

Does Science Advance One Funeral at a Time?

Pierre Azoulay

MIT and NBER

Sloan School of Management

100 Main Steet – E62-487

Cambridge, MA 02142

USA

Christian Fons-Rosen

Universitat Pompeu Fabra and CEPR

Barcelona GSE

Carrer Ramon Trias Fargas, 25-27

08005 Barcelona

Spain

Joshua S. Graff Zivin

UCSD and NBER

School of Global Policy & Strategy

9500 Gilman Drive

La Jolla, CA 92093

USA

September 20, 2016

Abstract

We study the extent to which eminent scientists shape the vitality of their fields by examining entry rates into the fields of 452 academic life scientists who pass away while at the peak of their scientific abilities. Key to our analyses is a novel way to delineate boundaries around scientific fields by appealing solely to intellectual linkages between scientists and their publications, rather than collaboration or co-citation patterns. Consistent with previous research, the flow of articles by collaborators into affected fields decreases precipitously after the death of a star scientist (relative to control fields). In contrast, we find that the flow of articles by non-collaborators increases by 8% on average. These additional contributions are disproportionately likely to be highly cited. They are also more likely to be authored by scientists who were not previously active in the deceased superstar’s field. Overall, these results suggest that outsiders are reluctant to challenge leadership within a field when the star is alive and that a number of barriers may constrain entry even after she is gone. Intellectual, social, and resource barriers all impede entry, with outsiders only entering subfields that offer a less hostile landscape for the support and acceptance of “foreign” ideas.

Keywords: economics of science, scientific fields, superstars, invisible college, cumulative knowledge production.

*Address all correspondence to pazoulay@mit.edu. Azoulay and Graff Zivin acknowledge the financial support of the National Science Foundation through its SciSIP Program (Award SBE-1460344). Fons-Rosen acknowledges financial support by the Spanish Ministry of Economy and Competitiveness (ECO2014-55555-P). Mikka Rokkanen provided additional research assistance. The project would not have been possible without Andrew Stellman’s extraordinary programming skills (www.stellman-greene.com). We thank Heidi Williams, Sameer Srivastava, Scott Stern, Bruce Weinberg, and seminar audiences at the NBER, UC Berkeley, National University of Singapore, and Stanford University for useful discussions. The usual disclaimer applies.

“A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it.”

MAX PLANCK

Scientific Autobiography and Other Papers

1 Introduction

Knowledge accumulation—the process by which new research builds upon ideas developed in prior research—has been long understood to be of central importance to scientific progress and economic growth (Mokyr 2002). In deference to Sir Isaac Newton, this cumulative process is often referred to as “standing on the shoulders of giants,” but is conceptualized more prosaically as the way in which researchers in one generation learn from and build upon prior research. Yet the literature is largely silent on the mechanisms that shape this slowly evolving process.¹ To borrow terminology from the economic pioneers in the field (Nelson 1962), we know far more about the determinants of the *rate* than that of the *direction* of scientific progress.

What guides researchers when choosing between various approaches to study a given problem? Does science evolve according to autonomous laws, or is the direction of science influenced by individuals, incentives, and institutions? Philosophers and historians have long debated the extent to which the pragmatic success of a scientific theory determines how quickly it gains adherents, or its longevity (e.g., Kuhn [1970], Laudan [1977], and their many detractors). The epigraph of this paper encapsulates the jaundiced view, attributed to Planck, that the idiosyncratic stances of individual scientists can do much to alter, or at least delay, the course of scientific advance. Yet, the proposition that established scientists are slower than younger ones in accepting novel theories has received little empirical support whenever it has been put to the test (Hull et al. 1978; Gorham 1971; Levin et al. 1995). Moreover, in contrast to technology development where market forces shape the direction of research effort (however imperfectly, cf. Acemoglu [2012]), the choice of a problem-solving

¹This stands in contrast to “paradigm shifts” (Kuhn 1970), which are exceedingly rare but garner far more scholarly attention. Bramoullé and Saint-Paul (2010) provide an equilibrium model of scientific revolutions with a Kuhnian flavor.

approach in basic research is less informed by market signals, and thus necessarily depends on a more nuanced system of non-pecuniary incentives (Feynman 1999; Foster et al. 2015).

In this paper, we test “Planck’s Principle” by examining how the death of 452 eminent academic life scientists alter the vitality (measured by publication rates and funding flows) of the subfields in which these scientists actively published in the years immediately preceding their premature passing. In contrast with prior work that focused on collaborators (Azoulay et al. 2010; Oettl 2012; Jaravel et al. 2015), our work leverages new tools to define scientific subfields in order to provide the first evidence on the response by non-collaborators. To our surprise, it is not competitors from within the field that assume the mantle, but rather the entry of outsiders that step in to fill the void created by a stars absence. Importantly, this surge in contributions from outsiders draws upon a different scientific corpus and is disproportionately likely to be highly cited. Thus, consistent with the contention by Planck, the loss of a luminary provides an opportunity for fields to evolve in new directions that advance the frontier of knowledge within them. The rest of the manuscript tries to elucidate the mechanisms responsible for this phenomenon.

While it is implausible that the stars who passed away exerted direct control over entry into their fields, since only a vanishingly small number were journal editors or members of NIH study sections, we do find evidence for several forms of indirect control. Deterrence appears to be largely driven by a reluctance to challenge particularly prominent or committed scholars in the field while they are alive. Even after a field has lost its shining star, entry can be regulated by key collaborators left behind. This is particularly true in fields that have coalesced around a narrow set of techniques or ideas or where collaboration networks are particularly tight-knit. Entry is also deterred when key collaborators of the star are in a position to channel resources (such as editorial goodwill or funding) to insiders. Though stars may have been a source of dynamism while alive, the turnover in leadership enables the injection of fresh ideas into the subfield, but only in those areas whose topology offers a less hostile landscape for the support and acceptance of “foreign” ideas.

To our knowledge, this manuscript is the first to examine the dynamics of scientific evolution using the standard empirical tools of applied microeconomics.² We conceptualize

²Considerable work outside of economics has examined the evolution of scientific fields through data visualization techniques (cf. Chavalari and Cointet (2013) for a recent example). While interesting, this work has been largely descriptive and mostly silent regarding the behavioral mechanisms that might explain the birth, fusion, split, or death of scientific fields.

the death of eminent scientists as shocks to the structure of the intellectual neighborhoods in which they worked several years prior to their death, and implement a procedure to delineate the boundaries of these neighborhoods in a way that is scalable, transparent, and does not rely on ad hoc human judgment. The construction of our dataset relies heavily on the *PubMed Related Citations Algorithm* [PMRA], which groups scientific articles into subfields based on their intellectual content using very detailed keyword information as well as the relative frequencies of these keywords in the scientific corpus.³ As such we are able to define circumscribed areas of scientific inquiry that are independent of training, personal relations, or self-proclaimed areas of expertise.

In addition to providing evidence regarding a central question for scholars studying the scientific process, our paper is a departure for the field of the economics of science in that it can attend to the ways in which scientists position themselves simultaneously in an intellectual space as well as a social space, whose boundaries do not overlap (Borjas and Doran 2015). As such, our work can be understood as integrating the traditional concerns of economists—understanding how incentives and institutions influence the rate of knowledge production or diffusion—with those of cognate disciplines such as sociology and philosophy, who have traditionally taken the direction of scientific change as the central problem to be explained.

The rest of the paper proceeds as follows. In the next section, we examine the institutional context and lay out our broad empirical strategy. In section 3, we then turn to data, methods and descriptive statistics. We report the results in section 4. Section 5 concludes by outlining the implications of our findings for future work.

2 Institutional Context and Empirical Design

Our empirical analyses are centered on the academic life sciences. The merits of this focus are several fold. First, the field has been an important source of scientific discovery over the past half century. Many modern medical therapies can trace their origins to research conducted in academic laboratories (Sampat and Lichtenberg, 2011). These discoveries, in turn, have generated enormous health and welfare gains for economies around the world.

³Unlike in economics, keywords for all publications indexed by *PubMed* (most of the life sciences) are assigned by staff at the National Library of Medicine and are drawn from a controlled vocabulary thesaurus. Thus, concerns about strategic or endogenous keyword choices are minimized in this setting.

