Series 2012-57, Board of Governors of the Federal Reserve System.

**Ihrig, Jane E., Ellen E. Meade, and Gretchen C. Weinbach.** 2015. "Rewriting Monetary Policy 101: What's the Fed's Preferred Post-Crisis Approach to Raising Interest Rates?" *Journal of Economic Perspectives* 29(4): 177–98.

**Kashyap, Anil K., and Jeremy C. Stein.** 1994. "Monetary Policy and Bank Lending." In *Monetary Policy*, edited by N. Gregory Mankiw, 221–261. University of Chicago Press.

**Kiley, Michael T.** 2014. "The Aggregate Demand Effects of Short- and Long-Term Interest Rates." *International Journal of Central Banking* 10(4): 69–104.

**Kim, Don H., and Jonathan H. Wright.** 2005. "An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates." Finance and Economics Discussion Series 2005-33, Board of Governors of the Federal Reserve System.

**Krishnamurthy, Arvind, and Annette Vissing-Jorgensen.** 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity*, no. 1: 215–87.

**Kurzman, Robert, Stephan Luck, and Tom Zimmerman.** 2017. "Did QE Lead Banks to Relax Their Lending Standards? Evidence from the Federal Reserve's LSAPs." Finance and Economics Discussion Series 2017-093, Board of Governors of the Federal Reserve System.

**Kuttner, Kenneth N.** 2001. "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market." *Journal of Monetary Economics* 47(3): 523–44.

**Li, Canlin, Andrew Meldrum, and Marius Rodriguez.** 2017. "Robustness of Long-Maturity Term Premium Estimates." FEDS Notes, April 3.

**Luck, Stephan, and Tom Zimmerman.** 2017. "Employment Effects of Unconventional Monetary Policy: Evidence from QE." Available at SSRN: https://ssrn.com/abstract=3002047.

**Meaning, Jack, and Feng Zhu.** 2011. "The Impact of Recent Central Bank Asset Purchase Programmes." *BIS Quarterly Review*, December, pp. 73–83.

**Modigliani, Franco, and Richard Sutch.** 1966. "Innovations in Interest Rate Policy." *American Economic Review* 56(1–2): 178–97.

**Nakamura, Emi, and Jón Steinsson.** 2018. "Identification in Macroeconomics." *Journal of Economic Perspectives* 32(3): 59–86.

**Neely, Christopher J.** 2015. "Unconventional Monetary Policy Had Large International Effects." *Journal of Banking and Finance* 52: 101–111.

**Powell, Jerome H.** 2017. "Low Interest Rates and the Financial System." Speech at the 77th Annual Meeting of the American Finance, Chicago I, January 7.

**Rajan, Raghuram G.** 2005. "Has Financial Development Made the World Riskier?" In *The Greenspan Era: Lessons for the Future*, Proceedings of the Federal Reserve Bank of Kansas City Jackson Hole Symposium, pp. 313–69.

**Rodnyansky, Alexander, and Olivier M. Darmouni.** 2017. "The Effects of Quantitative Easing on Bank Lending Behavior." *Review of Financial Studies* 30(11): 3858–87.

**Rudebusch, Glenn D., Brian P. Sack, and Eric T. Swanson.** 2007. "Macroeconomic Implications of Changes in the Term Premium." *Federal Reserve Bank of Saint Louis Review* 89(4): 241–70.

**Stein, Jeremy C.** 2012. "Evaluating Large-Scale Asset Purchases." Speech at the Brookings Institution, Washington, DC, October 11.

**Swanson, Eric T.** 2017. "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets." NBER Working Paper 23311.

**Swanson, Eric T., and John C. Williams.** 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review* 104(10): 3154–85.

**Tobin, James.** 1963. "An Essay on Principles

of Debt Management." *Fiscal and Debt Management Policies*, prepared for the Commission on Money and Credit, 143–218. Englewood Cliffs, NJ: Prentice-Hall.

**Vayanos, Dimitri, and Jean-Luc Vila.** 2009. "A Preferred-Habitat Model of the Term Structure of Interest Rates." NBER Working Paper 15487.

**Woodford, Michael.** 2012. "Methods of Policy Accommodation at the Interest-Rate Lower Bound." In *The Changing Policy Landscape,* Federal Reserve Bank of Kansas City Jackson Hole Symposium, pp. 185–288.

**Wright, Jonathan H.** 2012. "What Does Monetary Policy Do to Long-Term Interest Rates at the Zero Lower Bound?" *Economic Journal* 122(564): F447–F466.

**Wu, Jing Cynthia, and Fan Dora Xia.** 2016. "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound." *Journal of Money, Credit and Banking* 48(2–3): 253–91.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 147–172*

# Unconventional Monetary Policies in the Euro Area, Japan, and the United Kingdom

Giovanni Dell'Ariccia, Pau Rabanal, and Damiano Sandri

The global financial crisis hit hard in the euro area, the United Kingdom, and Japan. Real GDP from peak to trough contracted by about 6 percent in the euro area and the United Kingdom and by 9 percent in Japan. In all three cases, central banks cut interest rates aggressively and then, as policy rates approached zero, deployed a variety of untested and unconventional monetary policies (see Figure 1). In doing so, they hoped to restore the functioning of financial markets, and also to provide further monetary policy accommodation once the policy rate reached the zero lower bound.

In all three jurisdictions, the strategy entailed generous liquidity support for banks and other financial intermediaries and large-scale purchases of public (and in some cases private) assets. As a result, central banks' balance sheets expanded to unprecedented levels. Naturally, the specific measures employed by each jurisdiction varied in timing and detail. For example, the European Central Bank first focused on liquidity injections to restore frozen interbank activity and then on the targeted purchase of sovereign bonds to address the erupting sovereign debt crisis in Greece, Ireland, and Portugal in 2010. Only in mid-2014, in the context of anemic growth and below-target inflation, did the European Central Bank adopt its own large-scale asset purchase program. In the United Kingdom, the response to the crisis by the Bank of England was quite similar to that of the US Federal Reserve

■ *Giovanni Dell'Ariccia is a Deputy Director in the Research Department, Pau Rabanal is a Deputy Division Chief in the Research Department, and Damiano Sandri is a Senior Economist in the Western Hemisphere Department, all at the International Monetary Fund, Washington, DC. Their emails are gdellariccia@imf.org, prabanal@imf.org, and dsandri@imf.org.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.147

doi=10.1257/jep.32.4.147

*Figure 1*
**Real GDP and Central Bank Assets in the Euro Area, Japan, and the United Kingdom**

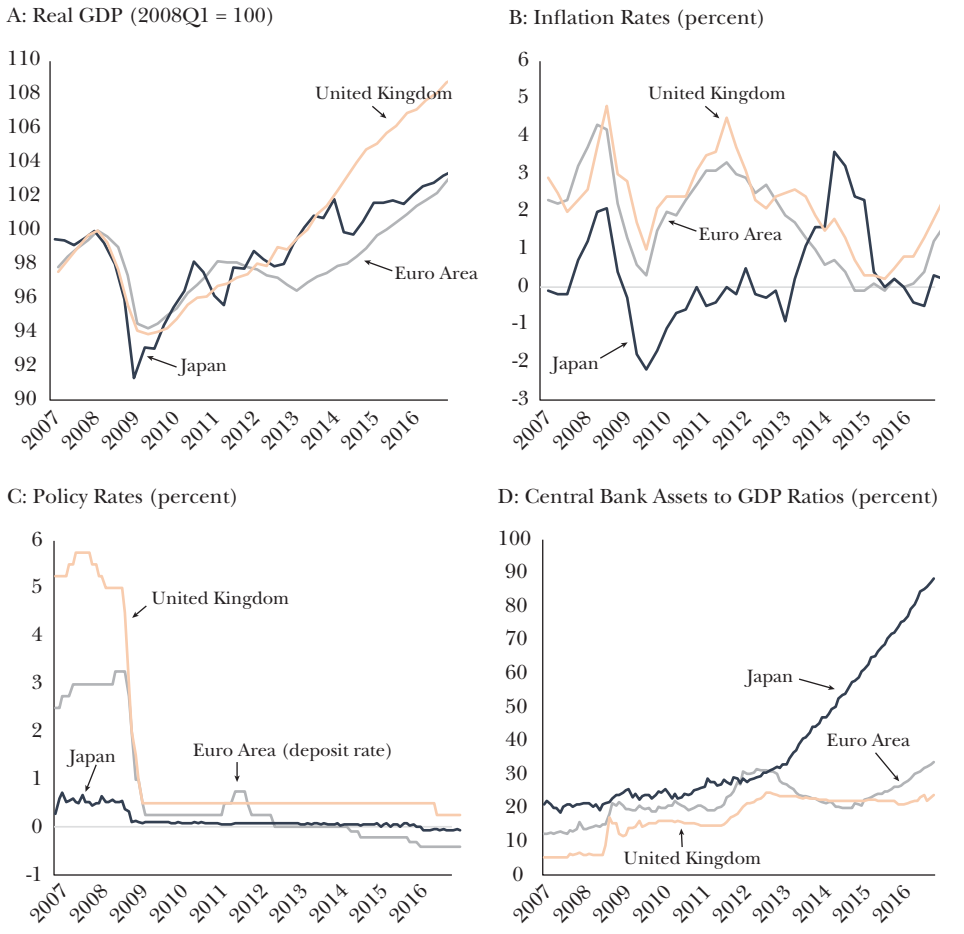


*Source:* CEIC, Haver Analytics, and authors' calculations.

in timing and style, including mostly via large-scale purchases of government debt. The Bank of Japan entered the crisis with more limited policy room, having experienced the combination of slow growth and near-zero policy interest rates since Japan's financial crisis in the early 1990s. After small initial purchases of government bonds, it became willing after 2012 to purchase a much wider and larger set of public and private securities while also adopting a numerical inflation target.

Overall, unconventional monetary policies have been quite effective in preventing further financial distress, restoring the functioning of financial markets, and providing additional monetary accommodation by compressing long-term interest rates. Furthermore, these policies likely had beneficial effects on macroeconomic variables such as real GDP growth and price stability, although these are more difficult to model and measure.

Looking ahead, to the extent that interest rates remain low over the medium and long run (as seems likely), the zero lower bound in nominal interest rates may be again binding in future recessions. With that in mind, this paper examines the experience with unconventional monetary policies in the euro zone, the United Kingdom, and Japan. The paper starts with a discussion of how quantitative easing, forward guidance, and negative interest rate policies work in theory, and some of their potential side effects. It then reviews the implementation of unconventional monetary policy by the European Central Bank, the Bank of England, and the Bank of Japan, including a narrative of how central banks responded to the crisis and the evidence on the effects of unconventional monetary policy actions. The conclusion summarizes the main empirical findings and offers some lessons for the use of unconventional monetary policies in the future. The appendix presents a detailed account of unconventional monetary policy actions taken by the three central banks, as well as the evolution of key macroeconomic indicators in the three economies. In a companion paper in this issue, Kuttner reviews the experience of the US Federal Reserve with unconventional monetary policies.

## Conceptual Framework for Unconventional Monetary Policy

This section introduces three forms of unconventional monetary policy: forward guidance, quantitative easing, and negative interest rate policies. More specifically, it focuses on the mechanisms through which such policies can lower and flatten the yield curve—that is, reduce the interest rate level and shrink the gap between long- and short-term interest rates. This in turn should lower borrowing costs for households and firms, increase credit, boost aggregate demand, and thus ultimately raise inflation and output. As in the case of conventional interest rate cuts, unconventional monetary policies can also support external demand by depreciating the exchange rate. This section concludes with a discussion of these policies' potential negative side effects.

### Forward Guidance

Forward guidance aims at providing market participants with information about the intentions of policymakers for the future path of the policy rate. It can come in two forms (Campbell, Evans, Fischer, and Justiniano 2012). In one version, the central bank aims to clarify how monetary policy will evolve in the future depending on its own expectations for economic activity or inflation. In the alternative and potentially more powerful version, the central bank commits to keeping interest rates low even if economic conditions improve in the future and warrant a monetary tightening.

Forward guidance can be made contingent on qualitative or quantitative criteria. As an example of qualitative guidance, the Bank of Japan announced in October 2010 that it would keep rates low until "price stability is in sight." As an example of quantitative guidance, the Bank of England announced in August 2013 that rates would stay low until the "unemployment rate has fallen to a threshold of

7 percent." Forward guidance targets can also be time dependent, for example by announcing that rates will stay low until a specific date in the future.

Several factors can hinder the effectiveness of forward guidance and complicate the empirical evaluation of its effects. First, forward guidance may fail to alter expectations. For example, guidance that interest rates will remain low for several years into the future may be ineffective if markets already expect such an interest path or do not find the commitment credible. Second, how the announcement of forward guidance is interpreted by markets can lead to counterintuitive results. For example, forward guidance that interest rates will remain low for longer than expected may signal that the central bank is more seriously concerned about future economic prospects than previously believed, and thus have negative economic effects by denting consumer and business confidence. Alternatively, if market participants believe that forward guidance will be successful in increasing inflation and real GDP growth, long-term rates may increase rather than decrease—although optimism about the future may offset the increase in rates.

**Quantitative Easing**

Quantitative easing involves the large-scale purchase of securities by the central bank. It is generally implemented through the acquisition of long-term government bonds (an asset in the central bank balance sheet) financed by an increase in the reserve accounts that commercial banks hold at the central bank (a liability for the central bank). The key idea is that, when the policy rate and thus the yield on short-term bonds are at zero, the central bank can still provide monetary stimulus by supporting long-term bond prices and thus lowering long-term yields. This mechanism has been formalized in macroeconomic models: for examples, see Gertler and Karadi (2011) and Chen, Curdia, and Ferrero (2012).

Former Federal Reserve chair Ben Bernanke (2014) once famously said that quantitative easing "works in practice, but it doesn't work in theory." His point was that in a theoretical model with no financial market frictions, and in which investors move freely across asset categories, central bank purchases of government bonds should not have any effects on bond yields. Instead, financial markets and arbitrageurs will reposition their portfolios offsetting the effects of central bank purchases. In practice, however, financial markets are segmented, for example because agents have a preference to hold specific securities or because they may find it difficult to short sell the bonds the central bank is buying. In this case, by reducing the net supply of government bonds on the market, central bank purchases raise bond prices and reduce the yield of those bonds. Note that for this to have an effect beyond government bond markets, segmentation cannot be too extreme. Indeed, agents that sell government bonds to the central bank are expected to rebalance their investment portfolios by purchasing other securities (like corporate bonds, equity, or real estate-backed securities), thus boosting prices for those assets, lowering yields and stimulating consumption and investment.

Quantitative easing can also decrease bond yields through a signaling channel. The purchase of large quantities of government bonds can help to convince markets

that the central bank is committed to keeping a loose policy stance. In this regard, quantitative easing operates as a useful complement to forward guidance.

Besides purchasing government securities, quantitative easing can be also implemented through the direct acquisition of privately issued securities. For example, the Bank of Japan has purchased not only government bonds, but also corporate bonds, exchange-traded funds, and real-estate investment funds. The European Central Bank purchased "covered bonds" (that is, collateralized bonds issued by banks or mortgage lenders) in three different phases between 2009 and 2017, and corporate bonds in the primary and secondary markets starting in June 2016. Purchases of private securities can reduce the borrowing costs faced by private agents and stimulate the economy more directly, but they expose the central bank to credit risk and potential losses.

Typically, quantitative easing has been implemented by announcing a specific timeline and amount of purchases. More recently, the Bank of Japan has adopted an alternative "yield curve control" approach, which sets targets for both short- and long-term yields and adjusts purchases to meet those targets. A possible advantage is that if the target is credible, market participants may coalesce around it without requiring purchases by the central bank. But if the target is not credible, the central bank may be forced to purchase bonds in very large quantities or further dent its credibility by revising the target.

**Negative Interest Rates**

The European Central Bank and the Bank of Japan, along with others, have implemented negative interest rates by charging, rather than paying, interest rates on the reserves that commercial banks hold at the central bank. The hope is that individual banks will reduce their excess reserves by increasing lending and purchasing other financial assets. In this way, the policy seeks to reduce lending rates, increase credit supply, and boost prices across financial markets.

The notion that policy interest rates cannot decline below zero derives from the idea that agents would rather hoard cash than deposit money in accounts that charge interest rates (that is, pay negative interest rates). However, using cash involves significant transaction costs and risks (it can be stolen), so that mildly negative rates are unlikely to generate major shifts into holding banknotes.

Several concerns have been raised about negative rates. First, banks appear reluctant to pass negative interest rates on to retail depositors (although there is evidence of pass-through to corporate deposits). Negative rates may thus harm bank profitability and possibly prevent a reduction in lending rates. However, banks seem to have supported profitability by increasing noninterest income through charging other fees.[1] The direct costs of negative reserves on bank reserves are also quantitatively small relative to banks' balance sheets. Central banks can further reduce these direct costs by charging negative rates only at the margin: for example, required

[1] See Cœuré (2016) and Hutchinson and Smets (2017) for the case of the euro area. In Japan, the evidence suggests that the profitability of small regional banks may have been reduced by the negative interest rate policy (IMF 2017b).

bank reserves are typically exempt from negative interest rates and the Bank of Japan charges negative rates only on a subset of other bank reserves. Second, there are concerns that negative rates can at a certain point lead to a major shift to cash, especially if they are perceived to be long lasting. Finally, an important impediment to negative rates may arise from the public perception that they are "unfair."

Of course, any costs from negative interest rates would need to be weighed against the positive effects that they might have on asset prices and the economic outlook. IMF (2017a) reviews the experiences of several countries with negative interest rates.

**Potential Side Effects**

In principle, unconventional monetary policies can contribute to financial stability if lower interest rates help to stimulate the economy and improve borrowers' ability to stay current with their loans. However, concerns have been raised about side effects that might endanger financial stability, even when unconventional monetary policy is successful in stimulating the economy. These concerns can be broadly divided into five sets of arguments.

First, when long-term securities purchases and forward guidance flatten the yield curve by compressing term premia, they put pressure on bank profitability (Borio, Gambacorta, and Hofmann, 2015; Borio and Gambacorta, 2017). After all, banks issue short-term liabilities such as deposits and commercial paper and invest in longer-term assets such as mortgages, asset-backed securities, and commercial loans. The profitability of this business model of credit and maturity transformation is proportional to the size of term premia.

Second, the compression in safe yields from monetary easing induces financial intermediaries to move toward riskier assets. For example, this situation can arise if institutional players such as life insurers and pension funds become unable to match promised yields on their long-term liabilities with safe assets (Rajan 2005). Taking additional risk is to some extent an intended effect of monetary easing (Chodorow-Reich 2014), but it can also become excessive from an aggregate welfare point of view.

Third, a low interest rate environment may reduce incentives for banks to recognize and write off nonperforming loans (Caballero, Hoshi, and Kashyap 2008).

Fourth, by increasing asset prices and reducing volatility, central bank purchases may lead to the build-up of asset-price deviations from their fundamentals and trigger a later sharp asset-price correction. This may also create liquidity risks in the nonbank financial sector, as investors may become too complacent (ECB 2017).

Fifth, critics of unconventional policies have also warned that when central banks stray from their traditional way of conducting monetary policy, they might jeopardize their hard-fought independence. The more a central bank becomes involved in multiple (and less measurable) objectives and policy instruments, critics say, the more monetary authorities become exposed to political interference (Taylor 2016). That said, central banks may also come under heavy political pressure if their reluctance to adopt unconventional monetary measures leads to a more severe financial and economic crisis.

## The Euro Area

The European Central Bank's response to the crisis and unconventional monetary policy implementation can be divided in three phases. First, between September 2008 and the end of 2009, the European Central Bank focused on supporting the banking sector, using instruments which can be categorized as part of a central bank's function as "lender of last resort." Second, during the sovereign-debt crisis of several euro area countries between early-2010 and late-2012, the European Central Bank purchased government bonds to restore market functioning and the transmission mechanism of monetary policy. Third, starting in mid-2013, the European Central Bank implemented a more aggressive combination of forward guidance, large-scale asset purchases, negative interest rates, and targeted credit supply policies.

**Economic Developments and Monetary Policy Responses in the Euro Area**

After reaching a pre-crisis peak in the first quarter of 2008, real GDP in the euro area contracted by 6 percent within the year. Consumer price inflation fell well below the 2 percent target, prompting the European Central Bank to lower its benchmark rate (the main refinancing rate) from 4.25 percent to 1 percent over the course of 2008. As the space for conventional monetary policy dwindled, the reaction of the European Central Bank to the global financial crisis can be divided into three main phases.

The initial strategy of the European Central Bank focused on addressing the increased credit and counterparty risk that had led to a sharp decline in interbank trading, thus impairing the monetary transmission mechanism in the euro area. Early in the financial crisis, then-president of the European Central Bank Jean-Claude Trichet (2009) pointed out the large differences in the composition of funding sources for nonfinancial corporations between the euro area, where bank financing accounted for roughly 70 percent of firms' total external financing, and the United States, where firms relied to a much larger extent on market-based sources (like issuing bonds) that made up about 80 percent of total external financing (Cour-Thimann and Winkler 2013).

Specifically, the policy response included expanding its main liquidity operations, and introducing several rounds of Longer-Term Refinancing Operations during 2008 and 2009. This de facto replaced the drop in interbank market activity with increased intermediation through the central bank (González-Parámo 2011).[2] Other policies included foreign-currency operations (swap lines with the Federal Reserve in US dollars), a broadening of the collateral framework allowing banks to use a broader range of assets in refinancing operations with the European Central

[2]The LTROs (Longer-Term Refinancing Operations) differed from the standard MROs (Main Refinancing Operations) because they were conducted at fixed rates and with full allotment, thereby making them unlimited, and their maturities were longer (3, 6, and 12 months, instead of the one-week MROs). The LTRO balance at the European Central Bank increased from €1.5 trillion to €2.8 trillion, a 90 percent increase, between early 2008 and early 2010.

Bank, and the launch of the Covered Bond Purchase Program. This last program was relatively small in size and was justified to "improve market liquidity" and "ease funding conditions for banks and enterprises." Broadly speaking, these policies can be understood as "lender of last resort" actions, in which a central bank makes credit available to financial intermediaries during financial stress.

But the euro-zone financial woes were soon compounded by a potential fiscal and sovereign-debt crisis in several member countries. In one notable event on November 5, 2009, the newly elected Greek Prime Minister Papandreou acknowledged that the Greek fiscal deficit would be 12.7 percent of GDP rather than the 3.7 percent reported by the outgoing government. The confidence shock associated with this announcement raised concerns about the high level of public debt in Italy and Portugal and the indebtedness of the private sector in Ireland and Spain. During the first half of 2010, government borrowing rates increased to unsustainable levels in Greece (see Figure 2B), which required financial assistance through an EU–IMF program.

Before the eruption of the sovereign debt crisis, the European Central Bank had hesitated to use large-scale asset purchases of government bonds in the euro area. Its legal authority was unclear since Article 123 of the Treaty on the Functioning of the European Union prohibits monetary financing of governments. But on May 10, 2010, one week after the IMF and the European Union had announced the first bailout package for Greece, the European Central Bank launched the Security Markets Program. The program involved purchasing government debt issued by Greece, Ireland, and Portugal and was implemented "to address the malfunctioning of securities markets and restore an appropriate monetary policy transmission mechanism" (ECB 2010). Because the program involved government bonds purchases in the secondary market, it was viewed as not breaching Article 123.

Later, the European Central Bank announced that it would also purchase Italian and Spanish government bonds. In total, the European Central Bank purchased €218 billion of Greek, Irish, Italian, Portuguese, and Spanish bonds as of end-2012 and held them until maturity. As shown in Figure 2, the launch of the Securities Market Program did not calm sovereign-debt markets in the euro area. Ireland, which had suffered a banking crisis in 2009–2010 amidst a collapsing housing bubble, requested an IMF–EU program, which was signed in December 16, 2010. Portugal also received a bailout package on May 20, 2011. Peripheral countries in the euro-zone entered a double-dip recession, and the real GDP in the euro area declined again during 2011–2012. Borrowing spreads for Italy and Spain kept creeping higher, unlike those in France and Germany (Figure 2). These countries were trapped in destructive self-fulfilling dynamics, where concerns about fiscal sustainability (because of large deficits and low growth prospects) increased the likelihood of sovereign debt default, which in turn increased borrowing costs, making it all more likely that countries would end up defaulting and perhaps having to leave the euro altogether. This "redenomination risk" (Cœuré 2013) increased borrowing costs further.

In this context, the European Central Bank strengthened its commitment to quantitative easing in two ways. Perhaps most critically, in a conference on July 26,

*Figure 2*
**Government Borrowing Costs and Policy Rates in the Euro Area**

A: Policy Rates and 10-Year Government Bond Yields (percent)

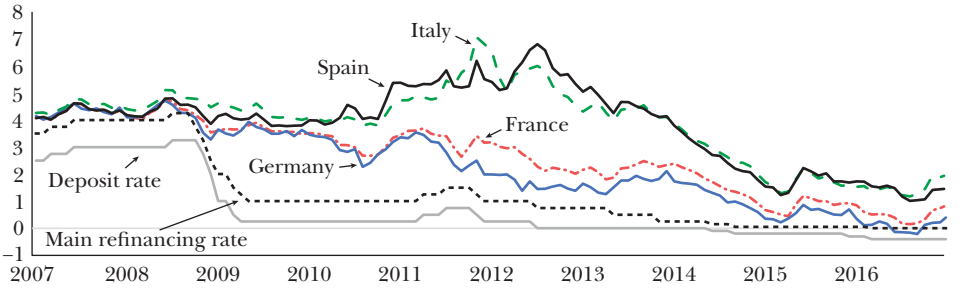


B: 10-Year Government Bond Yields in Program Countries (percent)

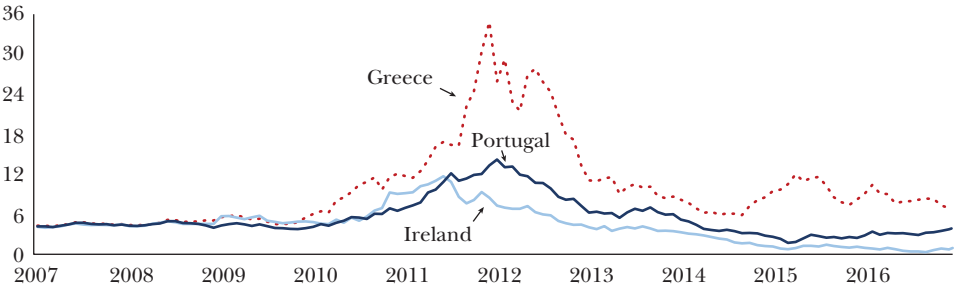


*Source:* Haver Analytics and the European Central Bank.

2012, in London, the European Central Bank President, Mario Draghi, made a statement that became famous: "Within our mandate, the ECB is ready to do whatever it takes to preserve the euro. And believe me, it will be enough." The European Central Bank soon followed up by announcing the Outright Monetary Transactions (OMT) program, which would include purchases of government bonds in secondary markets for member countries that requested its activation and accepted monitoring. The program was announced "to address severe distortions in government bond markets which originate from, in particular, unfounded fears on the part of investors of the reversibility of the euro." Together, the announcement of the OMT program and Draghi's "whatever it takes" speech reversed the sovereign-debt market self-destructing spiral. And this was accomplished without ever making purchases under the program, since to date, not a single member country has made a formal request.[3]

[3]The Outright Monetary Transactions program had some meaningful differences from the earlier Securities Market Program that enhanced its effectiveness. Under the Outright Monetary Transactions, the European Central Bank explicitly gave up seniority, which was instead retained under the Securities Market Program. Furthermore, differently from the Securities Market Program, the Outright Monetary Transactions included conditionality.

By 2013, the euro-zone economy had emerged from recession and the sharpest risk of a sovereign debt crisis had been averted, but growth remained anemic and inflation was stuck below its target. At this point, the European Central Bank employed more "standard" unconventional monetary policy tools to improve credit conditions and provide monetary stimulus. First, on July 4, 2013, the European Central Bank used forward guidance for the first time, with President Draghi stating: "The Governing Council expects the key ECB interest rates to remain at present or lower levels for an extended period of time." Second, the European Central Bank announced the introduction of a negative interest rate of -0.1 percent for its deposit facility on June 5, 2014. Subsequently, the European Central Bank announced further cuts to this rate up to -0.4 percent in March 2016. In parallel, the European Central Bank announced a new round of credit easing measures. The Targeted Longer-Term Refinancing Operations gave banks who lent to households and firms more favorable financing conditions from the European Central Bank, in order to reinforce the pass-through from negative interest rates to retail lending rates and increase credit supply in vulnerable countries (Hutchinson and Smets 2017).

Most important, the European Central Bank announced the introduction of its own large-scale asset purchase programs on September 4, 2014. President Draghi said that these programs "will have a sizeable impact on our balance sheet" and that their launching would "support our forward guidance on the key ECB interest rates." Under the umbrella of the asset purchase program, the ECB purchased asset-backed securities, covered bonds, corporate sector bonds, and government bonds. Total holdings of securities in the program as of May 2018 was about €2.4 trillion (or about 23 percent of euro area GDP).

**Effects of Unconventional Monetary Policy in the Euro Area**

Empirical researchers have used a variety of econometric tools to measure the effects of the European Central Bank policy measures. A central problem is identification: how to distinguish the effect of these policies from other macroeconomic news releases and policy initiatives taking place at the same time in the euro area. For instance, the Security Markets Program was announced on May 10, 2010, while the IMF board approved the first Greek program on May 9, 2010. These years have a constant overlap of changes in European Central Bank policy, national policy, IMF policy, all against an evolving backdrop of economic and market events.

One way to meet the identification challenge is to use high-frequency data and focus on a narrow window around policy changes. This "event study" methodology studies the reaction of financial variables such as government and corporate yields, and stock prices, which are available at high frequency (within minutes or hours) around a policy announcement. Table 1 summarizes the findings of a few studies in the euro area. In general, the literature finds significant effects on government bond yields from unconventional monetary policy actions undertaken by the European Central Bank. The largest effect is found by Krishnamurthy, Nagel, and Vissing-Jorgensen (2018), who estimate that the Security Markets Program and the

announcement of the Outright Monetary Transactions reduced two-year government bond yields by 200 basis points in Italy and Spain, 500 basis points in Portugal and Ireland, and 1,000 basis points in Greece. The decline can be attributed to a decline in default risk (which explains 37 percent of the total reduction), reduced redenomination risk (13 percent of the total reduction), and reduced market segmentation by increasing liquidity in distressed markets (50 percent of the total reduction). In contrast, estimates of the impact of the Long-Term Refinancing Operations are much smaller. These operations were intended to restore interbank liquidity, but in practice banks used the funds to purchase government debt, thereby reducing sovereign yields (Ongena, Popov, and van Horen 2016).

The effects on government yields in the euro area were generally highly persistent, as documented in the paper by Ghysels, Idier, Manganelli, and Vergote (2016) using a vector autoregressive model with daily data. Eser and Schwab (2016) conclude that the Security Markets Program improved liquidity conditions and reduced default-risk premia, but that the signaling of future low interest rates did not play a role. De Pooter, Martin, and Pruitt (2015) find that the liquidity premium declined between 32 and 40 basis points and helped reduce government yields. This channel is important since the European Central Bank justified its intervention due to a lack of liquidity and depth in certain securities markets.

Some studies confirm the effects of actions of the European Central Bank on peripheral Europe government bonds, although perhaps with smaller estimated effects (Fratszcher, Lo Duca, and Straub 2016). Other analyses have focused on the credit effects of the European Central Bank's programs. Ferrando, Popov, and Udell (2015) study the effect of the Outright Monetary Transactions announcement on small and medium enterprises' access to credit. They find that the probability of being credit-constrained in peripheral euro countries declined by 6.4 percent due to the announcement. Arce, Gimeno, and Mayordomo (2017) find that the Corporate Sector Purchase Program and Targeted Longer-Term Refinancing Operations increased bond issuance by large corporations, as well as bank credit to smaller corporations in Spain, thus providing support to the credit reallocation hypothesis.

The unconventional monetary policies of the European Central Bank clearly affected financial variables. But what about the ultimate objective of affecting macroeconomic aggregates? This connection is obviously harder to establish because macroeconomic variables are slower-moving and do not lend themselves to event studies. Methodologies use macroeconomic models, typically either a time series vector autoregressive or dynamic stochastic general equilibrium model at monthly or quarterly frequency. With these approaches, the effects of unconventional policies on government yields are fed through the model to obtain implications for prices, real activity, and other macroeconomic variables.

As one example, Altavilla, Giannone, and Lenz (2014) find that the announcement of the Outright Monetary Transactions program lowered 2-year Italian bond yields by 199 basis points and Spanish bonds yields by 234 basis points. In a second stage, using a multicountry vector autoregression model, the authors find that the announcement raised Italian GDP by 1.5 percent and consumer price index by

*Table 1*
**Unconventional Monetary Policy Effects in the Euro Area**

| *Study* | *Notes* | *Government bond yields* | *Real GDP* | *Prices* | *Other* |
|---|---|---|---|---|---|
| Darracq-Paries and De Santis (2015) | 3-year LTRO effects using a VAR model | | +0.8% | +0.3% | +3% credit, -0.2% lending spreads |
| Cahn, Matheron, and Sahuc (2014) | Effects of an LTRO of 2% of GDP. DSGE model with financial frictions. | | +1% | | |
| De Pooter, Martin, and Pruitt (2015) | Effect of the SMP on peripheral bonds liquidity premia | −32 to −40 bps on impact, −13 to −17 bps are lasting | | | |
| Ghysels, Idier, Manganelli, and Vergote (2016) | Effects of SMP with VAR model with high-frequency data. | −320 bps (Italy 2y), −180 bps (Spain 2y), −230 bps (Italy and Spain 10y). Similar results for Ireland and Portugal, not significant for Greece. | | | |
| Eser and Schwab (2016) | Cumulative SMP effects of purchases looking at high-frequency data | −10 bps ( 5y), −170 bps (Portugal .5y), −190 bps (Spain 5y), −210 bps (Italy 5y), −330 bps (Greece 5y) | | | |
| Ferrando, Popov, and Udell (2015) | OMT effects of SME access to credit in euro area distressed countries | | | | Probability of being credit constrained was reduced by 6.4% |
| Altavilla, Giannone, and Lenz (2014) | Effects of OMT announcements using event studies and VAR models | −199 bps (Italy 2y) −234 bps (Spain 2y), no effects in Germany and France | +1.5% (Italy), +2 % (Spain) | +1.2% (Italy), +0.74% (Spain) | +3.6% (credit, Italy), +2.3% (credit, Spain) |
| Fratszcher, Lo Duca, and Straub (2016) | Effects of LTRO, SMP, and OMT announcements using high-frequency data | −25 bps to −121 bps (Italy and Spain 10y) | | | +4.1 to +8.7% (equity prices) |
| Krishnamurthy, Nagel, and Vissing-Jorgensen (2018) | Effects of OMT, SMP, and LTROs | −200 bps (Italy and Spain 2Y), −500 bps (Portugal and Ireland 2y), −1,000 bps (Greece 2y). | | | +4% to +13% (stock prices) |
| Koijen, Koulischer, Nguyen, and Yogo (2016) | Effects of APP on portfolio holdings by institutional investors | Average −13 bps. Range −2 to −60 bps (higher in distressed countries) | | | |

*(continued on next page*

*Table 1 (continued)*
**Unconventional Monetary Policy Effects in the Euro Area**

| *Study* | *Notes* | *Government bond yields* | *Real GDP* | *Prices* | *Other* |
|---|---|---|---|---|---|
| Andrade, Breckenfelder, De Fiore, Karadi, and Tristani (2016) | Effects of APP using time series and DSGE models | −45 bps | +1.1% | +0.4% (actual), +0.45% (expectations) | |
| Mouabbi and Sahuc (2016) | Effects of APP and TLTRO using a DSGE model with an estimated shadow rate | | +0.56% (average of 2014–2016) | +0.25% (average of 2014–2016) | -400 bps (shadow rate) |
| Cova, Pagano, and Pisani (2015) | Effects of APP in a DSGE model | | +1.4% | +0.8% | |
| Hutchinson and Smets (2017) | Effects of NIRP, TLTRO, and APP | −155 bps (Average euro area 10y bond) | +1.7% (accum. 2016–2019) | +0.5% (accum. 2016–2019) | -70 bps (lending rate), 13 percent euro depreciation |

*Notes:* LTRO is Longer-Term Refinancing Operations; SMP is Security Markets Program; OMT is Outright Monetary Transactions; APP is Asset Purchase Program; TLTRO is Targeted Longer-Term Refinancing Operations; NIRP is Negative Interest Rate Policy. VAR means "vector autoregression"; DSGE means "dynamic stochastic general equilibrium." SME stands for "small and medium-sized enterprises"; "bps" are basis points.

