Working Paper Series
Congressional Budget Office
Washington, D.C.

Comparing the Effects of Current Pay and Defined Benefit Pensions on Employee Retention

Justin Falk and Nadia Karamcheva
Congressional Budget Office
justin.falk@cbo.gov
nadia.karamcheva@cbo.gov

Working Paper 2018-06

June 2018

To enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work, CBO's working paper series includes papers that provide technical descriptions of official CBO analyses as well as papers that represent independent research by CBO analysts. Papers in this series are available at http://go.usa.gov/ULE.

The authors thank Holly Boras, William Carrington, Molly Dahl, Nikolay Doudchenko, Guido Imbens, Joseph Kile, Jeffrey Kling, Amber Marcellino, James Poterba, Daniel Ready, Stephanie Ruiz, and David Torregrosa for their comments and suggestions.

www.cbo.gov/publication/54056

Abstract

Federal, state, and local governments continue to consider reducing the cost of their defined benefit pensions by decreasing annuity payments or having employees contribute a larger portion of their salaries toward them, thus reducing those workers' current pay. Such reductions to compensation can decrease the human capital of a workforce through lower employee retention. Using data that span more than 30 years and reflect substantial policy changes to federal workers' salary schedules and pension structure, we estimate that the average elasticity of job tenure with respect to the employer's cost is 1.5 for changes in current pay and 0.8 for changes in pension benefits. The magnitude of the estimated effects is different because cuts to the defined benefit pension cause many workers to delay retirement and also lead to fewer resignations than do similar cuts in current pay.

Keywords: public-sector compensation, pensions, employee retention

JEL Classification: J26, J33, J45

Contents

Introduction and Summary

Since 2009, the federal government, every state government, and many local governments have made substantial changes to the defined benefit (DB) pensions that they provide their employees. Those changes have typically combined increases in employee contributions and decreases in pension benefits, although some governments have relied almost exclusively on higher employee contributions (Brainard and Brown 2016).[1] Further changes are likely because the pension obligations of many of the plans continue to far exceed the value of the assets that have been set aside to fund them (Munnell and Aubry 2016). Cuts to pension benefits and reductions in current pay caused by higher employee contributions can affect people's incentive to continuing working for their current employer and thus decrease retention. Lower retention in turn can reduce the human capital of the workforce. This paper compares the effect on retention of changes in those forms of compensation by examining changes in resignation rates, retirement rates, and job tenure (measured as length of service at the time of separation). In doing so, it enhances the transparency of the Congressional Budget Office's work by providing a technical description of the analysis underlying a report that the agency published for the Congress in 2017 (CBO 2017).

Numerous studies have examined the effects of current pay and DB pensions on retention. For example, researchers have used the rigidities in teachers' pay schedules to estimate the effects of salaries on retention. Many others have examined the retention effects of prospective pension accruals. And a complementary literature has estimated teachers' and soldiers' willingness to pay for DB pensions.

This paper contributes to the literature on current pay, augments a separate literature on DB pensions, and provides a direct comparison of the ways in which those forms of compensation affect retention. With regard to current pay, we use a sudden adjustment in pay schedules to estimate the effect of changes in salary and mandatory employee contributions on retention. That sudden adjustment occurred at the beginning of 1991, when federal workers in New York City, San Francisco, and Los Angeles received an additional 8.2 percent increase in their salaries as the government began to adjust its pay schedules to accommodate local labor market conditions. That abrupt change allows us to measure the effect of current pay on retention for a workforce that includes a wide array of white-collar occupations. In contrast, most previous studies have focused on teachers, who account for only about 10 percent of state government employees and 30 percent of local government employees (Gittleman and Pierce 2012). We also make a methodological contribution to the literature by using the synthetic control method to validate our estimates of average effects through a systematic series of placebo tests.

[1] Unlike most DB plans in the private sector, government plans are typically financed with required contributions from both employees and employers.

With regard to the effect of DB pensions on retention, we study adjacent cohorts of federal workers who were in different retirement systems to estimate the parameters of a model that we then use to conduct policy simulations. Federal workers hired after 1983 were placed in a new retirement system, which had a less generous DB component but provided employer contributions to workers' defined contribution (DC) accounts. We use differences in retention between workers hired in 1983 and those hired in 1984 to estimate the parameters of an option value model. Because we estimate that model using panel data covering 31 years, we can predict retention over employees' entire careers with few extrapolations.

We can directly compare the effects of current pay and DB pensions on retention because we have administrative data that span 31 years and encompass both the abrupt salary increase and the reduction in the DB pension. After using those policy changes to estimate the parameters of our models, we simulate the average retention effects of a 2 percent cut in current pay and a 10 percent cut in the DB pension for the entire federal workforce. We selected those scenarios because they would decrease the cost of compensation by a similar amount and because they are similar in magnitude to recent reductions policymakers have made to those forms of compensation.

We find that the effect of current pay on retention is about twice the magnitude of the effect of DB pensions. We estimate that a 2 percent cut in current pay would decrease average job tenure by 2.3 quarters, primarily through an increase in resignations. In contrast, a 10 percent cut to DB pensions would increase resignations less and substantially decrease the separation rate among workers who are already eligible to draw a pension upon retirement, resulting in an average decrease in job tenure of 0.9 quarters. Combining those changes in separation rates, we estimate that the job tenure elasticity with respect to the employer's cost is 0.8 for DB pensions compared with 1.5 for current pay.

Using information on workers' annual performance ratings, we explore whether the effect of changes in compensation on retention is heterogeneous. For example, if the drops in retention caused by cuts to compensation are concentrated among workers with lower performance ratings, then reductions in the firm-specific human capital of the workforce from lower tenure might be offset by increases in the share of workers who are more skilled in general. However, point estimates for the job tenure elasticities suggest that cuts in compensation reduce retention more for better performers, thus potentially intensifying the effect of reduction in compensation on the average productivity of the workforce.

Background

We examine a period during which new policies led to substantial changes in the two main sources of compensation for federal employees: salaries and pensions. Those changes allow us to contribute to the literature that examines the effects of current pay and DB pensions on retention, as well as a closely related literature on workers' willingness to pay for DB pensions.

Compensation of Federal Employees

The federal government spends about $215 billion a year on salaries, retirement benefit obligations, health insurance, and payroll taxes for its roughly two million workers. Those employees receive most of their compensation through their salary, which is typically determined by their rank in a pay schedule (a series of rising salaries). Workers advance through those salary steps based on job tenure and performance (Falk 2012). They also have a substantial portion of their compensation deferred until retirement, primarily in the form of a DB pension, which provides retirees annuity payments until death.[2] The amounts of those payments are a function of length of service and the salaries they earn during their final few years of service. To be eligible to start drawing those pensions, workers need to meet age and years-of-service requirements. Mostly because of the rigidity of those eligibility criteria, pension benefits accrue unevenly over workers' careers. As a result, pension benefits are expected to affect retention differently depending on workers' ages and job tenure.

We use a sudden increase in the salaries of certain federal workers to estimate the effects of current pay on retention. For decades, the main federal pay schedule did not account for the location of employees. Then the Federal Employees Pay Comparability Act of 1990 (FEPCA) stipulated that the salaries in that General Schedule, as well as the schedule for the Department of Veterans Affairs, should be comparable to nonfederal salaries in the same location. It took over three years to design and apply a system of locality-based payments for those two pay schedules. In the meantime, the President used authority provided by FEPCA to raise those pay schedules in three metropolitan areas where private-sector salaries were particularly high. That locality raise was announced on December 12, 1990, and took effect on January 1, 1991.

We use data that captures changes in retention following a major overhaul of the retirement system for federal workers. The Civil Service Retirement System (CSRS) was established in 1920, before Social Security was instituted; consequently, employees who participate in CSRS do not participate in the Social Security program. However, the Social Security Amendments of 1983 required such participation from federal employees hired after 1983. Thus, policymakers created the Federal Employees Retirement System (FERS), which consists of Social Security, a smaller DB pension, and a DC plan. Workers in CSRS are also allowed to contribute to the DC plan, but only workers in FERS receive employer contributions, which are capped at 5 percent of a worker's salary. In both systems, most workers can choose to continue receiving subsidized health insurance until death if they are eligible to begin drawing a pension immediately upon retirement.

[2] Surviving spouses and dependents continue to receive annuity payments after the retiree's death if the retiree agreed to receive reduced payments during his or her retirement.

Related Literature

Several studies have estimated the effect of current pay on retention, but they have focused on the salaries of public school teachers.[3] Most of those studies measure the relationship between retention and salaries while allowing for school-district fixed effects and controlling for the observed traits of the teachers and their students. Other potential sources of endogeneity are less likely to be important for teachers than for most groups of workers because teachers' salaries generally just depend on the degrees they have earned, the district they work in, and their years of teaching experience. Most studies have found that higher salaries lead to substantial gains in retention (see Murnane and Olsen 1990 and Hendricks 2014 for examples). However, Hanushek, Kain, and Rivkin (2004) found that the gains are modest and that retention is more closely related to the characteristics of the students.

Numerous studies have examined the effect of DB pensions on retention. Unlike Social Security and DC plans, DB pensions often provide large incentives for workers to stay with their current employer until they meet the age and job tenure requirements for an immediate annuity. Accordingly, most studies of DB pensions recognize that the financial incentives extend beyond the benefits that the workers have accrued already and those they would accrue over the current period of employment. By staying for that current period, the workers also retain the option to stay through future periods, which provide much larger accruals in some cases. A few of those studies have embodied that option value in recursive models of indirect utility over the life cycle (for example, Gustman and Steinmeier 1986). In those models, the worker's choice is based on the expected utility from all potential job spells. Presumably because of the computational burden of that approach, most researchers have instead calculated the job spell that will maximize expected utility prior to estimation and then used the difference between that value and the utility the employees would receive from the benefits they have already accrued as the primary explanatory variable. Studies of older workers have generally found that that option value explains more variation in separations than the amount of benefits already accrued or the amount that would be accrued in the current period (for example, Stock and Wise 1990; Asch, Haider, and Zissimopoulos 2005; and Ni and Podgursky 2016). In contrast, studies of younger workers have had mixed results (for example, Warner and Goldberg 1984 and Black, Moffitt, and Warner 1990).[4]

---

[3] One exception is the 2017 report by the Congressional Budget Office, which examines the consequences of the federal government's cutting current pay by increasing the portion of salaries that its employees must contribute to their DB pensions. That change applied to workers hired after 2012, and workers hired in 2013 had much lower retention rates than workers hired in 2012. However, those results were not explored in depth because they were not the focus of that report.

[4] A related literature has examined the effect of Social Security on the timing of retirement. In particular, Coile and Gruber (2007) find that workers tend to retire when the option value of Social Security is low.

A complementary literature estimates workers' willingness to pay for DB pensions. Fitzpatrick (2015) finds that teachers in Illinois are willing to pay about 20 cents for an increase in future pension benefits that cost their employer a dollar. Similarly, Warner and Pleeter (2001) find that about 57,000 of the 66,000 eligible, separating soldiers chose an immediate lump sum over a DB pension that would have cost the government about twice as much. That literature also discusses why public-sector compensation packages are inefficiently tilted toward DB pensions.

Data

We analyze a panel of administrative data that covers nearly all federal employees from 1983 through the third quarter of 2014.[5] Those data contain quarterly records for each employee that provide point-in-time data on their salaries and demographics. In addition, the data include the day of each worker's hiring and separation and classify the voluntary separations as retirements—for workers who are eligible to draw a pension immediately—and as resignations otherwise. Involuntary separations are also specified in the data, which allows us to focus on employees that the government wants to retain by dropping workers from the sample if they separated involuntarily during the quarter.[6]

We limit the sample to workers who were in the pay systems that were altered the most by the policy changes. That includes the two pay schedules subject to the locality raise, which cover about 65 percent of the workers. Similarly, we limit the sample to the roughly 80 percent of workers who are eligible for the standard retirement benefits provided through CSRS or FERS.[7] We also drop workers if we are uncertain how much service they will credit to the DB pension because they have past military service or were rehired.

We extract two samples from the remaining data. For the analysis of DB pensions, we compare retention between the last cohort hired into CSRS and the first cohort hired into FERS to minimize differences between those workers' traits and the economic environments they faced. Because using only the 1983 and 1984 cohorts to estimate the effect of the locality raise in 1991 would limit us to workers who had only seven or eight years of job tenure, we construct a second sample for the analysis of current pay. That sample consists of observations from 1988 through 1993 for all cohorts. We do not use data past 1993 because other areas started receiving locality-

[5] The Office of Personnel Management provided those data from the Enterprise Human Resources Integration Data Warehouse Statistical Data Mart.

[6] Involuntary separations account for about 15 percent of the separations in the sample. The most common types of involuntary separations are retirements attributable to disability, death, termination, or resignation instead of involuntary action.

[7] Generally, workers in non-full-time positions receive reduced retirement benefits, workers in temporary positions do not receive retirement benefits, and workers in positions such as law enforcement receive augmented retirement benefits.

based raises in 1994. Constructing the sample that way allows us to measure the effect of the locality raise for up to three years after the raise. For symmetry, we include three years of pretreatment data in the main analysis.