Second, the life science research workforce is exceedingly large and specialized. Academic medical centers in the United States employ 150,000 faculty members. Moreover, scientific discoveries over the past half-century have greatly expanded the knowledge frontier, necessitating increasing specialization by researchers and a greater role for collaboration (Jones 2009). If knowledge and techniques remain at least partially tacit long after their initial discovery, tightly-knit research teams may be able to effectively control entry into intellectual domains. The size and maturity of this sector, including its extensive variety of narrowly-defined subfields, makes it an ideal candidate for an inquiry into the determinants of the direction of scientific effort in general, and how it is influenced by elite scientists in particular.

Third, the academic research setting also offers the practical benefits of an extensive paper trail of research inputs, outputs, and collaboration histories. On the input side, reliance of researchers on one agency for the majority of their funding raises the possibility that financial gatekeeping by elite scientists could be used to regulate entry into scientific fields. Data on NIH funding at the individual level, as well as membership in “study sections” (the peer-review panels that evaluate the scientific merits of grant applications) will allow us to examine such concerns directly. Most importantly for our study, the principal output of researchers—publications—are all indexed by a controlled vocabulary of keywords managed by the National Library of Medicine. This provides the raw material that allows us to define scientific subfields in a way that is stripped of “social baggage” (the specifics of this process will be described in detail in Section 2.2).

Lastly, while accounts by practicing scientists indicate that collaboration plays a large role in both the creation and diffusion of new ideas (Reese 2004), historians of science have long debated the role of controversies and competition in shaping the direction of scientific progress and the process through which new subfields within the same broad scientific paradigm are born and grow over time (Hull 1989; Morange 1999; Shwed and Bearman 2010). Our study presents a unique opportunity to test some of their insights in a way that is more systematic and can yield generalizable insights on the dynamics of field evolution.

3 Empirical Design, Data, and Descriptive Statistics

Below, we provide a detailed description of the process through which the matched scientist/subfield dataset used in the econometric analysis was assembled. We begin by describing the criteria used to select our sample of superstar academics, with a particular focus on “ex-

inction events”; the set of subfields in which these scientists were active prior to their death and the procedure followed to delineate their boundaries. Finally, we discuss the matching procedure implemented to identify control subfields associated with eminent scientists who did not pass away but are otherwise similar to our treatment group.

3.1 Superstar sample

Our basic approach is to rely on the death of “superstar” scientists as a lever to estimate the extent to which the production of knowledge in the fields in which they were active changes after their passing. The study’s focus on the scientific elite can be justified both on substantive and pragmatic grounds. The distribution of publications, funding, and citations at the individual level is extremely skewed (Lotka 1926; de Solla Price 1963) and only a tiny minority of scientists contribute, through their published research, to the advancement of science (Cole and Cole 1972). Stars also leave behind a corpus of work and colleagues with a stake in the preservation of their legacy, making it possible to trace back their careers, from humble beginnings to wide recognition and acclaim.

The elite academic life scientist sample includes 12,935 individuals, which corresponds to roughly 5% of the entire relevant labor market. In our framework, a scientist is deemed elite if they satisfy at least one of the following criteria for cumulative scientific achievement: (1) highly funded scientists; (2) highly cited scientists; (3) top patenters; and (4) members of the National Academy of Sciences or of the Institute of Medicine. Since these four criteria are based on extraordinary achievement over an entire scientific career, we augment this sample using additional criteria to capture individuals who show great promise at the early and middle stages of their scientific careers. These include: (5) NIH MERIT awardees; (6) Howard Hughes Medical Investigators; and (7) early career prize winners. Appendix I provides additional details regarding these seven metrics of “superstardom.”

For each scientist in the sample, we reconstruct their career from the time they obtained their first position as independent investigators (typically after a postdoctoral fellowship) until 2006. Our dataset includes employment history, degree held, date of degree, gender, and department affiliations as well as complete list of publications, patents and NIH funding obtained in each year by each scientist.⁴

⁴Though we apply the term of “superstar” to the entire group, there is substantial heterogeneity in intellectual stature within the elite sample (see Table 1).

The 452 scientists who pass away prematurely, and who are the particular focus of this paper, constitute a subset of this larger pool of 12,935. To be included in our sample, their deaths must intervene between 1975 and 2003 (this allows us to observe at least 3 years' worth of scientific output for every subfield after the death of a superstar scientist). Although we do not impose any age cutoff, the median and mean age at death is 61 with 85% of these scientists having passed away before the age of 70 (we will explore the sensitivity of our results to the age at death later). We also require evidence, in the form of published articles and/or NIH grants, that these scholars were still in a scientifically active phase of their career in the period just preceding their death (this is the narrow sense in which we deem their deaths to have occurred prematurely).

Within this sample, 229 (51%) of these scientists pass away after a protracted illness, whereas 185 (41%) die suddenly and unexpectedly. We were unable to ascertain the particular circumstances of 37 (8.20%) death events.⁵ Appendix G provides the full list of extinct superstars, together with their year of birth, year of death, institutional affiliation at the time of their passing, and a short description of their research expertise.

Table I provides descriptive statistics for the extinct superstar sample. The median star received his degree in 1957, died at 61 years old and was associated with 4 distinct subfields in the five years leading up to his/her death. On the output side, the stars each received an average of roughly 16.6 million dollars in NIH grants, and published 138 papers that garnered 8,347 citations over the course of their careers (as of early 2014).

3.2 Delineating Research Fields

The source of the publication data is *PubMed*, an online resource from the National Library of Medicine that provides fast, free, and reliable access to the biomedical research literature. *PubMed* indexes more than 40,000 journals within the life sciences.

To delineate the boundaries of the research fields in which each deceased star was active, we develop an approach based on topic similarity as inferred by the overlap in keywords between each article the star published in the five years prior to his/her death, and the rest of the scientific literature. Specifically, we use the *PubMed Related Citations Algorithm*

⁵We exclude from the sample one scientist who took his own life, and a further two for whom suicide could not be ruled out.

(PMRA) which relies heavily on Medical Subject Headings (MeSH). MeSH terms constitute a controlled vocabulary maintained by the National Library of Medicine that provides a very fine-grained partition of the intellectual space spanned by the biomedical research literature. Importantly for our purposes, MeSH keywords are assigned to each scientific publication by professional indexers and not by the authors themselves.⁶ We then use the “Related Articles” function in *PubMed* to harvest journal articles that are intellectually proximate to star scientists’ own papers.⁷

To fix ideas, consider “The transcriptional program of sporulation in budding yeast” [PubMed ID 9784122], an article published in the journal *Science* in 1998 originating from the laboratory of Ira Herskowitz, an eminent UCSF biologist who died in 2003 from pancreatic cancer. As can be seen in Figure I, PMRA returns 72 original related journal articles for this source publication.⁸ Some of these intellectual neighbors will have appeared before the source to which they are related, whereas others will have only been published after the source. Some will represent the work of collaborators, past or present, of Herskowitz’s, whereas others will represent the work of scientists in his field he may never have come in contact with during his life, much less collaborated with. The salient point is that nothing in the process through which these related articles are identified biases us towards (or away from) articles by collaborators, frequent citers of Herskowitz’s work, or co-located researchers. Rather, the only determinants of relatedness are to be found in the overlap in MeSH keywords between the source and its potential neighbors.

Consider now the second most-related article to Herskowitz’s *Science* paper listed in Figure I, “Phosphorylation and maximal activity of *Saccharomyces cerevisiae* meiosis-specific transcription factor Ndt80 is dependent on Ime2.” Figure C1 in Appendix C displays the MeSH terms that tag this article along with its source. As a byproduct, PMRA also provides a cardinal dyadic measure of intellectual proximity between each related article and its

⁶The algorithm also uses as inputs title and abstract words, which are obviously selected by authors, rather than by NLM staff. However, neither the choice of MeSH keywords nor the algorithm depend on cited references contained in publications.

⁷To facilitate the harvesting of *PubMed*-related records on a large scale, we have developed an open-source software tool that queries *PubMed* and PMRA and stores the retrieved data in a MySQL database. The software is available for download at <http://www.stellman-greene.com/FindRelated/>.