1.2 percent, while it raised Spanish GDP by 2 percent and consumer price index by 0.74 percent over a 3-year horizon. Credit also increased in both countries. The same study finds that France and Germany's yields did not fall due to the Outright Monetary Transactions announcement, but these countries still benefitted from higher growth in the periphery through trade linkages. Other studies summarized in Table 1 also suggest that the effects of the Longer-Term Refinancing Operations (LTRO) and asset purchase program (APP) were positive and stimulated the euro area economy (especially, Mouhabbi and Sahuc 2015). But as Burriel and Galesi (2016) show, there can be substantial heterogeneity in the effects of unconventional monetary policies on euro area member countries.

The effects of the negative interest rate policies (NIRPs) are harder to quantify because the European Central Bank and other macroeconomic policymakers undertook other measures during the same period. However, evidence on the behavior of government yields and lending rates (see Figure 2) suggests that negative interest rates operate very much like interest rate cuts when these are in positive territory. Using a simple event analysis after the four policy announcements that placed the deposit facility rate in negative territory (the initial announcement was done on June 5, 2014, with subsequent cuts of 0.1 percent in three different meetings), we found that deposit rate cuts were typically followed by downward shifts in

the yield curves, especially for Germany on the first announcement date, and for Italy and Spain in the subsequent rate changes.

The combined effect of all the measures implemented since the summer of 2013 reduced benchmark lending rates for households and corporations in the euro area from 3 percent in the summer of 2013 to 1.7 percent in the fall of 2017. In addition, the euro depreciated by about 30 percent vis-à-vis the US dollar during the same period. Hutchinson and Smets (2017) find that the combined effect of all these measures provided a boost to real GDP of 1.7 percent (on a cumulative basis between 2016–2019) and to inflation of 0.5 percent (over the same period).

## The United Kingdom

Since the beginning of the global financial crisis, unconventional monetary stimulus in the United Kingdom has passed through three main phases. First, large-scale quantitative easing programs between 2009 and 2012 sought to halt the recession and support the economic rebound. Second, forward guidance announcements in 2013 and 2014 clarified the intention of the Bank of England not to raise policy rates. Third, an additional round of quantitative easing occurred when the United Kingdom voted to leave the European Union in 2016.

### Economic Developments and Monetary Policy Responses in the United Kingdom

The UK economy was severely affected by the global financial crisis, especially because of the reliance of its large financial sector on wholesale funding. GDP started to contract in the second quarter of 2008 and fell by about 6 percent in real terms within a year. The Bank of England responded by sharply cutting interest rates from 5.75 to 0.5 percent by March 2009. Despite this large rate cut and the inflationary pressures arising from the sterling depreciation, the Bank of England feared that inflation would fall below the 2 percent target. To provide additional monetary stimulus, the Bank of England launched the first round of quantitative easing (QE1), announcing the purchase of £75 billion of government bonds financed with the expansion of central bank reserves. Purchases were later increased to £200 billion (about 13 percent of 2009 GDP) and completed in January 2010.

GDP started to recover in the third quarter of 2009, but at a subdued pace. In 2011, the intensification of the euro-area sovereign debt crisis threatened the UK economic outlook. Despite inflation rising toward 5 percent because of an increase in the value-added-tax and rising energy prices, in October 2011 the Bank of England announced a new round of quantitative easing (QE2). This was followed by a third round of asset purchases (QE3) in July 2012 that brought the total size of quantitative easing to £375 billion, about 25 percent of GDP and 35 percent of the outstanding stock of British government bonds.[4]

[4]The Bank of England also launched several programs to support the banking sector, including the Special Liquidity Scheme in April 2008 allowing banks to swap high-quality mortgage-backed securities

Despite economic growth remaining weak, the persistence of inflation above target raised concerns that the Bank of England could soon raise interest rates. In August 2013, the Bank of England decided to clarify the expected path of monetary policy by introducing forward guidance with explicit quantitative targets. It announced that it would not raise policy rates or reduce the stock of bonds purchased under quantitative easing at least until unemployment declined below 7 percent (from the ongoing 8 percent level). This guidance would cease to hold if medium-term inflation was expected to increase by 0.5 percent above target; if inflation expectations became de-anchored; or if the stance of monetary policy posed significant threats to financial stability.

Economic growth finally gained strength in late 2013. As the unemployment rate fell toward the 7 percent threshold faster than expected, in February 2014 the Bank of England provided additional forward guidance. It argued that considerable spare capacity remained in the economy, and so it expected to keep rates low for longer and eventually raise them only gradually.

Even though the UK economy was growing at a steady pace, the Bank of England deployed a new round of monetary stimulus in August 2016 following the "Brexit" vote to leave the European Union. To sustain economic growth, the Bank of England launched a fourth round of asset purchases (QE4) including £60 billion of government bonds and £10 billion of corporate bonds. The purchase of corporate bonds was expected to provide more stimulus than government bonds because corporate bonds have higher yields, and so investors selling them could be more likely to invest in other risky securities. Furthermore, the Bank of England cut the policy rate from 0.5 to 0.25 percent and launched the Term Funding Scheme to provide banks with funding close to the policy rate and encourage the pass-through of the policy rate cut to lower lending rates.

**Effects of Unconventional Monetary Policy in the United Kingdom**

Several papers, summarized in Table 2, have sought to measure the impact of quantitative easing on UK government bond yields using event studies. This approach analyzes the responses of bond yields right after Bank of England announcements of new asset purchases. Results show that the £200 billion purchases under QE1 had a sizable impact on yields, lowering them between 50 and 100 basis points.

The effects on bond yields seem to have weakened over the subsequent quantitative easing announcements. This does not necessarily imply that quantitative easing is subject to decreasing returns, so that it becomes less effective in larger volumes. Instead, yields may not have moved as much in response to later announcements because market participants started to anticipate new rounds

with UK Treasury bills; the Extended Collateral Term Repo in June 2012 to provide banks with liquidity in case of exceptional market-wide stress; and the Fund for Lending Scheme in July 2012 to incentivize banks to increase domestic lending. For a review of the effects of liquidity facilities during the global financial crisis in advanced economies, see Gagnon and Hinterschweiger (2013).

*Table 2*
**Unconventional Monetary Policy Effects in the United Kingdom**

| *Study* | *Notes* | *Government bond yields* | *GDP* | *Inflation* |
|---|---|---|---|---|
| Joyce, Lasaosa, Stevens, and Tong (2011); Joyce and Tong (2012) | Event study of £200b QE1. Portfolio rebalancing channel is predominant. | -100 bps | | |
| Christensen and Rudebusch (2012) | Event study of £275b QE1/QE2. Portfolio rebalancing is predominant. | -47 bps | | |
| Bridges and Thomas (2012) | Money demand/supply model to analyze impact of £200b QE1. | -150 bps | +2% | +1% |
| Kapetanios, Mumtaz, Stevens, and Theodoridis (2012) | VAR responses to 100 bps reduction in government bond yields | | +1.5% | +1.25% |
| Baumeister and Benati (2013) | VAR responses to 50 bps reduction in government bond yields | | +3% | +2% |
| Churm, Joyce, Kapetanios, and Theodoridis (2015) | Event study of £175b QE2/QE3 and VAR estimates of macro impact | -45 bps | +0.5/0.8% | +0.6% |
| Weale and Wieladek (2016) | VAR responses to QE of 1% of GDP, estimated over QE1/QE2/QE3 | negligible | +0.25% | +0.32% |

*Note:* QE1, QE2, and QE3 are three consecutive rounds of quantitative easing. VAR is "vector autoregression."

of quantitative easing depending on the inflation and growth outlook (Joyce, McLaren, and Young 2012). For example, McLaren, Banerjee, and Latto (2014) analyze unexpected announcements by the Bank of England about changes in the maturity distribution of bond purchases and find that asset purchases within a certain maturity segment were associated with similar reductions in yields under both QE1 and QE2.

The literature has also tried to shed light on the channels through which quantitative easing in the UK affected government bond yields. Event studies (Joyce, Lasaosa, Stevens, and Tong 2011; Christensen and Rudebusch, 2012) suggest a minor role for the signaling channel since the expectation of long-term rates did not decline much in response to quantitative easing announcements. Therefore, the reduction in bond yields seemed mostly due to a compression in term premia. In contrast, the vector autogression analysis of Weale and Wieladek (2016) found that quantitative easing had a negligible effect on yields, while it stimulated the economy by reducing uncertainty.

The effects of quantitative easing on financial markets went beyond the effect on government bond yields. Joyce et al. (2011) find evidence that pension funds and insurance companies responded to the reduction in government bond yields

by increasing their holdings of corporate bonds, leading to a reduction in corporate yields. Furthermore, quantitative easing seems to have led to a moderate depreciation of the British pound. The effects on equity prices have been instead more difficult to identify (Joyce et al. 2011). Quantitative easing also had minimal effects on bank lending because UK banks were poorly capitalized and trying to deleverage (Butt, Churm, McMahon, Morotz, and Schanz 2014; Joyce and Spaltro 2014).

Regarding the implications of quantitative easing for growth and inflation, various studies find sizable positive effects. Most of the papers listed in Table 2 rely on vector autoregression estimates that trace out the economy responses to an increase in asset purchases by the central bank or to a given reduction in government bond yields. Estimates vary considerably across studies, with QE1 increasing real GDP between 1.5 and 3 percent and raising inflation between 1.5 and 4 percent.

Finally, evidence about the effect of forward guidance in the United Kingdom is very limited. Filardo and Hoffman (2014) show that the forward guidance announcements in 2013 and 2014 did not lead to a reduction in expected future rates or in government bond yields. However, forward guidance seems to have reduced the volatility of expected future interest rates and has likely contributed to keeping expected interest rates low, despite sustained GDP growth in 2014.

## Japan

In Japan, the initial monetary response to the global financial crisis was relatively weak, involving forward guidance announcements between 2010 and 2012 supported by limited asset purchases. The Bank of Japan delivered much stronger monetary stimulus after the election of Prime Minister Shinzo Abe in 2012, by adopting a 2 percent inflation target and launching very large quantitative easing programs in 2013 and 2014. In 2016, the Bank of Japan entered a third phase of monetary stimulus by introducing the "yield curve control" framework and charging negative interest rates on central bank reserves.

### Economic Developments and Monetary Policy Responses in Japan

At the onset of the global financial crisis in 2008, Japan had already gone through a long period of low growth and inflation dating back to the early 1990s. The crisis made things much worse. Japan's GDP started to contract in the second quarter of 2008, falling in a year by about 8.5 percent. The Bank of Japan responded by marginally lowering the policy interest rate from 50 to 10 basis points and providing liquidity to the banking sector. Economic growth resumed at a moderate pace in the second half of 2009, but the recovery stalled in late 2010.

On October 5, 2010, the Bank of Japan entered a first phase of unconventional monetary stimulus based on forward guidance and modest asset purchases. It started by clarifying its intention not to raise rates until "price stability is in sight" if no major financial risk materialized. Furthermore, the Bank of Japan announced the purchase of ¥5 trillion assets (later increased to ¥20 trillion, equal to about 4

percent of GDP) under a newly established asset purchase program. However, the economic recovery remained feeble because of the headwinds from weak global demand and the appreciation of the yen. With inflation stubbornly anchored around zero, in February 2012 the Bank of Japan reiterated its intention not to raise rates and to use asset purchases (further increased by ¥10 trillion) until inflation is expected to reach the "1 percent goal" in so far as this does not raise "significant risk" especially regarding financial imbalances.

The strength of monetary stimulus increased significantly after the 2012 election of Prime Minister Abe who called for fiscal stimulus, structural reforms, and much greater monetary accommodation. The Bank of Japan responded by introducing an explicit inflation target of 2 percent and announcing that it would use asset purchases and keep rates low to "achieve this target at the earliest possible time." In February 2013, it launched the first round of "quantitative and qualitative monetary easing" (QQE1). The program involved the open-ended purchases of ¥50 trillion Japanese government bonds and ¥1 trillion exchange-traded funds per year. After a sizable but temporary increase in inflation, deflationary pressures reemerged in late 2014, leading the Bank of Japan to increase asset purchases (QQE2) up to ¥80 trillion Japanese government bonds and ¥3 trillion in exchange-traded funds per year. QQE1 and QQE2 also involved the purchase of Japan real estate investment trusts in the amount of ¥30 and ¥90 billion, respectively.

Despite the large size of asset purchases, Japanese consumer prices remained broadly flat in 2015 partly due to a decline in oil prices and weak external demand because of the economic slowdown in China. Thus, the Bank of Japan provided additional monetary stimulus in 2016 through several major announcements. In January 2016, it announced the introduction of a negative interest rate of -0.1 percent on a portion of the reserve deposits that financial institutions held at the central bank. In September 2016, the Bank of Japan launched the "yield curve control" framework under which the Bank aims to control both short and long-term interest rates. Specifically, it announced that it would keep short-term rates on central bank reserves at –0.1 percent and continue to purchase Japanese government bonds to keep the 10-year yield around the current zero percent. Furthermore, the Bank of Japan introduced an "inflation-overshooting commitment" to keep expanding the monetary base until inflation exceeds the 2 percent target on an ongoing basis.

**Effects of Unconventional Monetary Policy in Japan**

To understand the challenges faced by monetary policy in Japan during the global financial crisis, it is helpful to review Japan's prior economic history. In the early 1990s, Japan faced a collapse in real estate and stock prices. This led to strong deflationary pressures that the Bank of Japan tried to offset by pioneering unconventional monetary policy tools. After bringing policy rates to zero in 1999, in 2001 the Bank of Japan embarked on a ¥35 trillion quantitative easing program (described in McCauley and Ueda 2009). Faced with an improved economic outlook, in 2006 the Bank of Japan exited quantitative easing and increased policy interest rates. Nonetheless, inflation remained around zero raising concerns that the Bank of Japan had

*Table 3*
**Unconventional Monetary Policy Effects in Japan**

| *Study* | *Notes* | *Government bond yields* | *GDP* | *Inflation* |
|---|---|---|---|---|
| Lam (2011); Ueda (2012) | Event study of BoJ announcement on Oct 5, 2010. | -10 bps | | |
| De Michelis and Iacoviello (2016) | VAR analysis of 2% inflation target announcement in January 2013 | | | +0.8% |
| Arai (2017) | Event study of announcements around QQE1 | -14 bps | | |
| Hausman and Wieland (2014) | Announcement effects and VAR analysis of QQE1 | -11 bps | up to +1% | |
| Kan, Kishaba, and Tsuruga (2016) | QQE1 effects using BoJ's large-scale macroeconomic model | | +0.6/4.2% | +0.3/1.5% |
| Michaelis and Watzka (2017) | VAR analysis of QQE, response to 1% increase in bank reserves | | Not significant | +0.2% |

*Note:* QQE is "quantitative and qualitative easing." QQE1 is the first round of QQE. VAR is "vector autoregression." BoJ is Bank of Japan.

not provided enough monetary stimulus and was not committed to reinflating the economy (Ito and Mishkin 2006; Bernanke 2000).

The Bank of Japan's announcement in 2010 to start a new asset purchase program and keep rates at zero led to a mild reduction in government bond yields, by about 10 basis points at the 10-year maturity, as shown in Table 3. However, mild deflation persisted, inflation expectations remained largely unchanged, and the yen exchange rate did not depreciate. The muted response of inflation was at least in part a reflection of the limited credibility in fighting deflation that the Bank of Japan had acquired during the previous decades. This perception was probably strengthened by the fact that asset purchases were scaled up in small amounts of about ¥5 trillion each during several subsequent announcements (Shirai 2017).

Forward guidance to keeping rates low until inflation resumed was also largely ineffective because of the entrenched deflationary pressures. If an economy is expected to experience deflation for a prolonged period, promises to keep policy rates low until a distant point in the future when inflation will eventually increase are largely ineffective to stimulate present output and inflation (Katagiri 2016).

The election of Prime Minister Abe reinvigorated the efforts of the Bank of Japan to provide monetary stimulus. The decision in January 2013 to adopt the 2 percent inflation target had significant effects on inflation (De Michelis and Iacoviello 2016), strengthened by the announcement of the QQE1 program. Despite market participants already anticipating new monetary easing measures, the announcement of such a large-scale program led to a further reduction in long-term yields, by about 11–14 basis points (Arai 2017; Hausman and Wieland 2014).

Regarding the ultimate effects of QQE1 on output and inflation, identification is particularly challenging because the government enacted a concomitant fiscal stimulus as part of the so-called "Abenomics" program. However, analyses using vector autoregressions tend to find positive, albeit moderate, effects on both inflation and output (Hausman and Wieland 2014; Kan, Kishaba, and Tsuruga 2016; Michaelis and Watzka 2017). QQE1 was also associated with an increase in stock prices and decline in corporate bond yields (Arai 2017). The impact on the exchange rate was marginal, probably because the yen had already depreciated sharply a few months earlier when Abe called for extraordinary monetary easing during his electoral campaign. Despite the extraordinary size of asset purchases, QQE1 and QQE2 failed to deliver a sustained increased in inflation. Long-term inflation expectations increase to only about 1 percent (De Michelis and Iacovello 2016). Furthermore, inflation declined again to zero in 2015, thus defying the Bank of Japan's intention to meet the 2 percent target by then (as announced at the launch of QQE1). This has renewed concerns about the inflation target not being credible, given the entrenched deflationary mindset.

Regarding the impact of negative interest rates, Japan's announcement in 2016 led to a downward shift in the yield curve that pushed 10-year yields below zero (IMF 2017a). Negative rates led also to a reduction in deposit and lending rates. Nonetheless, the impact on GDP and inflation has been hard to identify. Furthermore, negative rates appear to have put further strain on the profitability of several of Japan's regional banks.

## Lessons of Unconventional Monetary Policy: What Worked and What Did Not

The experience with unconventional monetary policy in the euro area, Japan, and the United Kingdom provides abundant evidence that these measures contributed to easing monetary policy conditions after the policy rate hit the zero lower bound. Most studies find significant cumulative effects of quantitative easing and forward guidance on the yields of long-term government bonds. Negative interest rates have also been effective in lowering bond yields. These effects are clearly visible on the short-end of the yield curve, which dipped below zero in the euro area and Japan after the implementation of negative policy rates. Furthermore, unconventional monetary policy has contributed to reducing corporate yields, raising stock prices, and depreciating the exchange rate. The experience with unconventional monetary policy of other countries not covered in this paper (such as Denmark, Sweden, and Switzerland) is qualitatively in line with this evidence (IMF 2017a; De Graeve and Lindé 2015).

Evidence on how unconventional monetary policy affected output and inflation is more limited. While the effects on financial variables can be assessed using high-frequency indicators, the impact on macro variables can only be observed over longer lags and is thus confounded by possible concomitant shocks. Constructing a

counterfactual scenario to tease out the effects of unconventional monetary policy is particularly difficult because these measures were often adopted in periods of exceptional financial distress—which are hard to model. That said, existing studies suggest positive impacts of unconventional monetary policies on both GDP and inflation.

The analysis also suggests that unconventional monetary policy tends to be more effective under three specific conditions.

First, the effects appear stronger in periods of heightened financial distress, which is consistent with economic theory that quantitative easing should have stronger price effects under segmented-market conditions. For example, in the United Kingdom the first round of quantitative easing at the peak of the crisis in 2009 seems to have had stronger effects on bond yields than subsequent rounds. Furthermore, unconventional monetary policy measures in the euro area have been particularly effective when they contributed to the reestablishment of market liquidity and elimination of redenomination risks. For instance, the reduction in spreads in peripheral European countries was very large after the "whatever it takes" statement and the announcement of the Outright Monetary Transactions program.[5]

Second, unconventional monetary policy is less effective when deflationary pressures are entrenched, so that the economy is expected to remain at the zero lower bound for a long period of time. The experience of Japan is emblematic. When the yield curve is already quite flat and close to zero even at long maturities, quantitative easing can have only a limited impact on further reducing yields. Forward guidance also tends to lose effectiveness. The promise to keep rates low in a distant future when the economy exits the zero lower bound has limited effects on current spending and investment decisions. Negative interest rates can provide some additional stimulus, but they are bounded below by concerns over their effects on bank profitability and on the threshold at which economic agents switch to cash.

Third, unconventional monetary policy requires that the central bank is credible in its attempt to provide sustained monetary accommodation. For example, the unconventional monetary policy measures adopted by the Bank of Japan between 2010 and 2012 had a muted impact on inflation probably because the central bank's commitment to deliver sustained inflation was undermined by decades of mild deflation. Stronger effects were associated with the unconventional monetary policy measures adopted in 2013 when Prime Minister Abe provided political backing for the Bank of Japan to provide aggressive monetary stimulus. Similarly, the Long-Term Refinancing Operations undertaken by the European Central Bank in the early years of the crisis had limited impact on lending rates and aggregate demand, partly because markets were skeptical about the determination and legal powers of the European Central Bank. The open-ended commitment under Outright

[5] In addition, the size of the quantitative easing surprise was essentially larger in the euro area during heightened financial stress, and the effects were probably commensurate with the size of the surprise. The difference may be due as much to the extent of the surprise as to the overall size of quantitative easing (including both its surprise and expected components). Indeed, Altavilla, Carboni, and Motto (2015) suggest that low financial distress can actually reinforce some monetary policy transmission channels (specifically, the duration and credit channel).

Monetary Transactions was much more effective in reducing sovereign yields and laying the foundations for the economic recovery.

So far, there is little evidence that the undesired side effects of unconventional monetary policy, such as the risk of runaway inflation and financial instability, have materialized. Inflation has been below target in Japan and the euro area for quite some time. In the United Kingdom, inflation was above target in several years, but this was mostly due to one-off shocks (such as value-added tax increases) and the effect of the exchange rate depreciation. Bank profitability has been under pressure but this seems more a byproduct of stagnant economic conditions than of unconventional policies. If anything, banks seem to have benefitted from the stimulative effects of unconventional monetary policy and the associated reflation of asset prices and improvement in asset quality. There was concern over banks and nonfinancial institutions taking on too much risk because of a "search for yield" strategy when interest rates are low for too long, but this does not appear to have materialized.

Concerns about central bank independence may have been more pertinent. The activism of central banks during the crisis has increased political scrutiny (for examples, see Davies 2016). This happened in part because the expansion of central bank balance sheets has increased the potential for capital losses that may in turn have significant fiscal implications. Furthermore, unconventional monetary policy, with its effects on asset prices and bank balance sheets, is often perceived to have greater distributional effects than traditional interest rate policy. Finally, greater political scrutiny can also arise simply because unconventional monetary policy at this scale is new and not fully understood. While politicians should retain their prerogative to amend central banks' goals and instruments as our understanding of monetary policy and economic circumstances evolve, central banks should retain operational independence.

Looking ahead, unconventional monetary policy measures should remain part of the monetary toolkit because they can provide crucial relief if a future severe recession pushes economies against the zero lower bound. However, unconventional monetary policy is far from a panacea—in particular, it faces limits once the yield curve is close to zero even at long maturities. Therefore, it is crucial to design policy responses that can complement unconventional monetary policy, such as fiscal stimulus. It would also be useful to think about how to reduce the likelihood that the zero lower bound may become binding in the first place. Possible measures include increasing the inflation target to leave more room for conventional cuts in nominal interest rates, or promoting structural reforms to boost growth.

The experience with unconventional monetary policies also leaves several questions unanswered. Two in particular deserve further research. First, there is the question of the relationship between the amount of securities purchased by the central bank and the effectiveness of quantitative easing. In that context, how does the total amount announced and the periodic flow of purchases interact in affecting market yields? The second question is whether central bank purchases should be limited to certain classes of assets, and what are the risks of widening those limits.

Put differently, when is it that quantitative easing ceases to be a monetary measure and enters the realm of fiscal policy?

■ *The views expressed in this paper are those of the authors and do not necessarily represent the views of the IMF, its Executive Board, or IMF management. The authors thank Ben Bernanke, Olivier Blanchard, Paul Cashin, Joseph Gagnon, Mark Gertler, Donald Kohn, Srobona Mitra, Maury Obstfeld, Louise Sheiner, Fergal Shortall, Frank Smets, Timothy Taylor, David Wessel, and participants at a seminar convened by the Hutchins Center on Fiscal and Monetary Policy at Brookings for helpful comments and discussions, and Hala Moussawi for excellent research assistance.*

## References

**Altavilla, Carlo, Giacomo Carboni, and Roberto Motto.** 2015. "Asset Purchase Programmes and Financial Markets: Lessons from the Euro Area." ECB Working Paper 1864, European Central Bank.

**Altavilla, Carlo, Domenico Giannone, and Michele Lenza.** 2014. "The Financial and Macroeconomic Effects of OMT Announcements." ECB Working Paper 1707, European Central Bank.

**Andrade, Philippe, Johannes H. Breckenfelder, Fiorella De Fiore, Peter Karadi, and Oreste Tristani.** 2016. "The ECB's Asset Purchase Programme: An Early Assessment." ECB Working Paper 1956, European Central Bank.

**Arai, Natsuki.** 2017. "The Effects of Monetary Policy Announcements at the Zero Lower Bound." *International Journal of Central Banking* 13(2): 159–196.

**Arce, Oscar, Ricardo Gimeno, and Sergio Mayordomo.** 2017. "Making Room for the Needy: The Credit-Reallocation Effects of the ECBs Corporate QE." Bank of Spain Working Paper 1743.

**Baumeister, Christiane, and Luca Benati.** 2013. "Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound." *International Journal of Central Banking* 9(2): 165–212.

**Bernanke, Ben S.** 2000. "Japanese Monetary Policy: A Case of Self-Induced Paralysis?" Chap. 7 in *Japan's Financial Crisis and Its Parallels to US Experience*, edited by Ryōichi Mikitani and Adam Simon Posen. Institute for International Economics.

**Bernanke, Ben S.** 2014. Remarks made in a conversation moderated by Liaquat Ahamed, January 16, 2014, in Washington, DC. In "Central Banking After the Great Recession: A Discussion with Federal Reserve Chairman Ben Bernanke on the Fed's 100th Anniversary." https://www.brookings.edu/wp-content/uploads/2014/01/20140116_bernanke_remarks_transcript.pdf.

**Borio, Claudio, and Leonardo Gambacorta.** 2017. "Monetary Policy and Bank Lending in a Low Interest Rate Environment: Diminishing Effectiveness?" BIS Working Paper 612.

**Borio, Claudio, Leonardo Gambacorta, and Boris Hofmann.** 2015. "The Influence of Monetary Policy on Bank Profitability." BIS Working Paper 514.

**Bridges, Jonathan, and Ryland Thomas.** 2012. "The Impact of QE on the UK Economy—Some Supportive Monetarist Arithmetic." Bank of England Working Paper 442.

**Burriel, Pablo, and Alessandro Galesi.** 2016. "Uncovering the Heterogeneous Effects of ECB Unconventional Monetary Policies across Euro Area Countries." Bank of Spain Working Paper 1631.

**Butt, Nick, Rohan Churm, Michael F. McMahon, Arpad Morotz, and Jochen F. Schanz.** 2014. "QE and the Bank Lending Channel in the United Kingdom." Bank of England Working Paper 511.

**Caballero, Ricardo J., Takeo Hoshi, and Anil K. Kashyap.** 2008. "Zombie Lending and Depressed Restructuring in Japan." *American Economic Review* 98(5): 1943–77.

**Cahn, Christophe, Julien Matheron, and Jean-Guillaume Sahuc.** 2014. "Assessing the Macroeconomic Effects of LTROS." Banque de France Working Paper 528.

**Campbell, Jeffrey, Charles Evans, Jonas Fischer, and Alejandro Justiniano.** 2012. "Macroeconomic Effects of FOMC Forward Guidance." *Brookings Papers on Economic Activity*, no. 1, pp. 1–80.

**Chen, Han, Vasco Cúrdia, and Andrea Ferrero.** 2012. "The Macroeconomic Effects of Large-Scale Asset Purchase Programmes." *Economic Journal* 122(564): 289–F315.

**Chodrow-Reich, Gabriel.** 2014. "Effects of Unconventional Monetary Policy on Financial Institutions." Brooking Papers on Economic Activity, no. 1, pp. 155–227.

**Christensen, Jens H. E., and Glenn D. Rudebusch.** 2012. "The Response of Interest Rates to US and UK Quantitative Easing." *Economic Journal* 122(564).

**Churm, Rohan, Michael Joyce, George Kapetanios, and Konstantinos Theodoridis.** 2015. "Unconventional Monetary Policies and the Macroeconomy: The Impact of the United Kingdom's QE2 and Funding for Lending Scheme." Bank of England Working Paper 542.

**Cœuré, Benoît.** 2013. "Outright Monetary Transactions, One Year On." Speech at the conference "The ECB and its OMT Programme," organized by the Center for Economic Policy Research, German Institute for Economic Research, Berlin, and KfW Bankengruppe, September 2, 2013. http://www.ecb.europa.eu/press/key/date/2013/html/sp130902.en.html.

**Cœuré, Benoît.** 2016. "Assessing the Implications of Negative Interest Rates." Speech at Yale Financial Crisis Forum, Yale School of Management, July 28. https://www.ecb.europa.eu/press/key/date/2016/html/sp160728.en.html.

**Cour-Thimann, Philippine, and Bernhard Winkler.** 2013. "The ECB's Non-Standard Monetary Policy Measures: The Role of Institutional Factors and Financial Structure." ECB Working Paper Series 1528.

**Cova, Pietro, Patrizio Pagano, and Massimiliano Pisani.** 2015. "Domestic and International Macroeconomic Effects of the Eurosystem Expanded Asset Purchase Programme." Banca d'Italia Working Paper 1036.

**Darracq-Paries, Matthieu, and Roberto A. De Santis.** 2015. "A Non-Standard Monetary Policy Shock: The ECB's 3-year LTROs and the Shift in Credit Supply." *Journal of International Money and Finance* 54: 1–34.

**Davies, Howard.** 2016. "Central Bank Independence and Taking Monetary Policy to the People." World Economic Forum. https://www.weforum.org/agenda/2016/10/central-bank-independence-and-taking-monetary-policy-to-the-people.

**De Graeve, Ferre, and Jesper Lindé.** 2015. "Effects of Unconventional Monetary Policy: Theory and Evidence." *Sveriges Riksbank Economic Review* no. 1, pp. 43–74.

**De Michelis, Andrea, and Matteo Iacoviello.** 2016. "Raising an Inflation Target: The Japanese Experience with Abenomics." *European Economic Review* 88: 67–87.

**De Pooter, Michiel, Robert F. Martin, and Seth Pruitt.** 2015. "The Liquidity Effects of Official Bond Market Intervention." International Finance Discussion Papers 1138, Board of Governors of the Federal Reserve System.

**Draghi, Mario.** 2014a. Remarks by President of the European Central Bank, at a press conference in Frankfurt am Main, September 4, 2014. https://www.ecb.europa.eu/press/pressconf/2014/html/is140904.en.html.

**Draghi, Mario.** 2014b. Remarks by the President of the ECB, at a press conference in Frankfurth am Main, November 6, 2014. https://www.ecb.europa.eu/press/pressconf/2014/html/is141106.en.html.

**Eser, Fabian, and Bernd Schwaab.** 2016. "Evaluating the Impact of Unconventional Monetary Policy Measures: Empirical Evidence from the ECB's Securities Markets Programme." *Journal of Financial Economics* 119(1): 147–67.

**European Central Bank (ECB).** 2017. *Financial Stability Review,* November 2017.

**Ferrando, Annalisa, Alexander A. Popov, and Gregory F. Udell.** 2015. "Sovereign Stress, Unconventional Monetary Policy, and SME Access to Finance." ECB Working Paper 1820.

**Filardo, Andrew, and Boris Hofmann.** 2014. "Forward Guidance at the Zero Lower Bound." *BIS Quarterly Review*, March.

**Fratszcher, Marcel, Lo Duca, Marco, and Roland Straub.** 2016. "ECB Unconventional Monetary Policy: Market Impact and International Spillovers." *IMF Economic Review* 64(1): 36–74.

**Gagnon, Joseph E., and Marc Hinterschweiger.** 2013. "Responses of Central Banks in Advanced Economies to the Global Financial Crisis." Chap. 3 in *Responding to Financial Crisis: Lessons from Asia Then, The United States and Europe Now,* edited by Changyong Rhee and Adam S. Posen. Columbia University Press.

**Gertler, Mark, and Peter Karadi.** 2011. "A Model of Unconventional Monetary Policy." *Journal of Monetary Economics* 58(1): 17–34.

**Ghysels, Eric, Julien Idier, Simone Manganelli, and Olivier Vergote.** 2016. "A High-Frequency Assessment of the ECB Securities Markets Programme." *Journal of the European Economic Association* 15(1): 218–43.

**González-Parámo, José Manuel.** 2011. "The ECB's Monetary Policy during the Crisis." Speech given at the Tenth Economic Policy Conference,

Málaga, October 21, 2011. Available at: https://www.ecb.europa.eu/press/key/date/2011/html/sp111021_1.en.html.

**Hausman, Joshua K., and Johannes F. Wieland.** 2014. "Abenomics: Preliminary Analysis and Outlook." *Brookings Papers on Economic Activity* no. 1, pp. 1–63.

**Hutchinson, John, and Frank Smets.** 2017. "Monetary Policy in Uncertain Times: ECB Monetary Policy Since June 2014." *The Manchester School* 85: e1–e15.

**International Monetary Fund (IMF).** 2017a. "Negative Interest Rate Policies—Initial Experiences and Assessments." IMF Policy Paper, March.

**International Monetary Fund (IMF).** 2017b. "Low Growth, Low Interest Rates, and Financial Intermediation." Chap. 2 of *Global Financial Stability Report: Getting the Policy Mix Right*, April 2017.

**Ito, Takatoshi, and Frederic S. Mishkin.** 2006. "Two Decades of Japanese Monetary Policy and the Deflation Problem." In *Monetary Policy with Very Low Inflation in the Pacific Rim*, pp. 131–202. NBER Book Series East Asia Seminar on Economics (NBER-EASE), vol. 15.

**Joyce, Michael A. S., Ana Lasaosa, Ibrahim Stevens, and Matthew Tong.** 2011. "The Financial Market Impact of Quantitative Easing in the United Kingdom." *International Journal of Central Banking* 7(3): 113–61.

**Joyce, Michael A. S, Nick McLaren, and Chris Young.** 2012. "Quantitative Easing in the United Kingdom: Evidence from Financial Markets on QE1 and QE2." *Oxford Review of Economic Policy* 28(4): 671–701.

**Joyce, Michael A. S., and Marco Spaltro.** 2014. "Quantitative Easing and Bank Lending: A Panel Data Approach." Bank of England Working Paper 504.

**Joyce, Michael A. S., and Matthew Tong.** 2012. "QE and the Gilt Market: A Disaggregated Analysis." *Economic Journal* 122(564): F348–F384.

**Kan, Kazutoshi, Yui Kishaba, and Tomohiro Tsuruga.** 2016. "Policy Effects Since the Introduction of Quantitative and Qualitative Monetary Easing (QQE)—Assessment Based on the BoJ's Large-Scale Macroeconomic Model (Q-JEM)." Bank of Japan Working Paper no. 16-E-15.

**Kapetanios, George, Haroon Mumtaz, Ibrahim Stevens, and Konstantinos Theodoridis.** 2012. "Assessing the Economy-wide Effects of Quantitative Easing." *Economic Journal* 122(564): F316–F347.

**Katagiri, Mitsuru.** 2016. "Forward Guidance as a Monetary Policy Rule." Bank of Japan Working Paper no. 16-E-6.

**Koijen, Ralph S. J., Francois Koulischer, Benoît Nguyen, and Motohiro Yogo.** 2016. "Quantitative Easing in the Euro Area: The Dynamics of Risk Exposures and the Impact on Asset Prices." Banque de France Working Paper 601.

**Krishnamurthy, Arvind, Stefan Nagel, and Annette Vissing-Jorgensen.** 2018. "ECB Policies Involving Government Bond Purchases: Impact and Channels." *Review of Finance* 22(1): 1–44.

**Lam, W. Raphael.** 2011. "Bank of Japan's Monetary Easing Measures: Are They Powerful and Comprehensive?" IMF Working Paper 11/264.

**McCauley, Robert N., and Kazuo Ueda.** 2009. "Government Debt Management at Low Interest Rates." In *BIS Quarterly Review, June 209: International and Financial Market Developments*, pp. 35–52.