In our base specifications, we limit both samples to workers with at least three years of job tenure. That restriction removes workers from the current pay sample who were hired after the locality raise, which allows us to separate the effect on retention from any effect on recruitment. We drop workers with fewer than three years of job tenure from the DB pension sample because the terms of FERS were not finalized until the end of 1986. Thus, workers in the 1984 cohort did not know what retirement incentives they would face for the first three years.

Although there is considerable overlap in the characteristics of the workers in the two samples, there are also noticeable differences (see Table 1). In both samples, the vast majority of workers are in professional, administrative, or technical occupations. The most common professional and administrative occupations are contract administrator, accountant, and engineer. Technical occupations are typically filled with less-educated workers who support the professionals and administrators. Nearly all the other workers were in clerical occupations. The number of those jobs has been declining over the past 35 years as records and communications have been switched to electronic formats. Clerical occupations are less common in the DB pension sample because those workers tend to have been hired more recently. That difference contributes to the following: Workers in the DB pension sample are more educated, more highly paid, and more likely to be men. We will address those differences by estimating average effects for both sets of analysis by simulating outcomes for the 1984 cohort.

Effect of Current Pay on Retention

The first part of this section describes our empirical approach, and the second part presents the estimates for the model and the results of the policy simulation.

Approach

We use difference-in-difference methods to compare retention rates in the three metropolitan areas where all employees received a large pay raise with retention rates across the rest of the United States. On December 12, 1990, the President declared that federal employees in the metropolitan areas of New York, Los Angeles, and San Francisco would receive an 8.2 percent raise at the beginning of 1991 (see Figure 1).[8] That locality raise was in addition to a 3.9 percent increase given to all employees on one of the two affected pay schedules. Although it was the only nonuniform change in those schedules until the government began introducing other locality-based raises in 1994, other factors could obscure the effect that the raise had on retention. Notably, unemployment was rising rapidly in 1991, particularly in the areas that

[8] Exec. Order No. 12736, 55 Fed. Reg. 51,385 (December 14, 1990).

received the locality raise. In addition, federal salaries were growing a bit faster in those areas before the locality raise went into effect, which could be a product of managerial discretion. Although most federal workers progress through their pay schedule on the basis of their years of federal employment, managers do have a few tools, such as promotions and performance-based raises, that they can use to retain some employees.

We use a variety of difference-in-difference methods to control for any secular differences in retention trends between the three treated areas and the rest of the country. In our base specification, we use a probit model to estimate the probability that worker *i* in area *j* voluntarily separates in quarter *t*.

$$\Pr\left(Separates_{i,j,t}\right) = \Phi\left(\beta D_{i,j,t} + X'_{i,j,t}\gamma + \delta_j + \tau_t\right)$$

D equals one in the quarters following the locality raise for the workers who received it, and zero otherwise. In addition to that indicator of treatment, all specifications include the state's unemployment rate and the worker's sex in the vector of control variables *X*, as well as area-fixed effects δ and quarter-fixed effects τ. Areas are generally based on boundaries for the states and the District of Columbia, but the metropolitan areas of the three cities that received the raise are separated from the rest of California, New York, New Jersey, and Connecticut. Lastly, we suspect that the values for some parameters vary on the basis of whether the worker is eligible to retire, so we estimate each specification separately for retirement-eligible and noneligible workers.

To separate the effect on retention from any effects on recruitment, we exclude workers who were hired after 1990 by dropping workers with fewer than three years of job tenure from the inference sample.[9] In doing so, we exclude from the main analysis the career stage in which the employees are most likely to leave. In supplementary analysis, we limit the sample to 1990 through 1991 and find that the nonlinearity in the probit model allows the base specification to accurately capture the effect of the locality raise on workers with one or two years of tenure (see Table A-1 in the appendix).

We interpret the parameter estimates by calculating the average retention effects of permanently reducing the current pay of all federal workers by 2 percent at the beginning of their careers. For comparability to the analysis of pensions, those effects are estimated for the 1984 cohort. The penultimate section of this paper provides a more detailed description of those simulations.

Results

We begin with the graphical evidence before turning to the estimates of the parameters and average effects. Before the locality raise went into effect, resignation rates were much higher in

[9] For consistency, we also drop workers with fewer than three years of job tenure from the earlier cross sections.

the areas that went on to receive them, but after the raise, the treated and control areas had similar resignation rates, on average (see Figure 2). In case those areas' averages were already converging before the locality raise went into effect, we compare the resignations rates in the treated areas with a weighted average of the resignation rates in the control areas that exhibits similar trends before treatment. Placing all of the weight on Virginia, Georgia, and the untreated portion of New York provides the best fit over the pretreatment period (see the appendix for a more detailed description of how the weights are selected and applied). After treatment, resignation rates in that synthetic control area fell much less than in the treated areas, indicating that the locality raise cut the resignation rate by about a third.

In contrast, we cannot detect an effect of the locality raise on retirement rates whether we compare simple averages or use the synthetic control method (see Figure 3). That is surprising because raises increase the amount of salary that workers would forgo by retiring and also provide the opportunity for them to earn larger pensions by staying. One explanation for why the raise did not delay retirements is that workers might base their retirement decision on the portion of their salaries that they would be able to replace in retirement, instead of dollar amounts. If their reference point remains the salary they were accustomed to, then the relevant replacement rate would not change immediately after the raise. If they continue working, they might become accustomed to the higher salary, but the raise would also cause their future pension to increase, and thus the replacement rate could remain unchanged.

We provide estimates for the parameters of our models on the basis of the locality raise and, when practical, use those models to simulate the average effect of a 2 percent reduction in current pay because such a change is similar to policies that are currently under consideration. Under our base specification, we simulate that a 2 percent reduction in current pay would increase the quarterly probability of resignation by an average of 0.04 percentage points, which amounts to 10 percent (see Table 2, specification 1). In contrast, the point estimate for the average effect on retirements is smaller and not statistically significant. Combining the estimates for resignations and retirements indicates that cutting the current pay of the FERS cohort by 2 percent would have decreased their average job tenure by about 2.3 quarters.

The four alternative specifications listed in Table 2 give very similar estimates of the effect of current pay on resignations. The estimated effect barely changes when area-specific, quadratic time trends are added in the second probit specification even though higher unemployment rates are associated with lower resignation rates in all the specifications, which suggests the unemployment rate is capturing any confounding trends in local economic conditions. Our two linear probability specifications and synthetic control specification also indicate that a 2 percent cut in current pay increased the probability of resignation by about 0.04 percentage points on

average.[10] The synthetic control specification further validates the results because the test of its statistical significance captures uncertainty from the possibility that assumptions underlying the estimate are errant. Specifically, the p-values are constructed using a systematic series of placebo tests (see Abadie, Diamond, and Hainmueller 2015), in which we estimate the difference in retention by replacing the treated area with one of the control areas and then by estimating the weights for that area's synthetic control group using the same procedure that we implemented for the treated area. Only one of the 46 control areas produced a larger estimated effect than the treated area.[11]

Only one of the four alternative specifications yields a statistically significant estimate of the effect of raises on retirements. Specifically, we estimate that a 2 percent cut in current pay decreases the probability of retirements by 0.28 percentage points under the linear probability model with area-specific trends (specification 3). However, that specification appears to conflate the effects of the locality raise with those of a temporary change in postemployment rules that occurred in the fourth quarter of 1990, and the estimate becomes statistically insignificant when that quarter is dropped from the sample (specification 4).[12] The other coefficient estimates have the anticipated sign, as eligibility to continue receiving employer-provided health insurance in retirement, the early retirement age for Social Security, and the full retirement age for that program are all associated with higher retirement rates.

Effect of Defined Benefit Pensions on Retention

The first part of this section describes our empirical model and the option value theory behind it. The second part presents the estimates for the model and the results of the policy simulation.

Approach

In this subsection, we describe the option value formulation we use to represent the financial incentives from the pension, illustrate how that option value and retention differ between CSRS and FERS, and then lay out the empirical specifications that we use to determine the effect of the option value on retention.

---

[10] The linear probability model is commonly used for difference-in-difference analysis but frequently predicts negative separation rates. Thus, we do not use it to estimate changes in job tenure.

[11] The synthetic control method does not take into account the amount of sampling variation underlying the area's separation rate, so we drop five areas that have a small number of federal employees. The appendix provides a more detailed description of our synthetic control estimates. It also illustrates that the estimates change little when a longer pretreatment period is used to select the synthetic control weights.

[12] During the buildup to the first Gulf War, the Congress suspended for six months a restriction that required former federal employees to wait a year before working for contractors from whom they procured services while employed by the federal government. Thus, some retirement-eligible employees had a onetime opportunity to immediately start drawing a federal pension while earning a salary from a contractor.

Option Value of DB Pensions. Most of the recent literature has focused on models of pension incentives in which workers realize that by staying for the current period, they retain the option to stay through future periods in which they often can accrue a large pension. Thus, the workers decide whether to stay in each period by considering the utility that they expect to receive for job spells of various lengths. Typically, studies have used the indirect utility function from Stock and Wise (1990), which we generalize to allow for the possibility that the employees resign instead of retire.

$$V(l)_t = \sum_{s=t}^{l-1} p_{s|t} d^{s-t} \left(y_{fed,s}\right)^g + \sum_{s=l}^{S} p_{s|t} d^{s-t} (ky(l)_s)^g$$

When deciding whether to leave in period *t*, the workers consider the expected utilities they would receive instead from leaving in future periods, which are indexed by *l*. Utility is derived from income *y* and summed over the life cycle, which is indexed by *s* with death a certainty following period *S*. Prior to workers leaving at *l*, their income is denoted y_{fed} and consists of their federal salary less the contributions they make to their pension. After they leave, *y* includes the DB pension benefits they had accrued, which depend on when they left, and earnings from nonfederal employment. The parameter *g* allows workers to prefer smooth income profiles, and the parameter *k* allows the value of income to generally differ after they leave the federal government. The value of future income is reduced by a constant discount factor *d* and the probability of surviving to the period in which the income will be received *p*.

To the canonical model described above, we apply two refinements from the recent literature. Like Coile and Gruber (2007; see Table 4), we set *g* and *k* to 1 so that we can separate the influence of the DB pension from other income sources. Data limitations described below probably cause inaccuracies in our incentive measures for Social Security and DC plans. Instead of having that measurement error directly contaminate our determinant of interest, we impose a linear utility function so that the coefficients can differ according to the source of income. That simplification is unlikely to substantially reduce the explanatory power of the DB pension because forward-looking workers will optimize consumption over the life cycle by adjusting their contributions to the DC plan.[13] We also follow Coile and Gruber (as well as Friedberg and Webb 2005) in minimizing the effect of differences in salary on the option value. As discussed earlier, managers can target raises at the workers they believe have the best nonfederal job prospects. To mitigate such endogeneity, we exclude salaries from the option value.

We incorporate the refinements discussed above and follow the bulk of the literature in asserting that workers base their decisions on the utility they expect to receive if they stay until the future

[13] Stock and Wise's (1990) estimate for *g* differs from 1, but the employer they studied did not offer a DC plan to its workers.

period l^* that maximizes V less the utility they expect to receive if they leave in the current period. That option value can be viewed as having three components:

$$OV(l^*)_t = V(l^*)_t - V(t)_t = \sum_{s=t}^{l^*-1} p_{s|t} d^{s-t} c_s + \sum_{s=l^*}^{S} p_{s|t} d^{s-t} b_{s|l^*} - \sum_{s=t}^{S} p_{s|t} d^{s-t} b_{s|t}$$

The first two sums represent the utility of staying until l^*. In that case, workers must contribute c toward the DB pension and later receive b_{l^*} in benefits from it. If they leave in the current period instead, they stop making contributions and receive b_t in benefits.[14] To calculate the option value, we need estimates of the discount factor and mortality rates. Our empirical model fits the data best with a discount rate of 6 percent, so we use a discount factor of about 0.94. The mortality rates are taken from the Office of Personnel Management (2017), which uses them to calculate the obligations of CSRS and FERS. Lastly, we measure the option value as a portion of the worker's salary at age 62. That replacement-rate-like approach focuses identification on the aspects of the pension formula that are common to all workers in the cohort, instead of differences in their salaries.

We also incorporate the option values for Social Security benefits and the FERS supplement, whereas the DC plan has only a trivial option value because federal workers receive their matching contributions immediately. The FERS supplement provides additional benefits to certain annuitants until they become eligible for Social Security at age 62. To receive the FERS supplement, workers must remain in federal service until they are eligible to retire with an immediate, full annuity. Those benefits are subject to the same earnings test that is applied to Social Security benefits prior to the full retirement age. Thus, precise measurement of the FERS supplement and Social Security benefits requires data on federal employees' earnings in other jobs, which we do not have (see the appendix for a description of how we calculate benefits for Social Security, the FERS supplement, and the DC plan). So we relegate the option value for those benefits to robustness tests.

Effect of the Overhaul to the DB Pension on the Option Value. For background, both FERS and its predecessor, CSRS, provide final-average-salary (FAS) pensions, which are also the norm for state and local governments. In such pension plans, the amount of the initial annuity is based on the workers' years of service, their salary over the last few years of service, and a multiplier.