⁸Why exactly 72? In fact, PMRA lists 152 “intellectual neighbors” for PubMed ID 9784122. But once we exclude articles published after 2006 (the end of our observation period), purge from the list reviews, editorials and other miscellaneous “non-original” content, and drop a handful of articles that appeared in minor journals not indexed in Thomson-Reuter’s *Web of Science*, the number of publications associated with this source article indeed drops to 72. Appendix C provides more details on the rules that govern the cut-off for the number of articles returned by PMRA for any given source article.

associated source article. In this particular instance, the relatedness score of “Phosphorylation...” is 94%, whereas the relatedness score for the most distant related article in Figure I, “Catalytic roles of yeast...” is only 62%.

In the five years prior to his death (1998-2002), Herskowitz was the last author on 12 publications.⁹ For each of these publications, we treat the set of publications returned by PMRA as constituting a distinct subfield, and we create a star/field panel dataset by counting the number of related articles in each of these subfields in each year between 1975 and 2006. An important implication of this data construction procedure is that the subfields we delineate are quite limited in scope. One window into the degree of intellectual breadth for subfields is to gauge the overlap between the articles that constitute any pair of subfields associated with the same star. In the sample, the 452 deceased stars account for 3,074 subfields, and 21,633 pairwise combination of subfields (we are only considering pairs of subfields associated with the same individual star). Figure II displays the histogram for the distribution of overlap, which is extremely skewed. A full half of these pairs exhibit exactly zero overlap, whereas the mean of the distribution is 0.06. To find pairs of subfields that display substantial amounts of overlap (for example, half of the articles in subfield 1 also belong in subfield 2), one must reach far into the right tail of the distribution, specifically, above the 98th percentile.

As such, the subfields we delineate are relatively self-contained. Performing the analysis at the level of the subfield-star combination—rather than lumping together all the subfields of an individual star—will provide us with an opportunity to exploit variation in the extent of participation of the star within each of his/her subfields. We will also check the validity of the main results when rolling the data up from the subfield-star level to the star-level. Finally, since even modest amounts of overlap entail that the observations corresponding to the subfields of individual stars will not be independent in a statistical sense, we will cluster standard errors at the level of the star scientist.

⁹A robust social norm in the life sciences systematically assigns last authorship to the principal investigator, first authorship to the junior author who was responsible for the conduct of the investigation, and apportions the remaining credit to authors in the middle of the authorship list, generally as a decreasing function of the distance from the extremities of the list.

3.3 Identification Strategy

Given our interests in the effect of superstar death on entry into scientific subfields, our empirical strategy is focused on changes in published research output after the superstar passes away, relative to when s/he was still alive. To ensure that we are estimating the effect of interest and not some other influence that is correlated with the passage of time, our specifications include age and period effects, as is the norm in studies of scientific productivity (Levin and Stephan 1991). These temporal controls are tantamount to using subfields that lost a superstar in earlier or later periods as an implicit control when estimating entry into subfields that currently experienced the death of a superstar. If the death of a superstar only represented a one-time shift in the level of entry into the relevant subfields, this would not be problematic. But if these unfortunate events affect trends—and not simply levels—of scientific activity, this approach may not suffice to filter out the effect of time-varying omitted variables, even when flexible age and calendar time controls are included in the econometric specification. One tangible concern about time-varying effects relates to the life-cycle of subfields, where productive potential may initially increase over time before peaking and then slowly declining.

To mitigate this threat to identification, our preferred empirical strategy relies on the selection of a matched scientist/subfield for each treated scientist/subfield. These control observations are culled from the universe of subfields in which superstars who do not die are active (see Section 2.1 and Appendix D). Combining the treated and control samples enables us to estimate the effect of superstar death in a difference-in-differences framework. Figure III illustrates the procedure used to identify control subfields in the particular case of the Herskowitz publication highlighted above.

We begin by looking at all the articles that appeared in the same journal and in the same year as the treated source articles. From this set of articles, we keep only those that have one of the still-living superstars in the last authorship position. Then, using a “coarsened exact matching” procedure detailed in Appendix D, the control source articles are selected such that (1) the number of authors in the treated and control are approximately similar; (2) the age of the treated and control superstars differ by no more than five years; and (3) the number of citations received by the treated and source article are similar. For the Herskowitz/“sporulation in budding yeast” pair, we can select 10 control articles in this way. All of these controls were also published in *Science* in 1998, and have between five and

seven authors. One of these controls is “Hepatitis C Viral Dynamics in Vivo...,” whose last author is Alan Perelson, a biophysicist at Los Alamos National Lab. Perelson and Herskowitz obtained their PhD only a year apart. The two papers had received 514 and 344 citations respectively by the end 2003. Though this is a large difference, this places both well above the 99th percentile of the citation distribution for 5-year old articles published in 1998.

One potential concern with the addition of this “explicit” control group is that control subfields could be affected by the treatment of interest. What if, for instance, a control source article happens to be related (in a PMRA sense) with the treated source? Because the subfields identified by PMRA are narrow, this turns out to be an infrequent occurrence. Nonetheless, we remove all such instances from the data. We then find all the intellectual neighbors for these control source articles using PMRA; a control subfield is defined by the set of related articles returned by PMRA, in a manner that is exactly symmetric to the procedure used to delineate treated subfields. When these related articles are parsed below to distinguish between those published by collaborators vs. non-collaborators of the star, or between those by intellectual outsiders vs. insiders, treated and control observations will always be defined with perfect symmetry.

3.4 Descriptive Statistics

The procedure described above yields a total of 34,216 distinct subfields; 3,074 subfields correspond to one of the 452 dead scientists, whereas 31,142 subfields correspond to one of 5,809 still-living scientists. Table II provides descriptive statistics for control and treated subfields in the baseline year, i.e., the year of death for the deceased scientist.¹⁰

Covariate balance. In the list of variables displayed in Table II, it is important to remember that a number of covariates are balanced between treated and control subfields solely by virtue of the coarsened exact matching procedure—for instance, (star) investigator year of degree, the source article number of authors, or the source article number of citations at baseline.

However, there is nothing mechanical to explain the balance between treated and control subsamples with respect to the stock of our main outcome variable: the number of articles

¹⁰We can assign a counterfactual year of death for each control subfield, since each control subfield is associated with a particular treated subfield through the matching procedure described above.

in the star’s field. Figure IV compares the corresponding distribution and also shows a great deal of overlap between the two histograms. Of course, balance in the levels of the outcome variable is not technically required for the validity of the empirical exercise.¹¹ Yet, given the ad hoc nature of the procedure used to identify control subfields, this degree of balance is reassuring.

Another happy byproduct of our matching procedure is that treated and control scientists also appear quite similar in the extent of their eminence at the time of (counterfactual) death, whether such eminence is measured through NIH funding, the number of articles published, or the number of citations these articles received.

Collaborators vs. non-collaborators. One critical aspect of the empirical analysis is to distinguish between collaborators and non-collaborators of the star when measuring publishing activity in a subfield. It is therefore crucial to describe how this distinction can be made in our data. Information about the superstars’ colleagues stems from the Faculty Roster of the Association of American Medical Colleges (AAMC), to which we secured licensed access for the years 1975 through 2006, and which we augmented using NIH grantee information (cf. Azoulay et al. [2010] for more details).

An important implication of our reliance on these sources of data is that we can only identify authors who are faculty members in U.S. medical schools, or recipient of NIH funding. We cannot systematically identify trainees (at least not until they secure a faculty position), scientists working for industrial firms, or scientists employed in foreign academic institutions. The great benefit of using AAMC data, however, is that they ensure we have at our disposal both demographic and employment information for every individual in the relevant labor market: their (career) age, type of degree awarded, place of employment, gender, and research output, whether measured by publications or NIH grants.

To identify authors, we match the authorship roster of each related article in one of our subfields with the AAMC roster.¹² We tag as a collaborator any author who appeared as an author on a publication prior to the death with the star associated with the subfield. Each

¹¹What is required is that the trends in publication activity be comparable between treated and control subfields up until the death of the treated scientist. We verify that this is the case below.