**McLaren, Nick, Ryan N. Banerjee, and David Latto.** 2014. "Using Changes in Auction Maturity Sectors to Help Identify the Impact of QE on Gilt Yields." *Economic Journal* 124(576): 453–79.

**Michaelis, Henrike, and Sebastian Watzka.** 2017. "Are There Differences in the Effectiveness of Quantitative Easing at the Zero-Lower-Bound in Japan over Time?" *Journal of International Money and Finance* 70: 204–233.

**Mouabbi, Sarah, and Jean-Guillaume Sahuc.** 2016. "Evaluating the Macroeconomic Effects of the ECB's Unconventional Monetary Policies." Unpublished paper, Banque de France.

**Ongena, Steven, Alexander Popov, and Neeltje Van Horen.** 2016. "The Invisible Hand of the Government: 'Moral Suasion' during the European Sovereign Debt Crisis." ECB Working Paper 1937, European Central Bank.

**Rajan, Raghuram.** 2005. "Has Financial Development Made the World Riskier?" Proceedings of the Economic Policy Symposium, Jackson Hole, Federal Reserve Bank of Kansas City, 313–69.

**Shirai, Sayuri.** 2017. *Mission Incomplete: Reflating Japan's Economy.* Asian Development Bank Institute.

**Taylor, John.** 2016. "Independence and the Scope of the Central Bank's Mandate." *Sveriges Riksbank Economic Review*, no. 3 (Special issue: Rethinking the Central Bank's Mandate), pp. 96–103.

**Trichet, Jean-Claude.** 2009. "The ECB's enhanced credit support." Speech delivered at the University of Munich on July 13, 2009. https://www.ecb.europa.eu/press/key/date/2009/html/sp090713.en.html.

**Ueda, Kazuo.** 2012. "The Effectiveness of Non-Traditional Monetary Policy Measures: The Case of the Bank of Japan." *Japanese Economic Review* 63(1): 1–22.

**Weale, Martin, and Tomasz Wieladek.** 2016. "What are the Macroeconomic Effects of Asset Purchases?" *Journal of Monetary Economics* 79: 81–93.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 173–200*

# Ending Global Poverty: Why Money Isn't Enough

Lucy Page and Rohini Pande

The share of the world's population living below the global extreme poverty line ($1.90 in consumption per day) has plunged dramatically in recent decades, from 42 percent in 1981 to 11 percent in 2013 (PovcalNet 2018). This remarkable decline has buoyed hopes of continued reductions and created expectations about where future reductions will take place. In 2015, the international community enshrined the aim of ending extreme poverty by 2030 in the Sustainable Development Goals. The current literature talks of passing the "baton" of poverty reduction from China to India, and then to nations of Africa (Chandy, Ledlie, and Penciakova 2013; Commission on State Fragility, Growth, and Development 2018).

Historically, the quest to reduce poverty has relied on two levers: economic growth (the idea that "a rising tide lifts all boats") and the intentional redistribution of resources to the poor, either by the domestic state or foreign aid. In this essay, we argue that growth and aid, at least as currently constituted, are unlikely to suffice to end extreme poverty by 2030. To end extreme poverty sustainably and as quickly as possible, the states governing the world's poor need to be strengthened such that they are both accountable to the needs of the poor and have the capacity to meet those needs. The international development community should recalibrate the allocation of resources to increase accountability and state capacity.

■ *Lucy Page is a PhD student in Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Rohini Pande is Rafik Hariri Professor of International Political Economy, Harvard Kennedy School of Government, Cambridge, Massachusetts. Their email addresses are lucypage@mit.edu and Rohini_Pande@hks.harvard.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.173

doi=10.1257/jep.32.4.173

Underlying our argument is the changing global geography of need. Table 1 describes a dramatic shift in the concentrations of extreme poverty over the last 30 years. Panels A and B of Table 1 list the 20 countries that were home to the highest shares of the world's poor in 1987 and 2013, respectively. In 1987, 90 percent of the world's poor lived in low-income countries, while only 6.5 percent lived in middle-income countries. Only five of the 20 countries with the most people in poverty were middle-income. By 2013, over 60 percent of the world's poor lived in middle-income countries, and nine of the 20 countries with the highest concentrations of extreme poverty were middle-income. The eight middle-income countries that each have 1 percent or more of the world's poor are India, Nigeria, China, Indonesia, Pakistan, the Philippines, South Africa, and Zambia. In 2013, just under half of the world's extreme poor (49.3 percent) lived in these eight countries, which we refer to as the high-poverty middle-income countries.[1]

As the countries where the poor live have grown richer, the world's poorest people are increasingly split between two country groupings: low-income, fragile states like Afghanistan, Liberia, and the Democratic Republic of Congo (DRC); and the set of fast-growing but increasingly unequal high-poverty middle-income countries. Countries in these two groupings have often seen diverging growth trajectories over the last three decades. In 1987, China and the DRC had similar GDP. That year, China was home to more than one-third of the world's extreme poor, and DRC was home to 1.1 percent. By 2013, China had become a middle-income country and its share of the world's extreme poor had fallen tenfold, to just over 3 percent. Meanwhile, the share of the world's poor in DRC increased roughly sixfold. Low-income fragile countries are often trapped in cycles of erratic growth and misdirected aid, while high-poverty middle-income countries typify a global trend of falling cross-country inequality accompanied by greater within-country inequality (Hammar and Waldenström 2017). While a poor person in Liberia might live in a village where nearly everyone else is destitute, a growing share of the poor live in places like Dharavi in Mumbai—Asia's largest slum—in view of a high-rise reported to be the most expensive private residence in the world (Crabtree 2018).

What does this changing geography suggest about how to reduce poverty? In low-income countries, steady economic growth likely remains the most important tool for improving the lives of the poor. Yet instigating and sustaining such growth has often proven hard. Instead, the pattern seems to be one of erratic economic

[1]We use data on extreme poverty from PovcalNet (2018). Ferreira et al. (2016) provides a useful summary of PovcalNet's methods for estimating extreme poverty and of the $1.90 per day poverty line. To be consistent with the 2013 poverty data, we classify countries as low-, middle-, or high-income using the World Bank's country income classifications from FY2015, which are based on data from calendar year 2013. We do not classify any low-income countries that transitioned to middle-income status since FY2015, like Bangladesh and Kenya, as high-poverty middle-income countries. We continue to use FY2015 income classifications throughout the text and figures. The World Bank's PovcalNet released revised data on global poverty through 2015 in September 2018. These estimates suggest that the increasing concentration of the poor in relatively wealthy countries held true through 2015, when the World Bank estimates that 62.1 percent of the world's extreme poor lived in middle-income countries (using FY2017 income classifications).

*Table 1*
**Global Geographies of Extreme Poverty, 1987 and 2013**

| | *Millions in extreme poverty* | *Poverty headcount (%)* | *Share of world's poor (%)* | *Ranking in # of world's poor* | | *Millions in extreme poverty* | *Poverty headcount (%)* | *Share of world's poor (%)* | *Ranking in # of world's poor* |
|---|---|---|---|---|---|---|---|---|---|
| **A: 1987** | | | | | | | | | |
| **Low-income countries:** | | | | | **Middle-income countries:** | | | | |
| Total | 1,587 | 57.3 | 90.3 | | Total | 115.2 | 11.9 | 6.5 | |
| China | 659.5 | 60.8 | 37.5 | 1 | Brazil | 25.1 | 17.8 | 1.4 | 9 |
| India | 391.1 | 47.9 | 22.2 | 2 | Philippines | 15.4 | 26.9 | 0.9 | 12 |
| Indonesia | 122.5 | 71.4 | 7.0 | 3 | South Africa | 8.9 | 25.8 | 0.5 | 17 |
| Pakistan | 61.1 | 62.2 | 3.5 | 4 | Thailand | 8.3 | 15.4 | 0.5 | 19 |
| Nigeria | 56.8 | 64.5 | 3.2 | 5 | Mexico | 7.8 | 9.7 | 0.4 | 20 |
| Vietnam | 42.3 | 68.5 | 2.4 | 6 | | | | | |
| Myanmar | 36.5 | 94.4 | 2.1 | 7 | | | | | |
| Bangladesh | 33.4 | 33.9 | 1.9 | 8 | | | | | |
| Ethiopia | 24.6 | 56.6 | 1.4 | 10 | | | | | |
| Dem. Rep. of the Congo | 19.6 | 62.3 | 1.1 | 11 | | | | | |
| Tanzania | 15.0 | 64.7 | 0.9 | 13 | | | | | |
| Nepal | 12.7 | 72.6 | 0.7 | 14 | | | | | |
| Mozambique | 11.7 | 89.5 | 0.7 | 15 | | | | | |
| Uganda | 10.7 | 68.2 | 0.6 | 16 | | | | | |
| Sudan | 8.4 | 45.7 | 0.5 | 18 | | | | | |
| **B: 2013** | | | | | | | | | |
| **Low-income countries:** | | | | | **Middle-income countries:** | | | | |
| Total | 284.3 | 36.9 | 36.3 | | Total | 478.1 | 9.6 | 61.1 | |
| Dem. Rep. of the Congo | 54.1 | 75.9 | 6.9 | 3 | India | 210.4 | 16.5 | 26.9 | 1 |
| Ethiopia | 27.8 | 29.3 | 3.6 | 4 | Nigeria | 85.2 | 49.6 | 10.9 | 2 |
| Bangladesh | 26.5 | 16.8 | 3.4 | 5 | China | 25.2 | 1.9 | 3.2 | 6 |
| Tanzania | 23.3 | 45.9 | 3.0 | 8 | Indonesia | 23.6 | 9.4 | 3.0 | 7 |
| Madagascar | 17.9 | 77.8 | 2.3 | 9 | Pakistan | 12.7 | 7.0 | 1.6 | 13 |
| Mozambique | 16.9 | 63.9 | 2.2 | 10 | Philippines | 10.7 | 10.8 | 1.4 | 15 |
| Kenya | 15.1 | 33.7 | 1.9 | 11 | South Africa | 9.3 | 17.5 | 1.2 | 16 |
| Uganda | 13.5 | 35.8 | 1.7 | 12 | Zambia | 8.9 | 58.8 | 1.1 | 17 |
| Malawi | 11.7 | 70.4 | 1.5 | 14 | South Sudan | 7.8 | 69.5 | 0.9 | 20 |
| Mali | 8.6 | 52.0 | 1.1 | 18 | | | | | |
| Niger | 8.5 | 46.3 | 1.1 | 19 | | | | | |

*Note:* Panels A and B include the twenty countries with the highest share of the world's extreme poor in 1987 and 2013, respectively. Note that Panel B includes the full list of eight high-poverty middle-income countries in 2013, which we define as middle-income countries with at least one percent of the world's poor in 2013: India, Nigeria, China, Indonesia, Pakistan, Philippines, South Africa, and Zambia. We classify countries as low- or middle-income in 1987 and 2013 based on the World Bank's list of economies for FY1989 and FY2015, respectively; classifications for these years use income data from calendar years 1987 and 2013. We use data on extreme poverty from PovcalNet (2018).

growth episodes in which the periods of prosperity reached few (Acemoglu and Robinson 2012) or evaporated or reversed in periods of conflict (Jones and Olken 2008). In the absence of sustained growth, direct provision of cash and services to the poor is a critical, immediate way to alleviate poverty in low-income countries. Foreign aid will likely play a key role in providing these services.

In the second cluster of countries, growth has lifted millions out of poverty, but has also left millions behind amid increasing inequality (Alvaredo, Chancel, Piketty, Saez, and Zucman 2018). Continued growth may ultimately lift up those remaining millions, but it may do so much more slowly than is necessary. Ending poverty by 2030 in this second group of countries will require not just growth of the economy, but redistribution of new domestic resources to the poorest. Such redistribution must come in the form of services and institutions that the poor need for economic mobility. Because these countries receive relatively little foreign aid, domestic states will bear most of the responsibility for providing these services to the poor.

Perhaps because we typically identify the poor as those living below a certain income or consumption level, providing the poor with resources to exit poverty is often characterized in terms of cash transfers: that is, give the poor money and they will stop being poor. But poverty is more than just a lack of money, and escaping it requires more than cash. A variety of studies have shown that extreme poverty can be reduced by providing poor households with health, education, and access to a secure financial system and credit services, and by creating and enforcing regulation to ensure they are not exploited by shopkeepers, landowners, and employers.

The effective use of resources targeting extreme poverty, therefore, requires a complementary focus on investments in what we term "invisible infrastructure." We conceive of invisible infrastructure as the social and human systems that enable citizens to realize their capabilities and escape poverty. This comprises traditional elements of social infrastructure like health care and education but also, importantly, the incentive and information structures that bring the actions of those who control resources in line with the needs of the poor.

In advocating for investment in invisible infrastructure, we emphasize that the domestic state is the inevitable regulator, if not always the provider, of these services and institutions for the poor. First, the state is the only body with the mandate to provide certain critical institutions, like property rights and a monopoly of violence. Second, even where for-profit businesses and nongovernmental organizations are best-placed to provide specific services, such as micro-credit, the state alone can regulate the provision of these services to the poor. Third, the state has a role to play in spotting gaps in service provision and intervening in the absence of viable private sector providers. The final reason is pragmatic: the size of the state in each high-poverty middle-income country dwarfs foreign aid. While aid may play a role in providing invisible infrastructure and relieving immediate suffering in low-income countries, these countries too will graduate out of foreign aid as they grow richer; as they do, the state will increasingly bear responsibility for providing the invisible infrastructure and will likely still house large poor populations.

Therefore, enabling the provision of invisible infrastructure requires building capable and accountable domestic states. How can the international development community best deploy its resources to help?

A key part of the task at hand is to ensure that aid policies strengthen domestic institutions rather than undermine them. Especially in low-income countries, aid agencies often bypass messy, corrupt states and instead channel funds through a cadre of nongovernment organizations, contractors, and other nonstate actors. There are reasons for this. Doing so may be necessary on occasion, as, for instance, when delivering humanitarian aid after a natural disaster. Also, donor-country politicians may find it hard to justify working with governments seen as corrupt or compromised. But in the long term, aid transfers that bypass the state may fail to improve—and in some cases may even harm—the state's capacity to provide invisible infrastructure to its citizens. Even in the short term, cutting out the domestic state inhibits the use of two vital tools: local information about what works in context, and mechanisms for taking citizen preferences into account. The loss of these tools can damage long-term prospects for poverty reduction, because people who feel they have no voice in development may be less willing to support it by paying taxes.

We argue, therefore, that a sustainable end to global poverty will require that the international development community and civil society organizations invest resources in interventions that can help build capable, democratic state institutions. Some guidance on successful interventions comes from recent empirical contributions in the political economy of development literature, which support an agency perspective on government functioning: governments comprise individuals interacting along a *human chain* of command. Governance failures like corruption and leakage of funds reflect failures to resolve misaligned incentives and informational asymmetries along this human chain (for an overview, see Finan, Olken, and Pande 2017). Designing such reforms requires insights from the fields of political economy and mechanism design, as well as a theory of government that allows the disempowered to act as principal. Ultimately, it is democracy, done right, that best allows citizens to demand what they need to end poverty.

## Can We Rely on Growth to End Poverty?

Economic growth has significantly lowered global poverty (Kraay 2006; Dollar, Kleineberg, and Kraay 2016). China alone was home to three-quarters of the 1.12 billion people lifted out of extreme poverty worldwide between 1981 and 2013, when it grew at an average rate of 10 percent per year. India grew at an average annual rate of 6.2 percent over the same period, and it had about 190 million fewer people in extreme poverty in 2013 than in 1981; Indonesia, which saw average growth of 5 percent, had 92 million fewer.[2]

[2]Authors' calculations using poverty data from PovcalNet (2018) and data on GDP growth from World Development Indicators (2018b).

One would hope for growth to produce similar gains in today's low-income countries, lifting their citizens up the income ladder. Yet freeing low-income countries from cycles of conflict, natural disasters, and recession has proved challenging, and it is not clear when and how sustained economic growth will arrive as a driver of substantial declines in absolute poverty (Jones and Olken 2008).

In today's high-poverty middle-income countries, economic growth will certainly continue to reduce extreme poverty. But poverty can have a long half-life in the presence of inequality. In India, which in 2013 contained the largest share of the world's extreme poor, over 100 billionaires lived alongside 210.4 million people in extreme poverty in 2013. This imbalance arises from unequal growth: India's top 10 percent of incomes captured 66 percent of growth between 1980 and 2016, while the bottom 50 percent captured only 11 percent (Alvaredo et al. 2018). Furthermore, growth often discriminates: in India, disadvantaged social groups (Hindu lower castes and Muslims) came to represent 55 percent of the poor in 2011—up from 44 percent in 1983.[3] Assouad, Chancel, and Morgan (2018) provide congruent evidence for high-poverty middle-income Brazil and South Africa.

At minimum, these trends in inequality suggest that growth does not reduce poverty as quickly as the equitable distribution of resources might permit. A stronger conjecture is that as the poor are increasingly drawn from socially disadvantaged groups, discrimination and inequality-fueled conflict will weaken growth's ability to raise the incomes of the poor (for instance, Mitra and Ray 2014). In either case, we argue that ending extreme poverty as quickly as possible in both low-income and high-poverty middle-income countries will require coupling growth with mechanisms to directly redistribute resources to the poor in the forms that they need.

**Can Physical Infrastructure or Cash Suffice to End Poverty?**

Consider people living in a remote rural village separated from the nearest city by a river, a forest, and steep mountains (Castañeda et al. 2018). What would it take for them to gain enough income to exit poverty? A traditional model of economic development that focuses on raising earnings might call for investments in physical infrastructure—perhaps the construction of a road to allow them to sell goods or make their way to the city for work.

How would this road get built? A bridge across the river isn't enough, nor is a tunnel under the mountain, nor is a way through the forest. Rather, all of these things need to be constructed and linked into a viable path from the village to the city. Private companies might build some of these elements—but the road as a whole is expensive and difficult enough that the state will need to coordinate and, likely, subsidize parts of construction.

Rural roads, by themselves, may not bring jobs to the village (Asher and Novosad 2018). But for many male villagers, the road might still enable an escape from poverty. They can get on the bus, go to the town, and find manual work there,

[3] Authors' calculations using rounds 61 and 68 of India's National Sample Survey (NSS 2016).

perhaps enabled by a free bus ticket (Bryan, Chowdhury, and Mobarak 2015). But neither the road nor a free bus ticket may suffice for a poor female villager, since social norms and safety concerns may prevent her from getting on the bus in the first place. Perhaps she can make some money selling vegetables at the bus stand, but beyond that, the road will do very little for her directly.

A more modern model of development might advocate sending the woman cash. Indeed, modern digital technology has made it possible to transfer the equivalent of a $1.90 a day directly to a poor woman living in a remote rural area, bypassing corrupt and ineffective intermediaries. For now, assume that this cash could be funded either by domestic redistribution from the wealthy or through foreign aid, though we will document later that the poor's access to these funding pots varies substantially with country income status.

Many cash transfer programs have been shown to make life significantly better for the poor. For example, GiveDirectly is a transfer program that allows individuals in rich countries to send money directly—typically in the range of several hundred dollars—to poor Africans. Haushofer and Shapiro (2016) find that these grants significantly increased household consumption nine months after they were granted, and a longer-run study suggests that the gains in assets persist three years later (Haushofer and Shapiro 2018).[4] But if the woman's child becomes ill or her house is flooded, she may be knocked back down into poverty. Cash transfers can help her to pay her children's school tuition, but what if the higher-quality academy run by a nongovernment organization is full or too far away? Furthermore, if we expand our perspective beyond the one woman to all of the world's extreme poor, then cash transfers likely become too expensive to be a sustainable answer for single-handedly ending global poverty.[5] (Relatedly, Hanna and Olken in this symposium discuss the tax implications of implementing a universal basic income program versus a cheaper targeted transfer program.)

Thus, a woman's road out of poverty is different from a man's, but similarly includes many components. She needs a way to save money and smooth consumption, receive remittances, hide money from friends and family when she needs to, and provide them with informal insurance when she can. Even if she is provided a bank account, if she isn't trained in its use she may well cash out any transfers, keep the account at zero balance, and leave the benefits of that account unrealized (Field, Pande, Rigol, Schaner, and Moore 2016). She needs to have quality education for her children so they can earn more than she does, escape poverty themselves, and

[4] Blattman, Fiala, and Martinez (2014, 2018) find similar mid-term evidence from Uganda's Youth Opportunity Program, but that in the longer run (nine years) the control and treatment groups converge in employment, earnings, and consumption.

[5] Asset transfer programs, or "ultra-poor graduation" programs, get closer to filling the broad set of needs necessary for a permanent escape from poverty, and have been shown to have very significant positive impacts on household income (BRAC 2013; Banerjee, Duflo, Chattopadhyay, and Shapiro 2016; Bandiera et al. 2017). However, a back-of-the-envelope calculation based on Banerjee et al. (2016) suggests that scaling up graduation programs to reach the 783 million people in extreme poverty worldwide would cost between $288 billion and $864 billion in US purchasing-power-parity 2014 dollars.

care for her in old age (Montenegro and Patrinos 2014). She needs access to health care in emergencies and the ability to invest early in her children's nutrition (Currie and Vogl 2013). Furthermore, she needs protection—from relatives who might cheat her from her land, from industrial pollution that might destroy her health and her ability to work, from warlords who would forcibly recruit her children. She needs the freedom to use all of these facilities strategically and at her own volition. She needs far more than either a road or $1.90 a day.

**The Need for Invisible Infrastructure**

Coordination and provision of these services will require investments in physical infrastructure like roads, schools, healthcare centers, and cellphone towers. It will also require a set of institutions. North (1990) famously characterized institutions as "humanly devised constraints that structure political, economic and social interactions," setting the *rules of the game* for an economy. He argued that well-functioning institutions enable prosperity by ensuring secure property rights, facilitating complex contractual arrangements, aligning norms to allow markets to function effectively, and so on.

But physical infrastructure and the rules of the game, as implemented, often fail to create the incentives for state and nonstate actors to ensure that the poor receive the services they need to escape poverty. Ultimately, the efficacy of a country's institutions and physical infrastructure depends on how people within each body choose to allocate resources, implement policies, regulate private sector providers, and respond to citizen grievances.

Helping poor and isolated people out of poverty will require more than physical infrastructure and the setting of formal rules of the game. We also need ways to ensure that these basic components translate into the services that the poor need for economic mobility, and that those services work for them and are not coopted by the powerful or derailed by traditional social structures. For this, we need to ensure that the human infrastructure that undergirds service provision selects qualified and motivated individuals to staff these institutions and then gives them the incentives and information they need to do their jobs well. This invisible human and social infrastructure is critical for enabling the poor to realize their capabilities for economic mobility. Ending extreme poverty as quickly as possible will require coupling economic growth with the direct provision of this invisible infrastructure to the poor.

## Providing Invisible Infrastructure: Aid, the State, and Private Players

Who are the providers of invisible infrastructure? Some components of invisible infrastructure, like the monopoly of violence and security of property rights, fall squarely within the ambit of the domestic state in all but the most fragile countries. But other components—such as access to health care, education, and financial services—are often provided by a range of players.

**A Role for the Domestic Private Sector**

Today, various private social sector organizations support the delivery of services for the poor, including for-profit companies, nongovernment organizations, and social enterprises. In some cases, the client is a domestic government that explicitly contracts out service provision.

Consider the provision of financial services, such as loans, savings, and insurance. In the 1960s and 1970s, governments of many developing countries created large-scale social banking programs to provide credit and bank accounts to poor citizens. While state-led financial inclusion efforts did often reduce poverty (Burgess and Pande 2005), these programs were plagued by low repayment rates (Besley and Coate 1995) and elite capture (Cole 2009). As a result, governments began disbanding many of these programs in the 1980s. Lending to the poor was considered a low-profit and risky activity, so private sector financial institutions failed to step in. Nongovernment organizations then played a key role in developing viable financial products for the poor: in the 1980s, Bangladeshi nongovernment organizations—BRAC[6] and Grameen Bank—innovated by introducing the group lending microfinance model. By 2010, private sector microfinance institutions reported about 26.7 million clients, particularly women and the poor (Srinivasan 2010; Khandker 1998).

While these private sector initiatives largely avoided the pitfalls of corruption and inefficiency, concerns about unregulated lending by microfinance institutions grew. Critics warned that a for-profit drive was incentivizing frontline agents to overload the poor with loans (CGAP 2010). When Compartamos, a for-profit microlender in Mexico, became publicly traded in 2007 and created enormous profits for its private investors, several commentators raised concerns of "mission drift" among private microfinance providers (Ashta and Hudon 2012).

In India, these concerns came to a head in October 2010 following news reports linking a series of suicides to allegedly coercive loan collection policies in the state of Andhra Pradesh. The state government responded with an ordinance imposing a set of restrictive regulations on microfinance institutions. This brought the microfinance industry to a sudden halt; the poor were left with no access to credit and suffered large decreases in both household earning and consumption (Breza and Kinnan 2018). The experience of Andhra Pradesh shows that, while private microfinance could offer large benefits to the unbanked poor, sensible government regulation is important, too.

In education—another key pillar of invisible infrastructure—the private sector may provide better service to the poor at reduced cost in the short run (Muralidharan and Sundararaman 2015). Educating the poor is unlikely to be profitable, however, and so private schools typically require state funding to serve the poor. Studies of such efforts have shown promise, but have also revealed some perverse effects. Under the Partnership Schools for Liberia (PSL) program, the Liberian government outsourced management of 93 public schools to eight private contractors, including Bridge International, a for-profit company operating over 500 schools across Africa and India.

[6] BRAC originally stood for Bangladesh Rehabilitation Assistance Committee, but now stands for Building Resources Across Communities. However, the organization is usually referred to by its acronym.

Under the outsourcing scheme, school admission is free, PSL teachers are paid by the government, and contractors cannot screen students by ability or other characteristics. Romero, Sandefur, and Sandholtz (2017) conducted a randomized evaluation of the project and found that contracted schools performed significantly better than regular public schools one year after the intervention, with higher teacher attendance and better student performance in English and math. However, one provider, Bridge International Academies, pushed excess students and worse-performing teachers to government-run schools, subverting policymakers' efforts to maximize access to quality education. This provider was the only one whose funding was not linked to the number of students enrolled and whose contract did not forbid direct dismissal of teachers.

In sum, while private players can often play a substantive role in providing invisible infrastructure, ensuring that those services reach the poor will require that the state remains as an active regulator.

**A Role for Foreign Aid**

Now consider foreign aid, or official development assistance. The total volume of aid has increased substantially over time, rising nearly fivefold between 1960 and 2016, from about $32 billion to $158 billion in 2016—both in constant 2016 US dollars (OECD 2018). The decline of poverty in the same period has raised the prospect of aid as a dominant force in ending deprivation. Indeed, if the cost of ending poverty were simply the dollar value of the shortfall between the poor's daily consumption and $1.90, then the problem would appear to have been solved; official development assistance has exceeded this value since 2006 (Chandy, Noe, and Zhang 2016).

While early aid flows focused almost exclusively on promoting economic growth, donors began targeting a significant fraction of aid to social sectors in the 1970s (Streeten 1979). Figure 1 plots the distribution of aid by its purpose over time; we focus, in particular, on the fluctuations of "economic" aid (aid for growth) compared to "social" aid (aid for basic social services like education, health care, water sanitation, and food assistance). Social aid made up about 20 percent of average annual aid spending during the 1970s, half the percentage going to economic aid over the same period. Social aid stagnated at around 20 percent of total flows during the structural adjustment era of the 1980s, when conditions from the World Bank and IMF stipulated aid only if borrower countries tightened social spending. Social aid began rising again in the mid-1990s and since 1996 has typically surpassed economic aid as a share of total official development assistance, at between 22 and 31 percent each year.

Some of the aid investments in social infrastructure have been successful. Consider global health, for instance. The world has seen an unprecedented improvement in health outcomes since World War II (Deaton 2013). Average life expectancy worldwide rose from 46 to 69 between 1950 and 2011, and child and infant mortality rates fell in every single country in the world during that period (Bloom 2011). Global health inequality has fallen faster than income inequality (Becker, Philipson, and Soares 2005), as low-income countries see falling child mortality driven by public health advances in access to clean water, immunization, and sanitation. In several prominent cases, foreign assistance has contributed to these improvements.

*Figure 1*
**Official Development Assistance (ODA) by Purpose over Time**
*(in billions of constant 2011 US dollars)*

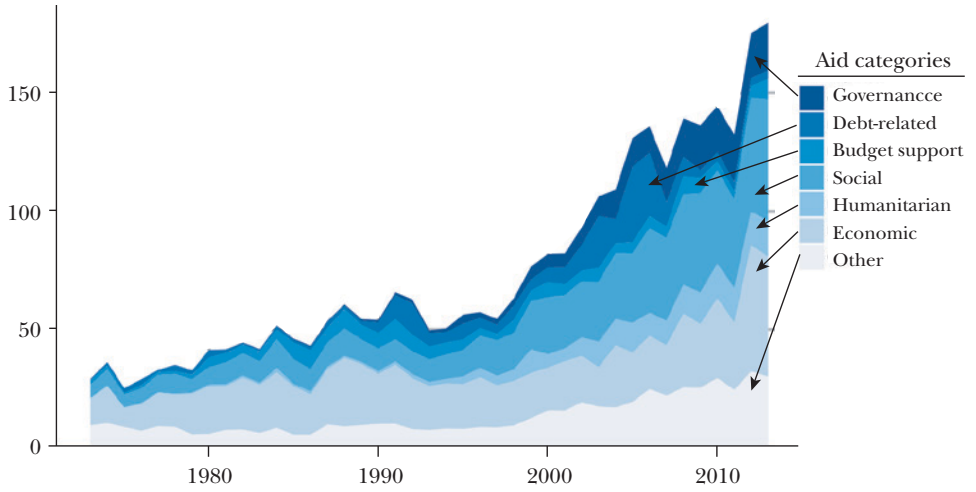


*Source:* Authors use data on official development assistance (ODA) flows from AidData (2017).
*Note:* We classify aid purpose according to AidData's assignment of OECD Creditor Reporting System (CRS) purpose codes. Economic aid includes aid for productive sectors like agriculture, mining, construction, transport and storage, communications, energy generation and supply, and banking and financial services (1-digit CRS codes 2 and 3). Social aid includes aid for education, health, population policies and reproductive health, water supply and sanitation, and other social infrastructure and services (2-digit CRS codes 11, 12, 13, 14, 16, 42, and 52). Humanitarian aid includes aid for emergency response and preparedness (1-digit CRS code 7). Debt-related aid includes debt forgiveness, rescheduling, and refinancing (1-digit CRS code 6). Governance aid includes institutional capacity building, public sector financial management, civil service reform, and conflict prevention and resolution (2-digit CRS code 15).

The eradication of smallpox—the only human disease ever successfully eradicated—offers a vivid example of the possible gains from aid for global health. In the mid-1960s, smallpox still infected 10 to 15 million people each year (Crosby 1993). In 1967, the World Health Organization established the Intensified Smallpox Eradication Program, which began a massive campaign of vaccination and case surveillance. Outbreaks tapered off, the last endemic case of smallpox was recorded in Somalia in 1977, and in May 1980 the World Health Assembly declared smallpox the first disease ever eradicated. In total, international donors provided $98 million, with about $200 million provided by recipient countries themselves (WHO 2011).[7]

[7]There are also some, though arguably fewer, success stories in aid for education. Take the case of Ghana—in the decade and a half after 1986, the World Bank provided the Rawlings government with technical assistance and loans totaling $260 million for primary education. With the Bank as catalyst, other donors joined the effort, more than doubling the Bank's contribution. Between 1987 and 2000, primary enrollment increased by over 60 percent and was accompanied by genuine learning gains: in identical English tests, two-thirds of primary school graduates in 1988 could not outperform random guessing, but in 2003, the figure was less than 20 percent (World Bank 2004).

Aid-funded health campaigns like the one against smallpox are often designed as "vertical" programs, a type of campaign that targets a particular need and is funded and overseen by external donors. Vertical initiatives may get rapid results by working outside of weak public systems, which may suffer from shortages of trained staff, funding, and equipment or other bureaucratic delays (Atun, Bennett, and Durán 2008).

Moreover, these campaigns may help to satisfy donors' preferences to safeguard aid from corrupt intermediaries or recipients (Dietrich 2013; Acht, Mahmoud, and Thiele 2015). Less than half of social aid in 2013 (47 percent) was channeled through recipient states (AidData 2017).

However, vertical public health programs may not contribute to the strengthening of domestic "horizontal" primary healthcare systems (Oliveira-Cruz, Kurowski, and Mills 2003). In some cases, vertical programs, or more generally, delivering social services through a cadre of nonstate actors, may even weaken public service delivery by diverting civil servants, funding, and political interest away from state structures and into parallel systems.

We hypothesize that while vertical programs are well-suited to solving problems that can be addressed with short-term and targeted attention, like inoculating children against smallpox or polio, when it comes to more diffuse projects that require working across systems, success or failure can depend on whether aid complements, or substitutes for, the state.[8] In these cases, aid interventions are more likely to have long-run success if they are designed and applied with state buy-in and eventually turned over to domestic actors.

**Why Aid Should Not Bypass the State**

A concrete example, focusing on the global philanthropic initiative to eradicate hookworm at the start of the 20th century, can help fix the ideas developed above. Unlike smallpox, which requires a single vaccination, eradicating hookworm requires both treating the infected and preventing reinfection by constructing modern sanitation systems and changing people's habits. John D. Rockefeller established the Rockefeller Sanitary Commission (RSC) for the Eradication of Hookworm Disease in 1910, when hookworm infections were widespread across the southern United States. The Rockefeller campaign treated hookworm disease in about 400,000 people across the South and ran large public education campaigns on the importance of hygiene and the symptoms of infection. As the campaign wound down, state and local governments took over responsibility for sponsoring construction of latrines, as well as dispensing hookworm medication. The campaign produced large

[8]The fact that many public health and education problems require continued attention, rather than a one-off fix, could explain why we observe mixed results in sector-level estimates of the impact of aid in health and education. While Mishra and Newhouse (2009) find that health aid significantly decreases infant mortality, Williamson (2008) finds no significant impact of health-related aid on infant mortality, and neither Williamson (2008) nor Wilson (2011) finds that health aid reduces mortality in general. Birchler and Michaelowa (2016) and Dreher, Nunnenkamp, and Thiele (2008) find that education-sector aid increases primary school enrollment.

and lasting reductions in hookworm infection rates, thereby increasing school attendance, literacy, and the returns to education among school children (Bleakley 2007).

Following the success of the Sanitary Commission, the Rockefeller Foundation was created in 1913 with the specific intent of developing a global health program. Its initial efforts were concentrated in Latin America and the British Caribbean, though it quickly expanded throughout the tropical world. But its international efforts on hookworm eradication had mixed results.

The Foundation's work in Costa Rica was a success. Between 1914 and 1921, it tested over 300,000 Costa Ricans and treated over 65,000. A primary reason for the success of this campaign was that it learned from the mistakes of an earlier government-led effort implemented by local doctors. The head of Rockefeller's program and governmental partners centralized control, worked through the public school system, and broadened the set of implementing actors to include schoolteachers, community leaders, and priests. Palmer (2003) argues that this use of the public school network was central to the program's success: by the end of Rockefeller Foundation's operations, more of the foundation's funds were being channeled through the country's School Health Department, rather than directly, to hookworm relief. The campaign contributed to centralizing and modernizing health care in Costa Rica and paved the way for the foundation of a national health ministry in 1927, the third in Latin America.

In contrast, the results in India were disappointing. In the 1920s, the Rockefeller Foundation initiated a large hookworm treatment campaign in Madras. Ten years after the campaign ended, the hookworm infection rate remained at about 90 percent (Kavadi 2007). While the campaign made substantial short-term progress by dispensing medications, it could not enduringly reduce hookworm infections without large-scale improvements in domestic sanitation systems. John F. Kendrick, one of the leaders of the Madras campaign, acknowledged that any such shift would rely on the domestic state, noting that "sanitation would have never reached its present state of perfection even in England had government not taken a hand in the matter" (as quoted in Kavadi 2007).

Vertical initiatives can beat smallpox, but building invisible infrastructure to reduce poverty is more akin to conquering hookworm: it requires the poor person to interact with a broad set of different systems. If we are to provide the poor with the systematic array of services needed to escape poverty, then we must invest in building domestic states that have the capacity to monitor and coordinate provision of services by nonstate actors and, when necessary, to provide services directly to the poor.