Initial Annuity Payment = YOS x *FAS* x *Multiplier*

[14] CSRS, FERS, and Social Security are the only sources of benefits included. Thus, we are asserting that workers will not receive a DB pension from their next employer. That assertion is supported by the fact that federal employees who leave for another job typically move to the private sector where DB pensions are uncommon.

To be eligible for those payments, employees must leave federal employment and meet age and years-of-service requirements. In most instances, subsequent annuity payments are boosted by a cost-of-living adjustment.

Both the CSRS and FERS pensions provide incentives for workers to leave in their late 50s or early 60s, but those incentives are typically much larger under CSRS (see Table 3). One reason is that CSRS pensions tend to be larger, primarily because of a higher multiplier. Another reason is that CSRS has more stringent eligibility rules. For example, under CSRS, employees must continue working for the government until they reach the minimum retirement age in order to receive a full annuity before age 62. In contrast, under FERS, workers can begin receiving a full annuity at the minimum retirement age as long as they have 30 years of service.

Illustrative comparisons of retention rates to the option value indicate that the reduction in the option value under FERS had a much larger effect on retirements than resignations. The difference in option value between the systems is particularly stark for workers who are nearing 30 years of service as they approach the minimum retirement age (see Panel A of Figure 4). By waiting to leave until they have served for 120 quarters, a FERS employee who was hired at age 25 accrues four more years of annuity payments. The same worker under CSRS would accrue seven more years of payments because they would not have received an annuity until age 62 if they had not served until the minimum retirement age. In addition, the initial payments would be nearly two times larger under CSRS than FERS, and only under CSRS would the annuitant receive cost-of-living adjustments before age 62. However, separation rates are similar between the plans prior to retirement eligibility, suggesting that the smaller pensions available through FERS are still sufficient to stop most workers from resigning. In contrast, the retirement rates are noticeably lower under FERS—an average of 2.5 percent instead of 6.8 percent over the first three quarters of retirement eligibility, which suggests that workers are more responsive to financial incentives after they have locked down an immediate pension. In many of those instances, the pension payments that workers forgo while they stay are larger than the additional future payment they accrue. Among workers hired at other ages, retirements are also more responsive to the option value than resignations are (see Panel B of Figure 4).

Empirical Specifications. We use a probit model to estimate the probability that worker *i* voluntarily separates in quarter *t*.[15]

$$\Pr\left(Separates_{i,t}\right) = \Phi\left(\theta_1 \frac{OV(l^*)_{i,t}}{Salary\ at\ Age\ 62_{i,t}} + \theta_2 \frac{Accrued\ Value_{i,t}}{Salary\ at\ Age\ 62_{i,t}} + X'_{i,t}\lambda\right)$$

[15] Workers are dropped from the sample once they separate, which is comparable to how hazard functions are typically estimated in duration analysis. Whereas duration analysis often uses the proportional hazard specification, the option value literature typically uses the probit specification, which we find fits our data better.

The DB pension is allowed to affect separations through its option value and through the amount of future benefits the worker has accrued by the current quarter. That accrued value is intended to capture the wealth effects of the pension, and thus we expect θ_2 to be positive. The vector X contains the control variables. All specifications include quadratic functions of age and job tenure that are interacted with a dummy for the worker's retirement plan so that the effects of the DB pensions are identified from the kinks and discontinuities in their option values and accrued values. For the same reason, all specifications include dummies for the first year the worker is eligible for employer-provided health insurance in retirement and the early and full retirement ages under Social Security.[16] Lastly, we estimate the specifications separately for men and women.[17]

We exclude from the inference sample workers in their first three years of service because the 1984 cohort did not learn the terms of their retirement benefits until they were finalized at the end of 1986. That restriction has little effect on the estimates for the pension incentives because accruals and option values are smooth over those values of job tenure.

We interpret the parameter estimates by calculating the average retention effects of a 10 percent decrease in the DB annuity for the FERS cohort. The savings from that decrease are similar to the savings from cutting salaries by 2 percent.

Results

Under our base specification, we find that a 10 percent decrease in the FERS annuity is unlikely to have a substantial effect on job tenure for men (see Table 4, specification 1). Our estimate for the average effect is statistically insignificant, and the lower bound of the 95 percent confidence interval is a one-quarter decrease in job tenure. That change in job tenure is the product of a small and statistically insignificant increase in the resignation rate and a reduction in the retirement rate that is about 70 times larger. However, the point estimate for the change in job tenure is close to zero because most workers are eligible for retirement for only a small portion of their careers. The coefficient estimates show that those results are a combination of lower option values that encourage some workers to leave earlier and lower pension wealth that forces others to leave later. Another factor that appears to play a role in the retirement decision is whether the worker has recently become eligible to continue receiving employer-provided health insurance in retirement. The coefficient on that indicator implies that such eligibility causes 5.6 percent of the remaining workforce to retire earlier.

---

[16] A few studies have found that workers tend to stay until they are eligible for employer-provided health insurance in retirement and that that behavior can be misinterpreted as a response to pension incentives if eligibility for those health benefits is ignored (for example, Fitzpatrick 2014).

[17] We also estimated the current pay model separately for men and women, but we presented the result for the joint estimation because the parameter estimates were similar.

The results are fairly robust to the other specifications we estimated. The inclusion of a quadratic in salaries (specification 2) has little effect on the estimates, which suggests that measuring the pension incentives as a percentage of salary at age 62 is sufficient to handle any endogeneity in salaries. The inclusion of indicators for education, indicators for occupation, and the state unemployment rate also has little effect on the estimates (specification 3). In contrast, the inclusion of the incentives from Social Security and the FERS supplement does cause changes in the estimates for the DB pension that are on the margin of statistical significance (specification 4). But those differences are not large enough to alter our main results. Lastly, the coefficient on DC balances is statistically insignificant, probably because workers who place more value on retirement income contribute more to the DC plan and wait longer to retire (specification 5). However, the inclusion of DC balances does not change the estimates for the DB pension because DC balances do not exhibit discontinuities in the same places as the DB incentives.

The estimates for women are quite robust across specifications and differ from the estimates for men is some important ways (see Table 5). In particular, women tend to more closely follow the option value, while adhering less to the accrued value, so the pension cut reduces their retention more. We predict that a 10 percent decrease in the FERS annuity would increase the percentage of women resigning in an average quarter by a statistically significant 0.05 percentage points. In addition to the larger increase in their resignation rates, women also respond to the pension cut with a smaller increase in their retirement rate. Overall, we estimate that the more generous pension would increase their job tenure by 1.4 quarters.

Comparing the Effects of Current Pay and Defined Benefit Pensions on Retention and Performance

To compare the effects of current pay and DB pensions on retention, we estimate the average elasticity of job tenure with respect to the employer's cost. Those elasticities—as well as the underlying average effects on resignation and retirement rates—are simulated using the characteristics of the 1984 cohort because those workers are in the current retirement system and we can follow them for more than 30 years. In addition, we project those workers' future salaries and pension incentives so that we can simulate differences in retention over the cohort's life cycle. Our measure of the employer's cost consists of the FERS basic annuity and current pay, which we define as salary less the contributions workers make toward their annuity. The costs of future salaries and pension payments are discounted at a rate of 5 percent to reflect their current

market value.[18] We use a similar approach to compare the effects of current pay and DB pensions on performance, but we cannot quantify the magnitude of those effects.

We find that cuts to current pay cause job tenure to drop about twice as much as cuts in DB pensions that save the employer the same amount. That is because the savings from a 2 percent reduction in salary are similar to those of a 10 percent reduction in the FERS pension, but the effect of the former on retention is substantially larger (see Table 6).[19] Specifically, the cut in current pay would cause average job tenure to fall by 2.3 quarters, or 3.1 percent, which implies an elasticity with respect to the employer's cost of 1.5. In contrast, we estimate an elasticity of 0.8 for the decrease in the FERS pension. We suspect that our estimates understate the difference between those elasticities because we infer the effect of a reduction in current pay from a raise, although research suggests reductions have larger effects on retention. Specifically, the evidence of downward wage rigidity can be interpreted as indicating a substantial jump in workers' dissatisfaction as changes in pay switch from positive to negative.

Further analysis suggests that cuts in current pay reduce the workforce's human capital more than cuts in pension benefits. Note that the larger job tenure elasticity for current pay is not sufficient support for that finding because decreases in tenure can increase human capital if they are concentrated among less-productive employees. To examine that possibility, we divide the samples between workers who received above-median performance ratings from their managers and workers whose ratings were below the median. Resignation and retirement rates are then estimated separately for each subsample using the base specifications from the analyses above. For neither form of compensation are the retention losses concentrated among lower-rated workers (see Table 7). In fact, the point estimates for the job tenure elasticities indicate that both cuts would reduce retention more for higher-rated workers. For DB pensions, the difference between the elasticity for higher- and lower-rated workers is statistically significant and driven by a tendency for higher-rated women to be more likely to remain in federal employment after they are eligible for immediate, large pensions. (See Table A-2 in the appendix.)

Conclusion and Discussion

We use a sudden increase in salaries for some federal workers and a shift to a smaller DB pension to estimate the parameters of models that can simulate changes in retention from cuts to current pay and DB pensions. We find that a 2 percent cut in current pay would reduce average

[18] Novy-Marx and Rauh (2009) explain the importance of using market-based discount rates that reflect the risk profile of the pension liabilities. Note that the discount rate we use in the workers' utility function is 1 percentage point higher. Using a higher rate for workers is consistent with the literature's findings on their time preference and results in the option value model fitting our data better. See Table A-3 in the appendix for results under alternative discount rates for workers.

[19] Note that a 2 percent cut in salaries would reduce the employer's cost by 2 percent because—in addition to reducing salaries—it would reduce DB pensions through its effect on final average salaries.

job tenure by 3.1 percent, whereas a cut to the DB pension that saved the employer the same amount of money would reduce job tenure by about half that percentage. Those reductions in tenure appear to be concentrated among workers who had received better performance ratings, which suggests that either cut would harm the workforce's performance, with the cut to current pay having a larger effect.

By focusing on retention, this paper provides an informative analysis of a salient consequence of changing public-sector compensation. For decades, researchers have evaluated public-sector compensation by comparing it to the compensation of private-sector workers who have similar observable characteristics (for recent examples, see Gittleman and Pierce 2012 and Falk 2015). Interpreting those results is difficult because self-selection, omitted variables bias, and flaws in private-sector compensation are substantial concerns. By examining differences in retention between similar public-sector workers under different compensation policies, this paper complements those efforts by credibly measuring a tangible cost of cuts in compensation. Specifically, a substantial portion of federal workers would respond to reductions in their current pay or DB pension by separating earlier, even though past research has shown they are paid more than similar private-sector workers on average.

We are able to disentangle the effect of current pay on retention from its effect on recruitment, and we find no evidence that the change from CSRS to FERS caused confounding changes in recruitment. In general, reductions in compensation can limit an employer to hiring workers who have worse job prospects because they are less skilled, which can lead to a rise in retention being misinterpreted as an increase in the quality of the workforce. For the analysis of current pay, we avoid effects on recruitment by dropping from the inference sample workers who were hired after the locality increase went into effect, along with comparable workers in the pretreatment period. We cannot apply the same approach to the change in the DB pension because it was not imposed on incumbent workers. Instead, we examine whether retention rates differ between the CSRS and FERS cohorts during the first two years of service. As the terms of FERS had not yet been determined, differences in retention rates during that period are unlikely to result from differences between the retirement systems but could result from differences in recruitment. However, we find that the retention rates are similar between the cohorts during the first two years of service.[20] One explanation for the similarity is that recruits expected the new system would provide about the same amount of benefits as CSRS, which it does because the smaller DB pension is offset by Social Security benefits and employer contributions to workers' DC accounts (Martin 2003/2004).

[20] In contrast, Ippolito (2002) finds that early-career retention is higher under CSRS than FERS. He compares retention between civilians employed by the Air Force in 1987 who participated in CSRS and civilians employed by the Navy in 1996 who participated in FERS, whereas we compare adjacent cohorts of employees who were sampled from the same agencies.

Additional research in this area could provide further valuable insights. In particular, we examined just two of several levers that policymakers have at their disposal with regard to DB pensions. Other cost-saving changes, such as increasing the minimum retirement age, could affect retention differently.