¹²We limit ourselves to authors with relatively infrequent names. Though this may create some measurement error, there is no reason to suspect that the wrongful attribution of articles to authors will impact treated and control subfields in a differential way.

related article is therefore assigned to one of two mutually-exclusive bins: the “collaborator” bin comprises the set of publications with at least one identified author who coauthored with the star prior to the year of death (or counterfactual death); the “non-collaborator” bin comprises the set of publications with no identified author who coauthored with the star prior to the year of death (or counterfactual death). As can be seen in Table II, roughly 15% of the publication activity at baseline can be accounted for by collaborators. Moreover, this proportion is very similar for control and treated subfields.¹³

4 Results

The exposition of the econometric results proceeds in stages. After a brief review of methodological issues, we provide results that pertain to the main effect of superstar death on subfield growth, measured by publication rates and funding flows. Second, we attempt to elucidate the mechanism (or set of mechanisms) at work to explain our most robust finding, that of relative subfield growth in the wake of a star’s passing, a growth entirely accounted for by contributions from non-collaborators. We do so by examining the characteristics of the articles published by non-collaborators, before turning to the characteristics of their authors. We also explore heterogeneity in the treatment effect through the interaction of the post-death indicator variable with various attributes of the stars.

4.1 Econometric Considerations

Our estimating equation relates publication or funding activity in subfield i in year t to the treatment effect of losing superstar j :

$$E[y_{it}|X_{ijt}] = \exp[\beta_0 + \beta_1 AFTER_DEATH_{jt} + f(AGE_{it}) + \delta_t + \gamma_{ij}] \quad (1)$$

where y is a measure of activity, $AFTER_DEATH$ denotes an indicator variable that switches to one in the year during which superstar j passes away, $f(AGE_{it})$ corresponds to a flexible function of the field’s age, the δ_t ’s stand for a full set of calendar year indicator variables, and the γ_{ij} ’s correspond to subfield/star fixed effects, consistent with our approach to analyze *changes* in activity within subfield i following the passing of superstar j .

¹³We define collaboration status by looking at the authorship roster for the entire corpus of work published by the star before or in the year of death, and not only with respect to the articles of the star that belong to the focal subfield.

The subfield fixed effects control for many time-invariant characteristics that could influence research activity, such as the need for capital equipment or the extent of disease burden (e.g., for clinical fields). A pregnant metaphor for the growth of scientific knowledge has been that of biological evolution (Hull 1989; Chavalarias and Cointet 2013): a field is born when new concepts are introduced, resulting in an accelerating production of “offsprings” (articles), until the underlying scientific community loses its thematic coherence, ushering in an era of decline (or alternatively, splitting or merging events). To flexibly account for such life cycle effects, we include subfield age indicator variables, where subfield age is computed as the number of years since the year of publication for the underlying article.¹⁴ The calendar year effects filter out the effects of the general expansion of the scientific enterprise as measured by the number of journals and articles published each year.¹⁵

Estimation. The dependent variables of interest, including publication counts and NIH grants awarded, are skewed and non-negative. For example, 31.40% of the subfield/year observations in the data correspond to years of no publication activity; the figure climbs to 56.70% if one focuses on the count of NIH grants awarded. Following a long-standing tradition in the study of scientific and technical change, we present conditional quasi-maximum likelihood estimates based on the fixed-effect Poisson model developed by Hausman et al. (1984). Because the Poisson model is in the linear exponential family, the coefficient estimates remain consistent as long as the mean of the dependent variable is correctly specified (Gouriéroux et al. 1984).

Inference. QML (i.e., “robust”) standard errors are consistent even if the underlying data generating process is not Poisson. In fact the Hausman et al. estimator can be used for any non-negative dependent variables, whether integer or continuous (Santos Silva and Tenreiro 2006), as long as the variance/covariance matrix is computed using the outer product of the gradient vector (and therefore does not rely on the Poisson variance assumption). Further, QML standard errors are robust to arbitrary patterns of serial correlation (Wooldridge 1997), and hence immune to the issues highlighted by Bertrand et al. (2004) concerning inference

¹⁴An alternative way to measure subfield age is to date its birth year as the year during which the first related article was published. Though our main results are robust to this alternative parametrization, this is not a desirable way to proceed since it will fail to distinguish subfields that are genuinely long-established from fields that are more recent but happen to have an ancient precursor that PMRA is able to recognize.

¹⁵It is not possible to separately identify calendar year effects from age effects in the “within subfield” dimension of a panel in a completely flexible fashion, because one cannot observe two subfields at the same point in time that have the same age but were born in different years (Hall et al. 2007).

in DD estimation. We cluster the standard errors around superstar scientists in the results presented below.

Dependent Variables. Our primary outcome variable is publication activity in a subfield. However, we go beyond this raw measure by assigning the related articles that together constitute the subfield into a variety of bins. For instance, we can decompose publication activity in the subfield into two mutually exclusive subfields: articles that appear in prestigious journals (Journal Impact Factor [JIF] higher than two) and those that appear in less prestigious journals (JIF lower than two); or articles with a superstar on the authorship roster vs. articles without a superstar; etc. Articles in each bin can then be counted and aggregated up to the subfield/year level.

Capturing funding flows at the field level is slightly more involved. *PubMed* systematically records NIH grant acknowledgements using grant numbers. Unfortunately, these grant numbers are often truncated and omit the grant cycle information that could enable us to pin down unambiguously the particular year in which the grant was awarded. When it is missing, we impute the award year using the following rule: for each related publication that acknowledges NIH funding, we identify the latest year in the three-year window that precedes the publication during which funding was awarded through either a new award or a competitive renewal. To measure funding activity in a subfield, we create a count variable that sums all the awards received in particular year, where these awards ultimately generate publications in the focal subfield.

4.2 Main effect of superstar death

Table III and Figure V present our core results. Overall, we find that publication activity increases slightly following the death of a star scientist who was an active contributor to it, but the magnitude of the effect is not large (about 2%) and imprecisely estimated (column 1). Yet, this result conceals a striking pattern that we uncover when we distinguish between publications by collaborators and non-collaborators. The decline in publication activity accounted for by previous collaborators of the star is massive, on the order of 40% (column 2). This evidence is consistent with our previous findings, which showed that coauthors of superstar scientists who die suffer a drop in output, particularly if their non-collaborative work exhibited strong keyword overlap with the star, i.e., if they were intellectually connected in addition to being coauthors (Azoulay et al. 2010, Table VI, column 2).

A limitation of the previous work focusing on the fate of collaborators after the loss of an eminent scientist always lied in the failure to distinguish between social and intellectual channels of influence, since every treated scientist was by definition a collaborator, even if merely a casual one. In this study, we can relax this constraint, and when we do, we find that publication activity by non-collaborators in the subfield increases by a statistically significant 8.00% (column 3).¹⁶

We also explore the dynamics of the effects uncovered in Table III. We do so by estimating a specification in which the treatment effect is interacted with a set of indicator variables corresponding to a particular year relative to the superstar’s death, and then graphing the effects and the 95% confidence interval around them (Panels A, B, and C of Figure V correspond to columns 1, 2, and 3 in Table III). Two features of the figure are worthy of note. First, the dynamics amplify the previous results in the sense that we see the effects increasing (in absolute value) monotonically over time—there is no indication that the effects we estimated in Table III are merely transitory. Five years after a star’s death, the increase in publication activity by non-collaborators is large enough in magnitude to fully offset the decline in activity by collaborators. Second, there is no discernible evidence of an effect in the years leading up to the death, a finding that validates *ex post* our identification strategy.¹⁷

The last three columns of Table III focus on funding flows from the National Institutes of Health (NIH) rather than publication flows. More precisely, the outcome variable in columns 4, 5, and 6 is the number of distinct NIH awards that acknowledge a publication in the subfield in the three-year window before the year of publication for the related article (counting grant amounts, as opposed to the number of grants, yields similar results). The patterns are very similar to those obtained in the case of publication activity, both in terms of magnitudes and in terms of statistical significance.¹⁸

¹⁶The number of observations varies ever so slightly across columns because the conditional fixed effects specification drops observations corresponding to subfields for which there is no variation in activity over the entire observation period. This is true as well for the results reported in Tables IV through VIII.