Yet foreign aid has historically not devoted much attention to building accountable and effective states. One prominent manifestation of this tendency—which we highlighted earlier using the example of vertical health programs—is that aid initiatives frequently bypass relatively weak states, instead delivering resources through a network of nonstate actors, like international and domestic nongovernmental organizations, multilateral organizations, public-private partnerships, and private contractors. State bypass is most common in humanitarian aid, where a fast response is key: only about 7 percent of humanitarian aid commitments were

implemented through recipient states in 2013.[9] In some cases, aid agencies may have no choice but to deliver aid through nonstate systems immediately following conflict or natural disasters. But if aid continues to take this route as countries transition from emergency to recovery, states will likely fail to develop the institutional capacity necessary to oversee service delivery in the long run.

Aid projects that bypass the state also lack built-in mechanisms for accountability to recipients; rather, aid projects may be accountable primarily to donors. Bypassing the accountability mechanisms built into the social contract of the state, especially democratic states, may leave projects ill-informed by on-the-ground realities and citizen preferences and, therefore, less likely to meet citizen needs.[10]

All of this assumes that aid is even reaching the poor where they live. With an increasing mismatch between the countries that contain large fractions of the world's poor and the countries that receive large amounts of aid, the role of the domestic state in building invisible infrastructure becomes even more crucial.

While aid remains a sizable share of GDP in many low-income countries, the 61 percent of the world's poor living in middle-income countries receive relatively little aid per capita. Figure 2 plots 2016 official development assistance per person in extreme poverty, assuming a constant distribution of global poverty between 2013 and 2016. We include all low- and middle-income countries that were home to at least one percent of the world's poor in 2013; for middle-income countries this is the set of eight high-poverty middle-income countries. Some middle-income countries, like Pakistan, received substantial aid in 2016, but China and Indonesia were net aid *donors* in 2016, and India, the Philippines, and Nigeria received nearly the lowest net aid per poor person among all aid-receiving countries. The politics of this aid allocation, with aid targeting poor countries rather than poor people, are unlikely to change. Particularly in times of austerity, citizens of rich countries are unlikely to stomach giving aid to countries that give aid themselves, or that have the resources to invest in "vanity" projects. After the state-run Indian Space Research Organization announced plans to launch a rocket carrying 103 satellites in January 2017, the United Kingdom's popular *Daily Mail* tabloid ran an article titled, "India Boasts of Satellite Launch (as We Hand Them £54m of Aid)."

In today's high-poverty middle-income countries, low aid flows mean that the domestic state already bears the bulk of responsibility for providing invisible

[9] Authors' calculations using data on official development assistance from AidData (2017). We classify aid as being channeled through the public sector if it has an OECD creditor reporting system channel code beginning with 1.

[10] A classic example was PlayPumps International's merry-go-round water pump, which was based on the idea that children play on a merry-go-round, causing water to be pumped from the ground into an elevated tank for storage. This project received widespread international coverage and attracted significant aid: for instance, in 2006 the US President's Emergency Plan for Aids Relief (Pepfar) announced a $60 million public-private partnership with PlayPumps International, with $10 million to directly come from the US government. However, in reality, installing the pumps was expensive, children were not always keen to volunteer their labor at times of high demand (early morning and evening), and the complexity of the mechanism rendered local maintenance impossible. In 2010, PlayPumps International shut down operation.

*Figure 2*
**Net 2016 Official Development Assistance per Person in Extreme Poverty**

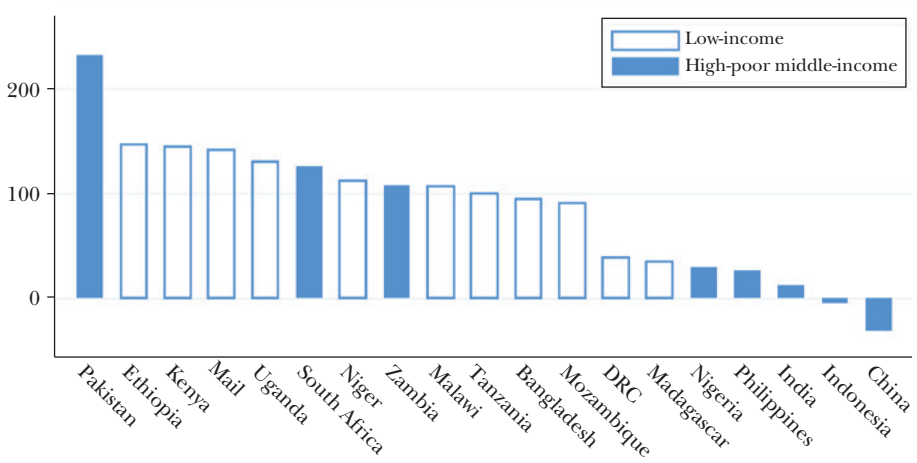


*Source:* Authors use data on official development assistance from World Development Indicators (2018a) and use data on extreme poverty from PovcalNet (2018).
*Note:* We use data on net official development assistance in 2016 but calculate aid per person assuming that the number of people in extreme poverty by country stayed constant between 2013 and 2016. Our sample comprises high-poverty middle-income countries (middle-income countries with at least one percent of the world's poor in 2013) and low-income countries that were home to at least one percent of the world's poor in 2013. To be consistent with our classification of high-poverty middle-income countries in Table 1, we classify country income status according to the World Bank's FY15 list of economies, which is based on data from calendar year 2013. DRC is the Democratic Republic of the Congo.

infrastructure to the poor. As low-income countries grow richer, they likely will also lose access to foreign aid. If aid today fails to contribute to building capable states in those countries, they may be left with weak institutions when aid dries up. Recent literature in economic history demonstrates how institutional persistence can influence a country's longer-term development (for instance, see Dell, Lane, and Querubin 2017). In addition, if citizens lack ownership of the process by which programs are decided, they may be less willing to pay taxes to fund those programs in the future.[11]

The domestic states of countries where the world's poor live will increasingly bear responsibility for provision of invisible infrastructure. Thus, building pathways out of poverty for the millions still in extreme poverty will require strong domestic states, and the role of aid should be to support, rather than substitute for, their institutions.

[11] Conversely, Weigel (2018) uses a field experiment in the Democratic Republic of Congo to show that citizens will respond to increased tax collection by participating more in politics. Citizens in neighborhoods where a door-to-door tax campaign took place increased political participation by 5 percentage points (28 percent), attending government-hosted townhall meetings and submitting suggestion cards evaluating government performance.

## What Are the Governance Challenges in Providing Invisible Infrastructure?

To propose ways to support developing states in the provision of invisible infrastructure, we must consider the governance challenges they face. We can approach them under two broad categories: problems of capacity and problems of will.

### Problems of Capacity

It takes money to run programs for the poor, and countries in both of our clusters of poverty typically lack the fiscal capacity to collect and spend resources at the scale needed to provide services to their populations. Weak fiscal capacity may be expected in the poorest countries, where low levels of economic activity, combined with the state's low ability to tax, result in a lack of funds to run programs for the poor. But Figure 3 shows that tax capacity in most high-poverty middle-income countries is as low as in many low-income countries. In 2015, high-poverty middle-income countries collected 16.9 percent of their GDP as tax revenue, versus 13.3 percent in low-income countries.[12]

To the extent that developing countries are successful in raising tax revenue, low state capacity causes them to rely extensively on a value-added tax and other more indirect taxation methods. These tend to fall on all consumers—particularly the poor, who spend a higher portion of their income on food and goods—in contrast to an income tax that can target the rich (Higgins and Lustig 2016).

Lacking both foreign aid and comprehensive tax nets, high-poverty middle-income countries constitute the "missing middle" of the global distribution of fiscal capacity, lagging between better-funded low-income and high-income countries (Kharas, Prizzon, and Rogerson 2014). This financing gap manifests in low government spending, especially on social services (Bastagli, Coady, and Gupta 2012).

Even when states can mobilize domestic funds for the provision of invisible infrastructure, corruption and leakage may hamstring their ability to reach the poor (Svensson 2005). The poor are particularly likely to face corruption in accessing services, either because wealthier households' connections and knowledge of the law equip them to resist corrupt officials or because the wealthy can opt out of free public services (Justesen and Bjornskov 2014; Peiffer and Rose 2016). As a result, a significant share of the resources that high-poverty middle-income countries earmark for social protection may not reach the right beneficiaries. Hanna and Olken discuss targeting in their companion paper in this issue.

### Problems of Will

Even where states have the capacity to deliver invisible infrastructure to the poor, they may lack the will to do so. This lack of will can find expression in spending resources on projects that seem frivolous given the high rate of poverty, and that can

[12] Authors' calculations using most recent taxation data since 2010 available from ICTD/UNU-WIDER (2017); 134 of 172 countries in our sample have data for 2015.

*Figure 3*
**Tax Revenue as a Share of GDP by Country Income**

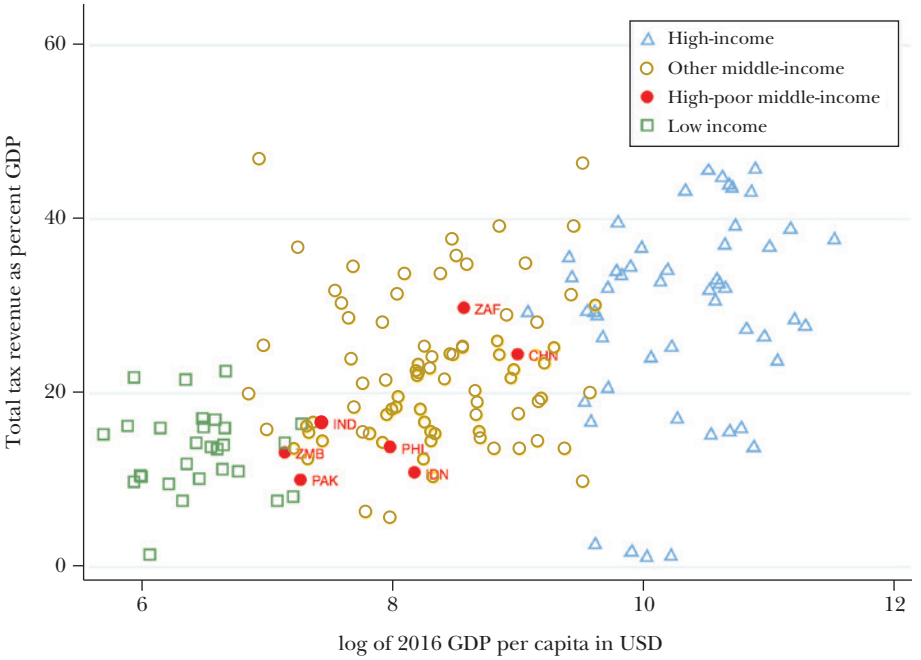


*Source:* Authors use data on tax revenue from ICTD/UNU–WIDER (2017) and data on GDP from World Development Indicators (2017).
*Note:* We use data for taxes including social contributions from the "merged" government revenue dataset available through ICTD. We exclude observations where ICTD flags that "accuracy, quality or comparability of data is questionable." We use the most recent taxation data available since 2010; data is from 2015 for 131 of 167 countries in the figure sample. No data is available for Nigeria, one of the eight high-poverty middle-income countries. To be consistent with our classification of high-poverty middle-income countries in Table 1, we classify country income status according to the World Bank's FY15 list of economies, which is based on data from calendar year 2013. We label the points for high-poverty middle-income countries with country codes from the World Bank as follows: IND is India, CHN is China, IDN is Indonesia, PAK is Pakistan, PHL is the Philippines, ZAF is South Africa, and ZMB is Zambia.

only be justified through complex trickle-down reasoning: the Indian state's space program is one example; Rwanda's £30 million sponsorship of Arsenal football club in 2018 is another. But it can also find expression through projects that ostensibly serve all, but that exacerbate poverty in pockets of the population. For example, a large dam construction policy in India increased regional inequality and aggregate poverty: while districts located downstream of the dam saw agricultural productivity rise and poverty fall, the districts where dams were built saw poverty rise. These rises were particularly pronounced in districts that had a history of extractive colonial institutions and, therefore, adversarial relationships between the elite and disadvantaged populations (Duflo and Pande 2007).

Lack of will can also find expression in states' choices to address extreme poverty in ways that do not fully account for poor people's values, preferences, and quality of life. For example, President Xi Jinping of China plans to relocate 9.8 million of the rural poor between 2016 and 2020 as part of a push to end extreme poverty in China by 2020 (as reported in Phillips 2018). But forced migration can disrupt valuable social networks and decrease quality of life in ways that are not captured by income measures (Barnhardt, Field, and Pande 2017).

Given the myriad vested interests in any society, governments—especially in the presence of resource constraints—will often only respond to clear demands from citizens. The poor are more likely to be sidelined both in economic development and in democratic processes. Recent evidence suggests that nondemocratic countries on average exhibit lower growth than democracies (Acemoglu, Naidu, Restrepo, and Robinson forthcoming). Moreover, the mechanisms that link growth with democracy include elements of invisible infrastructure. Democratic institutions also tend to be friendly to labor: they result in higher wages (Rodrik 1999).

In this light, problems of will on the part of government become problems of agency on the part of the poor: provision of good invisible infrastructure requires both that domestic states have the capacity to deliver it, and that poor citizens have the voice to demand it.

## Building Invisible Infrastructure that Delivers for the Poor

When a democracy functions as intended, there are two core positive consequences for invisible infrastructure. First, there is a systematic way for citizens to voice their needs (via voting on manifestos, or engaging in protests). Secondly, state bodies are accountable and incentives for delivering services are strong—ineffective or poor governance can be punished by removal from power at elections (or via impeachment) (Acemoglu and Robinson 2011).

Democracy is also increasingly the form of government the poor live under: We estimate that the proportion of the world's poor living in democracies rose from 25.8 to 47.1 percent between 1987 and 2013. Assuming that the distribution of the world's poor by country remained constant between 2013 and 2015, this figure would have risen to 60.7 percent by 2015—at which point, 11 of 33 low-income countries and all high-poverty middle-income countries but one (China) were classified as democracies by Polity IV measures.[13]

How do we further engage and empower citizens in developing countries to demand well-functioning invisible infrastructure?

[13]Authors' calculations using data on poverty from PovcalNet (2018) and Polity IV democracy data available through the Center for Systemic Peace (2016). We classify states as democracies if they have a Polity IV score of at least six. We continue to classify country income status according to the World Bank's FY2015 list of economies (data from calendar year 2013).

Free and fair elections in democracies are a critical first step. The international aid community has recognized their value: Between 1990 and 2013, annual official development assistance commitments for democracy and governance increased nearly twentyfold, from just above $1 billion to about $20 billion in 2011 US dollars.[14] While the literature on whether aid promotes or impairs democracy is mixed (Djankov, Montalvo, and Reynal-Querol 2008; Kersting and Kilby 2014; Knack 2004), a small literature finds that aid earmarked for democracy assistance does promote democratic institutions. The analysis of Finkel, Pérez-Liñan, and Seligson (2007) makes use of the 500 percent increase in US foreign assistance for democracy building between 1990 and 2003. Using program information for 165 countries, they find that democracy assistance helped build democratic institutions. Dietrich and Wright (2015) provide complementary data using all OECD democracy aid flows to 44 countries of sub-Saharan Africa in the 1990s and 2000s. They find that democracy aid stabilized multiparty regimes and decreased the incidence of electoral misconduct, which they interpret as increasing horizontal accountability.

Giving poor citizens the democratic tools to demand invisible infrastructure means not just giving them votes, but also establishing systems of broader accountability—where citizens have the tools and information to make demands of the various players involved in the provision of invisible infrastructure.

How can we design governance reforms to aid this? A first step is to model the behavior of actors engaged in the provision of invisible infrastructure in a way that can shed light on the root problems.

Principal-agent frameworks provide a natural way to model the provision of invisible infrastructure as involving a *human chain* of interlinked actors—upper tiers of management, or principals, delegate tasks to lower tiers of agents (Dixit 2002). In turn, agents at higher tiers often act as principals at lower tiers. When we so decompose the state from monolith to interlocking principal-agent relationships, we can see failures like widespread corruption as localized malfunctions of particular links in the human chain. The principal-agent framework allows us to model these malfunctions as agency problems: the principal and agent may have different preferences, and weak information limits the principal's ability to fully observe the agent's action. We can then approach policies to build invisible infrastructure for the poor as opportunities to solve a series of *mechanism design problems,* where efficient design requires understanding the political environment.

In the case of a democratic state providing services, the chain is often circular: citizens delegate policymaking to elected leaders, who delegate tasks to senior administrators, who in turn delegate tasks to lower-level bureaucrats, who, finally, direct the activities of frontline service providers. Ultimately, these providers often seek to influence the actions of citizens. It is useful to visualize a human chain forming a circle with two sides: a democratic side extending from the citizen up

[14] Authors' calculations using AidData (2017). Following Dietrich and Wright (2015), we classify governance and democracy aid as official development assistance to which AidData has assigned a Credit Reporting System purpose code beginning with 15.

to the politician, and an implementation side stretching from the politician back down to the citizen.

The failure to align the incentives of actors along the chain with the preferences of citizens is often at the root of weak delivery of invisible infrastructure. The human chain may be difficult to consider in the abstract, so we now give a series of concrete examples from a massive social protection program in a high-poverty middle-income country—the type of program states will need to execute successfully to draw their citizens out of extreme poverty.

**Seeing the State as a Human Chain: The Example of Workfare Program Reforms in India**

In 2005, India launched a federal workfare program—the Mahatma Gandhi National Rural Employment Guarantee Scheme (MGNREGS)—that seeks to guarantee employment to the rural poor when they need it. This large program (with a 2017–18 budget of over $7 billion) has historically been beset by significant corruption, though multiple program reforms appear to have reduced the leakage of funds over time (Imbert and Papp 2018).[15] In addition, the quality of MGNREGS implementation exhibits significant geographic variation. Research on the effects of this program and the trends in its implementation provides multiple lessons on how aligning incentives and improving information flows can help build states capable of delivering the invisible infrastructure.

First, it is crucial to ensure that information on reform reaches those with the will and ability to implement reforms. Working in Bihar, one of India's poorest states, Banerjee, Duflo, Imbert, Mathew, and Pande (2016) found that a new digital accounting system that cuts out administrative tiers lowered corruption and reduced MGNREGS program spending by 24 percent, with no detectable decline in payments to beneficiaries. This was a positive result, but scaling up the reform required dealing with the entire human chain of the state, not just tuning up a single malfunctioning link. The Bihar accounting reform first hit resistance from mid-tier administrators—those links in the chain who were being cut out—who lobbied against the reformed system. The state government repealed it. Eventually, the federal government—the program's funder (and so the primary principal within the administration)—overrode that decision and rolled out the reform nationwide because it saved them money. Thus, reform is more likely to be successful when actors higher up in the human chain with superior policy authority have both incentives to implement it well and access to independent information on program performance.

Second, it is necessary to align policy choices with the preferences and needs of the poor. Again, Bihar's reform experience is revealing. A continuing shortcoming

[15] Imbert and Papp (2018) compare MGNREGS employment in official reports to the estimated number of days spent by rural adults on any public works estimated based on National Sample Survey data. They find that in 2007–2008 only 51 percent of reported MGNREGS employment was independently confirmed by the survey data. They also find that this gap narrows over time to 71 percent in 2009–2010 and 80 percent in 2011–2012.

of the digital accounting reform was that citizens did not directly benefit from reduced corruption: the amount they earned from the workfare program remained unaffected, while implementation issues *increased* delays in their wage payments in the short term. Furthermore, since citizens had no way of knowing that the modified system had cut leakage, they could not lobby for the money saved to be spent on more job opportunities through the program.

In contrast, a different reform of the employment guarantee scheme did translate into higher wages for the poor, arguably because citizens were directly engaged in and, therefore, well informed of the reform. Muralidharan, Niehaus, and Sukhtankar (2016) demonstrate that investments in secure payment infrastructure for MGNREGS that directly included citizens—by altering how they obtain payments from banks—delivered a faster and less corrupt payment service while raising effective wages received by beneficiaries.

Third, leveraging the circular nature of the human chain in democratic settings can provide a powerful way of aligning incentives across the chain. Specifically, a well-designed human chain can ensure that elected politicians are incentivized to monitor administrators and verify that they provide services effectively. Gulzar and Pasquale (2017) compare MGNREGS performance in districts where bureaucrats are supervised by a single political principal with those supervised by multiple politicians and find that program performance is substantially better where bureaucrats answer to a single politician. They conclude that politicians face strong electoral incentives to motivate bureaucrats as long as they internalize the benefits from doing so.[16]

**Beyond the State: Enabling Citizens as the Ultimate Principal**

These examples focus on the circle of principal-agent interactions between the state and citizens. While this essay cannot do justice to the myriad ways in which state and nonstate actors engage with each other and with citizens to affect the provision of invisible infrastructure, we conclude by highlighting a few ways in which these interactions support an important tool of empowerment for citizens—the information they have on the provision of invisible infrastructure.

The media can play an important role in ensuring political accountability. To ensure impartiality, some of the institutions that provide information—like the media—need to exist outside the state (Besley and Prat 2006). Conversely, autocracies often use media censorship to reduce information available to citizens and, arguably, to lessen their will to engage with policymaking (Chen and Yang 2018). The role of a free media in highlighting situations of distress for the poor was famously argued by Sen (1999), who showed that famines in India disappeared with the establishment of democracy

[16] A famous historical example of how state capacity and accountability can go hand in hand comes from the simultaneous creation of the American modern welfare state and dramatic decline in corruption in the public service delivery in the 1930s. Wallis, Fishback, and Kantor (2006) provide evidence that Roosevelt's political interests were better served by ensuring that the American poor received services, rather than by using relief to aid small sets of politically connected contractors or to give jobs to loyalists.

and free media. Besley and Burgess (2002) show that calamity relief and public food distribution systems work better in Indian states with greater newspaper circulation.

Other institutions that empower citizens by providing information can be mandated by the state, but need to have operational autonomy and significant resources to be effective. A classic example is implementation of Freedom of Information Acts. Over the last half-century, these acts have spread from northern Europe to over 100 countries, rich and poor (according to data collected by Access Info Europe and the Centre for Law and Democracy). These acts are seen as embodying citizens' right to have access to information about the functioning of their government, and may also help citizens to access invisible infrastructure. In India, Banerjee, Enevoldsen, Pande, and Walton (2018) partnered with nongovernmental organizations to publish politician report cards in local newspapers, with performance data collected via right-to-information laws. These report cards moved politicians's spending allocations to more closely match citizen preferences for public goods delivery. In this case, as in many transparency initiatives, a nongovernment organization served as a vital intermediary between the poor and the state. Reinikka and Svensson (2011) report comparable evidence for Uganda.

Transparency initiatives—often organized by civil society organizations and nongovernment organizations—can empower citizens even in nondemocratic settings. After China passed regulations in 2008 giving the public access to certain types of environmental information, investigative journalist Ma Jun created an online public database that made information on water and air pollution violations easily available to citizens. By 2012, this portal had exposed over 90,000 private sector air and water pollution violations and was contributing to a swelling citizen-based environmental movement (Goldman Environmental Foundation 2012). Between 2000 and 2013, pollution was the largest driver of large public protests in China (Steinhardt and Wu 2016). While this example shows how transparency can enhance invisible infrastructure for all citizens, it is worth pointing out that a wide literature has shown that the costs of environmental damage fall most heavily on the poor through channels such as exposure to air pollution (Hajat, Hsia, and O'Neill 2015) and vulnerability to climate change (UN DESA 2016).

## Conclusion

Is it a realistic aim to end extreme poverty by 2030? We believe that achieving this goal within this timeframe will require substantial recalibration of efforts. While economic growth has fueled large reductions in poverty over recent decades, further reductions will also require providing the "invisible infrastructure" that the poor need for economic mobility. This provision can rely in part on aid and private players, but it will need to work principally through the domestic state. Thus, eliminating absolute poverty will require investing in—not circumnavigating—domestic states, both in low-income countries and high-poverty middle-income countries. It will also require empowering citizens to act as principals in demanding services from the state.

How can international actors help? On the grand scale, international democracy assistance groups who focus on fair and free elections can coordinate activities with groups that seek to strengthen state capacity via greater administrative efficiency and transparency. In low-income countries, aid should contribute to building effective, accountable state-run service delivery as much as possible. In high-poverty middle-income countries, well-targeted technical assistance intended to support transparency and accountability initiatives can yield high returns. Building sound invisible infrastructure will require working with individuals within the state who have the power and incentives to implement reform, and at the same time ensuring that poor citizens remain the ultimate principal.

In this spirit, we argue for research that unpacks the state and recognizes it as a chain of individuals, all acting on their own interest and responding to incentives. In recent decades, development economics has seen the emergence of an experimental literature that evaluates microeconomic policies one by one. This has led to a robust discussion on the relationship between single-program evaluations and system-level change—whether, for example, a program that improved learning in a small number of schools can guide reforms on a country's educational system (Alcott 2015; Bold, Kimenyi, Mwabu, Ng'ang'a, and Sandefur 2013). Some experts have suggested conducting experiments at the scale of the reform you want to effect (Muralidharan and Niehaus 2017), while others express concerns that the micro view may divert attention and resources from system-level improvement (Deaton and Cartwright 2018). We believe that micro-level evidence can inform system-level reforms, but to do so effectively requires engaging with the political economy of reform. That is, research must also examine whether policymakers have the means to monitor implementation, whether bureaucrats have the motivation to implement the policies, and whether citizens have effective mechanisms to make their voices heard. A focus on political economy allows us to develop hypotheses about how the incentives of different actors in the human chain of the state can be aligned and how information flows will influence their behavior—hypotheses that can be tested by rigorous evaluations, experimental or otherwise.

By 2030, we will likely be living in a much richer world. Whether it will be a world free of poverty will depend on whether we can reach the world's most isolated, disadvantaged, and demoralized people—those who remain untouched by record growth and unprecedented flows of aid. This task goes beyond money and into power: we must understand and restructure social and political institutions so that the powerful have reason to serve the powerless.

■ *We thank Dominic Leggett and Vestal McIntyre for extensive help in shaping this essay. We thank Michael Callen, Avinash Dixit, Edward Glaeser, Gordon Hanson, Adnan Khan, Enrico Moretti, Helena Roy, and Timothy Taylor for comments that substantially improved the draft. Pande thanks Angus Deaton and Jean Drèze for early insightful conversations on this topic. We are thankful to the Weatherhead Center for International Affairs at Harvard University and to the Women and Public Policy Program at Harvard Kennedy School for financial support.*

## References

**Acemoglu, Daron, Suresh Naidu, Pascual Restrepo, and James A Robinson.** Forthcoming. "Democracy Does Cause Growth." *Journal of Political Economy*.

**Acemoglu, Daron, and James Robinson.** 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Publishers.

**Acht, Martin, Toman Omar Mahmoud, and Rainer Thiele.** 2015. "Corrupt Governments Do Not Receive More State-to-State Aid: Governance and the Delivery of Foreign Aid through Non-State Actors." *Journal of Development Economics* 114: 20–33.

**AidData.** 2017. "AidData Core Research Release, Version 3.1." College of William and Mary. Available at: https://www.aiddata.org/data/aiddata-core-research-release-level-1-3-1 (accessed August 10, 2018). (Updated version of the data from Tierney, Michael J., Daniel L. Nielson, Darren G. Hawkins, J. Timmons Roberts, Michael G. Findley, Ryan M. Powers, Bradley Parks, Sven E. Wilson, and Robert L. Hicks. 2011. "More Dollars than Sense: Refining Our Knowledge of Development Finance Using AidData." *World Development* 39(11): 1891–1906.)

**Allcott, Hunt.** 2015. "Site Selection Bias in Program Evaluation." *Quarterly Journal of Economics* 130(3): 1117–65.

**Alvaredo, Facundo, Lucas Chancel, Thomas Piketty, Emmanuel Saez, and Gabriel Zucman.** 2018. "World Inequality Report 2018." World Inequality Lab.

**Asher, Sam, and Paul Novosad.** 2018. "Rural Roads and Local Economic Development." Policy Research Working Paper 8466, World Bank.

**Ashta, Arvind, and Marek Hudon.** 2012. "The Compartamos Microfinance IPO: Mission Conflicts in Hybrid Institutions with Diverse Shareholding." *Strategic Change* 21(7–8): 331–41.

**Assouad, Lydia, Lucas Chancel, and Marc Morgan.** 2018. "Extreme Inequality: Evidence from Brazil, India, the Middle East, and South Africa." *AEA Papers and Proceedings* 108: 119–23.

**Atun, Rifat A., Sara E. Bennett, and Antonio Durán.** 2008. "When Do Vertical (Stand-Alone) Programmes Have a Place in Health Systems?" WHO European Ministerial Conference on Health Systems, World Health Organization.

**Bandiera, Oriana, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman.** 2017. "Labor Markets and Poverty in Village Economies." *Quarterly Journal of Economics* 132(2): 811–70.

**Banerjee, Abhijit, Esther Duflo, Raghabendra Chattopadhyay, and Jeremy Shapiro.** 2016. "The Long-Term Impacts of a 'Graduation' Program: Evidence from West Bengal." https://economics.mit.edu/files/12031.

**Banerjee, Abhijit, Esther Duflo, Clement Imbert, Santhosh Mathew, and Rohini Pande**. 2016. "E-governance, Accountability, and Leakage in Public Programs: Experimental Evidence from a Financial Management Reform in India." NBER Working Paper 22803.

**Banerjee, Abhijit, Nils Enevoldsen, Rohini Pande, and Michael Walton.** 2018. "Information as an Incentive: Experimental Evidence from Delhi." August 25. https://scholar.harvard.edu/files/rpande/files/delhivoter_shared-14.pdf.

**Banerjee, Abhijit, Rema Hanna, Jordan C. Kyle, Benjamin A. Olken, and Sudarno Sumarto.** 2015. "The Power of Transparency: Information, Identification Cards and Food Subsidy Programs in Indonesia." NBER Working Paper 20923.

**Barnhardt, Sharon, Erica Field, and Rohini Pande.** 2017. "Moving to Opportunity or Isolation? Network Effects of a Randomized Housing Lottery in Urban India." *American Economic Journal: Applied Economics* 9(1): 1–32.

**Bastagli, Francesca, David Coady, and Sanjeev Gupta.** 2012. "Income Inequality and Fiscal Policy." IMF Staff Discussion Note SDN/12/08 (Revised), International Monetary Fund.

**Becker, Gary S., Tomas J. Philipson, and Rodrigo R. Soares.** 2005. "The Quantity and Quality of Life and the Evolution of World Inequality." *American Economic Review* 95(1): 277–91.

**Besley, Timothy, and Robin Burgess.** 2002. "The Political Economy of Government Responsiveness: Theory and Evidence from India." *Quarterly Journal of Economics* 117(4): 1415–51.

**Besley, Timothy, and Stephen Coate.** 1995. "Group Lending, Repayment Incentives and Social Collateral." *Journal of Development Economics* 46(1): 1–18.

**Besley, Timothy, and Andrea Prat.** 2006. "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability." *American Economic Review* 96(3): 720–36.

**Birchler, Kassandra, and Katharina Michaelowa.** 2016. "Making Aid Work for Education in Developing Countries: An Analysis of Aid Effectiveness for Primary Education Coverage and Quality." *International Journal of Educational Development* 48: 37–52.

**Blattman, Christopher, Nathan Fiala, and Sebastian Martinez.** 2014. "Generating Skilled Self-Employment in Developing Countries:

Experimental Evidence from Uganda." *Quarterly Journal of Economics* 129(2): 697–752.

**Blattman, Christopher, Nathan Fiala, and Sebastian Martinez.** 2018. "The Long Term Impacts of Grants on Poverty: 9-Year Evidence from Uganda's Youth Opportunities Program." NBER Working Paper 24999.

**Bleakley, Hoyt.** 2007. "Disease and Development: Evidence from Hookworm Eradication in the American South." *Quarterly Journal of Economics* 122(1): 73–117.

**Bloom, David E.** 2011. "7 Billion and Counting." *Science* 333(6042): 562–69.

**Bold, Tessa, Mwangi Kimenyi, Germano Mwabu, Alice Ng'ang'a, and Justin Sandefur.** 2013. "Scaling Up What Works: Experimental Evidence on External Validity in Kenyan Education." Working Paper 321, Center for Global Development.

**BRAC.** 2013. "An End in Sight for Ultra-Poverty: Scaling up BRAC's Graduation Model for the Poorest." Briefing Note #1: Ending Extreme Poverty. November 2013. http://www.brac.net/sites/default/files/BRAC%20Briefing%20-%20TUP.pdf (accessed November 6, 2017).

**Breza, Emily, and Cynthia Kinnan.** 2018. "Measuring the Equilibrium Impacts of Credit: Evidence from the Indian Microfinance Crisis." NBER Working Paper 24329.

**Bryan, Gharad, Shyamal Chowdhurym, and Ahmed Mushfiq Mobarak.** 2014. "Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh." *Econometrica* 82(5): 1671–1748.

**Burgess, Robin, and Rohini Pande.** 2005. "Do Rural Banks Matter? Evidence from the Indian Social Banking Experiment." *American Economic Review* 95(3): 780–95.

**Castañeda, Andrés, Dung Doan, David Newhouse, Minh Cong Nguyen, Hiroki Uematsu, and João Pedro Azevedo.** 2018. "A New Profile of the Global Poor." World Development 101: 250–67.

**Center for Systemic Peace.** 2016. "Polity IV Project, Political Regime Characteristics and Transitions, 1800–2016." At the website: http://www.systemicpeace.org/inscrdata.html (accessed January 5, 2017).

**CGAP.** 2010. "Andhra Pradesh 2010 : Global Implications of the Crisis in Indian Microfinance." Focus Note 67. Washington, DC: CGAP.

**Chandy, Laurence, Natasha Ledlie, and Veronika Penciakova.** 2013. "The Final Countdown: Prospects for Ending Extreme Poverty by 2030." Policy Paper 2013-04, Global Views, Brookings Institution.

**Chandy, Laurence, Lorenz Noe, and Christine Zhang.** 2016. "The Global Poverty Gap is Falling. Billionaires Could Help Close It." *Brookings Up Front*, January 20. https://www.brookings.edu/blog/up-front/2016/01/20/the-global-poverty-gap-is-falling-billionaires-could-help-close-it/.

**Chen, Yuyu, and David Y. Yang.** 2018. "The Impact of Media Censorship: Evidence from a Field Experiment in China." Job Market Paper, Stanford University.

**Cole, Shawn.** 2009. "Fixing Market Failures or Fixing Elections? Agricultural Credit in India." *American Economic Journal: Applied Economics* 1(1): 219–50.

**Commission on State Fragility, Growth, and Development.** 2018. *Escaping the Fragility Trap*. London, United Kingdom: LSE–Oxford Commission on State Fragility, Growth and Development.

**Crabtree, James.** 2018. *The Billionaire Raj: A Journey through India's New Gilded Age*. New York: Crown Publishers.

**Crosby, Alfred W.** 1993. "Smallpox." In *The Cambridge World History of Human Disease*, edited by Kenneth F. Kiple, 1008–1014. Cambridge University Press.

**Currie, Janet, and Tom Vogl.** 2013. "Early-Life Health and Adult Circumstance in Developing Countries." *Annual Review of Economics* 5: 1–36.

**Deaton, Angus.** 2013. *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton University Press.

**Deaton, Angus, and Nancy Cartwright.** 2018. "Understanding and Misunderstanding Randomized Controlled Trials." *Social Science and Medicine* 210: 2–21.

**Dell, Melissa, Nathaniel Lane, and Pablo Querubin.** 2017. "The Historical State, Local Collective Action, and Economic Development in Vietnam." NBER Working Paper 23208.

**Dietrich, Simone.** 2013. "Bypass or Engage? Explaining Donor Delivery Tactics in Foreign Aid Allocation." *International Studies Quarterly* 57(4): 698–712.

**Dietrich, Simone, and Joseph Wright.** 2015. "Foreign Aid Allocation Tactics and Democratic Change in Africa." *Journal of Politics* 77(1): 216–34.

**Dixit, Avinash.** 2002. "Incentives and Organizations in the Public Sector: An Interpretative Review." *Journal of Human Resources* 37(4): 696–727.

**Djankov, Simeon, Jose G. Montalvo, and Marta Reynal-Querol.** 2008. "The Curse of Aid." *Journal of Economic Growth* 13(3): 169–94.

**Dollar, David, Tatjana Kleineberg, and Aart Kraay.** 2016. "Growth is Still Good for the Poor." *European Economic Review* 81: 68–85.

**Dreher, Axel, Peter Nunnenkamp, and Rainer Thiele.** 2008. "Does Aid for Education Educate Children? Evidence from Panel Data." *World Bank Economic Review* 22(2): 291–314.