Figure 1. Changes in Two Determinants of Retention

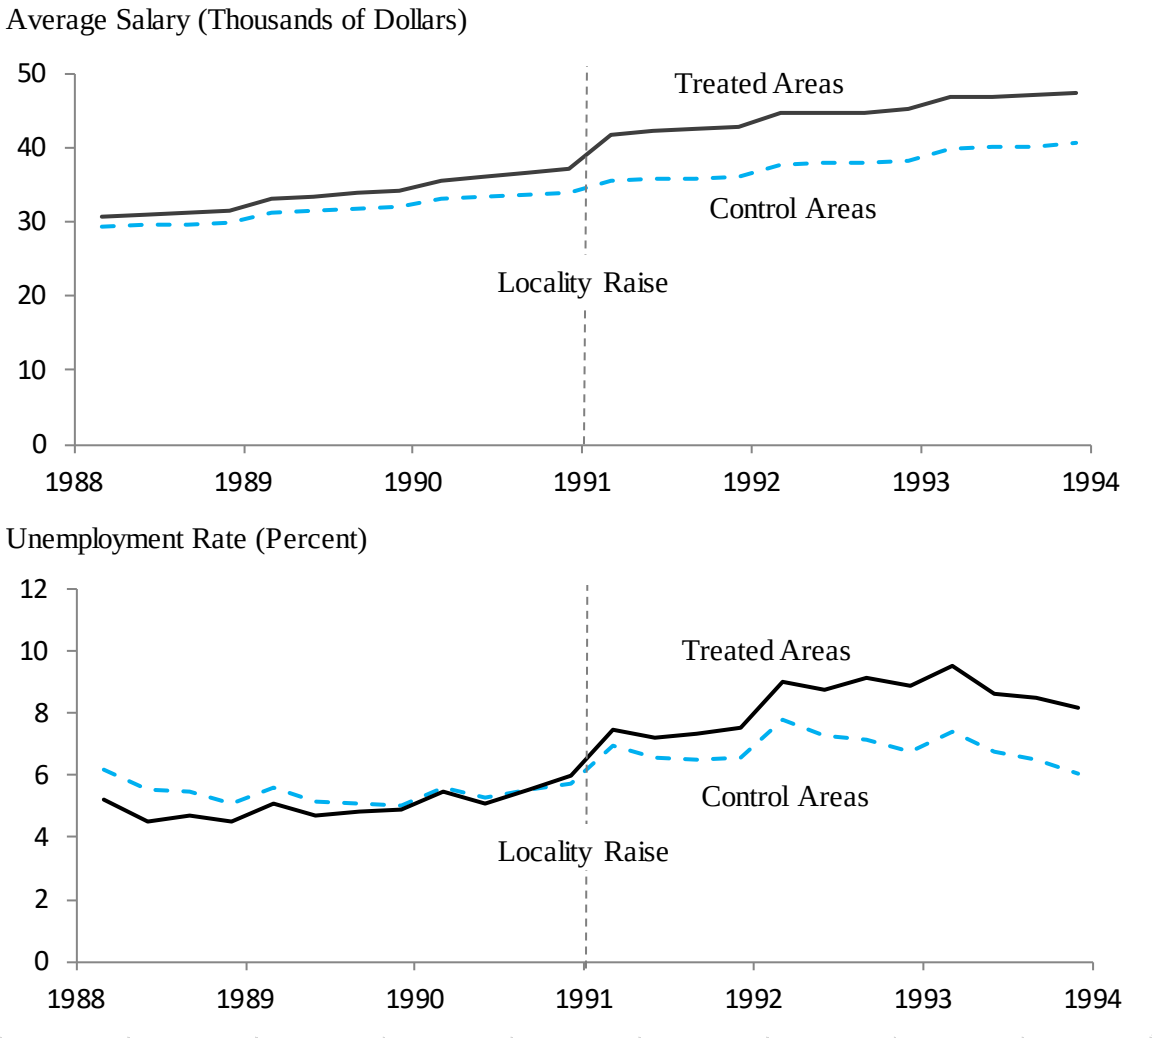

Source: Congressional Budget Office, using data from the Office of Personnel Management and the Bureau of Labor Statistics.

Workers in the metropolitan areas of New York City, San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. The figure shows average salaries for federal workers and the unemployment rates for the general workforces in the areas where federal employees work.

Figure 2. Resignation Rates by Receipt of Locality Raise

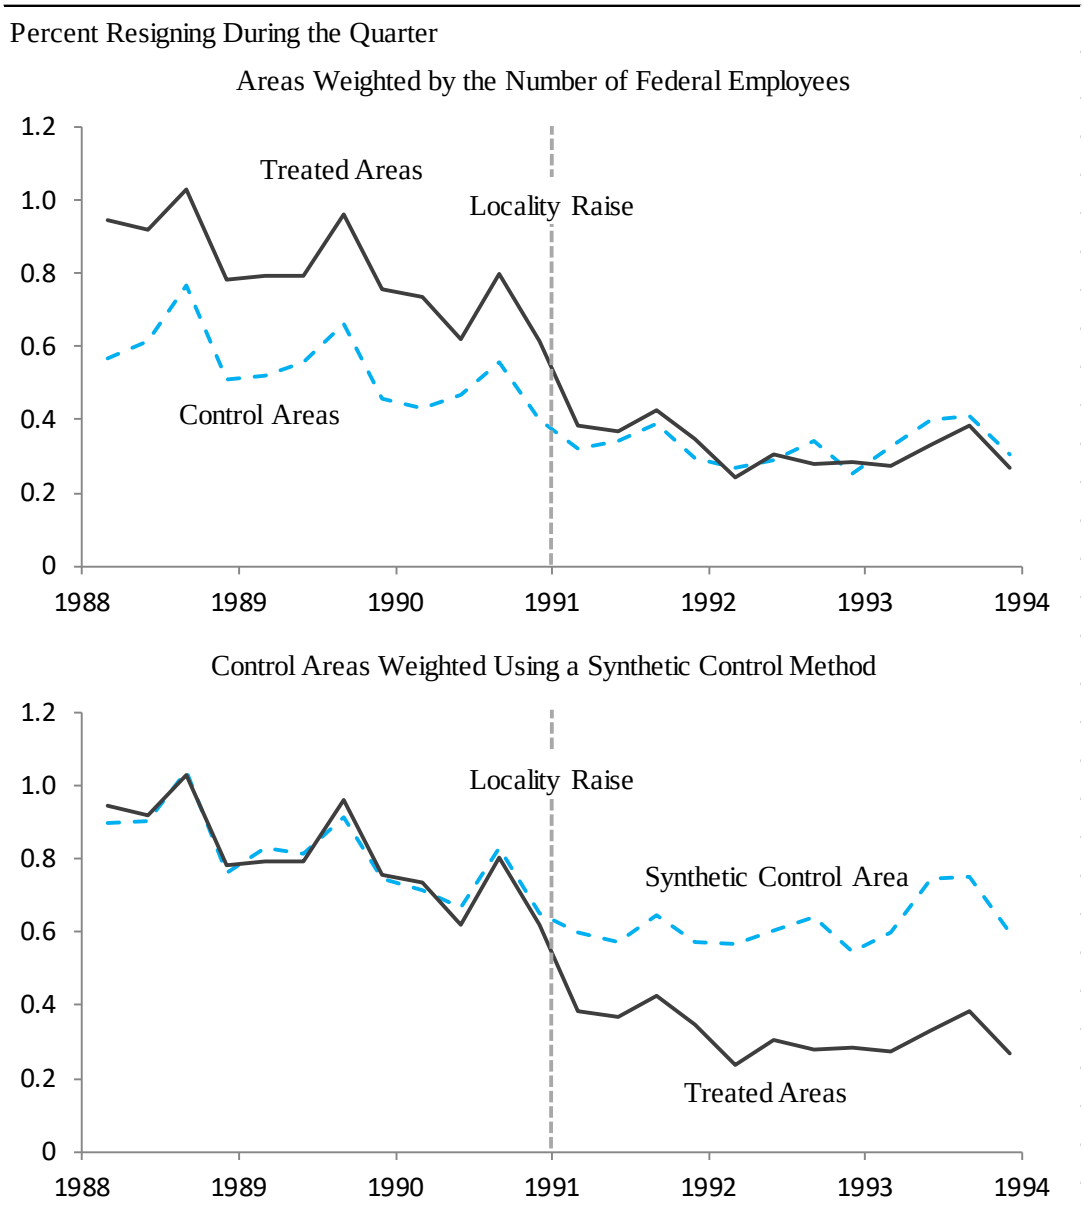

Area:	Virginia	Georgia	NY Excluding NYC
Synthetic Weight:	0.396	0.246	0.233

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Workers in the metropolitan areas of New York City (NYC), San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. Synthetic weights are estimated using an approach suggested by Doudchenko and Imbens (2016), which minimizes the differences between the treated areas and the synthetic control area in the pretreatment period subject to a penalty for overfitting. That approach allows for a constant difference between the treated and synthetic control areas, which we estimate to be 0.3 percentage points. Additional details are provided in the appendix.

Figure 3. Retirement Rates by Receipt of Locality Raise

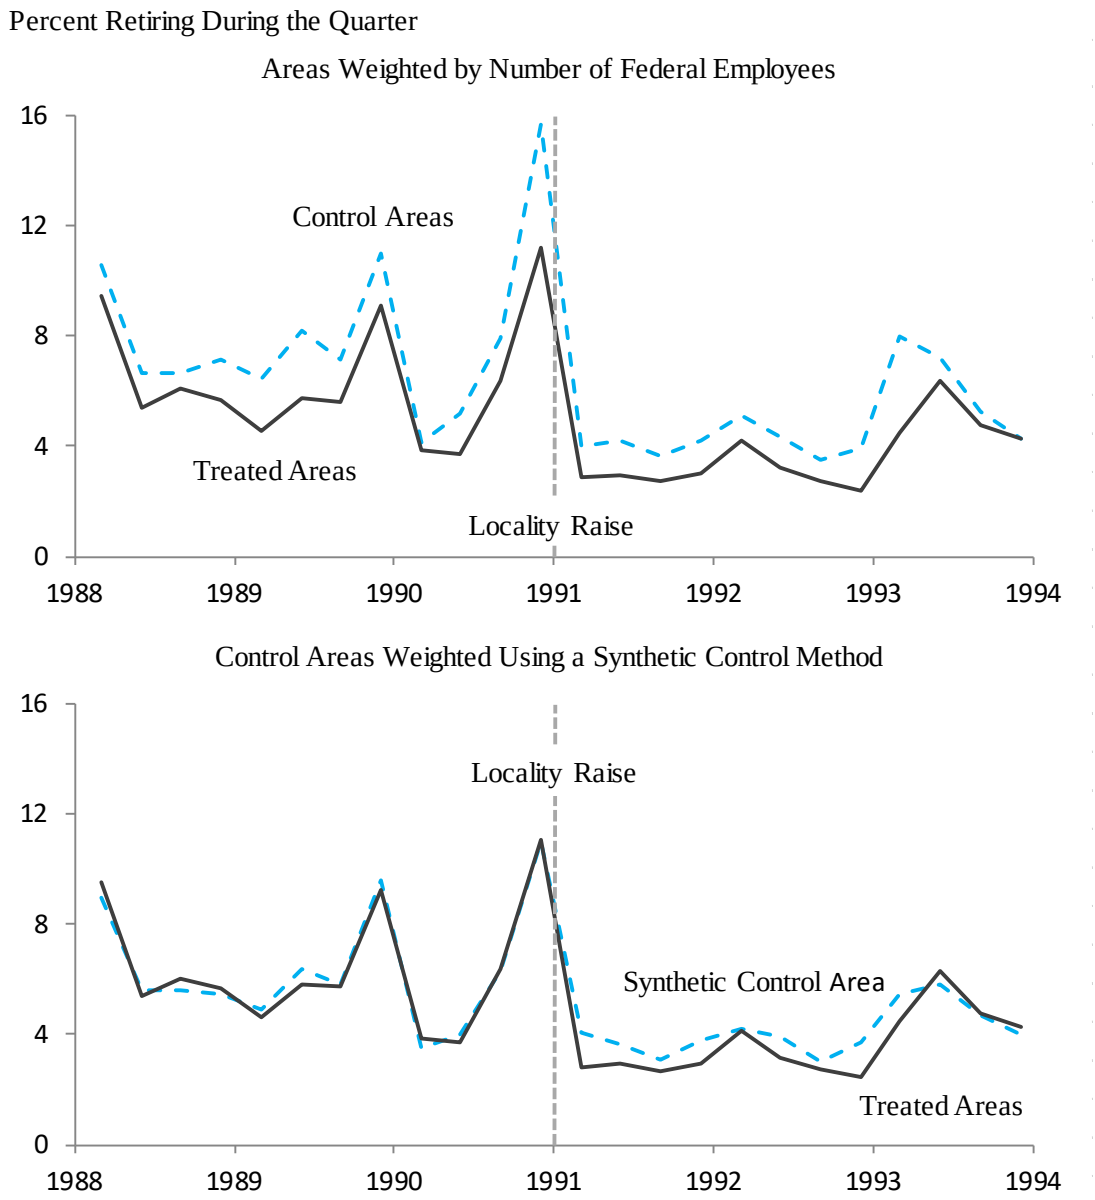

Area:	Maryland	Pennsylvania	Other Control Areas
Synthetic Weight:	0.515	0.127	0.197

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Workers in the metropolitan areas of New York City, San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. Synthetic weights are estimated using an approach suggested by Doudchenko and Imbens (2016), which minimizes the differences between the treated areas and the synthetic control area in the pretreatment period subject to a penalty for overfitting. That approach allows for a constant difference between the treated and synthetic control areas, which we estimate to be 0.6 percentage points. Additional details are provided in the appendix.

Figure 4. Comparisons of Option Values and Separation Rates Between FERS and CSRS

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Because our data go through the third quarter of 2014, we can follow some members of the 1984 cohort through 122 quarters of service. However, the small number of 40-year-old hires serving past age 67 leads to very imprecise separation rates beyond the 110th quarter.

CSRS = Civil Service Retirement System; FERS = Federal Employees Retirement System.