¹⁷This finding is reassuring as it suggests that death events are plausibly exogenous to the course of knowledge growth and decline within a subfield. The case for exogeneity is stronger in the case of sudden death than in the case of anticipated death, a distinction that we will examine in more detail below.

¹⁸The event study graphs corresponding to the dynamics of funding flows are available from the authors, but also show close similarity to those displayed in Figure V.

4.3 Robustness checks and extensions

We check the robustness of our main findings in Appendix E. In a difference-in-differences set up sharing many similarities with our own, Jaravel et al. (2015) raise the concern that individual fixed effects, age effects, and year effects might not fully account for the trends in publication flows around the year of a star’s death. Their recommended solution is the inclusion of a full set of leads and lags around star death for both treated and control subfields. These “common” effects can be separately identified from the leads and lags around star death that are specific to the treated fields because (i) death events are staggered over time in the data (rather than happening in a single year as in the typical DD setup); and (ii) control subfields can inherit the date of death of the treated subfield that caused them to enter the dataset. We implement their approach in the first three columns of Table E1 (analogous to the first three columns of Table III) and Figure E1 (analogous to Figure V). The point estimates are very similar to those obtained when not including effects common to both treated and control subfields.¹⁹

Our main results stem from a sample where the total number of articles in a given subfield-year includes the articles published by the star herself (the star is deemed to be her own collaborator). Clearly, part of the decrease in activity by collaborators is mechanically induced by the absence of the star in the post-death years. Yet, this is not enough to explain the decrease in activity by the star’s collaborators. In the last three columns of Table E1, and in Figure E2, activity in the subfield is computed without taking into account any paper that lists the star as an author. Relative to the second column of Table III and Panel B of Figure V, the magnitude of the treatment effect is attenuated (-20.39% vs. -33.77%), but remains large, statistically significant, and permanent, in line with the results presented in Azoulay et al. (2010).

The first three columns of Table E2 and Figure E3 drop from the sample all the control subfields. In these specifications, subfields who were treated in the past or will be treated in the future serve as implicit controls for the subfields currently experiencing the death of their associated star. The results are qualitatively similar to those displayed in Table III and Figure V, though a small upward trend can be discerned in Panel C of Figure E2. This

¹⁹Jaravel et al. (2015) examine the wages and patenting rates for co-inventors of deceased patent inventors—not necessarily eminent ones. As they point out, an inventor must necessarily have invented a patent before the year of death of their co-inventor and is therefore more likely to have been employed at that time, even conditional on a large set of fixed effects. In our setting, it is harder to see why the subfield associated with a departed superstar would mechanically be more active in the years prior to his/her death.

provides a clear reason to add to the specification an additional level of difference—that provided by control subfields. The last three columns of Table E2 and Figure E4 display coefficients estimated by ordinary least squares, rather than the fixed effects Poisson model of Hausman et al. (1984). The results indicate that in steady state, treated fields expand by one additional article per year on average, relative to control fields. The only anomaly presented by this change in estimation method is observed on Figure E4, Panel B. We fail to detect the pronounced downward trend in publication activity on the part of the star’s collaborators, whereas this was a salient feature in Figure V.

One last robustness check performs the main analysis after rolling up the data to the star-level. Because it is difficult to build a control group of deceased stars based solely on star-level covariates, the star-year level dataset does not include aggregates of fields associated with still-alive stars. Figure E5 presents the event-study graph corresponding to publication activity by non-collaborators. We observe a very pronounced upward trend, both before and after the death event (the pre-trend is not precisely estimated, but still relatively large in magnitude). As explained in Section 3.2, we strongly prefer the star/subfield level of analysis, primarily because the subfields delineated by the *PubMed* Related Citations Algorithm exhibit very limited overlap.

Impact of Star Age and Experience. As explained earlier, we do not impose a strict age cutoff for the deceased star, we merely insist that they exhibit tangible signs of research activity, such as publishing original articles (rather than simply reviews, editorials, or comments), obtaining NIH grants, and training students. Among our 452 departed superstars, the median age at death is 61, the seventy-fifth percentile 67, and the top decile 73. How do the core results change when the scientists who passed away at an advanced age are excluded from the sample? As can be observed on Table E3 (which focuses only on publication activity in subfields by non-collaborators of the star), the subfields of stars who passed away more prematurely are driving the bulk of the effect. The effect of the fields associated with older stars is still positive, but imprecisely estimated. We choose to keep these older stars in the sample because a larger sample size affords us opportunities to explore mechanisms without losing power to detect nuanced effects statistically. The last two columns of Table E3 investigate whether a star’s experience in the field (measured as the number of years between her first contribution in it and the year of death) moderates the core result. The median age in the field at the time of death is seven. We find no difference in the magnitude of the treatment effect along this dimension.

Displacement Effects. We find that non-collaborators of the star increase their publication activity in the fields in which the superstar was active prior to her death. Appendix F investigates whether there is evidence of commensurate declines in publication activity for these related authors in the fields where they were active but the star was not. These analyses entail a change in the level of analysis, from the subfield level to the related author level. A practical difficulty is that a related author can be—and is in fact frequently—related to more than a single star. To get around this issue and pin down for each related author a single year of treatment and a clear demarcation between in-field and out-of-field output, we build a panel dataset of related authors and their publication output using two different methods. In the first method, we associate each related author with the star who died (possibly counterfactually) in the earliest year of all possible years of treatment. In the second method, we associate each related author with their most-related star (i.e., the star for whom the cardinal relatedness score between her source article and the author’s related article is highest). Regardless of method, we divide each related author’s output according to whether it belongs to one of the fields of the star with whom s/he is associated, or whether it belongs to none of these fields. Table F1 then examines how these measures of output shift after the death event, relative to before, for treated authors, relative to control authors. We also distinguish between the overall number of publications, and the number of publications falling into various quantiles of the citation distribution.

We present OLS estimates, to ensure that the sample remains identical when examining in-field and out-of-field output. We also display elasticities, together with the mean of the dependent variable to help in comparing magnitudes. Panel A corresponds to the results obtained following the “earliest treating star” method. Panel B corresponds to the results obtained following the “most-related treating star” method. Regardless of the method employed, some stable patterns emerge. We can detect large effects on the rate of production of in-field articles, consistent with the results obtained when performing our analysis at the subfield level. Conversely, the magnitudes for the treatment effect on out-of-field output are typically much smaller, and sometimes imprecisely estimated. Figure F1 presents the corresponding event-study graphs (only for out-of-field publication output). The main takeaway is that we cannot detect any evidence of displacement. Non-collaborating related authors appear to increase their overall output modestly in the wake of a superstar’s premature passing.

4.4 Understanding subfield growth patterns induced by a star’s passing

In the remainder of the manuscript, we seek to understand the mechanisms that might explain the novel empirical regularity we uncovered: that of relative growth for subfields following the death of their superstar anchor, a phenomenon entirely accounted for by research activity undertaken by scientists who never collaborated with the star while alive. As a consequence, all the results below pertain to entry into the field by non-collaborators; any article with even one author who collaborated with the star is excluded from the count of articles that constitute the dependent variable.

Article Characteristics. What characterizes the additional contributions that together lead to increased activity in a subfield after a star has passed on? Are these in fact important contributions to the subfield? Do they focus on core issues, or should they be understood as taking the intellectual domain in a novel direction? Tables IV and V explore these issues. In Table IV, we parse every related article that constitute the subfields in our data to assign them into one of six mutually exclusive bins, based on long-run citation impact: articles that fall in the bottom quartile of the citation distribution; in the second quartile; in the third quartile; articles that fall above the 75th percentile, but below the 95th percentile; articles that fall above the 95th percentile, but below the 99th percentile; articles that fall above the 99th percentile of the citation distribution.²⁰

Panel A of Table IV produces a battery of estimates corresponding to each of these six bins in columns 2 through 7 (column 1 simply replicates the effect for all papers, regardless of impact, that was previously displayed in Table III, column 3). A startling result is that the magnitude of the treatment effect increases sharply as we focus on the rate of contributions with higher impact. In contrast, the number of lower-impact articles contributed by non-collaborators contracts slightly, though the effect is not precisely estimated.