**Duflo, Esther, and Rohini Pande.** 2007. "Dams." *Quarterly Journal of Economics* 122(2): 601–46.

**Ferreira, Francisco, Shaohua Chen, Andrew Dabalen, Yuri Dikhanov, Nada Hamadeh, Dean Jolliffe, Ambar Narayan, Espen Beer Prydz, Ana Revenga, Prem Sangraula, Umar Serajuddin, and Nobuo Yoshida.** 2016. "A Global Count of the Extreme Poor in 2012: Data Issues, Methodology and Initial Results." *Journal of Economic Inequality* 14(2): 141–72.

**Field, Erica, Rohini Pande, Natalia Rigol, Simone Schaner, and Charity Troyer Moore.** 2016. "On Her Account: Can Strengthening Women's Financial Control Boost Female Labor Supply?" http://scholar.harvard.edu/files/rpande/files/on_her_account.can_strengthening_womens_financial_control_boost_female_labor_supply.pdf.

**Finan, Frederico, Benjamin A. Olken, and Rohini Pande.** 2017. "The Personnel Economics of the Developing State." Chap. 6 in *Handbook of Field Experiments*, vol. 2, edited by Abhijit Banerjee and Esther Duflo. Elsevier.

**Finkel, Steven E., Aníbal Pérez-Liñan, and Mitchell A. Seligson.** 2007. "The Effects of U.S. Foreign Assistance on Democracy Building, 1990–2003." *World Politics* 59(3): 404–40.

**Gallup.** 2008–2016. "World Poll." https://www.gallup.com/analytics/232838/world-poll.aspx (accessed December 23, 2016).

**Goldman Environmental Foundation.** 2012. "Ma Jun: 2012 Goldman Prize Recipient Asia." https://www.goldmanprize.org/recipient/ma-jun/ (accessed September 9, 2018).

**Gulzar, Saad, and Benjamin J. Pasquale.** 2017. "Politicians, Bureaucrats, and Development: Evidence from India." *American Political Science Review* 111(1): 162–83.

**Hajat, Anjum, Charlene Hsia, and Marie S. O'Neill.** 2015. "Socioeconomic Disparities and Air Pollution Exposure: A Global Review." *Current Environmental Health Reports* 2(4): 440–50.

**Hammar, Olle, and Daniel Waldenström.** 2017. "Global Earnings Inequality, 1970–2015." CEPR Discussion Paper DP12019.

**Haushofer, Johannes, and Jeremy Shapiro.** 2016. "The Short-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya." *Quarterly Journal of Economics* 131(4): 1973–2042.

**Haushofer, Johannes, and Jeremy Shapiro.** 2018. "The Long-Term Impact of Unconditional Cash Transfers: Experimental Evidence from Kenya." Unpublished paper, Busara Center for Behavioral Economics, Nairobi, Kenya.

**Higgins, Sean, and Nora Lustig.** 2016. "Can a Poverty-Reducing and Progressive Tax and Transfer System Hurt the Poor?" *Journal of Development Economics* 122: 63–75.

**ICTD/UNU-WIDER.** 2017. "Government Revenue Dataset 2017: Merged." Available at https://www.wider.unu.edu/project/government-revenue-dataset (accessed January 18, 2018).

**Imbert, Clément, and John Papp.** 2018. "Short-Term Migration, Rural Public Works and Urban Labor Markets: Evidence from India." https://www.dropbox.com/s/d1w5qixdd-barf7o/2018-07-06%20Paper.pdf?dl=0.

**Jones, Benjamin F., and Benjamin A. Olken.** 2008. "The Anatomy of Start-Stop Growth." *Review of Economics and Statistics* 90(3): 582–87.

**Justesen, Mogens K., and Christian Bjørnskov.** 2014. "Exploiting the Poor: Bureaucratic Corruption and Poverty in Africa." *World Development* 58: 106–15.

**Kavadi, Shirish N.** 2007. "'Parasites Lost and Parasites Regained': Rockefeller Foundation's Anti-Hookworm Campaign in Madras Presidency." *Economic and Political Weekly* 42(2): 130–37.

**Kersting, Erasmus, and Christopher Kilby.** 2014. "Aid and Democracy Redux." *European Economic Review* 67: 125–43.

**Khandker, Shahidur R.** 1998. "Using Microcredit to Advance Women." PREM Note no. 8, World Bank.

**Kharas, Homi, Annalisa Prizzon, and Andrew Rogerson.** 2014. "Financing the Post-2015 Sustainable Development Goals: A Rough Roadmap." London, UK: Overseas Development Institute.

**Knack, Stephen.** 2004. "Does Foreign Aid Promote Democracy?" *International Studies Quarterly* 48(1): 251–66.

**Kraay, Aart.** 2006. "When Is Growth Pro-Poor? Evidence from a Panel of Countries." *Journal of Development Economics* 80(1): 198–227.

**Mishra, Prachi, and David Newhouse.** 2009. "Does Health Aid Matter?" *Journal of Health Economics* 28(4): 855–72.

**Mitra, Anirban, and Debraj Ray.** 2014. "Implications of an Economic Theory of Conflict: Hindu-Muslim Violence in India." *Journal of Political Economy* 122(4): 719–65.

**Montenegro, Claudio E., and Harry A. Patrinos.** 2014. "Comparable Estimates of Returns to Schooling around the World." Policy Research Working Paper 7020, World Bank Group.

**Muralidharan, Karthik, and Paul Niehaus.** 2017. "Experimentation at Scale." *Journal of Economic Perspectives* 31(4): 103–24.

**Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar.** 2016. "Building State Capacity: Evidence from Biometric Smartcards in India." *American Economic Review* 106(10): 2895–2929.

**Muralidharan, Karthik, and Venkatesh**

**Sundararaman.** 2015. "The Aggregate Effect of School Choice: Evidence from a Two-Stage Experiment in India." *Quarterly Journal of Economics* 130(3): 1011–66.

**National Sample Survey Office (NSS).** 2016. "National Sample Survey (NSS) data (unit level)." Ministry of Statistics and Program Implementation. https://doi.org/10.7910/DVN/K8BSDU (accessed via Harvard Dataverse, March 19, 2018).

**North, Douglass, C.** 1990. *Institutions, Institutional Change and Economic Performance.* Cambridge University Press.

**OECD.** 2018. "Aid (ODA) Disbursements to Countries and Regions [DAC2a]." Data. http://stats.oecd.org/Index.aspx?QueryId=42231&lang=en# (accessed August 13, 2018).

**Oliveira-Cruz, Valeria, Christoph Kurowski, and Anne Mills.** 2003. "Delivery of Priority Health Services: Searching for Synergies within the Vertical versus Horizontal Debate." *Journal of International Development* 15(1): 67–86.

**Palmer, Steven.** 2003. *From Popular Medicine to Medical Populism: Doctors, Healers, and Public Power in Costa Rica, 1800–1940.* Duke University Press.

**Peiffer, Caryn, and Richard Rose.** 2016. "Why Are the Poor More Vulnerable to Bribery in Africa? The Institutional Effects of Services." *Journal of Development Studies* 54(1): 18–29.

**Phillips, Tom.** 2018. "China to Move Millions of People from Homes in Anti-Poverty Drive." *The Guardian*, January 7.

**PovcalNet.** 2018. "Regional Aggregation Using 2011 PPP and $1.9/Day Poverty Line." World Bank. http://iresearch.worldbank.org/PovcalNet/povDuplicateWB.aspx (accessed June 22, 2018).

**Reinikka, Ritva, and Jakob Svensson.** 2011. "The Power of Information in Public Services: Evidence from Education in Uganda." *Journal of Public Economics* 95(7–8): 956–66.

**Rodrik, Dani.** 1999. "Democracies Pay Higher Wages." *Quarterly Journal of Economics* 114(3): 707–738.

**Romero, Mauricio, Justin Sandefur, and Wayne Sandholtz.** 2017. "Can Outsourcing Improve Liberia's Schools? Preliminary Results from Year One of a Three-Year Randomized Evaluation of Partnership Schools for Liberia." Center for Global Development Working Paper 462.

**Sen, Amartya.** 1999. *Development as Freedom.* New York: Anchor Books, Random House.

**Srinivasan, Narasimhan.** 2010. *Microfinance India: State of the Sector Report 2010.* SAGE Publications.

**Steinhardt, H. Christoph, and Fengshi Wu.** 2016. "In the Name of the Public: Environmental Protest and the Changing Landscape of Popular Contention in China." *China Journal* 75: 61–82.

**Streeten, Paul P.** 1979. "Basic Needs: Premises and Promises." *Journal of Policy Modeling* 1(1): 136–46.

**Svensson, Jakob.** 2005. "Eight Questions about Corruption." *Journal of Economic Perspectives* 19(3): 19–42.

**United Nations Department of Economic and Social Affairs (UN DESA).** 2016. *World Economic and Social Survey 2016: Climate Change Resilience: An Opportunity for Reducing Inequalities.* New York: United Nations.

**Wallis, John Joseph, Price V. Fishback, and Shawn Kantor.** 2006. "Politics, Relief, and Reform. Roosevelt's Efforts to Control Corruption and Political Manipulation during the New Deal." Chap. 11 in: *Corruption and Reform: Lessons from America's Economic History*, edited by Edward L. Glaeser and Claudia Goldin. NBER Conference Report. University of Chicago Press.

**Weigel, Jonathan L.** 2018. "Building State and Citizen: How Tax Collection in Congo Engenders Citizen Engagement with the State." https://static1.squarespace.com/static/59832fbcf9a61eb15deffb3c/t/5b6475e8758d469866a3e4336/1533310454595/weigel_jmp.pdf.

**Williamson, Claudia R.** 2008. "Foreign Aid and Human Development: The Impact of Foreign Aid to the Health Sector." *Southern Economic Journal* 75(1): 188–207.

**Wilson, Sven E.** 2011. "Chasing Success: Health Sector Aid and Mortality." *World Development* 39(11): 2032–43.

**World Bank Operations Evaluation Department.** 2004. "Books, Buildings, and Learning Outcomes: An Impact Evaluation of World Bank Support to Basic Education in Ghana." Operations Evaluation Department Report 28779, World Bank.

**World Development Indicators (WDI).** 2017. "GDP per capita (current US$)," with indicator ID: NY.GDP.PCAP.CD. World Bank. https://data.worldbank.org/indicator/ny.gdp.pcap.cd (accessed January 17, 2018).

**World Development Indicators (WDI).** 2018a. "Net Official Development Assistance Received (current US$)" with indicator ID: DT.ODA.ODAT.CD. World Bank. https://data.worldbank.org/indicator/ny.gdp.pcap.cd (accessed April 20, 2018).

**World Development Indicators (WDI).** 2018b. "GDP growth (annual %)," with indicator ID: NY.GDP.MKTP.KD.ZG. World Bank. https://data.worldbank.org/indicator/ny.gdp.mktp.kd.zg (accessed July 14, 2018)

**World Health Organization (WHO).** 2011. *Bugs, Drugs & Smoke: Stories from Public Health.* Geneva, Switzerland: World Health Organization.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 201–226*

# Universal Basic Incomes versus Targeted Transfers: Anti-Poverty Programs in Developing Countries

Rema Hanna and Benjamin A. Olken

Of the 17 Sustainable Development Goals articulated by the United Nations, number one is the elimination of extreme poverty by 2030. In recent years, the largest reductions in extreme poverty worldwide have resulted from substantial economic growth in many emerging economies, particularly China and India (Page and Pande, in this journal issue). While future growth should continue to reduce poverty, it will not solve the problem by itself—both because a substantial amount of poverty remains even in rapidly growing countries, and because sustained economic growth over decades of the type seen in China is the exception, not the norm (Jones and Olken 2008).

Given these realities, there is a potentially important role for national-level transfer programs that assist poor families in developing countries. Such programs are often run by developing country governments. For example, China's rural minimum living standard guarantee (Dibao) program reaches nearly 75 million individuals (Golan, Sicular, and Umapathi 2017), while Mexico's conditional cash transfer program reaches 32 million individuals (World Bank 2014). As countries become wealthier, a greater share of GDP usually goes to social transfer and insurance programs (Chetty and Looney 2006), suggesting that these types of safety net programs are likely to expand in many developing countries.

■ *Rema Hanna is the Jeffrey Cheah Professor of South-East Asia Studies, Harvard Kennedy School, Cambridge, Massachusetts. Benjamin A. Olken is Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Their email addresses are rema_hanna@hks.harvard.edu and bolken@mit.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.201

doi=10.1257/jep.32.4.201

Many countries have implemented transfer programs that seek to target beneficiaries: that is, to identify who is poor and then to restrict transfers to those individuals. In developed countries, because one rarely observes true income-earning ability, targeting is usually based on income (for example, Mirrlees 1971). But in developing countries, governments do not observe income for the vast majority of the population who work in the informal sector, which typically includes most of the poor. Imperfect targeting using various proxy measures for income leads to both *inclusion errors* (giving the transfer to those who are not poor) and *exclusion errors* (failing to deliver the transfer to poor individuals who slip through the cracks in the targeting protocol).

Given these challenges, some have begun to advocate for "universal basic income" programs, which dispense with trying to identify the poor and instead provide transfers to everyone. A universal basic income program is comparatively straightforward to implement—each individual receives a fixed transfer, regardless of income—so the main challenge is to ensure that each person receives the transfer only once. Although universal basic income programs distribute the same value of transfer to everyone, including the very rich, if they are financed through proportional or progressive taxation, they can still result in a substantial redistribution to the poor.

We begin by considering the universal basic income as part of the solution to an optimal income-taxation problem, focusing on the case of developing countries, where there is limited income data and inclusion in the formal tax system is low. We examine how the targeting of transfer programs is conducted in these settings, and provide empirical evidence on the tradeoffs involved between universal basic income and targeted transfer schemes using data from Indonesia and Peru—two countries that run nationwide transfer programs that are targeted to the poor. We conclude by linking our findings back to the broader policy debate on what tools should be preferred for redistribution, as well as the practical challenges of administering them in developing countries.

## Conceptual Connections and Underpinnings

### Universal Basic Income and the Tax Schedule

A universal basic income is usually conceived of as a fixed transfer given to everyone, regardless of income level. Such programs are fairly rare in practice, with prominent examples coming from places with substantial natural resource revenues. For example, Alaska's Permanent Fund transfers a fixed amount—usually between $1,000 and $2,000 per person, per year—to every citizen of the state. Iran implemented a similar program starting in 2011 (Salehi-Isfahani and Mostafavi-Dehzooei 2017). Common arguments that are made for universal basic income programs include their ease of implementation and low administrative costs, because the government does not have to verify income. Another common claim is that such

programs do not distort labor supply, because the payments do not decline if you work more, but as we discuss below, this conclusion is not as straightforward as it may at first appear.

For most developing economies, a substantial universal basic income would need to be financed via domestic taxation, because official development assistance is a small fraction of government budgets. For example, for countries ranked as "upper middle-income" by the World Bank, such as Peru, Lebanon, and the Dominican Republic, overall tax revenue was approximately 157 times the amount of net official development assistance received in 2010. This pattern also holds for "lower middle-income" countries, such as Indonesia, India, and Morocco, where overall tax revenue in 2010 was approximately 14 times the amount of development assistance received. It is only with "low-income" countries, such as Afghanistan, Ethiopia, and Mozambique, that official development assistance exceeded tax revenue—it was approximately 1.2 times the amount of tax revenue in 2010.[1]

Given that a universal basic income would need to be financed through taxation, one can recast universal basic transfers as a particular feature of a tax schedule, as described by Saez (2002) and illustrated by Figure 1. We assume that a universal basic income program is fully financed by a progressive income tax schedule. Pre-tax income $y$ for an individual appears on the horizontal axis, and after-tax income $\alpha(y)$ for that individual is on the vertical axis. The tax paid by a household is therefore given by $y - \alpha(y))$. The universal basic income—the transfer given to everyone, regardless of their income—is thus equal $\alpha(0)$ to on the vertical axis. This perspective suggests that any tax-and-transfer system that has the feature that $\alpha(0) > 0$ can be thought of as including a universal basic income.

Figure 1 plots a progressive tax schedule with and without a fully financed universal basic income. The universal basic income acts as an intercept shifter, raising post-tax income at $y = 0$. However, it will not raise everyone's after-tax income by that same amount, because to finance the universal basic income out of income taxes, marginal tax rates have to increase for *someone*. This illustrates three points about the universal basic income: 1) it can make use of income data collected through the tax system; 2) although a universal basic income acts as an intercept shifter, raising post-tax income at $y = 0$, it will not raise everyone's after-tax income by the same amount; and 3) as emphasized by Saez (2002), it is *not* the case that a universal basic income has no labor supply distortions, since there is still a tax on labor income, and in particular adding the universal basic income to the tax schedule requires adjusting the tax schedule elsewhere.

[1] Data on taxes and official development assistance come from the World Development Indicators. Data on social spending come from the International Labor Organization Social Security Inquiry, http://www.ilo.org/dyn/ilossi/ssimain.home?p_lang=en. For each comparison above, we include only countries with non-missing data for both variables (for example, we sum and compare tax revenue and official development assistance only among countries with non-missing data for both). Country classifications are from the World Bank, where low-income countries have 2015 gross national per-capita income of $1,025 or less; lower middle-income, $1,026–4,035; upper middle-income, $4,036–12,475; and high-income, $12,476 or more.

*Figure 1*
**Example of Progressive Post-Tax Income Schedules With and Without a Universal Basic Income (UBI)**

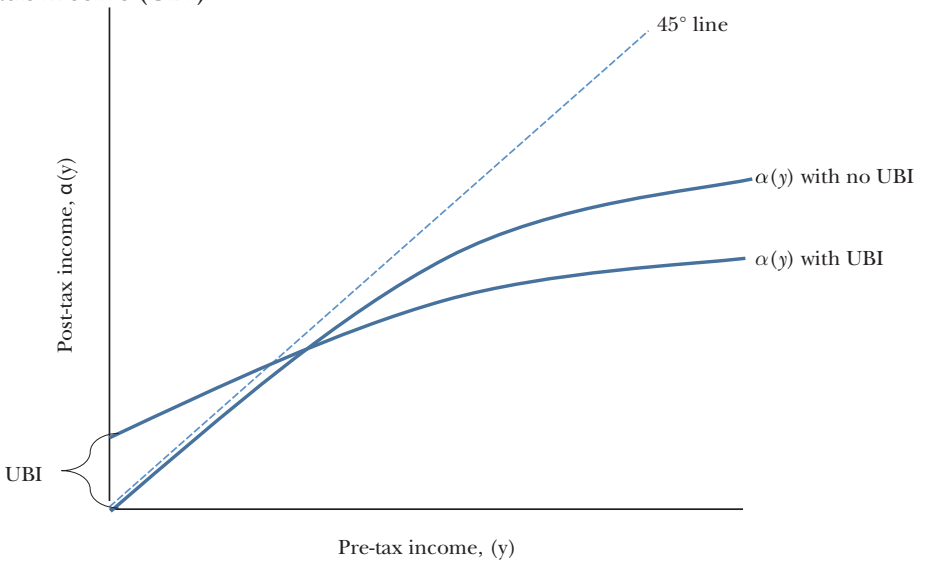


From this perspective, the tradeoffs involved in adding a universal basic income to an income tax schedule, therefore, depend on how the *rest* of the income tax schedule is adjusted to satisfy the government budget constraint—both in terms of overall redistribution and potential distortionary effects from the increased marginal tax rates that introducing a universal basic income will necessitate. Saez (2002) discusses these conditions in the context of a Mirrlees (1971) optimal income tax model in detail. He also discusses the conditions under which an optimal income tax schedule is likely to feature a universal basic income—essentially, when *intensive* labor supply elasticities (the effects of changes in wages on hours of existing workers) are larger than *extensive* labor supply elasticities (the effect of changes in wages on hours of workers entering or leaving the labor force).

Having a program that assures a transfer for those with zero or very low income could also be important for reasons outside the basic Mirrlees (1971)/Saez (2002) model. For one, the program could provide insurance for individuals to take risks: for example, choosing to plant a crop that could be high-return but has a nontrivial probability of zero income might be a more attractive option with a universal basic income in place. The same ideas could be applied to other types of risky decisions, such as human capital investment (Stantcheva 2014).

**What about Poor Countries, Where Tax Systems Are Less Developed?**

In a developed country, where the government observes income for most people and can redistribute through the tax system, this framework is broadly

*Figure 2*

**Example of Post-Tax Income Schedules With and Without a Universal Basic Income (UBI), with a Tax-Exempt Region**

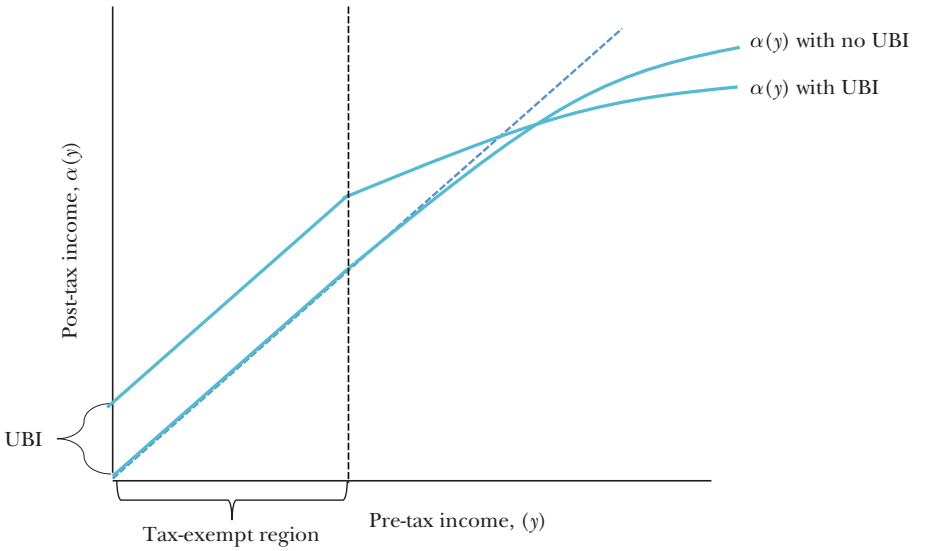


applicable. However, if we are interested in the poor in developing countries, it is more complicated.

In less-developed countries, substantial activity occurs in the "informal sector" that includes casual labor, undocumented firms, and small farms. Thus, most individuals are outside the tax net. Jensen (2016) documents that the share of employment for which people pay *any* income taxes (and therefore for which the government observes income) rises substantially with real per-capita national GDP. Conversely, in most poor countries, the government does not observe any information about income for most people, and in particular, for the poor. In Indonesia and Peru—the two countries that we explore in detail below—Jensen reports that 87.5 percent and 79 percent of the countries' employed populations, respectively, have incomes below the tax exclusion thresholds. Jensen argues that these thresholds emerge from the work environment—when most people work in large firms, the government can easily observe incomes (Kleven, Kreiner, and Saez 2016). But when there is a large informal sector, or many people work on their own or in tiny firms, as is true in most developing countries (Hsieh and Olken 2014), this is not possible.

Why does this matter? Figure 2 illustrates an example tax schedule with and without a universal basic income when most individuals—in developing countries, often upwards of 80 percent of people—fall into a tax-exempt region. In the tax-exempt region, everyone receives the same net transfer regardless of whether they are in the 5th or 70th percentile of income. However, to finance a universal basic income of a size similar to the case in Figure 1, one would need to increase marginal tax rates substantially more for those relatively few people inside the tax net. To the

extent that these individuals are particularly productive, or that, for a corporate income tax, the firms paying it are more productive, such a tax increase may have disproportionally larger efficiency consequences.

In this way, a universal basic income will function somewhat differently in a country with a large number of people outside the tax net. In particular, 1) the *net* transfer could be the same for the majority of the population (all those below the tax-exempt cutoff), which may not be the optimal resource allocation if we want to focus on the poorest of the poor; and 2) relatively few households will be able to contribute to the financing, which may restrict the overall level of funds available for the universal basic income. In countries with large numbers of people outside the tax net, any attempt to target resources toward the poorest of the poor requires using an alternative data collection approach beyond the income tax system.

The discussion thus far has assumed that a universal basic income is financed entirely through income taxes (individual or corporate), which in developing countries are paid primarily by the very top earners. These taxes combined with import duties (which are also likely to fall on the rich), make up about 42 percent of tax revenue for low- and middle-income countries (according to the most recent World Bank *World Development Indicators*), with the rest of tax revenue in these countries coming from consumption taxes, which are typically proportional to consumption. A universal basic income financed through a combination of an income and consumption tax will look closer to a combination of Figures 1 and 2. In this case, one can tax back some of the transfer on households below the income tax exemption through the consumption tax, but the fact that the tax is generally proportional to consumption limits the flexibility in doing so.

**How to Target the Poor in Developing Countries**

How can developing country governments target the poor when income is not observable? When the government observes a proxy for income and targets households based on that, one has what is termed a "proxy-means test." The implementation of a proxy-means test is typically based on large, periodic quasi-censuses of the population, focusing on those most likely to be poor. In these censuses, government enumerators go door-to-door, often visiting millions of households. Some countries, such as India, do target government assistance based on self-reported and unverified income, but this is the exception, not the rule, because people can easily lie if there is no way of verifying it. Instead, government census enumerators typically ask about assets, all of which are easy to observe directly. Examples might include ownership of items such as televisions and refrigerators, the type of material used in one's roof, floor, and walls, the number of rooms in one's house, and so on.

The government uses these assets to predict incomes (or per-capita consumption, which can be easier to measure in survey data). Thus, eligibility for benefits is based on *predicted*, rather than actual, income. This method is quite common, and is used in both large countries such as Indonesia, Pakistan, Nigeria, Mexico, and the Philippines, as well as a number of smaller countries, ranging from Burkina Faso to Ecuador to Jamaica (Fiszbein and Schady 2009).

There are a variety of methods to predict income (or consumption), but they all share basic features: The government takes another dataset that was collected in a low-stakes context (for example, one collected not for targeting, but rather just for research purposes), and therefore for which households have no strong incentive to lie. In this dataset, the government observes the same asset variables as in the proxy-means census and also observes a measure of poverty, such as a household's monthly income or per-capita expenditure. The government then estimates a regression with the measure of poverty as the dependent variable and the assets as explanatory variables. The proxy-means score is the predicted income or expenditure, which the government can calculate for any household using the coefficients from that regression.

The government then can set a threshold for eligibility and distribute benefits to all households with predicted incomes below the threshold.[2] If the government sets the threshold level for eligibility in such a way that it covers the entire population, then a "targeted" program effectively becomes a universal basic income. Because predictions of income, of course, happen with error, targeted systems based on a proxy-means test will feature both inclusion and exclusion errors.

## Comparing Targeted Programs to a Universal Basic Income: Examples from Indonesia and Peru

Depending on the policy context, either a universal basic income or a targeted program may be optimal. Besley (1990) provides an early discussion of tradeoffs between the two approaches. Here, we illustrate the tradeoffs using data from Indonesia and Peru, two countries that have targeted, nationwide transfer programs. The simulations described in this section (particularly Figures 4 and 5) are broadly similar to other analyses of targeting performance, such as Ravallion (2009), Alatas, Banerjee, Hanna, Olken, and Tobias (2012), and particularly Klasen and Lange (2016).

Indonesia has a number of targeted transfer programs, ranging from conditional cash transfers (Program Keluarga Harapan or PKH) to scholarships for poor students (Bantuan Siswa Miskin or BSM) to subsidized health insurance for the poor. In addition, the government has periodically conducted nationwide, unconditional cash transfers (previously called Bantuan Langsung Tunai or BLT, and more recently called Bantuan Langsung Sementara Masyarakat, or BLSM). The

[2]Once the eligibility threshold for a program is set, most developing countries do not vary the level of assistance among the eligible group. Part of the reason is that the imprecision of estimating income based on assets is well understood. Also, it is easier for reasons of politics and transparency just to explain that a given program provides all who are eligible with a certain benefit. However, countries can link benefits levels to predicted income levels by having programs with different eligibility cutoffs. For example, in Indonesia, the conditional cash transfer program called PKH targets roughly the bottom 7–10 percent of households (the "very poor"), but those who are in roughly the bottom 40 percent of the national income distribution (the "near-poor") can qualify for subsidized health insurance even if they cannot access PKH.

government has conducted nationwide targeting censuses approximately every three years since 2005, and then uses proxy-means testing.

Several different types of transfer programs also exist in Peru, ranging from nutritional subsidies (Vaso de Leche) to subsidized health insurance (Seguro Integral de Salud, or SIS) to conditional cash transfers for poor families (Programa Nacional de Apoyo Directo a los Más Pobres, or Juntos). Eligibility for these programs is determined by a proxy-means targeting system, the Sistema de Focalización de Hogares (SISFOH), which assigns households a numerical index to identify eligible households and then allows each individual program to select its final list of beneficiaries. Recertification occurs every three years via a new targeting survey.

For simplicity, we choose one targeted program per country to explore (although the basic results would continue to hold if we chose a basket of programs). In Indonesia, we focus on the Bantuan Langsung Tunai (BLT) program, a temporary and periodic unconditional cash transfer to poor households. The program was implemented first in 2005–2006, and then again in 2008–2009 and in 2014, to help offset shocks in fuel prices. Beneficiaries receive the equivalent of approximately $100 over the course of one year (World Bank 2012). For Peru, we focus on the Juntos program, which is a conditional cash transfer to mothers designed to subsidize child health and education. Beneficiary households receive a monthly transfer of 100 soles (approximately $30). Both programs target roughly one-third of the population.

**Predicting Income with Proxy Measures**

To simulate targeting in the programs, we obtained household-level data from the Indonesian National Socioeconomic Survey (SUSENAS) and the Peruvian National Household Survey (ENAHO), both for the years 2010–2011. Both datasets are used for targeting in the respective countries. Our sample contains 263,705 households in Indonesia and 46,305 households in Peru. Both surveys contain the complete set of asset variables used in targeting, as well a measure of *actual* per-capita consumption for the household, so we can use these datasets to examine targeting accuracy. For each country, we randomly divide the observations into equally sized "training" and "test" sets. In the training set, we regress monthly household per-capita consumption on the actual indicator variables used in each country's respective proxy-means test formula: 82 variables in Indonesia; 72 variables in Peru (regressions provided in the online Appendix). We then predict monthly per-capita consumption for each household in the test sets using the coefficients from the training regressions, and use this "predicted per-capita consumption" for targeting purposes in the simulations below.

The predictors used are not perfect. The typical fit we found of these regressions (the $R^2$) is between 0.53 and 0.66—so while the regressions have a good amount of explanatory power, they also lead to prediction errors: *exclusion errors* of excluding some who should have been eligible (households with true per-capita consumption below the poverty line) and *inclusion errors* of including some who should not have been eligible (households with true per-capita consumption above it).

We apply the formula to the half of our datasets that was not used for estimation (the "test" datasets). We can then use the actual per-capita consumption measure within the survey to assess whether the formulas accurately predict a household's place in the income distribution. Figure 3 shows the results, plotting actual per-capita consumption against predicted per-capita consumption. The figure shows four regions: correctly included (bottom-left region), correctly excluded (top-right region), inclusion error (top-left region), and exclusion error (bottom-right region).

For the exposition here, we are simplifying the actual process by which each country determines benefit eligibility. In modeling these two programs, we intend to target households below 1.5 times the poverty line in Indonesia (about 33 percent of our sample), and households below the poverty line in Peru (about 28 percent of our sample). However, in reality, the monetary poverty line varies by year, urban/rural zoning, and province (in Indonesia) or geographic region (in Peru). In Peru, instead of predicting income directly, the national household targeting system (SISFOH) actually produces a numerical index of household poverty using proxy means-testing variables; it then also incorporates certain other household circumstances (such as consumption of water and electricity) when determining final eligibility. In Indonesia, the proxy means-testing models are estimated separately by region of the country to allow the model to account for regional variation in how consumption of goods predicts income across different areas. In both countries, prediction using proxy-means tests is used to determine eligibility for various transfer programs, but each individual program may edit the list of ultimate beneficiaries.

**The Inclusion versus Exclusion Error Tradeoff**

Our first step is to explore the tradeoffs in the errors of inclusion and exclusion. One way to think about the government's problem is that by setting different cutoffs for program eligibility $c$, the government can choose the balance between the inclusion and exclusion errors that it makes. Figure 3 plots the results with one example $c$, but one can see how varying (that is, shifting the vertical line in Figure 3 to the left or to the right) would change the balance of inclusion and exclusion error.[3] Say that the government aims to assist those who are actually poor. Intuitively, not giving out the program to anyone (and setting $c = 0$) means no transfers, and hence very high exclusion error, because you are excluding everyone below the poverty line, but it also means no inclusion error, since no higher-income people who should not be receiving assistance are getting it. In contrast, a universal basic income (setting $c = \infty$) implies no exclusion error since all of the poor will get it, but very high inclusion error, because all the higher-income people are getting assistance, too. For values in between, varying the cutoff value $c$ allows us to trace out the tradeoffs between inclusion error and exclusion error available to the government.

[3]Note that these definitions of inclusion and exclusion error are with respect to a household's true poverty status (above or below poverty line based on true per-capita household consumption; that is, holding the horizontal line in Figure 3 fixed), not with respect to the program design (that is, the eligibility choice $c$, shown by the vertical line).

*Figure 3*
**Predicted versus Actual per-capita Consumption for Households in Test Set Data**

A: Indonesia

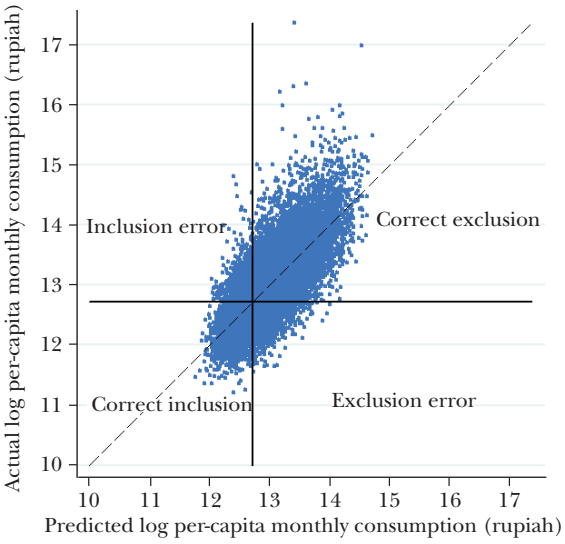


B: Peru

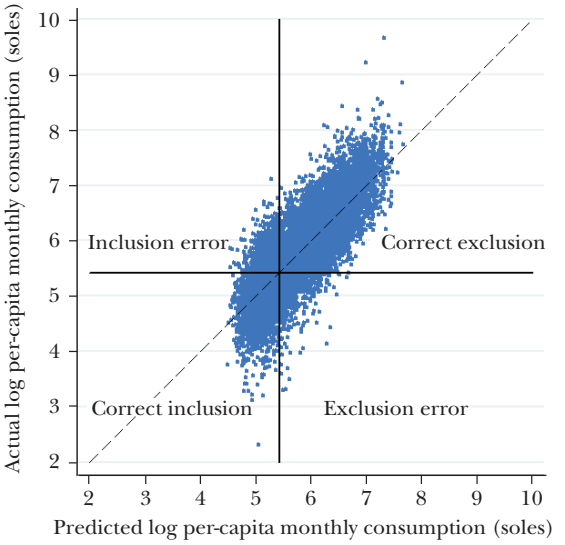


*Note:* The figures above plot actual log per-capita monthly consumption against predicted log per-capita monthly consumption for households in the test set data for both Indonesia and Peru (see text for details). The solid vertical and horizontal lines represent eligibility cutoffs for the programs we model here. (In Indonesia, we set the cutoff approximately at the 33rd percentile for consumption for the sample; and in Peru, we set the cutoff at the 28th percentile.) The dashed line is a 45° line. For legibility, the points plotted above represent a random sample of 10 percent of our full data from Indonesia and 50 percent of our full data from Peru.

These tradeoffs can be seen in Figure 4A, where we plot data from the proxy-means test regressions that we estimated for Indonesia and Peru. We define a target poverty level $y^*$ for each country, with the idea that the program is trying to target all individuals with household per-capita consumption below $y^*$.[4] We then define exclusion errors as failing to give benefits to those with household per-capita incomes below $y^*$, and inclusion errors as giving benefits to those with household per-capita incomes above $y^*$. These curves, also known as ROC curves,[5] plot inclusion error on the $x$-axis and 1 – exclusion error on the $y$-axis, and trace out the tradeoff between exclusion error and inclusion error obtained by varying the cutoff value $c$, holding the target poverty level for the program ($y^*$) fixed.