Table 1. Descriptive Statistics, by Sample

	Current Pay	DB Pension
Years in Sample	1988 – 1993	1983 – 2014
Years Hired	1934 – 1990	1983 – 1984
Sample Means		
Years of job tenure	16.1	14.4
Age	43.4	43.8
Annual salary[a]	76,600	85,100
Female	0.64	0.54
Occupation		
Professional and administrative	0.60	0.70
Technical	0.23	0.20
Clerical	0.17	0.10
Education		
High school or less	0.26	0.14
Some college	0.34	0.27
Bachelor's degree	0.27	0.43
Graduate degree	0.13	0.17
Region		
South	0.50	0.47
West	0.20	0.17
Midwest	0.17	0.21
Northeast	0.13	0.15
Observations	13,745,269	1,088,566

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Both samples are limited to workers in the following categories: those who are in one of the pay scales subject to the locality raise, those who are eligible for the standard defined benefit (DB) pension under the Civil Service Retirement System or the Federal Employees Retirement System, and those who have not served in the military or been rehired.

a. Salaries are adjusted to 2017 dollars using the employment cost index.

Table 2. The Effect of Current Pay on Retention

	Probit Specifications		Linear Probability Specifications		SCM
	(1)	(2)	(3)	(4)	(5)
Average Effects of a 2 Percent Decrease in Current Pay					
ΔQuarters of job tenure[a]	-2.31 ***	-2.54 ***			
	(0.31)	(0.58)			
ΔPr(Resigns) x 100	0.04 ***	0.05 ***	0.04 ***	0.04 ***	0.05 **
	(0.01)	(0.01)	(0.02)	(0.02)	n.a.
ΔPr(Retires) x 100	-0.02	0.02	-0.28 **	-0.06	0.21
	(0.04)	(0.07)	(0.14)	(0.11)	n.a.
Resignation Equation					
Coefficient estimates x 100					
1(Locality raise)	-8.72 ***	-9.42 ***	-0.16 ***	-0.15 ***	
	(3.12)	(1.77)	(0.06)	(0.06)	
State's unemployment rate	-3.77 ***	-1.37 ***	-0.02 ***	-0.02 **	
	(0.69)	(0.56)	(0.01)	(0.01)	
1(Female)	13.35 ***	13.38 ***	0.14 ***	0.14 ***	
	(1.46)	(1.46)	(0.02)	(0.02)	
Other statistics					
Log likelihood per observation	-0.0266	-0.0266	n.a.	n.a.	n.a.
Sample size	12,928,320	12,928,320	12,928,320	12,397,583	12,539,552
Retirement Equation					
Coefficient estimates x 100					
1(Locality raise)	0.73	-1.11	1.17 **	0.24	
	(2.39)	(4.85)	(0.59)	(0.47)	
State's unemployment rate	-1.00	-1.35	-0.14	-0.09	
	(0.83)	(1.29)	(0.15)	(0.17)	
1(Female)	12.29 ***	12.28 ***	1.40 ***	1.36 ***	
	(1.08)	(1.09)	(0.12)	(0.11)	
1(1st year eligible for employer-provided retiree health insurance)	38.67 ***	38.67 ***	4.65 ***	4.59 ***	
	(0.92)	(0.92)	(0.19)	(0.19)	
1(Early retirement age for Social Security)	8.85 ***	8.88 ***	0.82 ***	0.79 ***	
	(0.84)	(0.85)	(0.12)	(0.11)	
1(Full retirement age for Social Security)	20.83 ***	20.92 ***	2.56 ***	2.58 ***	
	(1.31)	(1.31)	(0.23)	(0.23)	
Other statistics					
Log likelihood per observation	-0.2217	-0.2214	n.a.	n.a.	n.a.
Sample size	816,949	816,949	816,949	785,212	783,565
Details of Specifications Common to Resignations and Retirements					
Other included regressors					
Area and quarter effects	Yes	Yes	Yes	Yes	n.a.
Area-specific time trends (quadratic)	No	Yes	Yes	Yes	n.a.
Excluded from the sample					
Quarter of change in postemployment rules	No	No	No	Yes	No
Areas with few federal workers	No	No	No	No	Yes

Source: Congressional Budget Office, using data from the Office of Personnel Management.

The specifications for resignations and retirement are estimated separately. Specifications 1 through 4 include quadratics in age and job tenure. For the probit specifications, the average effects are projected over the work lives of the 1984 cohort. The estimators for the standard errors are robust to heteroskedasticity and dependent sampling by area, and the estimates are reported in parentheses. We use the "in-space placebos" approach of Abadie,

Diamond, and Hainmueller (2015) to calculate a p-value for the synthetic control method. That approach does not yield standard errors. More details and results are in the appendix.

Significance levels: * = 10%, ** = 5%, *** = 1%.

n.a. = not available; SCM = synthetic control method specification.

a. We do not estimate changes in job tenure for the linear probability models or synthetic control method because those approaches frequently predict negative separation rates.

Table 3. Rules for the Basic Annuity, by Retirement System

	Civil Service Retirement System	Federal Employees Retirement System
Eligibility for a Full Annuity	Worker *stays until* age ≥ 55 with YOS ≥ 30, or age ≥ 60 with YOS ≥ 20, or *reaches* age ≥ 62 with YOS ≥ 5	Worker *reaches* age ≥ MRA[a] with YOS ≥ 30, or age ≥ 60 with YOS ≥ 20, or age ≥ 62 with YOS ≥ 5
Eligibility for a Reduced Annuity[b]	n.a.	Worker *reaches* age ≥ MRA with YOS ≥ 10
Multiplier	0.015 for first 5 YOS, 0.0175 for second 5 YOS, 0.02 after that	0.01 unless worker *stays until* age ≥ 62 with YOS ≥ 20, in which case it is 0.011
COLA	CPI-W	0 if age < 62, otherwise based on the CPI-W[c]

Source: Congressional Budget Office.

Deductions from the worker's salary are similar between the two systems with workers contributing either 7 percent to the Civil Service Retirement System (CSRS) or 0.8 percent to the Federal Employees Retirement System (FERS) and 6.2 percent to Social Security.

COLA = cost-of-living adjustment; CPI-W = consumer price index for urban wage earners; MRA = minimum retirement age; YOS = years of service; n.a. = not available.

a. The MRA for FERS is 55 for employees born before 1948. It gradually increases to 57 for employees born in later years. Workers in FERS only have to live until the relevant age (for example, the MRA) to receive the annuity, whereas workers in CSRS have to remain in federal employment until the relevant age.

b. Payments are permanently reduced by 5/12 percent for every month that the date of the first payment precedes the date at which the former employees would have received their first unreduced payment.

c. FERS annuitants who are at least 62 receive a COLA equal to the CPI-W if it is less than 2 percent, a COLA of 2 percent if the CPI-W is between 2 and 3 percent, and a COLA equal to the CPI-W minus 1 percent otherwise.

Table 4. The Effect of Defined Benefit Pensions on Retention for Men

	(1)	(2)	(3)	(4)	(5)
Average Effects of a 10 Percent Decrease in the FERS Annuity					
ΔQuarters of job tenure	-0.09	-0.06	-0.05	-1.12 **	-1.09 **
	(0.51)	(0.50)	(0.48)	(0.52)	(0.54)
ΔPr(Resigns) x 100	0.01	0.01	0.01	0.03 ***	0.03 ***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
ΔPr(Retires) x 100	-0.87 ***	-0.87 ***	-0.89 ***	-0.62 ***	-0.62 ***
	(0.17)	(0.17)	(0.17)	(0.21)	(0.21)
Probit Coefficient Estimates x 100					
Option value/earnings at 62					
Defined benefit pension	-9.71 ***	-9.29 ***	-8.90 ***	-12.61 ***	-12.56 ***
	(1.89)	(1.86)	(1.87)	(1.97)	(1.97)
Social Security and the FERS supplement				-27.96 ***	-27.97 ***
				(3.21)	(3.21)
Accrued value/earnings at 62					
Defined benefit pension	6.15 ***	6.19 ***	6.24 ***	3.78 ***	3.84 ***
	(1.19)	(1.18)	(1.19)	(1.25)	(1.26)
Social Security and the FERS supplement				2.17 ***	2.15 ***
				(0.36)	(0.36)
Defined contribution plan					-0.14
					(0.36)
1(1st year eligible for employer-provided retiree health insurance)	26.62 ***	27.19 ***	27.79 ***	28.85 ***	28.88 ***
	(3.96)	(3.94)	(3.94)	(4.06)	(4.06)
Other Included Regressors					
Base controls	Yes	Yes	Yes	Yes	Yes
Quadratic in earnings	No	Yes	No	No	No
Education and occupation	No	No	Yes	No	No
State's unemployment rate	No	No	Yes	No	No
Other Statistics					
Log likelihood per observation	-0.0324	-0.0323	-0.0323	-0.0322	-0.0322
Sample size	502,201	502,201	502,201	502,201	502,201

Source: Congressional Budget Office, using data from the Office of Personnel Management.

All specifications are estimated by probit and include quadratics in age and job tenure, indicators for the early and full retirement ages under Social Security, and an intercept, all of which are interacted with the worker's retirement plan. Specification 3 also includes sets of dummies for four categories of educational attainment and three occupational categories, as well as unemployment rates that are specific to the state and quarter. The average effects are based on a 10 percent decrease in annuity payments—which saves about as much as a 2 percent decrease in current pay—and calculated over the projected work lives of the cohort in the Federal Employees Retirement System (FERS). The estimators for the standard errors are robust to heteroskedasticity and dependent sampling by worker, and the estimates are reported in parentheses. The standard errors for the average effects are calculated using 1,000 bootstrapped samples.

Significance levels: * = 10%, ** = 5%, *** = 1%.

Table 5. The Effect of Defined Benefit Pensions on Retention for Women

	(1)	(2)	(3)	(4)	(5)
Average Effects of a 10 Percent Decrease in the FERS Annuity					
ΔQuarters of job tenure	-1.38 ***	-1.31 ***	-1.32 ***	-1.45 ***	-1.50 ***
	(0.31)	(0.31)	(0.30)	(0.29)	(0.29)
ΔPr(Resigns) x 100	0.05 ***	0.05 ***	0.05 ***	0.05 ***	0.05 ***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
ΔPr(Retires) x 100	-0.51 ***	-0.53 ***	-0.53 ***	-0.51 ***	-0.49 ***
	(0.13)	(0.13)	(0.13)	(0.15)	(0.15)
Probit Coefficient Estimates x 100					
Option value/earnings at 62					
Defined benefit pension	-12.90 ***	-12.58 ***	-12.64 ***	-12.99 ***	-13.13 ***
	(1.25)	(1.25)	(1.25)	(1.23)	(1.23)
Social Security and the FERS supplement				-17.81 ***	-17.58 ***
				(1.99)	(1.98)
Accrued value/earnings at 62					
Defined benefit pension	3.59 ***	3.73 ***	3.65 ***	3.35 ***	3.21 ***
	(0.81)	(0.81)	(0.81)	(0.83)	(0.83)
Social Security and the FERS supplement				0.56 ***	0.63 ***
				(0.17)	(0.18)
Defined contribution plan					0.58 *
					(0.30)
1(1st year eligible for employer-provided retiree health insurance)	25.23 ***	25.44 ***	25.31 ***	26.32 ***	26.12 ***
	(2.65)	(2.65)	(2.65)	(2.67)	(2.67)
Other Included Regressors					
Base controls	Yes	Yes	Yes	Yes	Yes
Quadratic in earnings	No	Yes	No	No	No
Education and occupation	No	No	Yes	No	No
State's unemployment rate	No	No	Yes	No	No
Other Statistics					
Log likelihood per observation	-0.0496	-0.0495	-0.0495	-0.0495	-0.0495
Sample size	586,365	586,365	586,365	586,365	586,365

Source: Congressional Budget Office, using data from the Office of Personnel Management.

All specifications are estimated by probit and include quadratics in age and job tenure, indicators for the early and full retirement ages under Social Security, and an intercept, all of which are interacted with the worker's retirement plan. Specification 3 also includes sets of dummies for four categories of educational attainment and three occupational categories, as well as unemployment rates that are specific to the state and quarter. The average effects are based on a 10 percent decrease in annuity payments—which saves about as much as a 2 percent decrease in current pay—and calculated over the projected work lives of the cohort in the Federal Employees Retirement System (FERS). The estimators for the standard errors are robust to heteroskedasticity and dependent sampling by worker, and the estimates are reported in parentheses. The standard errors for the average effects are calculated using 1,000 bootstrapped samples.

Significance levels: * = 10%, ** = 5%, *** = 1%.

Table 6. Comparing the Effects of Current Pay and Defined Benefit Pensions on Retention

	Current Pay	DB Pension
Job Tenure Elasticity With Respect to the Employer's Cost	1.5 ***	0.8 ***
	(0.2)	(0.2)
Underlying Estimates		
Percentage change in employer's cost	-2.0 ***	-2.0 ***
	(0.0)	(0.1)
Percentage change in job tenure	-3.1 ***	-1.6 ***
	(0.4)	(0.5)
Change in job tenure	-2.3 ***	-0.9 ***
	(0.3)	(0.3)

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Our measure of the employer's cost consists of expenditures for the defined benefit (DB) pension and current pay, which we defined as salary less employee contributions to the pension. Estimates for current pay are based on a 2 percent decrease in the 1984 cohort's salaries, and estimates for the DB pension are based on a 10 percent decrease in the same cohort's annuity payments. All estimates are calculated over the projected work lives of those employees. For current pay, the change in job tenure at exit is from specification 1 in Table 2. For the pension, the change in job tenure at exit is based on specification 1 in Tables 4 and 5. Standard errors are calculated using 1,000 bootstrapped samples.