Panels B and C break down these results further by examining separately the growth of subfields by cause of death (anticipated vs. sudden). As mentioned earlier, the case

²⁰Note that when we are referring to the citation distribution, we mean the vintage-specific citation distribution for the universe of articles simultaneously indexed by *PubMed* and the Web of Science. For example, the article by Sopko et al. highlighted on Figure C1 (in Appendix C) received 39 citations from other articles in *PubMed* by 2014. This puts this article above the 76th percentile of the citation distribution for articles published in 2002.

for exogeneity is stronger in the case of sudden death, since when the death is anticipated, it would be theoretically possible for the star to engage in “intellectual estate planning,” whereby particular scientists (presumably close collaborators) are anointed as representing the next generation of leaders in the subfield. The results in column 1 imply that there is an important difference between the two type of events—subfield growth is observed mostly when the death of the star was anticipated. Decomposing this effect across the quantile bins as above reveals that these differences can be accounted for by shifts in activity for low-impact contributions. In the right tail of the distribution, there is very little evidence that the manner of superstar death matters at all for the fate of their subfields. In both cases, non-collaborators increase their contribution sharply—on the order of 40%. Because of this convergence in the upper tail, the remainder of the manuscript will lump together anticipated and unanticipated events.²¹

Table V parses the related articles in each subfield to ascertain whether contributions by non-collaborators constitute a genuine change in intellectual direction. Panel A distinguishes between contributions that are proximate in intellectual space to the source article from those that are more distant (though still part of the subfield as construed by PMRA). Because we have at our disposal both a cardinal and an ordinal measure of intellectual proximity, we present four different estimates. In both cases, the magnitude of the treatment effect pertaining to publication activity by proximate articles is approximately twice as large as the magnitude corresponding to more distant articles. These differences, however, are not themselves statistically significant at conventional levels. But we can at least rule out the conjecture that non-collaborators enter the field from the periphery. Their contributions seem to lie smack-dab in the middle of the subfield as it existed when the star was still alive.

Panel B sheds light on the intellectual direction of the field, by examining the cited references contained in each related article. The first two columns separate related articles in two groups. The first contains only publications that cite at least some work which belongs to the subfield identified by PMRA for the corresponding source. The second contains publications that cite exclusively out of the PMRA subfield. Only articles in the second group appear to experience growth in the post-death era. The next two columns proceed similarly, except that the list of references is now parsed to highlight the presence of articles

²¹The most salient results reported below continue to hold when analyzed separately by cause of death. However, we gain statistical power from pooling these observations, and some empirical patterns would be estimated less precisely if we chose to focus solely on observations corresponding to subfields for which the star died suddenly and unexpectedly.

authored by the star, as opposed to all other authors. We find that subfield growth can be mostly accounted for by articles from non-collaborators who do not build on the work of the star. Finally, we investigate the vintage of the references cited by related articles. The last two columns in Panel B indicate that the new contributions are more likely to build on science of a more recent vintage.

Taken together, the results in Panels A and B of Table V paint a nuanced picture of directional change in the wake of superstar passing. The new contributions do not represent a radical departure from the subfield’s traditional questions—their MeSH keywords overlap with those of the source article even more than is typical for the “average” article in the subfield. At the same time, the citation evidence makes it clear that these additional contributions often draw from more recent and different sources of knowledge for inspiration.

Related Author Characteristics. The next step of the analysis is to investigate the type of scientists who publish the articles that account for subfield growth in the wake of a star’s death. Table VI reports these results. Perhaps the simplest author characteristic is age. For each related article in the subfield, we match the authorship roster to the AAMC Faculty Roster. Then, we compute the mean career age over matched authors for each related article. Since the median career age for matched authors turns out to be 16, we assign each article to one of two bins, the first comprising all related articles with an “older” authorship team (mean author career age greater than 16), the second comprising all related articles with a “younger” authorship team (mean author career age less than or equal to 16). We then compute publication activity separately for these two groups by aggregating these data up to the subfield/year level of analysis. As can be observed in the first two columns of Table VI, there really is not any difference in the magnitude of the post-death effect across these two groups.

The second step is to distinguish between the related articles with at least one eminent author from related articles for which none of the authors is particularly famous at the time of its publication. To do this, we use two distinct measures of eminence. The first is whether a matched author belong to our sample of 12,935 stars. The second is whether a matched author belongs to an even more elite set comprising Nobel prize winners, Howard Hughes Medical Investigators, and members of the National Academy of Sciences. In the final four columns to Table VI, we find that articles published by non-elite members of the profession appear to account for much of the relative growth for treated subfields. This is consistent

with the idea that elite scientists face weaker incentives to deviate from their existing research trajectory, relative to less-established scientists.

Finally, we probe the standing of the non-collaborators in the subfield. One possibility is that they are competitors of the star, with much of their publication activity in the subfield when the star was alive. Another possibility is that they are recent entrants into the subfield—not social outsiders but intellectual outsiders. To distinguish these different types of authors empirically, we create a metric of intellectual proximity for each matched author, by computing the fraction of their publications that belongs to the star’s subfields up to the year before the publication of each related article. Whenever we match more than one author on a single related article, we assign to that article the maximum proximity score. A full 50% of the related articles turn out to have authors with exactly zero intellectual overlap with the star’s subfield. In addition to the bottom two quartiles, we create 10 bins for every five percentiles above the median (50th to 55th percentile, 55th to 60th percentile, . . . , 95th to 99th percentile), as well as top percentile bin. We then compute the corresponding measures of subfield activity by aggregating the data up to the subfield/year level. This time, we opt to present the results graphically in Figure VI. Each dot corresponds to the magnitude of the treatment effect in a separate regression with the outcome variable being the number of articles in each subfield that belong to the corresponding bins.

A striking pattern emerges. The authors driving the growth in publication activity following a star’s death are largely outsiders. They do not appear to have been substantially active in the subfield when the star was alive. To borrow a term from industrial organization, they are new entrants into these subfields, though the evidence presented above also shows that they are not especially likely to be younger scientists overall.

4.5 The Nature of Entry Barriers

The evidence so far points to fields of deceased stars enjoying bursts of activity after the death event. The influx of outsiders documented above suggests that stars may be able to regulate entry into their field while alive. While it is tempting to envisage conscious effort by the stars to block entry through the explicit control of key resources, such as funding and/or editorial goodwill (Li 2015; Brogaard et al. 2014), this explanation appears inconsistent with the facts on the ground. In the five-year window before death, only three of our stars (out of 452) were sitting on study sections, the funding panels that evaluate the scientific merits of

NIH grant applications. Another three were journal editors in the same time window. This handful of individuals could not possibly drive the robust effects we have uncovered.²² If barriers to entry are not the result of explicit control by stars, what is discouraging entry?

Goliath’s Shadow. One possibility is that outsiders are simply deterred by the prospect of challenging a luminary in the field. The existence of a towering figure may skew the cost-benefit calculations from entry by outside scholars toward delay or alternative activities. Table VII examines this role of implicit barriers to entry by focusing on the importance and commitment of the star in terms of publications and NIH funding within the field. Importance is defined as the fraction of papers (respectively, NIH grant amounts) in the subfield that have the star as an author (principal investigator). Commitment to the field is defined as the fraction of a star’s entire corpus of publications (respectively, cumulative NIH grant awards) that falls in the focal subfield. Splitting the sample at the median of these measures reveals an interesting pattern of results.

Stars that were especially important to the field in terms of research output appear to be an important deterrent to entry, with their passing creating a larger void for non-collaborators to fill. In contrast, a star’s commitment to the field—the degree to which a star’s main research interests lay within the field—does not appear to play a similar role. The last four columns of Table VII underscore the importance of financial resources in regulating entry. When the star commandeers large amounts of funding under either of our measures, we see a surge in entry by outsiders after the star’s passing, when competition for these resources is presumed to be more vigorous. Together these results suggest that, rather than directly thwarting the efforts of would be entrants, it is the presence of a preeminent scholar that dissuades intellectual outsiders from engaging with the field.