Figure 4A confirms the fundamental tradeoff between inclusion and exclusion error. In both Indonesia and Peru, to reach 80 percent of the intended beneficiaries (exclusion error of 20 percent), the government would need to tolerate an inclusion error of between 22 and 31 percent. Further reductions in exclusion error come at the cost of much higher inclusion error.

While more sophisticated prediction methods (like more flexible prediction equations, machine learning methods, and so on) can improve on these predictions (McBride and Nichols 2016), the fundamental tradeoff between inclusion and exclusion error that we document here remains. Also, these graphs ignore the very real problems of incomplete take-up by those who are deemed eligible, which could affect both the targeted programs and universal transfers. We will revisit this topic below.

**Narrowly versus Widely Targeted Programs with a Fixed Budget**

If the government only cares about making sure the poor have access to transfers of some type, a universal basic income—which eliminates exclusion error entirely—may seem attractive. However, as we move to the right in Figure 4A, by raising the cutoff level $c$, more and more people become eligible for the transfer. If the total budget for the program is fixed, then the transfer given per-person falls as we move to the right on this curve. Therefore, moving to the right on the curve not only changes the tradeoff between exclusion error and inclusion error, but holding the total program budget fixed, also changes the amount given.

Figure 4B shows this other aspect of the tradeoff, graphing the level of cash transfers disbursed per beneficiary under Indonesia's BLT program and Peru's Juntos program under different targeting schemes, as defined by different levels of the cutoff income $c$. The transfer level to eligible households steeply falls as you allow more people into the program. Again, the universal basic income can

[4]For Indonesia, we plot the ROC curve with $y^*$ equal to the 1.5× the government's official poverty line, and for Peru, we plot the curve with $y^*$ equal to the government's official poverty line. In the Appendix, we compare curves setting $y^*$ equal to 0.75× the government's official poverty line, which we call the "extreme poor"; 1.0× the government's official poverty line; and 1.5× the government's official poverty line, or the "near poor." The tradeoff is similar in all three cases.

[5]ROC stands for "receiver operating characteristic." It refers to a situation in which there is binary classifier—in this case, the proxy-means score combined with different cutoff levels $c$—and shows the tradeoff between true positives and false positives as one varies the cutoff level $c$.

*Figure 4*

**Tradeoffs between Inclusion Error and Exclusion Error by Varying Eligibility Cutoff**

*(exclusion error is the portion of intended beneficiaries excluded; inclusion error is the portion not intended to be beneficiaries who are included)*

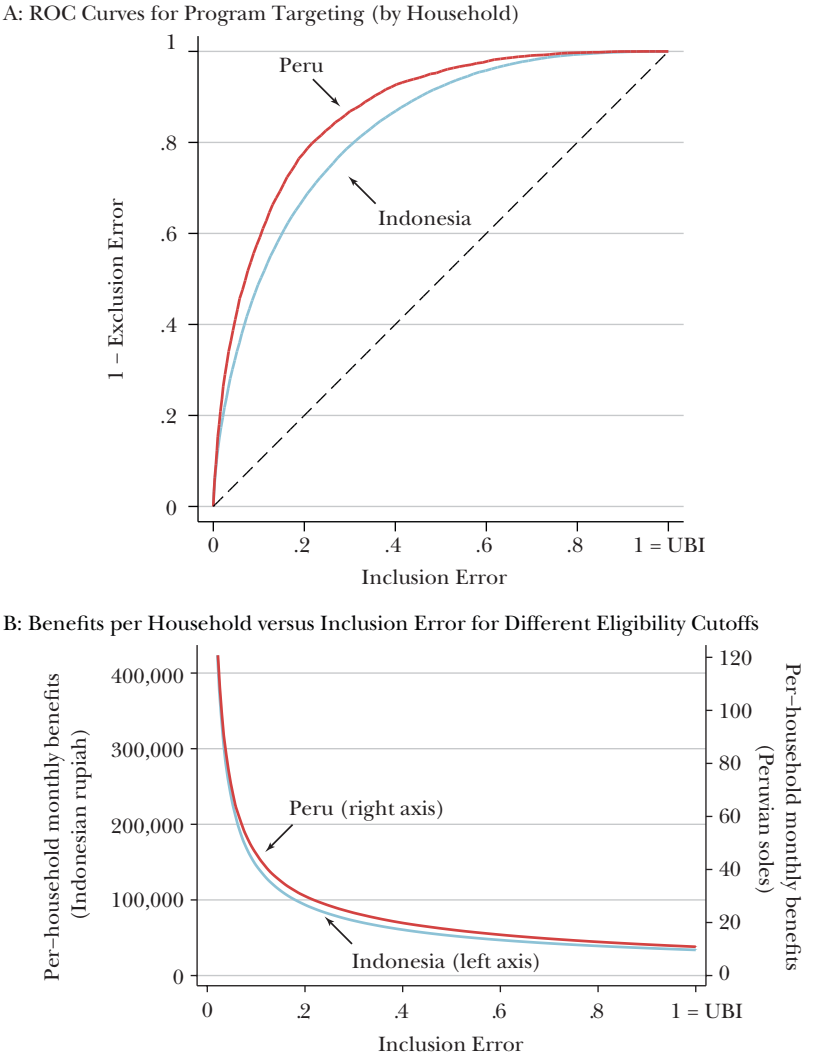


*Notes:* ROC stands for "receiver operating characteristic." It refers to a situation in which there is binary classifier and shows the tradeoff between true positives and false positives as one varies the cutoff level *c*. Figure 4A plots the ROC curves for both Peruvian and Indonesian test set data, using a target household per-capita consumption value equal to the government poverty line in Peru and 1.5 times the government poverty line in Indonesia. Figure 4B traces the per-household benefit amount as inclusion error increases, assuming a fixed transfer budget of approximately $1.83 billion per year in Indonesia and $274 million per year in Peru (modeled after the sizes of BLT and Juntos, respectively). We include cost-savings from administrative targeting at the universal basic income (UBI) point, where inclusion error equals 1. See text for details.

be thought of as the rightmost extreme point, with zero exclusion error and the smallest per-capita transfer for a given budget.

One argument for a universal basic income is that you can save on the administrative costs of targeting, because you do not need to do the periodic targeting census. To capture this, in Figure 4B we add back in the costs of targeting for the UBI, so the benefits per-capita ticks up very, very slightly at the universal basic income point ($c = \infty$). However, the administrative cost of targeting is very low relative to the benefits given out: Indonesia spends roughly $42 million every three years for updating its Unified DataBase, with annual operating costs of about $1.1 million (for further information, see Bah, Nazara, and Satriawan 2015). Peru spends roughly $10.8 million every three years with annual operating costs of about $1.1 million (Ministerio de Economía y Finanzas 2008). Per year, this translates to an additional 0.8 and 1.7 percent of the overall transfer budget in Indonesia and Peru, respectively. Thus, administrative costs are not a large driver of the benefits level, especially in comparison to the level of inclusion error one chooses, which is why the discontinuity at the universal basic income point is visually imperceptible in Figure 4B (although you will be able to see it in Figure 5).

**Welfare Comparisons**

For a given total budget, which point on the graph is "best" depends on various social welfare weights. For example, it is necessary to put a weight on inclusion errors, while remembering that just because someone is above the cutoff line does not imply that the social value of them receiving a transfer is zero—indeed, generically such a transfer still has a positive value. It is necessary to put a weight on exclusion errors: how costly is it to the government for some deserving people not to obtain the transfer, merely because of errors in the targeting formula? It is necessary to put a weight on the implied reduction in per-capita transfer to the very poorest that is required to accommodate the increased number of beneficiaries. It is also necessary to take a stand on how much more the poor value a given dollar than the rich (that is, the difference in marginal utilities for poor and rich), which is given by the curvature on the utility function. For a given amount of targeting, one can then determine the per-capita transfer.

Deciding how to parse the tradeoffs between inclusion error, exclusion error, and per-capita benefits requires specifying a social welfare function that allows the government to evaluate the social benefits from these different decision rules while holding the overall size of the transfer constant. Here, we will evaluate the total social welfare using what is commonly known as a constant relative risk-aversion CRRA-utility function:

$$U = \frac{\Sigma (y_i + b_i)^{1-\rho}}{1-\rho}$$

where $y_i$ is household $i$'s pre-tax per-capita income, $b_i$ is the per-capita benefits assigned to household $i$, and $\rho$ is a coefficient of relative risk-aversion, where higher values of $\rho$ put higher weights on transfers received by the very poor. (Note that

*Figure 5*
**Social Welfare versus Inclusion Error**

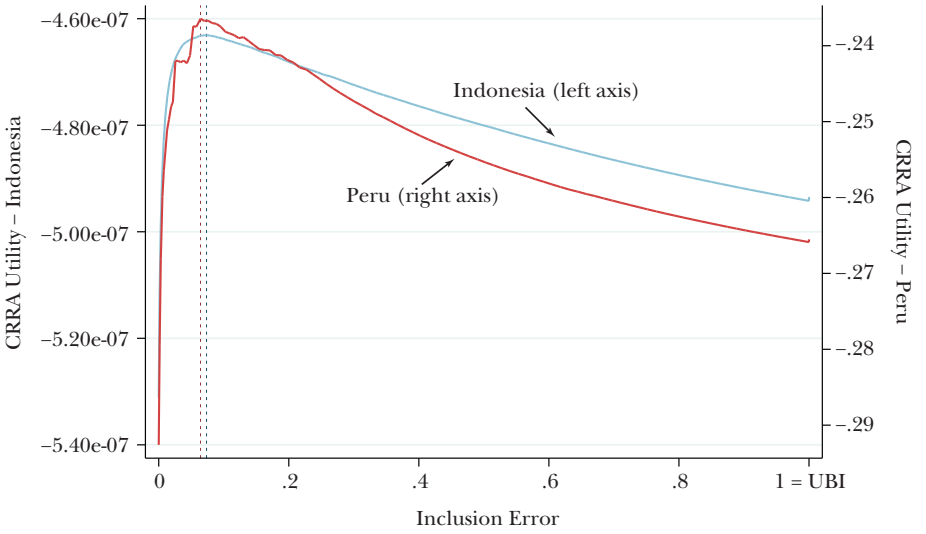


*Notes:* Figure 5 uses the same pooled household data as in Figure 4. For each cutoff value c, we calculate the per-capita benefit amount for included households, and then calculate constant relative risk-aversion utility (CRRA) using the formula presented in the text with $\rho = 3$. I. Utility is negative given the CRRA assumptions, with higher (that is, less negative) values indicating higher utility. If a household is not included in the program at a given value of $c$, we simply set $b_i = 0$, but their utility is still included. Dashed lines indicate the point of maximum social welfare in each country. We include cost-savings from administrative targeting at the universal basic income (UBI) point, where inclusion error equals 1.

utility is negative given the CRRA assumptions with less negative values indicating higher utility.) By comparing the sum of $U$ for different configurations of benefits, one can see which scheme produces the highest net social welfare judged by this metric.

Of course, conclusions will depend on the particular value of $\rho$ one chooses. Here, we calculate the social welfare from programs with different levels of targeting in Figure 5 for each country, using $\rho = 3$. (As a robustness check, we also calculated social welfare with $\rho = 1$ and $\rho = 5$, and the conclusions are qualitatively similar.) Specifically, in Figure 5, we plot social welfare evaluated against the amount of inclusion error for a fixed budget, for both countries. As before, social welfare will increase discontinuously for the universal basic income because no targeting costs are incurred in this program.

The key finding from this graph is that narrowly targeted programs—those focused on distributing large benefits-per-capita to the poorest of the poor—appear to achieve much higher utility levels than less narrowly targeted programs, including but not limited to a universal basic income. In Indonesia, the socially optimal program calculated in this way targets about 19 percent of the population, with inclusion error of 7.4 percent and exclusion error of 58.2 percent;

for Peru, the socially optimal program targets approximately 18 percent of the population, with inclusion error of 6.4 percent and exclusion error of 52.4 percent.

Even programs that appear somewhat badly targeted could still dominate a universal basic income. To understand where the programs in Indonesia and Peru fall in terms of social welfare, we can see where the current level of inclusion error roughly falls. Bah, Bazzi, Sumarto, and Tobias (2018) calculate that inclusion error in the 2008 Indonesian BLT program was roughly 34 percent, which is higher than the social optimum, but still substantially more narrowly targeted than a universal basic income. Peru's Juntos is at roughly 6.4 percent inclusion error, according to Robles, Rubio, and Stampini (2015), suggesting that the rate of inclusion error is close to the social optimum as calculated here.

In short, this analysis illustrates that targeting could be an effective mechanism in improving overall social welfare, even with high levels of targeting error. However, note that this discussion focuses only on tradeoffs involved in choosing whether to target transfers based on one survey—that is, how to identify the poor in the static situation of a given point in time. In reality, income status varies over time. Thus, the size of the relative errors will also depend on how frequently the government collects asset data from households and how much churn in and out of poverty occurs over time. Because targeting surveys are infrequent, actual targeting errors may be higher than what we report here, and so one may want to account for this factor in determining where to place the cutoff.

**Other Challenges with Targeted Transfers**

In choosing between a broader approach like a universal basic income and a more narrowly targeted approach, several other important considerations arise in addition to inclusion and exclusion errors. We consider four issues: breadth of political support, horizontal equity, transparency, and labor market distortions.

First, we have focused on the case where the government has a fixed budget available for the transfer program. However, political support for redistribution may be much higher if everyone gets a "piece of the pie" (Gelbach and Pritchett 2002). Thus, as inclusion error increases, the overall amount of funding available for transfers may increase. As the elasticity between the total budget and the number of beneficiaries increases, universal transfers become more attractive (Klasen and Lange 2016).

Second, the principle of horizontal equity holds that each person who has the same relevant conditions should be treated the same. Transfer programs using proxy-means tests will have errors in the prediction algorithm—and so they do not treat all people with the same utility of receiving the benefits identically. For programs where the income cutoff for providing benefits is near-zero, horizontal equity violations are (mechanically) close to zero, because almost no households receive the program. Violations of horizontal equity are also mechanically zero for universal basic income: by definition, if all households receive the transfer, then all households are treated the same.

*Figure 6*
**Horizontal Equity versus Inclusion Error in the Programs from Peru and Indonesia**

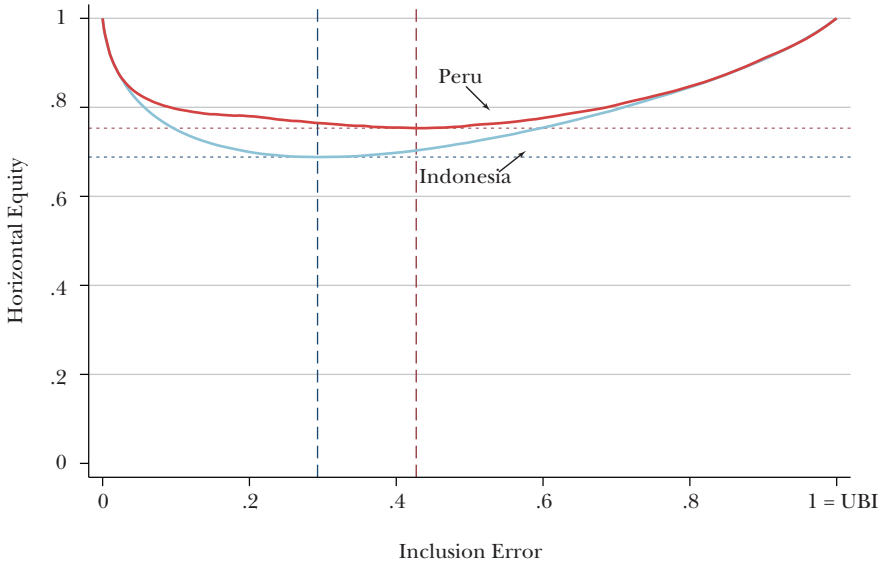


*Note:* Figure 6 traces horizontal equity as a function of inclusion error. At each cutoff $c$, we calculate, for each household, the percentage of households within ±5 income percentiles (based on actual income) that received the same benefit status—included or excluded—based on the results of proxy-means test prediction. In other words, for households that were included in the program at a given $c$, we calculate the percentage of similar households that were also included; for households that were excluded, we calculate the percentage of similar households that were also excluded. At each cutoff point, we average these percentages over all households to compute the horizontal equity variable graphed above. Dashed lines indicate the location and value of minimum horizontal equity.

To calculate a measure of horizontal equity and see how it varies with the cutoff value $c$, we do the following for both Peru and Indonesia. For each cutoff value $c$, we determine which households would receive the program and which would not. For each household income level, we then compute the percentage of households 5 percentiles above or below who have the same outcome (that is, for households receiving the program, the percentage of households 5 percentiles above and below who also would receive the program; for households not receiving the program, the percentage of households 5 percentiles above and below who would not receive the program). For each cutoff (and hence allocation rule) $c$, we average these two percentages over all households to get a measure of horizontal equity. We then repeat this exercise for each possible value of $c$.

Figure 6 shows the resulting measure of horizontal equity and how it varies as we vary the cutoff $c$, and hence the inclusion error, of the program. We find that for cutoff levels in an intermediate range, in which inclusion errors were roughly 30–40 percent, it is common to find that horizontal equity is violated about 20 percent of the time in Peru and 30 percent of the time in Indonesia. In addition to the

unfairness they represent, violations of horizontal equity can also lead to significant political problems.

A third important consideration involves the logistics of implementing targeting schemes in a low-capacity (and potentially corruption-prone) environment. A proxy-means test has an inherent lack of transparency, because eligibility is determined based on a weighted sum of many different variables, which would be hard for ordinary citizens to understand even if the weights were public information. Moreover, the proxy-means test formula (essentially the weight that the coefficients in the regression would give to each asset) are kept secret because if they are known, households (perhaps in cooperation with better-informed agents) may strategically misreport or hide assets to make sure they fall under the cutoff. Using data from Colombia, Camacho and Conover (2011) show that over time as the formula became known, there was substantial bunching of reporting right below the eligibility cutoff. Given this concern, governments not only keep the formula confidential, but they also tweak the coefficients every few years.

Thus, it is very difficult for citizens to verify that a proxy-means test scheme is being implemented properly. If a local official says to a citizen, "Sorry, you're not on the list of eligible beneficiaries," villagers have little recourse, because their true eligibility is hard to verify. On the other hand, a universal basic income is inherently transparent, in the sense that *everyone* is entitled to a transfer.

The evidence suggests that greater transparency of the lists of eligible beneficiaries may be important to effective implementation of a proxy-means test-based scheme. In Banerjee, Hanna, Kyle, Olken, and Sumarto (2018), we and our coauthors worked with the Indonesian government on a randomized experiment in which heads of all villages received a list of who was eligible to receive subsidized rice, but in a treatment group of villages, the central government also mailed out an "identification card" directly to citizens deemed eligible. Mailing identification cards to the beneficiaries resulted in treatment villages receiving 26 percent more subsidized rice compared to control villages—and reduced "leakage" (rice that, as far as we can tell, was not distributed to anyone) by between 33 and 58 percent.

Finally, targeting may introduce distortions. These distortions can, in theory, reduce the quantity of work, either because income effects decrease the incentive to work, or because phasing out the transfer over some income range decreases the after-tax-and-transfer effective marginal wage in that range. To study the income effects, we and our coauthors re-examined seven different randomized trials of government-run cash transfer programs throughout the developing world (Banerjee, Hanna, Kriendler, and Olken 2017), which provided between 4 and 20 percent of household consumption to beneficiaries for many years. We found no systematic evidence that any of these programs reduced labor supply.

Regarding the second channel (effects that come through changing the after-tax-and-transfer effective marginal wage), a vast literature in developed economies has studied the alteration in work incentives that arises either from the changes in the income tax rate formula needed to raise the revenue to pay for the transfer, or from the effective marginal tax rate imposed by the phase-out of the transfers.

In the US context, for example, the earned-income tax credit provides a negative marginal income tax rate at the very bottom of the tax schedule (thus subsidizing work), and was designed with these concerns in mind, given that there is evidence in the United States that labor supply elasticities are higher on the extensive margin (working at all versus staying out of the labor force) than on the intensive margin (how much you work) (for example, Eissa and Leibman 1996; Meyer and Rosenbaum 2001; Saez 2002). Relatedly, a major goal of the 1996 US welfare reform was to redesign the welfare system to reduce very high marginal tax rates that came from rapid phase-outs of benefits over certain income tax ranges. As discussed above, a universal basic income program does not eliminate this issue entirely, although the relevant distortions to think about for a universal basic income are the increases in marginal income tax rates required to finance the universal basic income.

These issues are related to, but somewhat different from, what happens with a proxy-means test in a developing economy. In this setting, the labor supply distortion will involve the implied tax rates from the phase-out of the program for those who are eligible, but will also be smoothed out by the error rate in the formula for those who have equivalent income (or consumption) but are not designated as eligible by the formula.[6] Perhaps ironically, a more accurate proxy-means test formula, with fewer errors, may actually induce *more* labor supply distortions, since fewer errors in the proxy-means formula will imply that some households face a steeper effective tax rate over some range. More noise in the proxy-means test formula, making the proxy-means formula less predictive, tends to smooth out the phase-out region, and will thus tend to reduce the implicit tax rate that a population faces from a targeted benefit.

For example, Figure 7 illustrates the implied marginal tax rates from actual proxy-means test formulas and from a counterfactual formula where we double the amount of noise, making the proxy-means test less predictive than it is in practice. As the figures show, doubling the amount of noise in the proxy-means test formula reduces the implied marginal tax rate for the poor, but it increases the implied marginal tax rate somewhat for the middle-class, who face some probability of receiving the benefits (due to inclusion error) that is declining in their income. If households are aware that increasing their income or consumption may reduce their chances of being eligible for benefits, they may reduce their labor supply accordingly, though documenting to what extent this happens empirically remains an important direction for future research.

The proxy-means test puts particularly heavy weight on certain assets. While the precise proxy-means test formula may not be known, households do observe which variables the government asks about—that is, they know that the poverty census asks about televisions, refrigerators, and so on. The potential for linkage from ownership

[6] To see this, recall that for each income level $y$, one can think of the proxy-means test as producing a probability of receiving the benefit $p(y)$. The fact that the proxy-means test has predictive power implies that $p$ will be decreasing with income ($p'(y) < 0$), but the fact that there are inclusion and exclusion errors means that $p(y)$ is continuous. Abstracting away from the rest of the tax system, one can then write a household's expected after-tax income as $\alpha(y) = y + p(y)b$. One can therefore think of the "expected tax rate" imposed by the proxy-means test system as $-p'(y)b$.

*Figure 7*

**Implied Tax Rate versus Household Income, with Noisy Proxy Means Test Formulas**

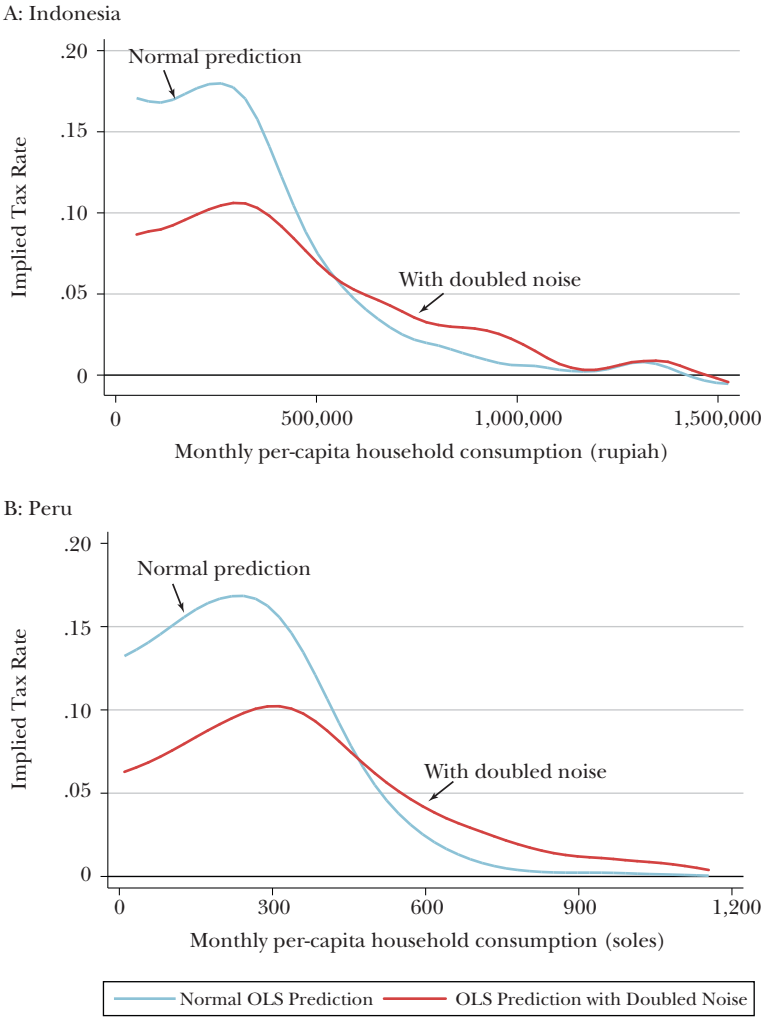


*Note:* Figure 7 illustrates the implied marginal tax rates from actual proxy-means test formulas and from a counterfactual formula where we double the amount of noise, making the proxy-means test less predictive than it is in practice. The implied marginal tax rates are plotted against actual per-capita household consumption for the Indonesian and Peruvian cases described in the text. We designated households as benefit recipients if *predicted* income was less than 1.5 times the poverty line (in Indonesia) or less than 1 times the poverty line (in Peru). Households above the 95th income percentile were dropped from both country samples. We then performed a local polynomial regression of benefit receipt status on *actual* income, using a bandwidth of 50,000 rupiah (Indonesia) and 75 soles (Peru). We took the first derivative of this function and multiplied it by (minus) the benefit amount $b$ in order to calculate the implied tax rate as a function of household income ($-p'(y)\,b$). We then applied the same process after adding a random, normally distributed noise term (with the same standard deviation as our original predicted incomes) to these predicted income values. OLS is ordinary least squares.

of such assets to reduced eligibility for benefits puts a particularly high implicit tax rate on these particular assets, which could lead households to reduce their consumption of these assets. For example, from 1696 to 1851, Britain imposed a tax on windows, which were easily observable to the taxman—but this led to the construction of buildings with very few windows, low light, and poor ventilation (Glantz 2008; Oates and Schwab 2015). It is not clear that households understand the formula in a proxy-means test enough to distort their purchases of assets, and we are currently studying this question in cooperation with the government of Indonesia.

## Alternative Methods of Targeting

Our discussion has assumed that the government produces an eligibility list for targeted transfers through a proxy-means test based on household-level data. However, we now consider two other methods of targeting: community-based targeting and a system that imposes costs on beneficiaries and lets people assist in targeting themselves. We also consider how conditional cash transfer programs, popular throughout the developing world, can be thought of in relation to the targeting challenge.

### Community-Based Targeting

In a community-based targeting process, a fixed number of program slots are allocated to a given community, such as a hamlet (a part of a village). The hamlet, through a participatory meeting of some type, decides who among them is most in need of the transfer.

One might think that such a system would be unpopular: after all, who wants to discuss their income in front of their neighbors and peers? (Not Americans, as discussed in Card, Mas, Moretti, and Saez 2012!) In addition, one might fear that such a system is prone to elite capture: for example, a powerful headman might manipulate the meeting to put his nephew on the list of beneficiaries. However, when we and our coauthors tried such an approach experimentally in Indonesia and compared it experimentally, we found that citizens vastly preferred the community based-approach to a data-based proxy-means test. In addition, the community approach did a much better job of identifying those households who self-assessed themselves to be poor, and was only slightly worse at identifying the poor based on per-capita consumption levels. Furthermore, there was no detectable elite capture (Alatas, Banerjee, Hanna, Olken, and Tobias 2012).

A community targeting approach avoids some key pitfalls of the proxy-means test related to transparency and horizontal equity. The process for allocating beneficiaries under the community approach was completely public and transparent—unlike the proxy-means test that relies on secret formulas to produce a list of eligible and ineligible households that villagers do not really understand. It also addresses the perception of a horizontal equity problem—from the perspective of the villagers, if it is broadly agreed that person *A* is poorer than person *B*, they can ensure that person *A* receives the transfer.

Community targeting also has the potential to be less distortionary. In Alatas et al. (2012), we showed that communities appear to target based on earning potential, rather than actual earnings. For example, widows—who have lower earning potential than households headed by men—appear to be more likely to be ranked as poor, even conditional on their actual income level. The same is true for education: those with high education, and hence high earning power, are ranked as wealthier than their actual income alone would imply. To the extent that communities target based on *earning potential* rather than actual income, such a system has the potential to be less distortionary not only than a proxy-means test, but also than a regular income tax.

Community targeting has important limitations—in particular, it does not readily allow for comparisons *across* locations, and if implemented at scale, could lead households to take actions to hide their income from their neighbors. But its potential advantages may mean that it is worth wider consideration as part of the targeting toolkit.

**Differential Costs of Take-Up and Self-Targeting**

A substantial fraction of households who are eligible for a government benefit program do not enroll in the program. Moffitt (2003) and Currie (2006a) summarize an extensive US-based literature studying the reasons for incomplete take-up, including the stigma of receiving benefits, incomplete information, and both the psychic and real costs of applying for programs.

A universal basic income may reduce the take-up problem, because the government can in theory send the checks to everyone automatically. This is, of course, easier said than done. Currie (2006b) shows that take-up is incomplete even in programs with universal eligibility in the United States, and a developing country setting raises additional issues. For example, the government would need to have a list of every individual without duplications. Most developing countries do not have such lists, although India has been working over the past decade to issue a unique, unduplicated and biometrically authenticated identity number through the Aadhaar program to each of its more than 1 billion citizens. Even with such a list, a country would need a way of automatically delivering the transfer. In most developing country contexts, most people do not have an address, let alone a bank account, so the idea of "mailing a check to the address on file" would not work. This suggests that even if enrollment is in theory automatic, in practice some take-up decisions and enrollment actions will need to occur.

However, having some type of costly take-up can also provide a benefit by helping in the targeting of a program. If the cost/benefit of taking up a program is higher for the rich than for the poor, the poor may be more likely to take up the program. This idea—that having an "ordeal" associated with receiving a program can improve the ability to screen among rich and poor—was developed by Nichols and Zeckhauser (1982), and further studied in the context of take-up decisions by Kleven and Kopczuk (2011). A key result is that one may want to offer a transfer program to all takers, but with the catch that signing up for the program entails a cost that is differentially costly for the rich.

Workfare programs, such as those offered by the US Works Progress Administration in the 1930s and through the National Rural Employment Guarantee Act (NREGA) in India today, operate on this principle. In the NREGA program, for example, any Indian in a rural area is entitled to 100 days of work at the official minimum wage, and currently about 180 million people per year are employed through the program, making it one of the largest (if not the largest) anti-poverty programs in the world today. This is a universal entitlement—there is no poverty screening, and there is a guarantee to everyone who wants the work. On the other hand, people with better job options tend to prefer doing something other than performing manual labor for minimum wage in the hot sun. Whether these types of programs—universal in nature, but achieving screening through so-called ordeal mechanisms—are better or worse than a universal basic income program depends on whether the substantial utility costs imposed on beneficiaries (like manual labor in the hot sun in order to receive benefits) are outweighed by the cost savings from improved targeting.

However, it may be possible to obtain many of the benefits from self-selection with much smaller utility costs to recipients if one marries self-selection to a proxy-means test. In Alatas et al. (2016), we and our coauthors study what happens when, instead of having the government come to everyone's house and conduct the poverty census automatically, households instead have to come and apply for the program, at which point they are screened using the proxy-means test.[7] Some fraction of households who pass the in-person eligibility test have their eligibility verified via a home visit, but this allows the government to skip the home visits for those not eligible and to check the eligible at home probabilistically, avoiding substantial costs. The proxy-means test is the same—what differs is the take-up step. Compared to manual labor in the sun, the cost of applying is relatively small—about three hours total in our context—and pales in comparison to the potential benefits, which average about $150 per year for six years in the context we studied. Yet, despite the relatively small costs, we show that self-selection yielded substantially improved screening; the beneficiaries selected by the application-based method were about 20 percent poorer than those selected through automatic enrollment.

The key seems to be that the application process substantially reduces the inclusion errors in the proxy-means test process while holding the official eligibility cutoff $c$ fixed. A relatively well-off household forecasts correctly that the probability they pass the proxy-means test is small, and hence even the relatively small application cost can be enough to discourage them from applying. Since there are so many more relatively well-off households than very poor households, inducing these households to self-select out can actually lead to a substantial improvement in

[7]Both the automatic door-to-door approach and the application-based approach are used in other contexts. For example, Colombia uses an automatic enrollment door-to-door proxy-means test; Costa Rica and Chile use application-based proxy-means tests; and both Mexico and Brazil use different approaches in different contexts.

program targeting. This logic suggests that the benefits from adding self-selection to a proxy-means test system may be largest when the nontarget population is large relative to the target population. Moreover, an additional benefit is that some of the poorest households may live at the very margins of society and therefore be off the government's radar screen, and may be missed by enumerators trying to do the universal door-to-door screen. When beneficiaries can come to apply, such households can make themselves known.

However, application processes need not necessarily always improve targeting. For example, a complicated application form may dissuade those who are less literate or comfortable with bureaucracy from filling it out, leading to worse targeting (Gupta 2017). The degree to which these approaches can improve social welfare depend on whether the selection they induce enables the government to improve the efficiency with which it can deliver assistance to those who need it most, and also on the extent to which the efficiency gains achieved thereby are sufficiently large (and hence can translate into larger benefits) to offset the costs they impose on beneficiaries.

**Conditional Transfers**

A number of transfer programs in developing countries have explicit conditions that beneficiaries must meet in order to receive assistance. For example, one of the first of these programs, Mexico's *Prospera* program (originally called *Progresa*), gives cash transfers to poor households (screened via a proxy-means test) who also meet basic maternal and child health and education conditions, including regular pre- and postnatal care, regular growth monitoring, immunizations, school enrollment, and school attendance. Conditional cash transfer programs have spread throughout the world, and now are present in more than 63 countries, reaching tens of millions of households annually (Bastagli et al. 2016).

A number of randomized trials of these programs, including the original launch of *Progresa*, showed that these programs led to substantial improvements on the conditioned indicators, both initially (for example, see Behrman and Todd 1999; Gertler 2004) and over the medium term (Behrman, Parker, and Todd 2011; Barham, Macours, and Maluccio 2017; Kugler and Rojas 2018; Cahyadi, Hanna, Olken, Prima, Satriawan, and Syamsulhakim 2018). These conditions also may help make the programs more politically palatable, since voters in many countries may prefer that individuals do something in return for receiving aid.

While much attention has been given to the *incentive* effects of the health and education conditions, they can also have *targeting* effects. In particular, if some very poor households are unable to meet the conditions, and if the conditions are actually enforced (which they are not always), then some very poor households may be excluded from the programs. Baird, McIntosh, and Özler (2011) studied this issue, randomly comparing unconditional and conditional transfer programs in Malawi. While they found that the conditions increased school enrollments, they noted that the unconditional program actually had a larger effect on reducing teen pregnancy. The reason was that the unconditional program (unlike the conditional program)

still provided cash transfers to girls who dropped out of school anyway and were at the highest risk of pregnancy, and the cash transfers reduced the likelihood of early pregnancy for these girls.

While in principle conditional cash transfers could be either universal (anyone who meets the education and health conditions receives the transfer) or targeted using a proxy-means test, almost all are targeted, most often using proxy-means tests (Fiszbein and Schady 2009). One reason for this is that, in general, richer households are more likely to meet the health and education conditions than the poor. Thus, in the absence of explicit additional restrictions limiting these programs to the poor, the fact that the conditions are more likely to be fulfilled by wealthier households would actually make a universal version of these programs regressive rather than progressive.

## Summing Up

Government-led anti-poverty programs have a crucial role to play in helping to eliminate extreme poverty worldwide. However, important questions remain about the form of these transfers, and in particular, whether they should be universal or targeted more narrowly to the poor. While some government programs are universal, such as publicly funded primary schools that are free to the families using them, most cash transfer programs today involve some type of targeting mechanism.

In this paper, we explore how to think about this tradeoff in developing countries, where incomes are unobserved for very large portions of the population. In this setting, universal transfers cannot be easily taxed back as one moves up the income distribution. As a result, universal transfer programs will give out the same *net* transfer quite high up the income distribution—and therefore the choice between a universal basic income program and transfers targeted through other means implies a very substantial tradeoff between eliminating exclusion error and giving much smaller transfers on a per-beneficiary basis.