Significance level: *** = 1%.

Table 7. Comparing the Effects of Current Pay and Defined Benefit Pensions on Retention, by Worker's Performance Rating

	Higher-Rated Workers		Lower-Rated Workers	
	Current Pay	DB Pension	Current Pay	DB Pension
ΔQuarters of Job Tenure	-2.495 ***	-1.312 ***	-1.987 ***	-0.393
	(0.491)	(0.387)	(0.430)	(0.413)
ΔPr(Resigns) x 100	0.042 ***	0.045 ***	0.037 ***	0.020 *
	(0.008)	(0.009)	(0.008)	(0.010)
ΔPr(Retires) x 100	0.007	-0.544 ***	-0.020	-0.845 ***
	(0.052)	(0.147)	(0.055)	(0.158)

Source: Congressional Budget Office, using data from the Office of Personnel Management.

The sample is divided between workers who receive performance ratings from their managers that were at the median or above and workers whose ratings are below the median. Estimates for current pay are based on a 2 percent decrease in the 1984 cohort's salaries, and estimates for the defined benefit (DB) pension are based on a 10 percent decrease in the same cohort's annuity payments. All estimates are calculated over the projected work lives of those employees. For current pay, the change in job tenure at exit is from specification 1 in Table 2. For the pension, the change in job tenure at exit is based on specification 1 in Tables 4 and 5. Standard errors are calculated using 1,000 bootstrapped samples.

Significance levels: * = 10%, ** = 5%, *** = 1%.

Appendix

This appendix provides additional details about benefit calculations for Social Security, the Federal Employees Retirement System (FERS) supplement, and the defined contribution (DC) plan; heterogeneity in the effect of current pay on retention; and details of the synthetic control analysis for the effect of current pay on retention. In addition, the appendix includes supplementary tables that are referenced in the main text of the working paper.

Benefit Calculations for Social Security, the FERS Supplement, and the DC Plan

Although we have extensive data on current pay and defined benefit (DB) pension accruals—data that span from the beginning of calendar year 1983 through the third quarter of 2014—we have limited information on benefits from the DC plan, Social Security, the FERS supplement, and other sources of retirement income. Because we control for age and job tenure for each retirement system using a smooth function of those variables, inaccuracies in our estimates of those other benefits are unlikely to bias our estimates of the effect of the DB pension on retention—unless discontinuities in the accruals of those other benefits coincide with discontinuities in DB pension accruals.

For the DC plan that the government provides its employees, we only have data from 2008 through 2014. For years before 2008, we extrapolate DC contributions and balances by combining the available data on workers' contribution rates, balances, and portfolio allocations with historic rates of return. However, we have no basis for such extrapolations for workers who separated before 2008. For them, we impute values using the contributions and balances of workers who did not separate before 2008 and who were similar in age and had similar levels of salary and education.

We cannot precisely estimate the income workers will receive through Social Security or the FERS supplement because we do not have data on their income from nonfederal employment. Instead, we project their expected retired worker benefits from Social Security by imputing earnings outside of the federal government using historic and projected growth rates in the employment cost index and then applying Social Security program rules.[21] Because Song and Manchester (2007) find that the retirement rate for the general population spikes at both the early retirement age and the full retirement age for Social Security, we include indicators for those ages in all specifications to avoid attributing those retirements to the DB pension. Accruals for the FERS supplement jump when workers gain eligibility for it by becoming eligible to receive

[21] Because we do not have information on spouses' earnings in the calculation, we consider only the Social Security benefits that are based on a worker's own earnings history. Also, we do not include cross-worker variation in the growth rates for nonfederal salaries because we lack relevant data.

an immediate, full DB pension before age 62. We do not apply the earnings test to the FERS supplement because we do not know whether eligible annuitants took other jobs after retiring from federal service.

We do not have data on other sources of retirement income, such as other DC plans, home equity, and spouses' investments. However, surveys indicate that few federal employees anticipate that those assets will be important sources of retirement income (Federal Thrift Retirement Investment Board 2008). Most federal employees might prefer the government-provided DC plan over those assets because the former has low administrative costs (an expense ratio of 0.03 percent per year) and offers a risk-free asset with above-market returns.

Heterogeneity in the Effect of Current Pay on Retention

Two sources of heterogeneity that we do not allow in the base specification could substantially alter our estimates of the effect of current pay on resignations. First, we assume that the nonlinearity in the probit model captures how that effect changes with job tenure. Second, we do not allow the effect of changes in current pay on resignations to change with the amount of time that has lapsed since the change in pay.

We conclude that the base probit specification accurately approximates the relationship between the effect of current pay on resignations and job tenure, even though it does not include an interaction term and data on workers with fewer than three years of tenure. We reach that conclusion by examining the sensitivity of our results to the inclusion of those elements (see Table A-1). By limiting the sample to observations within a year of the raise, we can include workers with one or two years of tenure without conflating resignation effects with recruitment effects. With those changes, our estimate of the interaction term's coefficient is marginally statistically significant and its inclusion leads to only a small increase in our estimate of the average effect of current pay on resignations. The average effect rises slightly because resignation rates are higher early in employees' careers, when a larger portion of them are still in federal employment. However, the point estimate is too small to cause a substantial increase in the average effect.

We find little evidence that the effect of current pay on retention changes with the amount of time that has lapsed since the change in pay. We examine that possibility by adding to the base specification an interaction between the treatment indicator and quarters since treatment. Our estimate of the coefficient on that indicator is not statistically significant.

Details of the Synthetic Control Analysis

Athey and Imbens (2016) call the synthetic control method "arguably the most important innovation in the policy evaluation literature in the last 15 years." Building on more arbitrary difference-in-difference methods, it systematically provides more valid counterfactuals. In this section, we describe our application of the synthetic control method and provide detailed results.

Following Doudchenko and Imbens (2016), we compare retention rates between the treated areas and a weighted average of the control areas that exhibits similar trends in those rates before treatment. Ignoring the overfitting penalty for the moment, we select the time-invariant weights ω and a constant difference γ_0 that minimize the squared distance between the retention rates for the three treated areas Y_1 and the rates for the weighted average of the J control areas before treatment, which occurs after T_0.

$$\min_{\gamma_0,\omega} \sum_{s=1}^{T_0} \left(Y_{1,s} - \gamma_0 - \sum_{j=1}^{J} \omega_i Y_{0,j,s} \right)^2$$

Thus, we could estimate the value of the parameters through a least squares regression of the retention rate for the aggregated treated area on the retention rates for the control areas with an observation for each pretreatment period. The constant is included to allow for parallel movements in the retention rates. We restrict the weights to zero for control areas with few federal workers to avoid conflating sampling-based fluctuations in the retention rates for the control areas with fluctuations in the population retention rates for the treated area, which includes many federal workers.

We add to the objective function above an overfitting penalty in order to maximize the precision with which the retention rates for the synthetic control area predict what those rates would have been for the treated areas in the absence of treatment. To understand the importance of the overfitting penalty, consider the case in which least squares perfectly fits the pretreatment retention rates for the synthetic control area to those of the treated area because there is only one less control area than pretreatment period. In that instance, the least squares fit captures not only the signal from the control areas' retention rates, but also the sampling-based noise, which is not useful for predicting the counterfactual rates for the treated group. We use the elastic net penalty suggested by Doudchenko and Imbens, which is a combination of the absolute values of the weights and the squares of the weights.

$$\min_{\gamma_0,\omega,\lambda,\alpha} \sum_{s=1}^{T_0} \left(Y_{1,s} - \gamma_0 - \sum_{j=1}^{J} \omega_j Y_{0,j,s} \right)^2 + \lambda \left[\frac{1-\alpha}{2} \sum_{j=1}^{J} |\omega_j| + \alpha \sum_{j=1}^{J} \omega_j^2 \right]$$

The parameter α gives the relative amount of weight given to the two penalty terms, and λ determines the amount of weight given to overfitting in total.

Before estimating the parameters, we drop areas with few federal employees from the analysis and then select which of the remaining control areas can have nonzero weight. We drop areas with few federal employees because the synthetic control procedure does not take into account

sampling variation underlying the area's separation rate.[22] In addition, to guarantee a unique solution, we set the maximum number of areas that can be included in the synthetic control group to three less than the number of pretreatment periods because we need to estimate γ_0, α, and λ in addition to the weights. Specifically, we only allow nonzero weights for the T_0-3 control areas whose average separation rate over the pretreatment period is closest to that average for the treated area.

Following Doudchenko and Imbens (2016), we use the values of the parameters that minimize the mean squared cross-validation error, after having selected potential values of λ using coordinate descent. The test period for the cross validation is four quarters (that is, we set the number of folds to $T_0/4$).

The point estimates for the effects on retention are based on the procedure described above, and the statistical significance of those estimates is determined by J placebo tests. In each of those placebo tests, we implement the procedure after replacing the treated area with one of the control areas.

We find that the locality raise—an 8.2 percent increase to the salaries of federal employees in three metropolitan areas—substantially reduced the resignation rate but did not have a discernible effect on the retirement rate. For 1991, the average difference between the resignation rates for the areas that received the locality raise and their synthetic control group was 0.2 percentage points compared with zero percentage points in 1990 (see Figure A-1). Thus the difference-in-difference estimate for the first year is a 0.2 percentage-point increase in the resignation rate, which is larger than the estimate for 45 of the 46 placebo tests. In a linear model, that point estimate implies that a 2 percent cut in current pay would lead to a 0.05 percentage-point increase in the resignation rate. Athey and others (2017) find that, under certain circumstances, a lengthy pretreatment period is needed for the synthetic control estimator to perform well. That does not appear to be the case in this application, as extending the pretreatment period to nine years has little effect on the estimates (see Figure A-2). (Detailed results for the effect of current pay on the retirement rate are shown in Figure A-3 and Figure A-4.)

[22] More specifically, we drop the five areas that had the fewest employees who were eligible to resign (that is, not eligible to retire) and the seven areas that had the fewest employees who were eligible to retire. Thus, the vast majority of the states and the District of Columbia remain in the analyses.

Figure A-1. Synthetic Control Estimates for the Effect of the Locality Raise on Resignations Using Three Years of Pretreatment Data

Difference in Percent Resigning

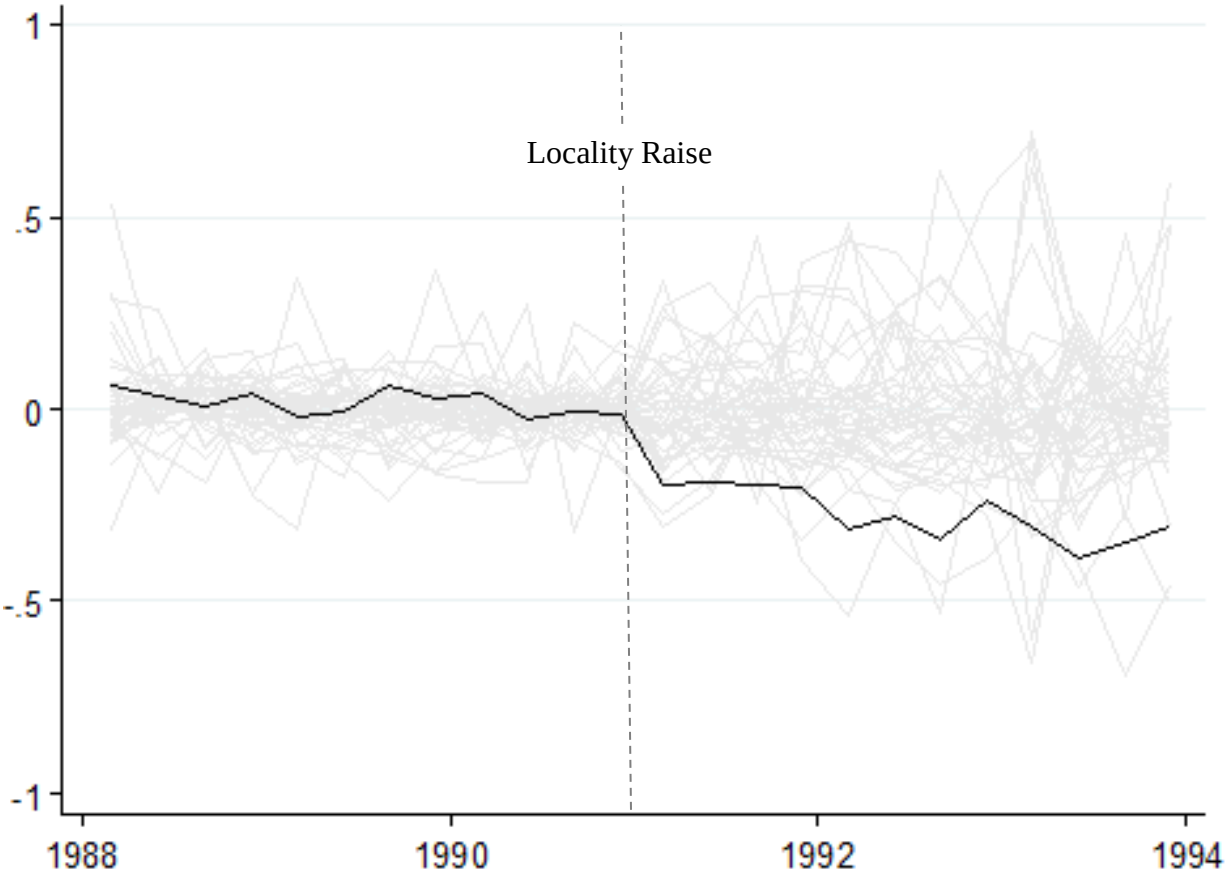

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Workers in the metropolitan areas of New York City (NYC), San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. The black line represents the difference in resignation rates between the areas that received that locality raise and their synthetic control group. Thus, it is the difference between the two lines shown in the bottom panel of Figure 2. The gray lines represent those differences for the 46 control areas.