Intellectual Closure. Entry into a field, even after it has lost its shining star, may be deterred if the subfield appears unusually coherent to outsiders. A subfield is likely to be perceived as *intellectually* coherent, when the researchers active in it agree on the set of questions, approaches, and methodologies that propel the field forward. To explore the notion of “paradigmatic closure” as a barrier to field entry we develop two measures of intellectual coherence.

²²We verified that omitting these scientists from the sample hardly change the core results.

The first index of intellectual coherence leverages PMRA to capture the extent to which articles in the subfield pack themselves into a crowded scientific neighborhood. Recall that for each article in a subfield, we have at our disposal both a cardinal and an ordinal measure of intellectual proximity with the source article from which all other articles in the subfield radiate. Focusing only on the set of articles published in the subfield before the year of death, we measure intellectual coherence as the cardinal ranking (expressed as a real number between zero and one) for the 25th most related article in the subfield.²³ According to this metric, subfields exhibit wide variation in their degree of intellectual coherence, with a mean and median equal to 0.62 ($sd = 0.13$). The second index of intellectual coherence exploits the list of references cited in each article in the subfield before the star’s death. We simply compute the proportion of these references that fall within the subfield. Our contention is that fields that are more self-referential will tend to dissuade outsiders from entering. Once again, we observe meaningful variation across subfields using this second index ($mean = 0.081$; $median = 0.067$; $sd = 0.059$).

Social Closure. Alternatively, a field might be perceived as *socially* coherent, when the researchers active in it form a tightly-knit clique, often collaborating with each other, and perhaps also reviewing each other’s manuscripts. To explore this barrier we develop two additional measures of coherence, only in this case those designed to capture social cohesion rather than paradigmatic closure.

A natural way to capture endogamy within a subfield is to focus on the extent to which the star trained a large number of the junior scientists within it. We conjecture that the fields of stars who produced many intellectual “offspring” would be less welcoming to outsiders than those in which the stars did not train many graduate students or postdoctoral fellows. To identify trainees, we focus on the subset of coauthors who occupy the first author position in articles where the star occupies the last position; with the added stipulation that the coauthored publication appears in a window of \pm three years around the year in which the collaborator’s highest degree was received. Our first index of social coherence at the subfield level is then simply the count of the number of investigators trained by the star before his/her (possibly counterfactual) death. Our second measure of social coherence

²³The choice of the twenty fifth-ranked article is arbitrary, and also convenient. After purging from each subfield reviews, editorials, and articles appearing in journals not indexed by *WoS*, 95% of the subfields contain 25 articles or more in the period that precedes the star’s death. In those rare cases where the number of articles is less than twenty-five, we choose as our measure of coherence the cardinal measure for the least-proximate article in the subfield.

summarizes the degree of subfield “cliquishness” by computing the clustering coefficient in its coauthorship network. The clustering coefficient is simply the proportion of closed triplets within the network, an intuitive way to measure the propensity of scientists in the field to choose insiders as collaborators.²⁴

Panel A of Table VIII investigates the role of these intellectual and social barriers in modulating the post-death expansion of fields. We find evidence of a large role for both types of barriers, no matter how they are measured. The treatment effect is systematically larger when the subfield is less intellectually coherent (we use a top quartile-split to contrast the effect in more coherent vs. less coherent subfields). The same is true when subfields are less socially coherent. In fact, in the subsample of unusually coherent subfields, we find no statistical evidence of a publication influx after the passing of a star.²⁵

Incumbent Resource Control. While we noted earlier that stars do not appear especially well positioned to directly block entry through the control of key resources, it is possible that those resources can be controlled indirectly through the influence of collaborators. If incumbent scholars within a field serve as gatekeepers of funding and journal access, they may be able to effectively stave off threats of entry from outsiders.

A practical challenge to assessing this indirect channel is that stars tend to have a large number of collaborators; which among them could be instrumental in shaping the intellectual direction of the field? To gain empirical traction on the concept of indirect control, we delineate two categories of “important” collaborators. The first comprises those individuals who coauthor frequently with the star: five coauthorships or more at the time of the star’s death (this corresponds to the top decile of collaborators when ranked by total number of coauthorships). The second uses “extreme” authorship positions, focusing on collaborators who were ever first author when the star was in last position, or last author when the star was in first position. Using information regarding the composition of NIH funding panels, we then tabulate, for each star, the number of important collaborators who were members of at least one of these committees in the five years preceding the death of the star.

²⁴The clustering coefficient is based on triplets of nodes (authors). A triplet consists of three authors that are connected by either two (open triplet) or three (closed triplet) undirected ties. The clustering coefficient is the number of closed triplets over the total number of triplets (both open and closed, cf. Luce and Perry [1949]).

²⁵A small caveat pertains to the measure based on the count of trainees. While the magnitudes of the coefficients are ordered in a similar way, the difference between them is not itself statistically significant.

We would like to proceed in a similar fashion using the composition of editorial boards, but these data are not easily available for the set of *PubMed*-indexed journals and the thirty-year time period covered by our sample. As an alternative, we develop a proxy for editorial position based on the number of editorials or comments written by every collaborator of the star.²⁶ We then sum the number of editorials written by important coauthors in the five years before the death. Together, the editorial and study section information allow us to distinguish between the stars whose important coauthors were in a position to channel resources towards preferred individuals or intellectual approaches from those stars whose important coauthors had no such power.

Panel B of Table VIII presents the evidence on the role of indirect control. The eight specifications paint a unified picture—subfield expansion is the rule, but is much more pronounced when stars have relatively few collaborators in influential positions. The differences between estimates in each pair of columns are large, and significantly different from zero at the 5% level of significance in one-tailed tests. Indirect control therefore appears to be a mechanism through which superstars can exert influence on the evolution of their fields, even from beyond the grave. Important coauthors, in their effort to keep the star’s intellectual flame alive, erect barriers to entry into those fields that prevent its rejuvenation by outsiders.

Taken together, these results suggest that outsiders are reluctant to challenge hegemonic leadership within a field when the star is alive. They also highlight a number of factors that constrain entry even after she is gone. Intellectual, social, and resource barriers all impede entry, with outsiders only entering subfields whose topology offers a less hostile landscape for the support and acceptance of “foreign” ideas.

5 Conclusion

In this paper, we exploit the applied economist’s toolkit, together with a novel approach to delineate the boundaries of scientific fields, to explore the effect that the passing of an eminent scientist exerts on the dynamics of growth—or decline—for the fields in which

²⁶We investigated the validity of this proxy as follows. In the sample of deceased superstars, every individual with five editorials or more was an editor. In a random sample of 50 superstars with no editorials published, only one was an editor (for a field journal). Finally, among the sixteen superstars who wrote between one and four editorials over their career, we found two whose CV indicate they were in fact editors for a key journal in their field. We conclude that there appears to be a meaningful correlation between the number of editorials written and the propensity to be an editor.

s/he was active while alive. We find that publications and grants by scientists that never collaborated with the star surge within the subfield, absent the star. Interestingly, this surge is not driven by a reshuffling of leadership within the field, but rather by new entrants that are drawn from outside of it. Our rich data on individual researchers and the nature of their scholarship allows us provide a deeper understanding of this dynamic.

In particular, this increase in contributions by outsiders appears to tackle the mainstream questions within the field but by leveraging newer ideas that arise in other domains. This intellectual arbitrage is quite successful—the new articles represent substantial contributions, at least as measured by long-run citation impact. Together, these results paint a picture of scientific fields as scholarly guilds to which elite scientists can regulate access, providing them with outsized opportunities to shape the direction of scientific advance in that space.

We also provide evidence regarding the mechanisms that enable the regulation of entry. While stars are alive, entry appears to be effectively deterred where the shadow they cast over the fields in which they were active looms particularly large. After their passing, we find evidence for influence from beyond the grave, exercised through a tightly-knit “invisible college” of collaborators (de Solla Price and Beaver 1966; Crane 1972). The loss of an elite scientist central to the field appears to signal to those on the outside that the cost/benefit calculations on the avant-garde ideas they might bring to the table has changed, thus encouraging them to engage. But this occurs only when the topology of the field offers a less hostile landscape for the support and acceptance of “foreign” ideas, and specifically when the star’s network of close collaborators is insufficiently robust to stave off threats from intellectual outsiders.