Our evidence from Indonesia and Peru shows that existing targeting methods in developing countries, while imperfect, appear to deliver substantial improvements in welfare compared to universal programs, because they can transfer much more on a per-beneficiary basis to the poor as compared with universal programs. The primary downside of these programs is horizontal equity—because targeting is imperfect, there will be a substantial number of poor households who slip through cracks and are excluded. Nevertheless, for many developing countries, our simulations suggest the welfare gains from targeting may be substantial.

■ *We thank Abhijit Banerjee, Amy Finkelstein, and Rohini Pande for helpful comments and discussions on these issues, and Aaron Berman and Samuel Solomon for helpful research assistance.*

## References

**Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, Ririn Purnamasari, and Matthew Wai-Poi.** 2016. "Self-Targeting: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 124(2): 371–427.

**Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias.** 2012. "Targeting the Poor: Evidence from a Field Experiment in Indonesia." *American Economic Review* 102(4): 1206–40.

**Bah, Adama, Samuel Bazzi, Sudarno Sumarto, and Julia Tobias.** 2018. "Finding the Poor vs. Measuring Their Poverty: Exploring the Drivers of Targeting Effectiveness in Indonesia." World Bank Policy Research Working Paper 8342.

**Bah, Adama, Suahasil Nazara, and Elan Satriawan.** 2015. "Indonesia's Single Registry for Social Protection Programmes." Research Brief 49, International Policy Centre for Inclusive Growth.

**Baird, Sarah, Craig McIntosh, and Berk Özler.** 2011. "Cash or Condition? Evidence from a Cash Transfer Experiment." *Quarterly Journal of Economics* 126(4): 1709–53.

**Banerjee, Abhijit, Rema Hanna, Jordan Kyle, Benjamin A. Olken, and Sudarno Sumarto.** 2018. "Tangible Information and Citizen Empowerment: Identification Cards and Food Subsidy Programs in Indonesia." *Journal of Political Economy* 126(2): 451–91.

**Banerjee, Abhijit, Rema Hanna, Gabriel E. Kreindler, and Benjamin A. Olken.** 2017. "Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs Worldwide." *World Bank Research Observer* 32(2): 155–85.

**Barham, Tania, Karen Macours, and John Maluccio.** 2017. "Are Conditional Cash Transfers Fulfilling Their Promise? Schooling, Learning, and Earnings after 10 Years." CEPR Discussion Paper DP11937.

**Bastagli, Francesca, Jessica Hagen-Zanker, Luke Harman, Valentina Barca, Georgina Sturge, and Tanja Schmidt.** 2016. *Cash Transfers: What Does the Evidence Say? A Rigorous Review of Programme Impact and the Role of Design and Implementation Features.* With Luca Pellerano. Overseas Development Institute Report, July.

**Behrman, Jere R., Susan W. Parker, and Petra E. Todd.** 2011. "Do Conditional Cash Transfers for Schooling Generate Lasting Benefits?: A Five-Year Follow-up of PROGRESA/Oportunidades." *Journal of Human Resources* 46(1): 93–122.

**Behrman, Jere R., and Petra E. Todd.** 1999. "Randomness in the Experimental Samples of PROGRESA (Education, Health, and Nutrition Program)." IFPRI Research Report.

**Besley, Timothy.** 1990. "Means Testing versus Universal Provision in Poverty Alleviation Programmes." *Economica* 57(225): 119–29.

**Cahyadi, Nur, Rema Hanna, Benjamin A. Olken, Rizal Adi Prima, Elan Satriawan, and Ekki Syamsulhakim.** 2018. "Cumulative Impacts of Conditional Cash Transfer Programs: Experimental Evidence from Indonesia." NBER Working paper 24670.

**Camacho, Adriana, and Emily Conover.** 2011. "Manipulation of Social Program Eligibility." *American Economic Journal: Economic Policy* 3(2): 41–65.

**Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez.** 2012. "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction." *American Economic Review* 102(6): 2981–3003.

**Chetty, Raj, and Adam Looney.** 2006. "Consumption Smoothing and the Welfare Consequences of Social Insurance in Developing Economies." *Journal of Public Economics* 90(12): 2351–56.

**Currie, Janet M.** 2006a. *The Invisible Safety Net: Protecting the Nation's Poor Children and Families.* Princeton University Press.

**Currie Janet.** 2006b. "The Take-up of Social Benefits." Chap. 3 in *Public Policy and the Income Distribution*, edited by Alan Auerbach, David Card, and John M. Quigley. New York: Russell Sage Foundation.

**Eissa, Nada, and Jeffrey B. Liebman.** 1996. "Labor Supply Response to the Earned Income Tax Credit." *Quarterly Journal of Economics* 111(2): 605–637.

**Fiszbein, Ariel, and Norbert Schady.** 2009. "Conditional Cash Transfers: Reducing Present and Future Poverty." With Francisco H. G. Ferreira, Margaret Grosh, Niall Keleher, Pedro Olinto, and Emmanuel Skoufias. World Bank Policy Research Report. World Bank.

**Gelbach, Jonah B., and Lant Pritchett.** 2002. "Is More for the Poor Less for the Poor? The Politics of Means-Tested Targeting." *Topics in Economic Analysis & Policy* 2(1).

**Gertler, Paul.** 2004. "Do Conditional Cash Transfers Improve Child Health? Evidence from PROGRESA's Control Randomized Experiment." *American Economic Review* 94(2): 336–41.

**Glantz, Andrew E.** 2008. "A Tax on Light and Air: Impact of the Window Duty on Tax Administration and Architecture, 1696–1851." *Penn History Review* 15(2): Article 3.

**Golan, Jennifer, Terry Sicular, and Nithin Umapathi.** 2017. "Unconditional Cash Transfers in China: Who Benefits from the Rural Minimum Living Standard Guarantee (Dibao) Program?" *World Development* 93: 316–36.

**Gupta, Sarika.** 2017. "Perils of the Paperwork: The Impact of Information and Application Assistance on Welfare Program Take-Up in India." Available at: https://scholar.harvard.edu/sarikagupta/publications/mediation-deficit-when-information-isn%E2%80%99t-enough-take-welfare-programs.

**Hsieh, Chang-Tai, and Benjamin A. Olken.** 2014. "The Missing 'Missing Middle.'" *Journal of Economic Perspectives* 28(3): 89–108.

**International Labor Organization.** No date. "ILO Social Security Inquiry." http://www.ilo.org/dyn/ilossi/ssimain.home?p_lang=en.

**Jensen, Anders.** 2016. "Employment Structure and the Rise of the Modern Tax System." Available at: https://www.semanticscholar.org/paper/Employment-Structure-and-the-Rise-of-the-Modern-Tax-Jensen/e7adac3d8699faca906b8f834be6b88cf870fc30.

**Jones, Benjamin F., and Benjamin A. Olken.** 2008. "The Anatomy of Start-Stop Growth." *Review of Economics and Statistics* 90(3): 582–87.

**Klasen, Stephan, and Simon Lange.** 2016. "How Narrowly Should Anti-poverty Programs Be Targeted? Simulation Evidence from Bolivia and Indonesia." Courant Research Center: Poverty, Equity and Growth—Discussion Papers, no. 213.

**Kleven, Henrik Jacobsen, and Wojciech Kopczuk.** 2011. "Transfer Program Complexity and the Take-Up of Social Benefits." *American Economic Journal: Economic Policy* 3(1): 54–90.

**Kleven, Henrik Jacobsen, Claus Thrustrup Kreiner, and Emmanuel Saez.** 2016. "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries." *Economica* 83(330): 219–46.

**Kugler, Adriana D., and Ingrid Rojas.** 2018. "Do CCTs Improve Employment and Earnings in the Very Long-Term? Evidence from Mexico." NBER Working Paper 24248.

**McBride, Linden, and Austin Nichols.** 2016. "Retooling Poverty Targeting Using Out-of-Sample Validation and Machine Learning." *World Bank Economic Review*. https://doi.org/10.1093/wber/lhw056.

**Meyer, Bruce D., and Dan T. Rosenbaum.** 2001. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." *Quarterly Journal of Economics* 116(3): 1063–1114.

**Ministerio de Economía y Finanzas.** 2008. *Marco Macroeconómico Multianual 2009–2011*. Lima, Peru: MFF.

**Mirrlees, James A.** 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38(2): 175–208.

**Moffitt, Robert A. (ed.)** 2003. *Means-Tested Transfer Programs in the United States*. University of Chicago Press.

**Nichols, Albert L., and Richard J. Zeckhauser.** 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review* 72(2): 372–77.

**Oates, Wallace E., and Robert M. Schwab.** 2015. "The Window Tax: A Case Study in Excess Burden." *Journal of Economic Perspectives* 29(1): 163–80.

**Ravallion, Martin.** 2009. "How Relevant is Targeting to the Success of an Antipoverty Program?" *World Bank Research Observer* 24(2): 205–231.

**Robles, Marcos, Marcela G. Rubio, and Marco Stampini.** 2015. "Have Cash Transfers Succeeded in Reaching the Poor in Latin America and the Caribbean?" Inter-American Development Bank Policy Brief no. IDB-PB-246.

**Saez, Emmanuel.** 2002. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *Quarterly Journal of Economics* 117(3): 1039–1073.

**Salehi-Isfahani, Djavad, and Mohammad H. Mostafavi-Dehzooei.** 2017. "Cash Transfers and Labor Supply: Evidence from a Large-Scale Program in Iran." Economic Research Forum Working Paper 1090.

**Stantcheva, Stefanie.** 2014. "Optimal Income Taxation with Adverse Selection in the Labour Market." *Review of Economic Studies* 81(3): 1296–1329.

**World Bank.** 2012. "BLT Temporary Unconditional Cash Transfer." Social Assistance Program and Public Expenditure Review 2.

**World Bank.** 2014. *The State of Social Safety Nets 2014*. World Bank.

**World Bank.** Various years. "World Development Indicators."

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 227–240*

# Retrospectives
# On the Genius Behind David Ricardo's 1817 Formulation of Comparative Advantage

Daniel M. Bernhofen and John C. Brown

This feature addresses the history of economic terms and ideas. The hope is to deepen the workaday dialogue of economists, while perhaps also casting new light on ongoing questions. If you have suggestions for future topics or authors, please contact Joseph Persky, Professor of Economics, University of Illinois, Chicago, at jpersky@uic.edu.

## Introduction

David Ricardo's formulation of comparative advantage is one of the oldest analytical results in economics. In a famous paragraph in the trade chapter of his *Principles of Political Economy and Taxation*, Ricardo (1817 [1966], pp. 134–135) employed what Paul Samuelson (1969) referred to as "Ricardo's four magic numbers":

> The quantity of wine which she [Portugal] shall give in exchange for the cloth of England, is not determined by the respective quantities of labour devoted to the production of each, as it would be, if both commodities were manufactured in England, or both in Portugal. England may be so circumstanced, that to produce the cloth may require the labour of 100 men for one year;

■ *Daniel M. Bernhofen is Professor of International Relations, School of International Service, American University, Washington, DC. John C. Brown is Professor of Economics, Clark University Worcester, Massachusetts. Bernhofen is the corresponding author at dbernhof@american.edu.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.227

doi=10.1257/jep.32.4.227

> and if she attempted to make the wine, it might require the labour of 120 men for the same time. England would therefore find it her interest to import wine, and to purchase it by the exportation of cloth. To produce the wine in Portugal, might require only the labour of 80 men for one year, and to produce the cloth in the same country, might require the labour of 90 men for the same time. It would therefore be advantageous for her to export wine in exchange for cloth.

As far back as Samuelson (1948), the four magic numbers have inspired the basis for numerous textbook treatments of the principle of comparative advantage. Following the lead of James Mill's (1821) reading of Ricardo's famous passage, these numbers have been interpreted as fixed unit labor requirements and have served as a powerful pedagogical device to illustrate the idea of comparative advantage and the gains two countries realize when they move from autarky to trading with each other.

But in a few key respects, the approach Ricardo followed in developing his argument appears at odds with the modern textbook treatment of comparative advantage. Instead of a comparison between autarky and trade, his canonical example starts with the evaluation of an existing trading relationship between England and Portugal. If the numbers are interpreted as labor requirements per unit of output, all four numbers are needed for a relative cost comparison between England and Portugal. However, Ricardo draws a conclusion about England's pattern of trade and its gains from trade based only on the magic numbers for England; he then draws a conclusion about Portugal's pattern of trade and its gains from trade based on the magic numbers for Portugal. His discussion makes no statement about the terms of trade between wine and cloth.[1] The contrast between Ricardo's presentation of comparative advantage and the labor unit requirement interpretation of his thinking has led to a debate in the history of thought literature discussed in Aldrich (2004) as to whether Ricardo really understood his principle.

The recent 200th anniversary of the publication of Ricardo's 1817 statement of comparative advantage offers an opportunity for revisiting the concerns of Ricardo and his contemporaries that motivated his development of the concept of comparative advantage. Our discussion offers a deeper reading of the fundamentals of Ricardo's logic, which viewed trade in goods as equivalent to workers moving across borders. From this perspective, the four magic numbers pertain to the amount of labor embodied in trade. Our account starts with mercantilism and the emergence of the "18th century rule" more than a century prior to the publication of the *Principles of Political Economy*.

[1] By contrast, if one assumes a terms of trade of one between cloth and wine, as Eaton and Kortum do in their recent discussion of Ricardo (in this journal, 2012), then the four numbers become both the labor value of trade and the unit labor coefficients.

## Identifying Why Trade Is Advantageous: From Mercantilism to the 18th Century Rule

The voyages of discovery of the late 15th and 16th centuries that opened up the European conquest of the Americas and the seaborne trade routes to Asia ushered in an era of unprecedented commercial rivalry among the great European powers of the era: England, the Dutch Republic, Spain, Portugal, and France. The trade was conducted either by large monopoly companies such as the British East India Company or by independent merchants. It focused on sugar, tobacco, and indigo from new colonies in the Caribbean and North America; gold and silver from Latin America; and products of the East such as spices, cotton cloth from India, and raw silk from China. The rivalry prompted one-and-a-half centuries of theorizing about how to maximize the benefit that states received from that trade—and finance the extraordinary military expenditures required to successfully compete. The doctrine that emerged and the policy recommendations of the period have become known as "mercantilism." Viner (1937) and Heckscher (1935) are classic attempts to characterize the theory and policy recommendations of mercantilism.

Mercantilist thought was pro-trade, but for mercantilists, the gains from trade emanated from an excess of exports over imports in enough branches of trade to allow the accumulation of the gold and silver bullion required to finance other areas (Irwin 1996, pp. 32–33).[2] The incentives of private traders were not always aligned with those of the state. Both as a response to concerns about overpopulation and a belief that exports of goods that embodied a substantial amount of domestic labor were the best means to secure a surplus, mercantilist trade policy favored exports that absorbed large amounts of labor (Grampp 1952, pp. 467–72; Tucker 1750). That meant promoting exports of manufactures that were among the most labor-intensive industries of pre-industrial economies. The main British export of the early mercantilist period—woolen cloth manufactured from English wool—fit the bill. Mercantilist policy focused on reducing imports of competing manufactures, unless they contributed to enhancing the productive capacity of the state. Instead, imports that used the least amount of labor, such as raw materials or foodstuffs, increased the balance of exported labor and the gains from trade: "When two countries are exchanging their produce or manufactures with each other, that nation which has the greatest number employed in this reciprocal trade; is said to receive a balance from the other; because the price of the overplus labor must be paid in gold and silver" (Tucker 1750, p. iii, as quoted in Viner 1937, p. 53). By contrast, imports of luxuries (French silks or lace) consumed only by the wealthiest classes were to be discouraged since they resulted in a *net* importation of embodied labor.

The idea that imports—not exports—could generate gains from trade first appeared in Henry Martyn's (1701) *Considerations of the East India Trade*. Martyn's

[2]Wilson (1949) notes that England's ongoing construction of a large naval fleet required importation of wood and iron from the Baltic for which it paid with gold and silver bullion.

argument, which appears to have gained few adherents at the time, would re-emerge a century later as a core element of Ricardo's conceptualization of the gains from trade. During the last decades of the 17th century, the British East India Company was importing cotton cloth (calicos) from Bengal, primarily for re-export to the continent and the American colonies. The cotton cloth imports threatened the domestic English woolen industry and prompted a backlash of tariffs and eventually a prohibition on the consumption of any all-cotton cloth in England. In the midst of a debate over how to limit these imports, Martyn argued that imports of cotton cloth freed up resources that could be deployed better elsewhere in the English economy. Martyn was the first to enunciate what Jacob Viner (1937, p. 440) termed the "eighteenth-century rule" (p. 440–41) that overturned the standard mercantilist view of how trade benefitted the national economy, a rule which Viner summarizes as: "[I]t pays to import commodities from abroad whenever they can be obtained in exchange for exports at a smaller real cost than their production at home would entail" (p. 440).

Seventy years after the 18th century rule appeared as part of an unsuccessful attempt to fend off an import prohibition, Adam Smith (1776) invoked it in the *Wealth of Nations*. Smith's perspective that trade could enhance the wealth of a country via specialization also reflected a world where the scarcity of productive means meant that savings (and gains) could be achieved with specialization and exchange: "If a foreign country can supply us with a commodity cheaper than we ourselves can make it, better buy it of them with some part of the produce of our own industry, employed in a way in which we have some advantage" (Smith 1776, Book IV, p. 185).

The ongoing British struggle against Napoleon in the first decade of the 19th century prompted the clearest expression of the 18th century rule by the early classical economists. The 1807 Orders of Council imposed a severe trade blockade on Napoleon's Europe and of course, on British exports to the continent. In response to arguments in some quarters that Britain should pursue a policy of self-sufficiency, James Mill (1808) and Robert Torrens (1808) argued forcefully for the advantages that accrued to a trading nation. In particular, Torrens (1808, p. 37) offers a remarkable and strikingly modern account that captures the gains that England had secured from its trading relationship with France:

> If I wish to know the extent of the advantage, which arises to England, from her giving France a hundred pounds of broad cloth, in exchange for a hundred pounds of lace, I take the quantity of lace which she might, at the same expense of labor and capital, have acquired by manufacturing it at home. The lace that remains, beyond what the labor and capital employed on the cloth might have fabricated at home, is the amount of the advantage which England derives from the exchange.

Torrens recognized that the advantage from trade is realized on the imports side and also that the evaluation involves a comparison between the (counterfactual) amounts of English labor and capital that would have been necessary to

produce the import (in this case, lace) in England with the (actual) amounts of English labor and capital embodied in the production of the export (in this case, woolen broad cloth). Robbins (1958, p. 22) argues that in his definitive articulation of the 18th century rule about when trade was beneficial, Torrens had presented "one-half of the principle of comparative [advantage]." But as Robbins notes, "it was only one-half," since Torrens and the 18th century rule were silent on the role cross-national differences in productivity played in trade and the gains from it.

## Ricardo's Labor Value Formulation of Comparative Advantage

It was left to David Ricardo to enlist the 18th century rule in the effort to develop a full statement of the sources of gains from trade and comparative advantage. One task of Ricardo's *Principles* was to elucidate the two additional steps required to achieve this goal: a theory of value and an explicit recognition of the distinction between domestic and international terms of trade.

In an insightful paper, Ruffin (2002) offers a narrative account of Ricardo's discovery of comparative advantage and brings attention to Sraffa's neglected interpretation of Ricardo's numbers as labor embodied in trade (Sraffa and Einaudi 1930).[3] Ruffin (2002, p. 736) reports on Ricardo's correspondence from the period when he was writing the *Principles*. In February 1816, Ricardo wrote "If I could overcome the obstacles in the way of giving a clear insight into the origin and law of relative or exchangeable value I should have gained half the battle." In the opening paragraph of his chapter seven on foreign trade, Ricardo recognized that the actual valuation of "foreign goods is measured by the quantity of the produce of our land and labor, which is given in exchange for them" (Ricardo 1817 [1966], p. 128). However, abstracting from other factors of production and focusing on labor alone offered him insight into the origin and law of relative and exchangeable values.[4] It can be argued that employing labor units alone as the sole metric for the relative valuations of goods opened the door to his comparative advantage statement.

As with the passage from Torrens (1808) quoted above, Ricardo's famous passage starts with an existing trading relationship between England and Portugal: "The quantity of wine which she (Portugal) shall give in exchange for the cloth of England is not determined by the respective quantities of labor devoted to the production of each, as

[3] Ruffin's rediscovery of Sraffa's insights has launched a reread of Ricardo among historians of economic thought. Recent writers like Maneschi (2004) and Faccarello (2015) are now in agreement that the textbook version of the one-factor fixed labor coefficient model is based on John Stuart Mill's (1852) misreading of Ricardo via his father James Mill (1821).

[4] In his seminal article, Stigler (1958, p. 361) points out that Ricardo's labor theory of value was based on empirical reasoning: "Ricardo believed that the changes brought about in the relative values of commodities by fluctuations in wages and profits were very small relative to those brought about by fluctuations in the quantity of labor (direct and indirect)."

it would be if both commodities were manufactured in England, or both in Portugal" (Ricardo 1817 [1966], pp. 134–35). Unlike Torrens, Ricardo's labor theory of value permitted him to employ labor units as a metric for evaluating the "real costs" in the context of the 18th century rule of the advantage of trade.

Ricardo next considers a configuration where the English cloth exported in exchange for the imported wine required the labor of 100 men; England in turn would have required the labor of 120 men to produce an equivalent amount of wine. Because the latter number is larger than the former, Ricardo applied the 18th century rule to conclude that England would "find it in her interest to import wine, and purchase it by the exportation of cloth."

After having established the advantage of trade for England, Ricardo follows up with a parallel 18th century rule statement for Portugal. He chooses numbers that indicate higher labor productivity in Portugal for both goods. The export of wine would require only the labor of 80 men; the cloth purchased with the exported wine would have required the labor of 90 men. Ricardo (1817 [1966], pp. 134–35) notes that Portugal's relative productivity in both goods poses a puzzle not resolved by simply invoking the 18th century rule: "This exchange might even take place, notwithstanding that the commodity imported by Portugal could be produced there with less labor than in England." Ricardo (p. 133) resolves the puzzle by asserting that "the same rule which regulates the relative value of commodities in one country, does not regulate the relative value of the commodities exchanged between two or more countries." His labor theory of value permits him to demonstrate the difference in the rules that govern international exchange: "[Within England] the labour of 100 Englishmen cannot be given for that of 80 Englishmen, but the produce of the labour of 100 Englishmen may be given for the produce of the labour of 80 Portuguese, 60 Russians or 120 East Indians" (p. 135).

By noticing that in domestic exchange, the labor of an Englishman in domestic cloth production is always given one for one for the labor of an Englishman in any other branch of production, Ricardo's formulation anticipates the ultimate source of the gains from trade in the neoclassical framework: the difference between the domestic and international terms of trade. In this formulation, the magnitude of the gains from trade stems from a benefit minus cost calculation: the benefits are on the import side and the costs are on the export side. Although Ricardo only considers the case of two goods and two countries, the logic extends in a straightforward manner to higher dimensions. In the many goods case, the labor needed to produce all import goods is subtracted from the labor needed to produce all export goods. Because the workers could be of different skill types, the formulation is not restricted to a single factor.

Table 1 summarizes Ricardo's logic. It also illustrates the separation property: that is, gains from trade can be calculated for each country separately. In addition, Ricardo's formulation does not require data on the technologies used by the trading partner. The amount of domestic labor actually embodied in a country's exports and imports—as a counterfactual—contains all the relevant information about the gains from trade for a given country.

*Table 1*
**Ricardo's Labor Value Formulation of Comparative Advantage**

| | *Exports (workers, actual)* | *Imports (workers, counterfactual)* | *Gains from trade (in workers)* |
|---|---|---|---|
| *England* | 100 Englishmen | 120 Englishmen | 20 Englishmen |
| *Portugal* | 80 Portuguese | 90 Portuguese | 10 Portuguese |

*Source:* Numbers are from Ricardo (1817 [1966], p. 135)
*Note:* Ricardo compares the number of a country's workers embodied in its exports with the counterfactual workers that would have been needed to produce its imports and implies the 18th century rule to determine the gains (in terms of workers) for each country separately.

## Ricardo meets Haberler

Because Ricardo's formulation of comparative advantage was so tied to his labor theory of value, it lost its prominence later in the 19th century and into the 20th century in tandem with the loss of prominence of the labor theory of value. In a path-breaking paper that reformulated comparative advantage in terms of opportunity costs, Haberler (1930) "revolutionized the theory of international trade… [and] laid the foundation for Ohlin's theory, as well as … Samuelson's" (Chipman 2008, p. 812). Although Haberler's formulation encountered some initial resistance, his diagrammatic representation became a major tool for the neoclassical synthesis between the classical emphasis on production costs and the utilitarian emphasis on consumer utility as a source of relative prices (or value).[5] Over time, Haberler's diagram established itself as what Robert Baldwin (1982, p. 142) has labelled the trade economist's "sacred diagram." It became a main textbook diagram for depicting the logic of comparative advantage and the corresponding gains from trade. This opportunity cost formulation of comparative advantage has allowed trade economists (and textbook writers) to view the Ricardian model and the Heckscher–Ohlin model as special cases pointing to alternative sources of comparative advantage: differences in relative productivities or endowments.

Figure 1 depicts Haberler's bowed-out production possibilities frontier diagram, with cloth on the horizontal and wine on the vertical axis, which serves as a useful tool for putting Ricardo's logic into a modern context. Because the standard textbook treatment of Ricardo is based on the assumption that Ricardo's numbers pertain to fixed unit labor requirements, the production possibility frontier of the canonical Ricardian model would have a linear slope, with the empirically stark prediction that international trade completely wipes out the comparative disadvantage sector

[5]Haberler's original paper was published in 1930 in German in *Weltwirtschafliches Archiv* and was not translated into English until 1985. Haberler's formulation reached the English-speaking audience in his 1936 English translation of his 1933 textbook on international trade. See Viner (1937, pp. 516–526) for his critical comparison between the opportunity cost approach and his favored real cost approach.

*Figure 1*
**Ricardo's Gains-from-Trade Formulation in the Modern Production Possibilities Frontier Diagram**

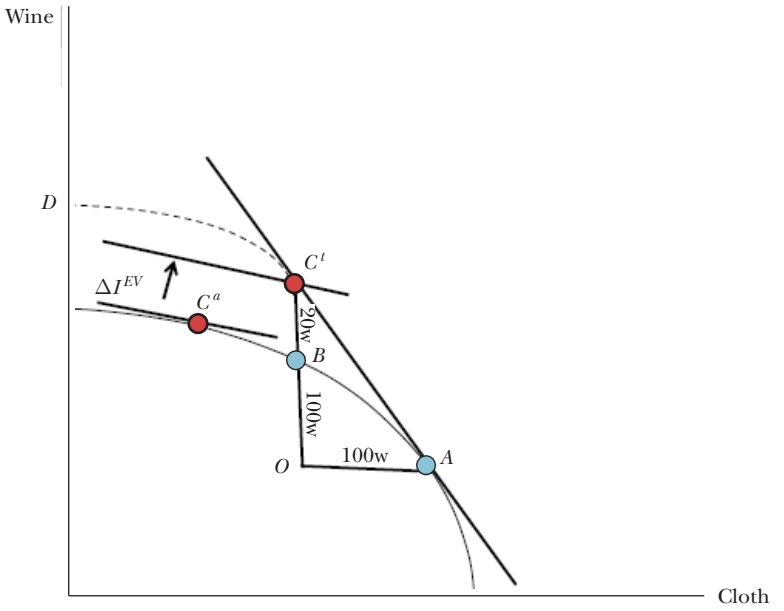


*Source:* Authors.
*Note:* Ricardo's gains-from-trade logic is depicted in the standard textbook production possibilities frontier. Starting from the trade equilibrium with production at point $A$ and consumption at point $C^t$, Ricardo compared a counterfactual domestic exchange of 100 workers along the segment $AB$ on the production possibilities frontier with the international exchange of 100 workers embodied in exports $OA$ for 120 workers embodied in imports $OC^t$. In Ricardo's formulation, trade is equivalent to a labor augmentation of 20 workers, depicted by the outward shift of the production possibilities frontier to $DC^t$. In Samuelson's formulation, autarky consumption $C^a$ is the starting point and trade is equivalent to an income increase of $\Delta I^{EV}$ to the economy's representative consumer allowing that consumer to afford the trade consumption point $C^t$ at autarky prices.

for at least one country. When we recognize that Ricardo's numbers pertain to units of labor embodied in a country's imports and exports, Ricardo's logic is compatible with a bowed-out production possibilities frontier, diminishing returns to labor, and incomplete specialization. Because of Ricardo's intuitive use of the separation property of comparative advantage, we can just focus on the trading equilibrium for England and compare his gains-from-trade formulation with the versions found in standard textbook presentations and in Torrens (1808).

The standard textbook formulation starts out with an autarky equilibrium and expresses the gains from trade as a comparison of (economy-wide) consumer utility or consumption between autarky and free trade. In Figure 1, England's autarky consumption point is depicted by point $C^a$ on its production possibilities frontier. Under autarky, the economy's consumption point must coincide with its

production point. The slope of the tangency at point $C^a$ measures the autarky price of cloth relative to wine. Assume now that when opening up to international trade, England faces a relative world price of cloth greater than the relative price of cloth in autarky. This provides an incentive for English resources (labor, capital, or land) to move into the production of cloth, which will drive up the relative price of cloth in England until it coincides with the world price of cloth at the free trade production point *A*. Exporting English cloth for imports of foreign wine will enable the English economy to reach a free trade consumption $C^t$ somewhere on the terms-of-trade line, but outside its autarky-based production possibilities frontier. Thinking about the line through $C^t$ as an income line under autarky prices (and invoking the assumption of a single consumer), the gains from trade can be expressed as the increase in income (denoted by $\Delta I^{EV}$) that must be given to English consumers in order to achieve the free trade consumption point under autarky prices. An attractive feature of this gains-from-trade formulation is that it holds under very general settings.[6] A major limitation is that an empirical assessment of this formulation requires data on autarky prices for a market economy.[7]

Unlike the textbook formulation, the starting point for both Torrens and Ricardo is a statement about the quantities of England's exports and imports, which are represented in Figure 1 by the trading vector $AC^t$. Their application of the 18th century rule to find the gains from trade means that both compare the resources embodied in the economy's exports with the domestic resources that would have been necessary to produce the economy's imports. In Ricardo's formulation, the exports of cloth ($OA$) correspond to a resource cost of 100 English workers (labeled 100w on the figure), while the imports of wine ($OC^t$) would have required 120 English workers (the labels 20w and 100w) if they were produced in England. In Ricardo's formulation, the issue of the gains from trade can be separated into two questions: 1) Are there gains from trade? And 2) if gains exist, how large are they? He addresses the existence question by noting that in domestic exchange, 100 workers embodied in cloth production must also be exchanged for 100 workers embodied in wine production. In terms of Figure 1, if the 100 English workers engaged in the production of cloth exports were domestically reallocated to domestic wine production, they could have produced only $OB$ units of wine. The movement between points *A* and *B* along an economy's production possibilities frontier (PPF) depicts this "no trade" scenario. Turning to the case of international trade, 120 workers is

[6] Paul Samuelson (1939), who was Haberler's student, formally proved the existence of the gains from trade. Although the production possibilities diagram is restricted to two goods, Samuelson showed that the underlying gains-from-trade logic holds for any number of goods or factors. Samuelson's gains-from-trade formulation assumes a representative consumer who follows the weak axiom of revealed preference; he introduced the axiom in his foundational consumer theory paper (Samuelson 1938).

[7] In Bernhofen and Brown (2005), we exploit data from Japan's 19th-century opening up from a market-based autarky economy to a trading economy as a natural experiment compatible with the theory's ceteris paribus assumptions. Using detailed autarky prices, we provide an upper bound on $\Delta I^{EV}$ that was the equivalent of about 6–9 percent of Japan's GDP at the time.

greater than 100 workers, so the gains from trade are positive. The gain of $BC^t$ units of wine corresponds to a gain of 20 English workers.[8]

We can embed Ricardo's formulation in the standard general equilibrium framework by recognizing that these 20 workers are the additional counterfactual workers that would be needed to attain the free trade consumption point $C^t$ through domestic production capabilities. In this formulation, the "compensation" measure that formally captures the gains from trade is in terms of the augmentation of the economy's labor endowment. Trade is equivalent to the economy gaining 20 fictitious workers who would enable the economy to produce the free trade consumption point $C^t$ through domestic production. These 20 workers can be thought of as extending the economy's production possibilities frontier in Figure 1 along the segment $DC^t$.

Ricardo's gains-from-trade formulation has several attractive features. First, the gains from trade can be illustrated without knowledge of the economy's autarky equilibrium consumption $C^a$ or autarky prices. The exchange of factor services embodied in the economy's actual trading vector is—to use modern terminology—a sufficient statistic for identifying the existence and the magnitude of the gains from trade.[9] Second, the gains from trade are formulated in terms of a cost–benefit comparison. The costs of trade are the workers lost to the production of exports; the gross benefits from trade are the domestic workers that would have been needed to produce the imports; and the net benefits (the gains from trade) are the difference between the gross benefits and the costs. Third, Ricardo's logic does not require any restrictive assumptions on the consumption side of the economy (like a representative consumer).

Finally, Ricardo's benefit–cost characterization of the gains can be helpful for illustrating the role trade costs play in the existence and magnitude of trade. For tractability, assume that each country incurs trade costs by using its own workers to ship its export good abroad. For example, if 25 English workers are needed to export the English cloth in Ricardo's example, trade is not expected to take place even if the relative autarky price of cloth is lower in England than in the rest of the world. Interestingly, under these circumstances a uniform increase in labor productivity across all activities (cloth manufacture, wine production, and shipping) will not make trade beneficial. For trade to occur would require an increase in labor productivity in shipping *relative* to other domestic production activities.[10]

[8] In Torrens's (1808) formulation of the 18th century rule, the gains are formulated in the units of imports, which would be $BC^t$ units of imported lace. Although the formulations of Ricardo and Torrens are isomorphic in the case of two goods, Ricardo's formulation has the considerable advantage that it can be generalized to multiple goods and factors.

[9] Fundamentally, the advantage of trade can be thought of as a refutable proposition. For example, if the domestic labor content of imports were revealed (by the data) to be lower than the domestic labor content of exports, the gains would be negative. Assuming that the "invisible hand" would not permit this to happen in a completely free market, this outcome could still occur from export subsidies distorting the law of comparative advantage.

[10] Krugman (2010) uses this line of reasoning to speculate that the productivity gains associated with the steam engine disproportionately influenced trade costs relative to costs of production and thus played a key role in ushering in the dramatic growth of world trade between 1870 and 1913. Given that piracy

## A Gains-from-Trade Formula Based on Ricardo's Numbers

Ricardo's framework suggests an algebraic gains-from-trade expression. The expression can be viewed as a special case of the influential gains-from-trade formula developed in Arkolakis, Costinot, and Rodríguez-Clare (2012), which we will refer to as the ACR formula. In their recent discussion in this journal, Costinot and Rodríguez-Clare (2018) provide a unifying perspective of the ACR formula by considering a trade equilibrium that is characterized by an international exchange of factor services.[11] Their approach is reminiscent of Ricardo's view of international trade as an exchange of labor services.

To see this, assume that England and Portugal are operating in a trading equilibrium where the terms of trade are normalized to be one, such that one bottle of wine is exchanged for one unit of cloth. The gains from trade for each economy can then be formulated as:

$$\text{Gains from trade} = \frac{IM(L_{IM} - L_{EX})}{\bar{L}}.$$

The numerator in this expression gives the factor augmentation that would be equivalent to the suspension of international trade and it depends on two factors: the quantity of imports $IM$ multiplied by the difference between the (counterfactual) average number of domestic workers needed to produce one unit of imports, $L_{IM}$, and the average number of domestic workers needed to produce one unit of exports, $L_{EX}$.[12] In Ricardo's example, $120 = IM \times L_{IM}$, $100 = IM \times L_{EX}$, and the labor augmentation equivalent of trade is 20. In order to evaluate whether 20 is a big or small number, the labor augmentation needs to be divided by the economy's labor force $\bar{L}$. This gains- from-trade formula measures the augmentation as a percentage of the economy's labor force.