Figure A-2. Synthetic Control Estimates for the Effect of the Locality Raise on Resignations Using Nine Years of Pretreatment Data

Difference in Percent Resigning

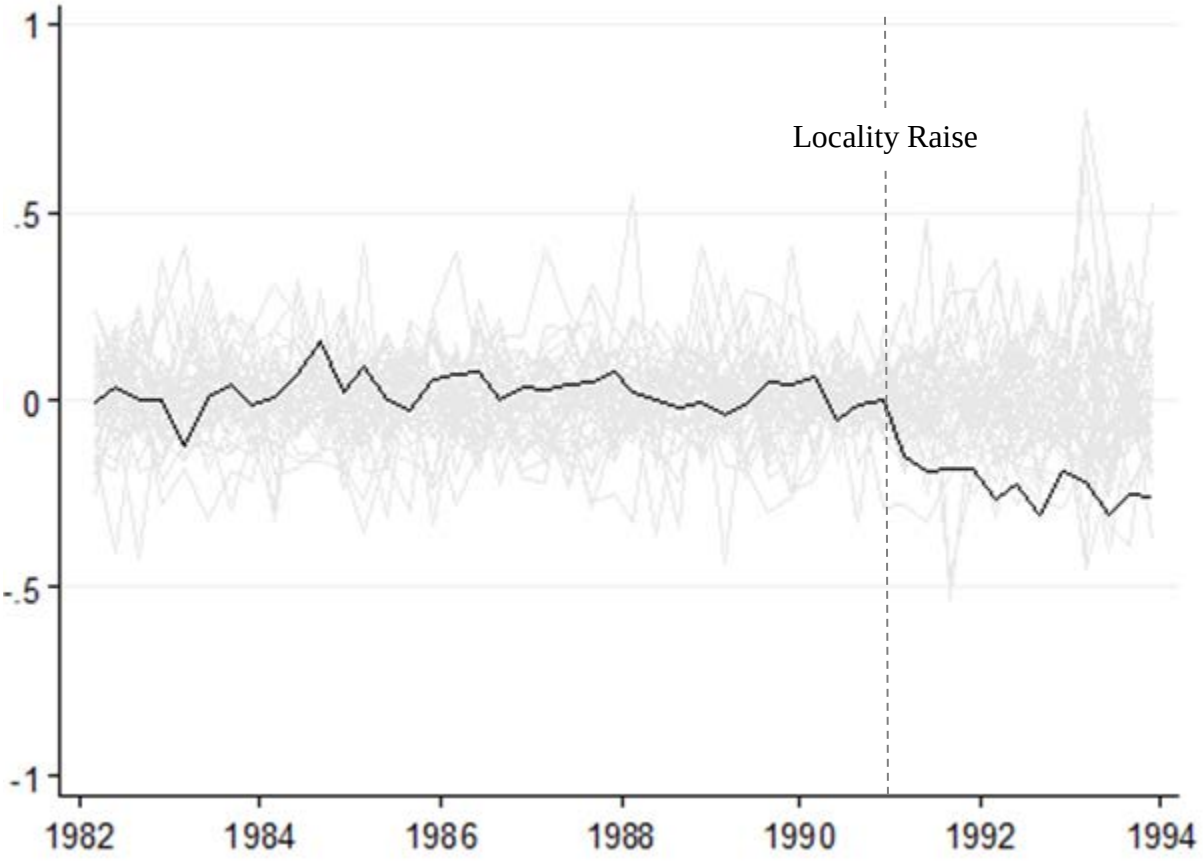

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Workers in the metropolitan areas of New York City, San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. The black line represents the difference in resignation rates between the areas that received that locality raise and their synthetic control group. The gray lines represent those differences for the 46 control areas.

Figure A-3. Synthetic Control Estimates for the Effect of the Locality Raise on Retirements Using Three Years of Pretreatment Data

Difference in Percent Retiring

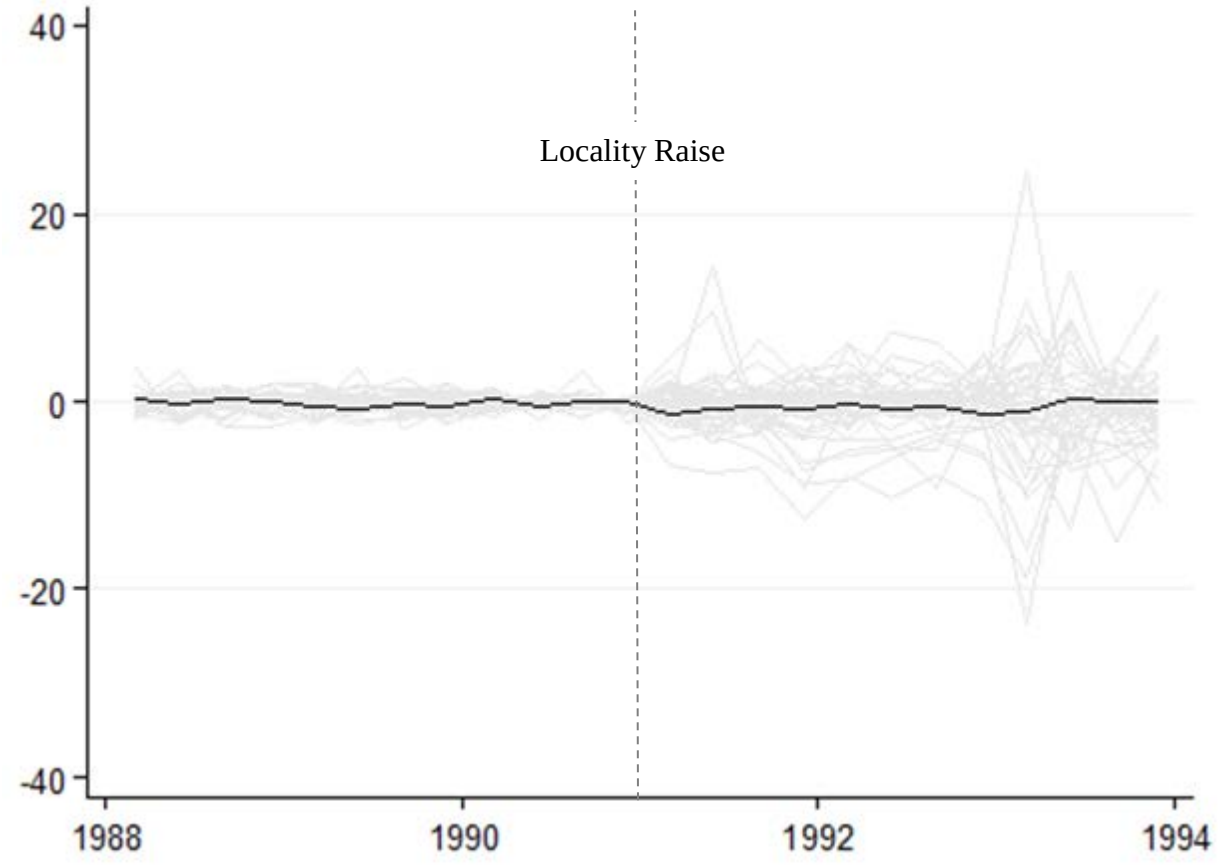

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Workers in the metropolitan areas of New York City, San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. The black line represents the difference in retirement rates between the areas that received that locality raise and their synthetic control group. Thus, it is the difference between the two lines shown in the bottom panel of Figure 3. The gray lines represent those differences for the 44 control areas.

Figure A-4. Synthetic Control Estimates for the Effect of the Locality Raise on Retirement Using Nine Years of Pretreatment Data

Difference in Percent Retiring

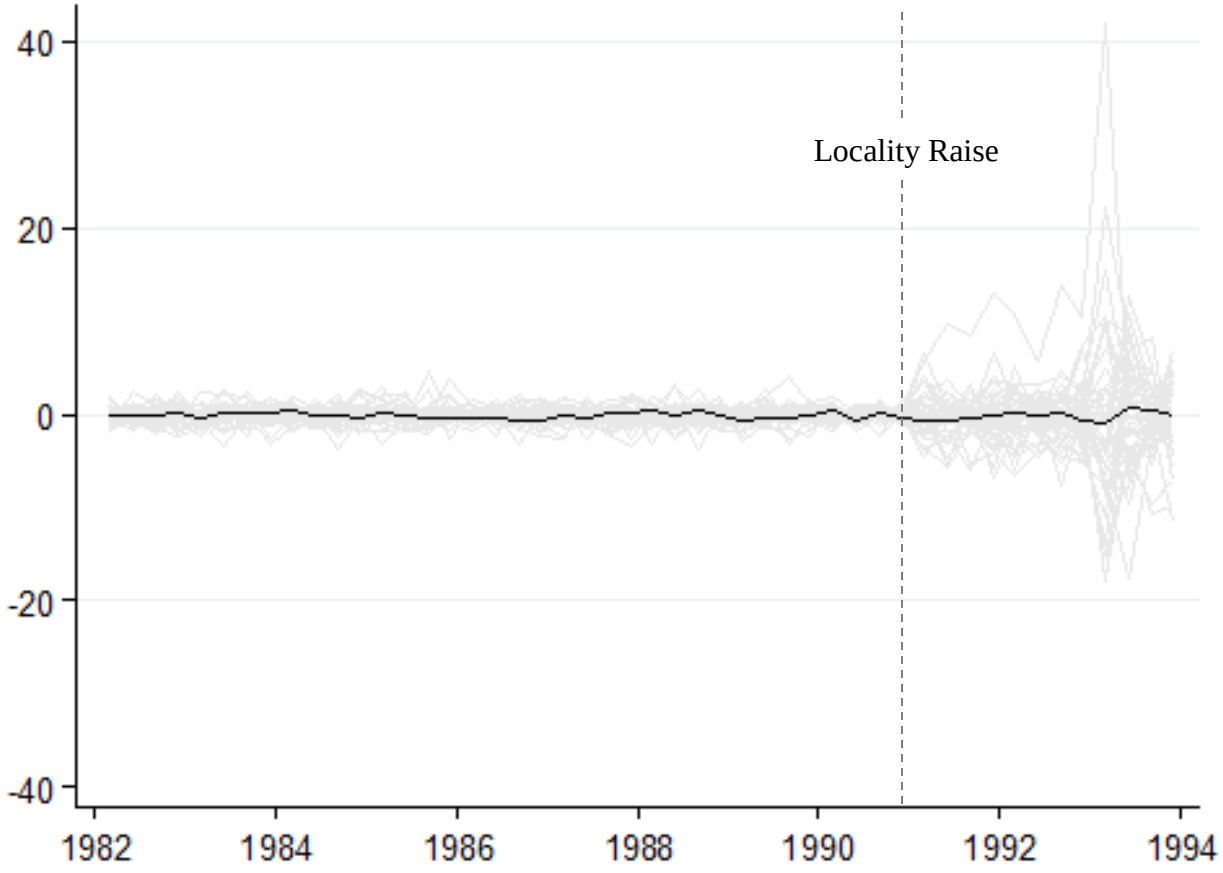

Source: Congressional Budget Office, using data from the Office of Personnel Management.

Workers in the metropolitan areas of New York City, San Francisco, and Los Angeles received an additional 8.2 percent salary increase in January 1991, whereas other areas did not begin receiving pay increases on the basis of local labor market conditions until 1994. The black line represents the difference in retirement rates between the areas that received that locality raise and their synthetic control group. The gray lines represent those differences for the 44 control areas.

Table A-1. Heterogeneity in the Effect of Current Pay on Resignation Rates

	Base Specification	Heterogeneity by Job Tenure	Heterogeneity by Lag
Probit coefficient estimates x 100			
1(Locality raise)	-8.78 ***	-11.59 ***	-6.69 ***
	(3.12)	(1.44)	(2.66)
1(Locality raise) x job tenure		0.06 *	
		(0.03)	
1(Locality raise) x quarters since raise			-1.62
			(1.06)
State's unemployment rate	-3.76 ***	-2.00 ***	-3.71 ***
	(0.69)	(0.51)	(0.71)
1(Female)	13.26 ***	10.40 ***	13.26 ***
	(1.45)	(1.35)	(1.45)
Other statistics			
Log likelihood per observation	-0.0267	-0.0295	-0.0267
Sample size	12,928,320	4,609,915	12,928,320
Characteristics of the sample			
Years included	1988 – 1993	1990 – 1991	1988 – 1993
Includes workers with 1 or 2 years of job tenure	No	Yes	No

Source: Congressional Budget Office, using data from the Office of Personnel Management.

All specifications were estimated by probit and include a quadratic in age, a quadratic in job tenure, area fixed effects, and quarter fixed effects. The estimators for the standard errors are robust to heteroskedasticity and dependent sampling by area, and the estimates are reported in parentheses.

Significance levels: * = 10%, ** = 5%, *** = 1%.

Table A-2. The Effect of Defined Benefit Pensions on Separation Rates, by Performance Ratings and Sex

	Men		Women	
	Higher Rated	Lower Rated	Higher Rated	Lower Rated
Average Effects of a 10 Percent Decrease in the FERS Annuity				
ΔQuarters of job tenure	-0.01	-0.43	-2.24 ***	-0.37
	(0.70)	(0.76)	(0.44)	(0.44)
ΔPr(Resigns) x 100	0.01	0.02	0.07 ***	0.02
	(0.01)	(0.02)	(0.01)	(0.01)
ΔPr(Retires) x 100	-1.03 ***	-0.73 ***	-0.20	-0.92 ***
	(0.26)	(0.25)	(0.19)	(0.20)
Probit Coefficient Estimates x 100				
Defined benefit pension				
Option value/earnings at 62	-9.92 ***	-9.02 ***	-17.19 ***	-7.91 ***
	(2.87)	(2.62)	(1.79)	(1.81)
Accrued value/earnings at 62	6.86 ***	5.26 ***	1.65	6.26 ***
	(1.74)	(1.73)	(1.08)	(1.25)
1(1st year eligible for employer-provided retiree health insurance)	26.85 ***	26.95 ***	22.57 ***	28.98 ***
	(6.01)	(5.39)	(3.82)	(3.75)
Other statistics				
Log likelihood per observation	-0.026	-0.030	-0.041	-0.047
Sample size	244,531	225,861	284,873	260,418

Source: Congressional Budget Office, using data from the Office of Personnel Management.

The sample is divided between workers who receive performance ratings from their managers that are at the median or above and workers whose ratings are below the median. Workers are dropped from the sample if they did not receive a performance rating or are not rated using the most common system. All specifications are estimated by probit and include quadratics in age and job tenure, indicators for the early and full retirement ages under Social Security, and an intercept, each of which is interacted with the worker's retirement plan. The average effects are based on a 10 percent decrease in annuity payments and calculated over the projected work lives of the cohort in the Federal Employees Retirement System (FERS). The estimators for the standard errors are robust to heteroskedasticity and dependent sampling by worker, and the estimates are reported in parentheses. The standard errors for the average effects are calculated using 1,000 bootstrapped samples.

Significance levels: * = 10%, ** = 5%, *** = 1%.

Table A-3. The Effect of Defined Benefit Pensions on Retention by Discount Rate

	Men			Women		
Worker's discount rate:	5%	6%	7%	5%	6%	7%
Average Effects of a 10 Percent Decrease in the FERS Annuity						
dQuarters of job tenure	-0.85	-0.09	0.39	-1.99 ***	-1.38 ***	-0.78 ***
	(0.62)	(0.51)	(0.45)	(0.36)	(0.31)	(0.25)
dPr(Resigns) x 100	0.03 ***	0.01	0.01	0.07 ***	0.05 ***	0.03 ***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
dPr(Retires) x 100	-0.78 ***	-0.87 ***	-0.92 ***	-0.47 ***	-0.51 ***	-0.61 ***
	(0.18)	(0.17)	(0.17)	(0.13)	(0.13)	(0.12)
Probit Coefficient Estimates x 100						
Defined benefit pension						
Option value/earnings at 62	-9.64 ***	-9.71 ***	-9.66 ***	-11.46 ***	-12.90 ***	-13.32 ***
	(1.62)	(1.89)	(2.19)	(1.06)	(1.25)	(1.44)
Accrued value/earnings at 62	5.15 ***	6.15 ***	7.07 ***	3.15 ***	3.59 ***	4.51 ***
	(1.12)	(1.19)	(1.29)	(0.77)	(0.81)	(0.85)
1(1st year eligible for employer-	26.20 ***	26.62 ***	27.21 ***	25.79 ***	25.23 ***	25.17 ***
provided retiree health insurance)	(3.95)	(3.96)	(3.95)	(2.63)	(2.65)	(2.66)
Other Statistics						
Log likelihood per observation	-0.032375	-0.032371	-0.032376	-0.049561	-0.049556	-0.049565
Sample size	502,115	502,201	502,115	586,411	586,365	586,411

Source: Congressional Budget Office, using data from the Office of Personnel Management.

All specifications are estimated by probit and include quadratics in age and job tenure, indicators for the early and full retirement ages under Social Security, and an intercept, each of which is interacted with the worker's retirement plan. The average effects are based on a 10 percent decrease in annuity payments and calculated over the projected work lives of the cohort in the Federal Employees Retirement System (FERS). The estimators for the standard errors are robust to heteroskedasticity and dependent sampling by worker, and the estimates are reported in parentheses. The standard errors for the average effects are calculated using 1,000 bootstrapped samples.

Significance levels: * = 10%, ** = 5%, *** = 1%.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method," *American Journal of Political Science*, vol. 59, no. 2 (April), pp. 495–510, http://dx.doi.org/10.1111/ajps.12116.

Asch, Beth, Steven J. Haider, and Julie Zissimopoulos. 2005. "Financial Incentives and Retirement: Evidence From Federal Civil Service Workers," *Journal of Public Economics*, vol. 89, no. 2–3 (February), pp. 427–440, http://dx.doi.org/10.1016/j.jpubeco.2003.12.006.

Athey, Susan, and others. 2017. *Matrix Completion Methods for Causal Panel Data Models*, Working Paper 3630 (Stanford Graduate School of Business, November), https://tinyurl.com/yacf5e74.

Athey, Susan, and Guido W. Imbens. 2017. "The State of Applied Econometrics: Causality and Policy Evaluation," *Journal of Economic Perspectives*, vol. 31, no. 2 (Spring), pp. 3–32, http://dx.doi.org/10.1257/jep.31.2.3.

Black, Matthew, Robert Moffitt, and John T. Warner. 1990. "The Dynamics of Job Separation: The Case of Federal Employees," *Journal of Applied Econometrics*, vol. 5, no. 3 (July–September), pp. 245–262, http://doi.org/10.1002/jae.3950050304.

Brainard, Keith, and Alex Brown. 2016. *Significant Reforms to State Retirement Plans* (National Association of State Retirement Administrators, June), www.nasra.org/content.asp?admin=Y&contentid=219.

Coile, Courtney, and Jonathan Gruber. 2007. "Future Social Security Entitlements and the Retirement Decision," *The Review of Economics and Statistics*, vol. 89, no. 2 (May), pp. 234–246, http://dx.doi.org/10.1162/rest.89.2.234.

Congressional Budget Office. 2017. *Options for Changing the Retirement System for Federal Civilian Workers* (August), www.cbo.gov/publication/53003.

Doudchenko, Nikolay, and Guido W. Imbens. 2016. *Balancing, Regression, Difference-In-Differences and Synthetic Control Methods: A Synthesis*, Working Paper 22791 (National Bureau of Economic Research, October), www.nber.org/papers/w22791.

Falk, Justin R. 2015. "Comparing Federal and Private-Sector Wages Without Logs," *Contemporary Economic Policy*, vol. 33, no. 1 (January), http://dx.doi.org/10.1111/coep.12061.

Falk, Justin R. 2012. "Comparing Benefits and Total Compensation Between Similar Federal and Private-Sector Workers," *The B.E. Journal of Economic Analysis & Policy*, vol. 12, no. 1, http://dx.doi.org/10.1515/1935-1682.3256.

Federal Thrift Retirement Investment Board. 2008. *2008 TSP Participant Survey Results* (prepared by Watson Wyatt Worldwide), https://go.usa.gov/xQfKA (477 KB).

Fitzpatrick, Maria D. 2015. "How Much Are Public School Teachers Willing to Pay for Their Retirement Benefits?" *American Economic Journal: Economic Policy*, vol. 7, no. 4 (November), pp. 165–188, http://dx.doi.org/10.1257/pol.20140087.

Fitzpatrick, Maria D. 2014. "Retiree Health Insurance for Public School Employees: Does It Affect Retirement?" *Journal of Health Economics*, vol. 38, no. 4 (December), pp. 88–98, http://dx.doi.org/10.1016/j.jhealeco.2014.03.009.

Friedberg, Leora, and Anthony Webb. 2005. "Retirement and the Evolution of Pension Structure," *The Journal of Human Resources*, vol. 40, no. 2 (Spring), pp. 281–308, http://dx.doi.org/10.3368/jhr.XL.2.281.

Gittleman, Maury, and Brooks Pierce. 2012. "Compensation for State and Local Government Workers," *Journal of Economic Perspectives*, vol. 26, no. 1 (Winter), pp. 217–242, http://dx.doi.org/10.1257/jep.26.1.217.

Gustman, Alan, and Thomas Steinmeier. 1986. "A Structural Retirement Model," *Econometrica*, vol. 54, no. 3 (May), pp. 555–584, http://dx.doi.org/10.2307/1911308.

Hanushek, Eric A., John F. Kain, Steven G. Rivkin. 2004. "Why Public Schools Lose Teachers," *The Journal of Human Resources*, vol. 39, no. 2 (Spring), pp. 326–354, http://dx.doi.org/10.3368/jhr.XXXIX.2.326.

Hendricks, Matthew D. 2014. "Does It Pay to Pay Teachers More? Evidence From Texas," *Journal of Public Economics*, vol. 109, no. 1 (January), pp. 50–63, http://dx.doi.org/10.1016/j.jpubeco.2013.11.001.

Ippolito, Richard A. 2002. "Stayers as 'Workers' and 'Savers': Toward Reconciling the Pension-Quit Literature," *The Journal of Human Resources*, vol. 37, no. 2 (Spring), pp. 275–308, http://dx.doi.org/10.2307/3069648.

Krueger, Alan B. 1988. "The Determinants of Queues for Federal Jobs," *Industrial and Labor Relations Review*, vol. 41, no. 4 (July), pp. 567–581, http://dx.doi.org/10.1177/001979398804100406.

Martin, Patricia P. 2003/2004. "Comparing Replacement Rates Under Private and Federal Retirement Systems," *Social Security Bulletin*, vol. 65, No. 1 (Social Security Office of Policy), https://go.usa.gov/xRk38.

Munnell, Alicia H., and Jean-Pierre Aubry. 2016. *The Funding of State and Local Pensions: 2015–2020*, Issue Brief 50 (Center for Retirement Research at Boston College, June), http://crr.bc.edu/briefs/the-funding-of-state-and-local-pensions-2015-2020.

Murnane, Richard J., and Randall J. Olsen. 1990."The Effects of Salaries and Opportunity Costs on Length of Stay in Teaching: Evidence From North Carolina," *The Journal of Human Resources*, vol. 25, no. 1 (Winter), pp. 106–124, http://dx.doi.org/10.2307/145729.

Ni, Shawn, and Michael Podgursky. 2016. "How Teachers Respond to Pension System Incentives: New Estimates and Policy Applications," *Journal of Labor Economics*, vol. 34, no. 4 (October), pp. 1075–1104, http://dx.doi.org/10.1086/686263.

Novy-Marx, Robert, and Joshua D. Rauh. 2009. "The Liabilities and Risk of State-Sponsored Pension Plans," *Journal of Economic Perspectives*, vol. 23, no. 4 (Fall), pp. 191–210, http://dx.doi.org/10.1257/jep.23.4.191.

Office of Personnel Management. 2017. *Civil Service Retirement and Disability Fund Annual Report, Fiscal Year Ended September 30, 2016* (February), https://go.usa.gov/xRQ56.

Song, Jae, and Joyce Manchester. 2007. "Have People Delayed Claiming Retirement Benefits? Responses to Changes in Social Security Rules," *Social Security Bulletin*, vol. 67, no. 2, pp. 1–23, https://www.ssa.gov/policy/docs/ssb/v67n2/v67n2p1.html.

Stock, James H., and David A. Wise. 1990. "Pensions, the Option Value of Work, and Retirement," *Econometrica*, vol. 58, no. 5 (September), pp. 1151–1180, http://dx.doi.org/10.2307/2938304.

Warner, John T., and Matthew S. Goldberg. 1984. "The Influence of Non-Pecuniary Factors on Labor Supply: The Case of Navy Enlisted Personnel," *The Review of Economics and Statistics*, vol. 66, no. 1 (February), pp. 26–35, http://dx.doi.org/10.2307/1924692.

Warner, John T., and Saul Pleeter. 2001. "The Personal Discount Rate: Evidence From Military Downsizing Programs," *The American Economic Review*, vol. 91, no. 1 (March), pp. 33–55, http://www.jstor.org/stable/2677897.