Costinot and Rodríguez-Clare (in this journal, 2018) argue that measuring the gains from trade in the ACR formula involves two questions: 1) How large are imports of factor services in the trade equilibrium? And 2) How elastic is the relative demand for imported factor services in the counterfactual move from trade to autarky? In an analogous manner in Ricardo's example, the gains from trade for England would be larger, or a counterfactual move to autarky would cost more, the higher the per capita consumption of foreign wine, $IM/L$. The difference $(L_{IM} - L_{EX})$ can be interpreted as capturing the degree of substitutability between

raises trade costs by requiring manpower to protect cargo, Ricardo's gains-from-trade characterization helps to illustrate how improvements in the rule of law on the high seas can stimulate trade. See also North (1968), who found a significant role for the decline in piracy in productivity improvements of ocean shipping prior to 1850.

[11] The discussion of Costinot and Rodriguez-Clare (2018) in this journal is based on the formal model and estimation procedure developed in Adao, Costinot, and Donaldson (2017).

[12] In order to guarantee that each country will produce both goods in a trading equilibrium, one needs to assume diminishing marginal products for labor which implies that $L_{IM}$ and $L_{EX}$ are not constant and will depend on $IM$.

foreign and domestic labor in the transition from trade to autarky, where domestic workers now produce the goods that would have been imported. A higher value of $L_{IM} - L_{EX}$ corresponds to a lower degree of substitutability between foreign and domestic workers and implies larger gains from trade. In a Ricardian world, the larger the relative productivity differences of workers, the lower will be their degree of substitutability.

## A Final Note on Ricardo's Genius

Ricardo's breakthrough formulation of comparative advantage and the gains associated with it stem from his insight that if countries ship goods across borders, it is as if their workers move across borders. Hence, Ricardo's reasoning anticipated the general idea of the factor content of trade, which has proven itself to be a useful analytical tool for generalizing the Heckscher–Ohlin model, among others. In fact, Deardorff and Staiger (1988) provide a formal proof that the reasoning behind a Haberler-style Figure 1 is a special case of a general equilibrium analysis of a neoclassical economy with an arbitrary number of goods and factors (like labor, capital, and land). Starting from a trade equilibrium, they show that if trade were suspended, but if the economy were given the labor, capital, and land embodied in its imports minus the labor, capital, and land embodied in its exports, then the economy would be able to obtain the same consumption level as through international trade.[13] In this way, our discussion provides further evidence for the continuity of economic thought from the classical economists onward and the decisive break they represented with mercantilist thinking. As Haberler (1977, p. 1) observed 160 years after the publication of the *Principles*: "[T]here is an unmistakable family likeness between the modern theories … and the early classical theories, just as there is between a modern jumbo jet and the Wright brothers' contraption."

■ *We thank Gordon Hanson and Timothy Taylor for insightful comments and suggestions.*

[13] Deardorff and Staiger (1988) use the term "equivalent autarky equilibrium" for the situation when trade is suspended but the economy is compensated through the factor augmentation. However, they do not link the factor content of trade to the aggregate gains from trade. Extending Ricardo's gains-from-trade logic to a multifactor framework and employing detailed data on 19th-century Japan, in Bernhofen and Brown (2012) we calculate the factor augmentation equivalents of trade in the case of five factors: female labor, skilled male labor, unskilled male labor, capital, and arable land.

## References

**Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson.** 2017. "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade." *American Economic Review* 107(3): 633–89.

**Aldrich, John.** 2004. "The Discovery of Comparative Advantage." *Journal of the History of Economic Thought* 26(3): 379–99.

**Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare.** 2012. "New Trade Models, Same Old Gains?" *American Economic Review* 102(1): 94–130.

**Baldwin, Robert E.** 1982. "Gottfried Haberler's Contributions to International Trade Theory and Policy." *Quarterly Journal of Economics* 97(1): 141–48.

**Bernhofen, Daniel M., and John C. Brown.** 2005. "An Empirical Assessment of the Comparative Advantage Gains from Trade: Evidence from Japan." *American Economic Review* 95(1): 208–25.

**Bernhofen, Daniel M., and John C. Brown.** 2012. "A Factor Augmentation Formulation of the Value of International Trade." CESifo Working Paper 3860.

**Chipman, John S.** 2008. "Haberler, Gottfried." *The New Palgrave Dictionary of Economics*, vol. 2, edited by Steven Durlauf and Lawrence E. Blume, 812–14. Macmillan.

**Costinot, Arnaud, and Andrés Rodríguez-Clare.** 2018. "The US Gains from Trade: Valuation Using the Demand for Foreign Factor Services." *Journal of Economic Perspectives* 32(2): 3–24.

**Deardorff, Alan V., and Robert W. Staiger.** 1988. "An Interpretation of the Factor Content of Trade." *Journal of International Economics* 24(1–2): 93–107.

**Eaton, Jonathan, and Samuel Kortum.** 2002. "Technology, Geography, and Trade." *Econometrica* 70(5): 1741–1779.

**Eaton, Jonathan, and Samuel Kortum.** 2012. "Putting Ricardo to Work." *Journal of Economic Perspectives* 26(2): 65–89.

**Faccarello, Gilbert.** 2015. "A Calm Investigation into Mr Ricardo's Principles of International Trade." *European Journal of the History of Economic Thought* 22(5): 754–90.

**Grampp, William D.** 1952. "The Liberal Elements in English Mercantilism." *Quarterly Journal of Economics* 66(4): 465–501.

**Haberler, Gottfried.** 1930. "Die Theorie der komparativen Kosten und ihre Auswertung für die Begründung des Freihandels." *Weltwirtschaftliches Archiv* 32: 349–70. (Reprinted 1985 in *Selected Essays of Gottfried Haberler,* edited by Anthony Y. C. Koo, 3–19. Cambridge, MA: MIT Press.)

**Haberler, Gottfried.** 1977. "Survey of Circumstances Affecting the Location of Production and International Trade as Analyzed in the Theoretical Literature." In *The International Allocation of Economic Activity,* Proceedings of a Nobel Symposium held at Stockholm, edited by Bertil Ohlin, Per-Ove Hesselborn, and Per Magnus Wijkman, 1–24. Springer**.**

**Heckscher, Eli F.** 1935. *Mercantilism.* 2 vols. London: G. Allen & Unwin.

**Irwin, Douglas A.** 1996. *Against the Tide: An Intellectual History of Free Trade.* Princeton.

**Krugman, Paul.** 2010. "A Globalization Puzzle." The Conscience of a Liberal, *New York Times*, February 21, 2010.

**Maneschi, Andrea.** 2004. "The True Meaning of David Ricardo's Four Magic Numbers." *Journal of International Economics* 62(2): 433–43.

**Martyn, Henry.** 1701. *Considerations on the East India Trade.* London, J. Roberts.

**Mill, James.** 1808. *Commerce Defended: An Answer to the Arguments by which Mr. Spence, Mr. Corbett, and Others, Have Attempted to Prove that Commerce is Not a Source of National Wealth.* London C. and R. Baldwin.

**Mill, James.** 1821. *Elements of Political Economy.* 1st edition. London: Baldwin, Cradock, and Joy.

**Mill, John Stuart.** 1852. *Principles of Political Economy with Some of their Applications to Social Philosophy.* 3rd edition. London. J. W. Parker.

**North, Douglass C.** 1968. "Sources of Productivity Change in Ocean Shipping, 1600–1850." *Journal of Political Economy* 76(5): 953–70.

**Ricardo, David.** 1817 [1966]. *On the Principles of Political Economy and Taxation.* London: John Murray. (Reprinted 1966 in *The Works and Correspondence of David Ricardo, Vol 1: Principles of Political Economy and Taxation,* edited by P. Sraffa. London: Cambridge University Press.)

**Robbins, Lionel.** 1958. *Robert Torrens and the Evolution of Classical Economics.* London and New York: Macmillan.

**Ruffin, Roy J.** 2002. "David Ricardo's Discovery of Comparative Advantage." *History of Political Economy* 34(4): 727–48.

**Samuelson, Paul A.** 1938. "A Note on the Pure Theory of Consumer's Behavior." *Economica* 5(17): 61–71.

**Samuelson, Paul A.** 1939. "The Gains from International Trade." *Canadian Journal of Economics and Political Science* 5(2): 195–205.

**Samuelson, Paul A.** 1948. *Economics: An Introductory Analysis.* New York, McGraw-Hill.

**Samuelson, Paul A.** 1969. "Presidential Address: The Way of an Economist." *International Economic*

*Relations: Proceedings of the Third Congress of the International Economic Association*, 1–11. London: Palgrave Macmillan.

**Smith, Adam.** 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations*. London and New York: W. Strahan and T. Cadell.

**Sraffa, Piero, and L. Einaudi.** 1930. "An Alleged Correction of Ricardo." *Quarterly Journal of Economics* 44(3): 539–545.

**Stigler, George J.** 1958. "Ricardo and the 93% Labor Theory of Value." *American Economic Review* 48(3): 357–67.

**Torrens, Robert.** 1808. *The Economists Refuted; Or, An Inquiry into the Nature and Extent of the Advantages Derived from Trade*. London: S.A. Ody.

**Tucker, Josiah.** 1750. *A Brief Essay on the Advantages and Disadvantages, which Respectively Attend France and Great Britain, with Regard to Trade: With Some Proposals for Removing the Principal Disadvantages of Great Britain*. London: Trye.

**Viner, Jacob.** 1937. *Studies in the Theory of International Trade*. New York and London: Harper & Brothers.

**Wilson, Charles.** 1949. "Treasure and Trade Balances: The Mercantilist Problem." *Economic History Review* 2(2): 152–61.

*Journal of Economic Perspectives—Volume 32, Number 4—Fall 2018—Pages 241–248*

# Recommendations for Further Reading

Timothy Taylor

This section will list readings that may be especially useful to teachers of undergraduate economics, as well as other articles that are of broader cultural interest. In general, with occasional exceptions, the articles chosen will be expository or integrative and not focus on original research. If you write or read an appropriate article, please send a copy of the article (and possibly a few sentences describing it) to Timothy Taylor, preferably by email at *taylort@macalester.edu*, or c/o *Journal of Economic Perspectives*, Macalester College, 1600 Grand Ave., St. Paul, MN 55105.

## Potpourri

Susan Lund, Jonathan Woetzel, Eckart Windhagen, Richard Dobbs, and Diana Goldshtein ask "Rising Corporate Debt: Peril or Promise?" "In a departure from the past, most of the growth in corporate debt has come from developing countries, in particular China. Companies in advanced economies accounted for just 34 percent or $9.9 trillion of the growth in global corporate debt since 2007, while developing countries accounted for 66 percent or $19.2 trillion. Since 2007, China's corporate debt has increased by $15 trillion, or more than half of global corporate debt growth. As a share of GDP, China's corporate debt rose from 97 percent of GDP in 2007 to 163 percent in 2017 … The growth in corporate debt

■ *Timothy Taylor is Managing Editor,* Journal of Economic Perspectives*, based at Macalester College, Saint Paul, Minnesota. He blogs at http://conversableeconomist.blogspot.com.*

† For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at
https://doi.org/10.1257/jep.32.4.241

doi=10.1257/jep.32.4.241

in China is mainly associated with a construction sector that increased its leverage as the housing market boomed. Today, 30 to 35 percent of corporate debt in China is associated with construction and real estate. … A relatively new feature of the debt landscape in recent years has been a shift in corporate borrowing from loans to bonds. Given the growing pressure on banks to meet new capital and liquidity standards, global nonfinancial corporate loans outstanding have been growing by only 3 percent annually on average since 2007 to stand at around $55 trillion in 2017. However, the share of global corporate debt in the form of bonds has nearly doubled, and the value of corporate bonds outstanding has grown 2.7 times since 2007. This is a positive trend, leading to a diversification of corporate financing. However, we also find risks." McKinsey Global Institute, June 2018, at https://www.mckinsey.com/business-functions/strategy-and-corporate-finance/our-insights/rising-corporate-debt-peril-or-promise.

The Joint Center for Housing Studies of Harvard University has published its 30th annual report: *The State of the Nation's Housing 2018*. "In 2017, the supply of for-sale homes averaged only 3.9 months—well below the 6 months considered a balanced market. … Lower-cost homes are especially scarce. … Why inventories are so tight is not entirely clear. CoreLogic data show that the number of owners underwater on their mortgages shrank from more than 12.1 million in 2011 to 2.5 million in 2017, so negative equity should no longer be a significant drag on sales. Still, conversion of 3.9 million single-family homes to rentals in 2006–2016 could be constraining the number of entry-level homes on the market. … Another factor is the low level of single-family construction. Despite six consecutive years of increases, single-family starts stood at just 849,000 units in 2017, well below the long-run annual average of 1.1 million. … The slow growth in single-family construction reflects in part homebuilder caution following the dramatic housing bust. But risk aversion aside, a significant constraint on new residential construction may be the dwindling supply of buildable lots. According to Metrostudy data, the inventory of vacant lots in the 98 metro areas tracked fell 36 percent in 2008–2017. Indeed, 21 of the nation's 25 largest metros reported inventories that would support less than 24 months of residential construction. Along with limited land, respondents to builder surveys cite rising input costs as adding to the difficulty of constructing entry-level homes. As a result, the share of smaller homes (under 1,800 square feet) built each year fell from 50 percent in 1988 to 36 percent in 2000 to 22 percent in 2017." June 19, 2018, http://www.jchs.harvard.edu/sites/default/files/Harvard_JCHS_State_of_the_Nations_Housing_2018.pdf.

The Bank for International Settlements devotes Chapter V of its *Annual Report 2017–18* to the topic "Cryptocurrencies: Looking Beyond the Hype." "[T]he essence of good money has always been trust in the stability of its value. And for money to live up to its signature property—to act as a coordination device facilitating transactions—it needs to efficiently scale with the economy and be provided elastically to address fluctuating demand. … The chapter then gives an introduction to cryptocurrencies and discusses the economic limitations inherent in the decentralised creation of trust which they entail. For the trust to be maintained, honest network participants

need to control the vast majority of computing power, each and every user needs to verify the history of transactions and the supply of the cryptocurrency needs to be predetermined by its protocol. Trust can evaporate at any time because of the fragility of the decentralised consensus through which transactions are recorded. Not only does this call into question the finality of individual payments, it also means that a cryptocurrency can simply stop functioning, resulting in a complete loss of value. Moreover, even if trust can be maintained, cryptocurrency technology comes with poor efficiency and vast energy use. Cryptocurrencies cannot scale with transaction demand, are prone to congestion and greatly fluctuate in value. Overall, the decentralised technology of cryptocurrencies, however sophisticated, is a poor substitute for the solid institutional backing of money. That said, the underlying technology could have promise in other applications, such as the simplification of administrative processes in the settlement of financial transactions. Still, this remains to be tested." June 24, 2018. https://www.bis.org/publ/arpdf/ar2018e5.htm.

John L. Mikesell and Sharon N. Kioko provide an overview of "The Retail Sales Tax in a New Economy." "The American retail sales tax emerged from a desperation experiment in Mississippi in the midst of the Great Depression. Revenue from the property tax, the largest single source of state tax revenue at the time, collapsed … Mississippi (followed by West Virginia) showed that retail sales taxes could produce immediate cash collections, even in low-income jurisdictions. … By 1938, twenty-two states (plus Hawaii, not yet a state) were collecting the tax; six others had also imposed the tax for a short time but had let them expire. … The national total retail sales tax collections exceeded the collections from every other state tax from 1947 through 2001. It was also the largest tax producer in 2003 and 2004 also (years in which individual income tax revenue was still impacted by the 2001 recession), but it was surpassed by state individual income tax revenues in other years since 2001. … [T]he history of mean retail sales tax breadth (implicit tax base/state personal income) across the states from 1970 to 2016 … is one of almost constant decline, from 49.0 percent in 1970 to 37.3 percent in 2016. … The typical state retail sales tax base has narrowed as a share of the economy of the state over the years and this has meant that, in order for states to maintain the place of their sales tax in their revenue systems, they have been required to gradually increase the statutory tax rate they apply to that base. … [L]ittle good can be said about a narrow base/high statutory rate revenue policy. … Unfortunately, many states got off to a bad start when they initially adopted their sales taxes and excluded all or almost all household service purchases from the tax base and it has proven to be difficult to correct that initial error." Brookings Institution, 7th Annual Municipal Finance Conference, July 16–17, 2018, https://www.brookings.edu/wp-content/uploads/2018/04/Mikesell-Kioko1.pdf; video of paper presentation, with comments and discussion, at https://www.youtube.com/watch?time_continue=452&v=mccFFLQ_Ydg.

The International Energy Agency has published "The Future of Cooling: Opportunities for Energy-Efficient Air Conditioning," "The world faces a looming 'cold crunch.' Using air conditioners and electric fans to stay cool accounts for

nearly 20% of the total electricity used in buildings around the world today. And this trend is set to grow as the world's economic and demographic growth becomes more focused in hotter countries. … Wider access to cooling is necessary, bringing benefits to human development, health, well-being and economic productivity. … If left unchecked, energy demand from air conditioners will more than triple by 2050, equal to China's electricity demand today. … The answer lies first and foremost in improving the efficiency of air conditioners … Measures such as strong minimum energy performance standards are well known and well proven to drive up equipment efficiencies quickly and cost-effectively. In the longer term, the underlying need for cooling can also be greatly reduced by better building design and tougher building codes, as well as by increased rates of energy efficiency improvements in existing buildings. … Household ownership of ACs varies enormously across countries, from around 4% in India and less than 10% in Europe, to over 90% in the United States and Japan, and close to 100% in a few Middle Eastern countries. In China, nearly 60% of households now have at least one AC … " May 2018 (free registration may be needed to access report), http://www.iea.org/publications/freepublications/publication/The_Future_of_Cooling.pdf.

## Symposia

The *Russell Sage Foundation Journal of the Social Sciences* has published a 10-paper symposium from a range of social scientists concerning "The Fiftieth Anniversary of the Kerner Commission Report." From the introductory essay by Susan T. Gooden and Samuel L. Myers Jr., "The Kerner Commission Report Fifty Years Later: Revisiting the American Dream" (pp. 1–17): "The Kerner report was the final report of a commission appointed by the U.S. President Lyndon B. Johnson on July 28, 1967, as a response to preceding and ongoing racial riots across many urban cities, including Los Angeles, Chicago, Detroit, and Newark. These riots largely took place in African American neighborhoods, then commonly called ghettos. On February 29, 1968, seven months after the commission was formed, it issued its final report. The report was an instant success, selling more than two million copies. … The Kerner report documents 164 civil disorders that occurred in 128 cities across the forty-eight continental states and the District of Columbia in 1967. Other reports indicate a total of 957 riots in 133 cities from 1963 until 1968, a particular explosion of violence following the assassination of King in April 1968 … President Johnson was enormously displeased with the report, which in his view grossly ignored his Great Society efforts. The report also received considerable backlash from many whites and conservatives for its identification of attitudes and racism of whites as a cause of the riots. 'So Johnson ignored the report. He refused to formally receive the publication in front of reporters. He didn't talk about the Kerner Commission report when asked by the media,' and he refused to sign thank-you letters for the commissioners." September 2018, https://www.rsfjournal.org/toc/rsf/4/6.

Two think-tanks, ThirdWay and the American Enterprise Institute, have published a set of five papers on the subject of "Elevating College Completion." Bridget Terry Long offers an overview in "The College Completion Landscape: Trends, Challenges, and Why it Matters": "The conventional way to measure graduation rates is to examine how many students complete a degree within 150 percent of the expected completion time—that is, six years for a bachelor's degree and three years for an associate degree. Using this metric, research suggests that about only half of students enrolled at four-year colleges and universities graduate within 150 percent of the expected completion time, and the completion rate is even lower for students enrolled at two-year colleges." Sarah Turner writes in "The Policy Imperative: Policy Tools Should Create Incentives for College Completion": "In 43 four-year public schools, the three-year cohort default rate is greater than the completion rate. This is also the case for 147 four-year private nonprofit schools and 98 for-profit schools. In other words, students in these schools who borrow face a greater likelihood of defaulting than completing a degree. It would seem, then, that college attendance at these schools leaves many students worse off—lacking a degree, defaulting on a student loan, or both." May 25, 2018, https://www.thirdway.org/series/elevating-college-completion.

The *World Happiness Report 2018*, edited by John F. Helliwell, Richard Layard, and Jeffrey D. Sachs, included seven chapters generally focused on happiness and migration. Here's a comment concerning the enormous internal migration in China from the overview chapter by Helliwell, Layard, and Sachs: "Over the years 1990–2015 the Chinese urban population has grown by 463 million, of whom roughly half are migrants from villages to towns and cities. By contrast, over the same period the increase in the number of international migrants in the entire world has been 90 million, less than half as many as rural to urban migrants in China alone. Thus internal migration is an order of magnitude larger than international migration. … Migrants [within China] have roughly doubled their work income by moving from the countryside, but they are less happy than the people still living in rural areas. … Could it be that many of the migrants suffer because of the remittances they send home? The evidence says, No. Could it be that the people who migrate were intrinsically less happy? The evidence says, No. Could it be that urban life is more insecure than life in the countryside—and involves fewer friends and more discrimination? Perhaps. The biggest factor affecting the happiness of [within China] migrants is a change of reference group: the happiness equation for migrants is similar to that of urban dwellers, and different from that of rural dwellers. This could explain why migrants say they are happier as a result of moving—they would no longer appreciate the simple pleasures of rural life." March 2018, https://s3.amazonaws.com/happiness-report/2018/WHR_web.pdf.

## Interview with Economists

Renee Haltom interviews Jesús Fernández-Villaverde. On the euro: "[I]f you ask me, 'Should I marry my friend X?' I may tell you, 'No, I don't think you are

compatible, you are going to end up divorced.' But that's a very different question from, 'Should I get a divorce now that we are married and have a mortgage, three kids in school, two cars, and a dog?' Like it or not, we got married to the Germans, and the Germans got married to the Spaniards. We need to make this work, because breaking up now would be way too costly. … There has to be a great bargain between those who point out the need for making financial and economic crises easier to go through and those who emphasize that, in the long run, rules are very important. That's the big question mark: Is the political process within Europe going to be able to deliver that solution?" On the state of macro: "If you take the best 20 macroeconomists of my generation, of course they don't agree on everything, but the things they talk about are very different from the type of things you will see on Twitter or the blogosphere. … Sometimes I see criticisms about the state of macro saying, 'Macroeconomists should do X,' and I'm thinking, 'Well, we have been doing X for 15 years.' … Many of the people who are currently very critical of macro are in another generation, and some of them may not be fully aware of where the frontier of research is right now. They also have plenty of free time, so it's much easier for them to write 20 pages of some type of exposé, if they want to use that word, on the state of macro." *Econ Focus*, published by the Federal Reserve Bank of Richmond, First Quarter 2018, pp. 22–27, at https://www.richmondfed.org/publications/research/econ_focus/2018/q1/interview.

Douglas Clement has an "Interview with Marianne Bertrand," subtitled "University of Chicago Economist on the Glass Ceiling, Implications of Growing Inequality and the Trouble with Boys." "We were talking about income inequality, and one of our colleagues said, basically, 'Well, at the end of the day, who cares? Yes, maybe we're growing apart economically, but on Sunday all we all do is watch TV. We are growing apart economically, but our lives may not be that different; they may, in fact, have converged.' … How much can we say about how the lives of the rich and the poor changed? … So we tried to assemble all the data sets we could; for example, time-use data, which go back to the 1960s. Another data set that a lot of social scientists use is the General Social Survey, which tells us something about views and opinions—views on abortion, gays, racial issues, government spending and the like. … [W]e had access to a marketing data set, which is truly remarkable. In that data set, we can see media consumption—what TV shows people watch, what movies they watch, what magazines they read. The data set also shows thousands of products that people may or may not buy, and thousands of brands that people may or may not buy or own. Then we built a metric of cultural distance between groups by income. . … The main headline result of the paper is that most of the trend lines are flat. Our ability to predict someone's income based on the consumption of particular goods and brands is essentially the same today as it was 25 years ago. There's no trend in our ability to predict people's income based on how they spend their time today, compared to close to 50 years ago. The only area where we see some slight evidence of divergence on income is with respect to social attitudes, where our ability to predict people's income based on what they think, their views, is slightly better today than it was in the early 1970s. … [N]ow we've done this exercise, as I said, for

race, gender and urbanicity. … We've also done it based on political attitudes, and the main result, which I just gave you for income—there's no big trend—essentially applies to, at a first-level of approximation, everything that we have looked at. The one really large exception quantitatively is our ability to predict whether someone is liberal or conservative/Democrat or Republican based on their social attitudes. That has been increasing over time. So liberals and conservatives haven't been diverging over time on TV consumption, brands or goods, but on social views they have been diverging a lot over time." *The Region*, Federal Reserve Bank of Minneapolis, June 19, 2018, https://www.minneapolisfed.org/publications/the-region/interview-with-marianne-bertrand.

The Institute for New Economic Thinking has posted video of a six-part series: "What Money Can't Buy." Most of the videos are a seminar-style discussion with philosopher Michael Sandel and 12 students, but there are also snippets from Greg Mankiw, Richard Posner, Joseph Stiglitz, Lawrence H. Summers, and others prominent economists. The lectures are: Episode 1: Sex Sells, But Should It? (Should We Be Able to Discriminate Based on Looks?); Episode 2: The Body Market (Should You Be Able to Sell Your Kidney?); Episode 3: The Walrus Quota (Should We Be Able to Sell Refugees?); Episode 4: Supply Shock (Should You Be Able to Sell Water In A Disaster?); Episode 5: The Golden Door (Should We Pay People to Vote?); Episode 6: The Death Pool (Should We Be Able to Profit Off of Death?). At https://www.ineteconomics.org/perspectives/videos/what-money-cant-buy.

## Discussion Starters

Amy L. Brooks, Shunli Wang, and Jenna R. Jambeck consider "The Chinese Import Ban and Its Impact on Global Plastic Waste Trade." "[U]pward of half of the plastic waste intended for recycling has been exported to hundreds of countries around the world. China, which has imported a cumulative 45% of plastic waste since 1992, recently implemented a new policy banning the importation of most plastic waste, begging the question of where the plastic waste will go now. We use commodity trade data for mass and value, region, and income level to illustrate that higher-income countries in the Organization for Economic Cooperation have been exporting plastic waste (70% in 2016) to lower-income countries in the East Asia and Pacific for decades. An estimated 111 million metric tons of plastic waste will be displaced with the new Chinese policy by 2030. As 89% of historical exports consist of polymer groups often used in single-use plastic food packaging (polyethylene, polypropylene, and polyethylene terephthalate), bold global ideas and actions for reducing quantities of nonrecyclable materials, redesigning products, and funding domestic plastic waste management are needed." *Science Advances*, June 20, 2018. http://advances.sciencemag.org/content/4/6/eaat0131/tab-pdf.

Peter J. Neumann and Joshua T. Cohen provide an overview of "QALYs in 2018—Advantages and Concerns." "A year in the hypothetical state of 'perfect health' is worth 1 QALY. Being deceased is worth 0 QALYs. Other health states

fall between these bounds, with less desirable states closer to 0. QALYs are useful because they combine mortality and morbidity into a single metric, reflect individual preferences, and can be used as a standard measure of health gains across diverse treatments and settings. … Typical value benchmarks in the United States have historically ranged from approximately $50 000 to, more recently, as high as approximately $150 000 per QALY. Those benchmarks purport to represent the 'value' of a QALY; ie, the 'willingness to pay' to gain 1 QALY of health. … The United Kingdom's National Institute for Health and Care Excellence, which is charged with assessing health technology value for that country's National Health Service, has used more stringent benchmarks. With a number of exceptions, favorable value has generally corresponded to cost-effectiveness ratios below £20 000 (about $28 000) per QALY, and unfavorable value has generally corresponded to ratios exceeding £30 000 (about $42 000) per QALY." *Journal of the American Medical Association,* June 26, 2018, vol. 319, no. 24, pp. 2473–74, at https://jamanetwork.com/journals/jama/fullarticle/2682917.

Colin Grabow, Inu Manak, and Daniel Ikenson discuss "The Jones Act: A Burden America Can No Longer Bear" "For nearly 100 years, a federal law known as the Jones Act has restricted water transportation of cargo between U.S. ports to ships that are U.S.-owned, U.S.-crewed, U.S.-registered, and U.S.-built. … While the law's most direct consequence is to raise transportation costs, which are passed down through supply chains and ultimately reflected in higher retail prices, it generates enormous collateral damage through excessive wear and tear on the country's infrastructure, time wasted in traffic congestion, and the accumulated health and environmental toll caused by unnecessary carbon emissions and hazardous material spills from trucks and trains. Meanwhile, closer scrutiny finds the law's national security justification to be unmoored from modern military and technological realities." Cato Institute Policy Analysis no. 845, June 28, 2018, https://object.cato.org/sites/cato.org/files/pubs/pdf/pa845.pdf.

# STUDY WITH PURPOSE

*"The 11-month **Master of Arts in International Economics and Finance** prepares early-career professionals to understand advanced economic theories, master professional quantitative and econometrics skills, and assess a wide range of international economic and financial scenarios. More than 90% of recent graduates are working in their chosen field within six months after graduation."*

— GORDON BODNAR, PhD
Director, International Economics Program and Morris W. Offit Professor of International Finance

SAIS JOHNS HOPKINS
SCHOOL *of* ADVANCED INTERNATIONAL STUDIES

**NOW ACCEPTING APPLICATIONS TO THE MASTER OF ARTS IN INTERNATIONAL ECONOMICS AND FINANCE PROGRAM**

**sais-jhu.edu/mief-aea**

**WASHINGTON, DC**

# Application/Renewal for Membership

**AMERICAN ECONOMIC ASSOCIATION**
2014 Broadway, Suite 305
Nashville, TN 37203
Ph. 615-322-2595 fax: 615-343-7590
Federal ID No. 36-2166945
www.vanderbilt.edu/AEA

| **RENEWING MEMBERS, ENTER ACCT. NUMBER & EXP. DATE** | **IF PAYING BY CREDIT CARD, PLEASE FILL OUT BELOW** |
|---|---|
| ACCOUNT NUMBER: | CARD NUMBER: |
| EXPIRATION DATE: | EXP DATE: CSC CODE: AMT. PAID: |

| | | | |
|---|---|---|---|
| FIRST NAME: | MI: | LAST NAME: | |
| ADDRESS: | | | |
| CITY: | | STATE/PROVINCE: | ZIP: |
| COUNTRY: | | ☐ Check here if non-US | |
| PHONE: | | FAX: | |
| PRIMARY FIELD OF SPECIALIZATION: | | | |
| SECONDARY FIELD OF SPECIALIZATION: | | | |
| EMAIL: | | ☐ Check here to exclude your email address from the public directory | |

Please include my email address to receive:
☐ Announcements about public policy affecting economists or the economics profession
☐ Surveys about economics or the economics profession

**MEMBERSHIP DUES** — Based on annual income. Please select one below.

| | | |
|---|---|---|
| ☐ Annual income of $70,000 or less | $20 | $ |
| ☐ Annual income of $70,000 to $105,000 | $30 | $ |
| ☐ Annual income over $105,000 | $40 | $ |

The AEA dues above include online access to all AEA journals and JSTOR archived issues.
For print or CD subscription(s) indicate preference below and add appropriate charge(s).

| **Journal** | | **Print** | **Int'l Postage*** | **CD*** | |
|---|---|---|---|---|---|
| AER | (12 issues, includes the AEA P&P) | ☐ Add $25 | ☐ Add $35 | ☐ Add $25 | $ |
| AEA Papers & Proceedings ONLY | **(1 May issue)** | ☐ Add $10 | n/a | n/a | $ |
| JEL | (4 quarterly issues) | ☐ Add $15 | ☐ Add $15 | ☐ Add $15 | $ |
| JEP | (4 quarterly issues) | ☐ Add $15 | ☐ Add $15 | ☐ Add $15 | $ |
| AEJ: Applied | (4 quarterly issues) | ☐ Add $15 | ☐ Add $15 | n/a | $ |
| AEJ: Policy | (4 quarterly issues) | ☐ Add $15 | ☐ Add $15 | n/a | $ |
| AEJ: Macro | (4 quarterly issues) | ☐ Add $15 | ☐ Add $15 | n/a | $ |
| AEJ: Micro | (4 quarterly issues) | ☐ Add $15 | ☐ Add $15 | n/a | $ |

* Int'l postage applies only to print journals mailed outside of the U.S. No additional postage is required for CDs or the AEA Papers and Proceedings.

| | | | | | |
|---|---|---|---|---|---|
| | | | | Sub Total | $ |
| Check One: | ☐ 1 Year | ☐ 2 Years | ☐ 3 Years | TOTAL AMOUNT | $ |

Make checks payable to: American Economic Association.
Must be drawn on a US bank.
Apply online at http://www.aeaweb.org/membership.php

Payments must be made in advance. We accept checks (in US dollars only, with correct coding for processing in US banks) and credit cards; online or by faxing or mailing the application. Please choose one method;
it is the Association's policy NOT TO REFUND dues.

New Journal!

# American Economic Review: INSIGHTS

**ABOUT THIS JOURNAL**

*AER: Insights* is designed to be a top-tier, general-interest economics journal publishing papers of the same quality and importance as those in the *AER*, but devoted to publishing papers with important insights that can be conveyed succinctly.

The aim of *AER: Insights* is to provide a high-quality outlet for important yet concise contributions to economics, both empirical and theoretical.

## CALL FOR PAPERS

### CONCISE PAPERS ⇄ QUICKER REVISION PROCESS

Turnaround time for first response is a maximum of three months. First responses will be limited to either Conditional Accept OR Reject with no lengthy responses to referees or second round comments. All requested revisions must be returned to the editor within eight weeks. **Length limits on submissions and revisions will be strictly enforced.**

Open Now! Submissions must be <=6,000 words and have a maximum of five exhibits (figures or tables).

New Journal!

**ANTICIPATED LAUNCH**
**2019**

AMERICAN ECONOMIC ASSOCIATION

**EDITOR**
**Amy Finkelstein**
Massachusetts Institute of Technology

**COEDITORS**
**Peter Klenow**
Stanford University

**Larry Samuelson**
Yale University

Full Submission Guidelines
**www.aeaweb.org/journals/aeri/submissions**

## *DON'T MISS...*

# CTREE 2019

## The Ninth Annual Conference on Teaching & Research in Economic Education

### *Plenary Speakers include:*

Bob Frank
Cornell University

Anne Winkler
University of Missouri, St Louis

Stephen Chew
Samford University

**May 29 – 31, 2019**
**St. Louis, Missouri**
**Hyatt Regency St. Louis at the Arch**

AMERICAN ECONOMIC ASSOCIATION

CTREE is hosted by the AEA Committee on Economic Education in conjunction with the *Journal of Economic Education*.

***For more info go to***
***www.aeaweb.org/ctree/2019***

# The American Economic Association

Correspondence relating to advertising, business matters, permission to quote, or change of address should be sent to the AEA business office: aeainfo@vanderbilt.edu. Street address: American Economic Association, 2014 Broadway, Suite 305, Nashville, TN 37203. For membership, subscriptions, or complimentary access to *JEP* articles, go to the AEA website: http://www.aeaweb.org. Annual dues for regular membership are $20.00, $30.00, or $40.00, depending on income; for an additional fee, you can receive this journal, or any of the Association's journals, in print. Change of address notice must be received at least six weeks prior to the publication month.



The following Statement of Ownership, Management and Circulation is provided in accordance with the requirements, as contained in 39 U.S.C. 3658. *Journal of Economic Perspectives* is owned, managed and published by the American Economic Association, a nonprofit educational organization, located at 2014 Broadway, Suite 305, Nashville, Davidson County, TN 37203-2425. The Editor is Professor Enrico Moretti. The Managing Editor is Timothy Taylor: *Journal of Economic Perspectives*, 2403 Sidney Street, Suite 260, Pittsburgh, PA 15203. The tax status of the American Economic Association has not changed during the preceding twelve months. During the preceding twelve months, the average number of copies printed for each issue was 5,972; the average total paid and/or requested circulation, 5,002; the average total non-requested distribution, 19; the average number of copies not distributed 951. Corresponding figures for May 2018, the issue published nearest to filing date: total number of copies printed, 5,943; total paid and/or requested circulation, 4,826; number of copies not distributed 1,040; total distribution, 4,903. During the preceding twelve months, the average number of requested and paid electronic copies each issue was 962; the total average requested and paid print and electronic copies, 5,964. Corresponding figures for May 2018, the issue published nearest to filing date: number of requested and paid electronic copies, 955; the total requested and paid print and electronic copies, 5,781. Certified by Barbara Fiser, Director of Finance and Administration.

*The Journal of*

# *Economic Perspectives*

Fall 2018, Volume 32, Number 4

## Symposia



## Features



AMERICAN ECONOMIC ASSOCIATION