In the end, our results lend credence to Planck’s infamous quip that provides the title for this manuscript. Yet its implications for social welfare are ambiguous. While we can document that eminent scientists restrict the entry of new ideas and scholars into a field, gatekeeping activities could have beneficial properties when the field is in its inception; it might allow cumulative progress through shared assumptions and methodologies, and the ability to control the intellectual evolution of a scientific domain might, in itself, be a prize that spurs much *ex ante* risk taking. Because our empirical exercise cannot shed light on these countervailing tendencies, we must remain guarded in drawing policy conclusions from our results. Yet, the fact that the presence of a tutelar figurehead can freeze patterns of participation into a scientific field increases the appeal of policies that bolster access to less

established or less well-connected investigators. Example of such policies include caps on the amount of funding a single laboratory is eligible to receive, “bonus points” for first-time investigators in funding programs, emeritus awards to induce senior scientists to wind down their laboratory activities, and double-blind refereeing policies (Kaiser 2011, Berg 2012, Deng 2015).

Our work leaves many questions unanswered. What is the fate of the fields that these new entrants departed? Do they decay, or instead “merge” with those whose star departed prematurely? Given a finite supply of scientists and the adjustment costs involved in switching scientific focus, one would expect some other field to contract on the margin in the wake of a superstar’s passing. Is this marginal field more novel, or already established? We are pursuing these questions in ongoing work.

Appendix A: Criteria for Delineating the Set of 12,935 “Superstars”

Highly Funded Scientists. Our first data source is the Consolidated Grant/Applicant File (CGAF) from the U.S. National Institutes of Health (NIH). This dataset records information about grants awarded to extramural researchers funded by the NIH since 1938. Using the CGAF and focusing only on direct costs associated with research grants, we compute individual cumulative totals for the decades 1977-1986, 1987-1996, and 1997-2006, deflating the earlier years by the Biomedical Research Producer Price Index. We also recompute these totals excluding large center grants that usually fund groups of investigators (M01 and P01 grants). Scientists whose totals lie above the 95th percentile of either distribution constitute our first group of superstars. In this group, the least well-funded investigator garnered \$10.5 million in career NIH funding and the most well-funded \$462.6 million.ⁱ

Highly Cited Scientists. Despite the preeminent role of the NIH in the funding of public biomedical research, the above indicator of “superstardom” biases the sample towards scientists conducting relatively expensive research. We complement this first group with a second composed of highly cited scientists identified by the Institute for Scientific Information. A Highly Cited listing means that an individual was among the 250 most cited researchers for their published articles between 1981 and 1999, within a broad scientific field.ⁱⁱ

Top Patenters. We add to these groups academic life scientists who belong in the top percentile of the patent distribution among academics—those who were granted 17 patents or more between 1976 and 2004.

Members of the National Academy of Science and of the Institute of Medicine. We add to these groups academic life scientists who were elected to the National Academy of Science or the Institute of Medicine between 1970 and 2013.

MERIT Awardees of the NIH. Initiated in the mid-1980s, the MERIT Award program extends funding for up to 5 years (but typically 3 years) to a select number of NIH-funded investigators “*who have demonstrated superior competence, outstanding productivity during their previous research endeavors and are leaders in their field with paradigm-shifting ideas.*” The specific details governing selection vary across the component institutes of the NIH, but the essential feature of the program is that only researchers holding an R01 grant in its second or later cycle are eligible. Further, the application must be scored in the top percentile in a given funding cycle.

Former and current Howard Hughes Medical Investigators (HHMIs). Every three years, the Howard Hughes Medical Institute selects a small cohort of mid-career biomedical scientists with the potential to revolutionize their respective subfields. Once selected, HHMIs continue to be based at their institutions, typically leading a research group of 10 to 25 students, postdoctoral associates and technicians. Their appointment is reviewed every five years, based solely on their most important contributions during the cycle.ⁱⁱⁱ

ⁱWe perform a similar exercise for scientists employed by the intramural campus of the NIH. These scientists are not eligible to receive extramural funds, but the NIH keeps records of the number of “internal projects” each intramural scientist leads. We include in the elite sample the top five percentiles of intramural scientists according to this metric.

ⁱⁱThe relevant scientific fields in the life sciences are microbiology, biochemistry, psychiatry/psychology, neuroscience, molecular biology & genetics, immunology, pharmacology, and clinical medicine.

ⁱⁱⁱSee Azoulay et al. (2011) for more details and an evaluation of this program.

Early career prize winners. We also included winners of the Pew, Searle, Beckman, Rita Allen, and Packard scholarships for the years 1981 through 2000. Every year, these charitable foundations provide seed funding to between 20 and 40 young academic life scientists. These scholarships are the most prestigious accolades that young researchers can receive in the first two years of their careers as independent investigators.

Appendix B: Linking Scientists with their Journal Articles

The source of our publication data is *PubMed*, a bibliographic database maintained by the U.S. National Library of Medicine that is searchable on the web at no cost.^{iv} *PubMed* contains over 14 million citations from 4,800 journals published in the United States and more than 70 other countries from 1950 to the present. The subject scope of this database is biomedicine and health, broadly defined to encompass those areas of the life sciences, behavioral sciences, chemical sciences, and bioengineering that inform research in health-related fields. In order to effectively mine this publicly-available data source, we designed PUBHARVESTER, an open-source software tool that automates the process of gathering publication information for individual life scientists (see Azoulay et al. 2006 for a complete description of the software). PUBHARVESTER is fast, simple to use, and reliable. Its output consists of a series of reports that can be easily imported by statistical software packages.

This software tool does not obviate the two challenges faced by empirical researchers when attempting to accurately link individual scientists with their published output. The first relates to what one might term “Type I Error,” whereby we mistakenly attribute to a scientist a journal article actually authored by a namesake; The second relates to “Type II error,” whereby we conservatively exclude from a scientist’s publication roster legitimate articles:

Namesakes and popular names. *PubMed* does not assign unique identifiers to the authors of the publications they index. They identify authors simply by their last name, up to two initials, and an optional suffix. This makes it difficult to unambiguously assign publication output to individual scientists, especially when their last name is relatively common.

Inconsistent publication names. The opposite danger, that of recording too few publications, also looms large, since scientists are often inconsistent in the choice of names they choose to publish under. By far the most common source of error is the haphazard use of a middle initial. Other errors stem from inconsistent use of suffixes (Jr., Sr., 2nd, etc.), or from multiple patronyms due to changes in spousal status.

To deal with these serious measurement problems, we opted for a labor-intensive approach: the design of individual search queries that relies on relevant scientific keywords, the names of frequent collaborators, journal names, as well as institutional affiliations. We are aided in the time-consuming process of query design by the availability of a reliable archival data source, namely, these scientists’ CVs and biosketches. PUBHARVESTER provides the option to use such custom queries in lieu of a completely generic query (e.g, "azoulay p"[au] or "graff zivin js"[au]). As an example, one can examine the publications of Scott A. Waldman, an eminent pharmacologist located in Philadelphia, PA at Thomas Jefferson University. Waldman is a relatively frequent name in the United States (with 208 researchers with an identical patronym in the AAMC faculty roster); the combination "waldman s" is common to 3 researchers in the same database.

^{iv}<http://www.pubmed.gov/>

A simple search query for "waldman sa"[au] OR "waldman s"[au] returns 377 publications at the time of this writing. However, a more refined query, based on Professor Waldman's biosketch returns only 256 publications.^v

The above example also makes clear how we deal with the issue of inconsistent publication names. PUB-HARVESTER gives the end-user the option to choose up to four *PubMed*-formatted names under which publications can be found for a given researcher. For example, Louis J. Tobian, Jr. publishes under "tobian l", "tobian l jr", and "tobian lj", and all three names need to be provided as inputs to generate a complete publication listing. Furthermore, even though Tobian is a relatively rare name, the search query needs to be modified to account for these name variations, as in ("tobian l"[au] OR "tobian lj"[au]).

Appendix C: *PubMed* Related Citations Algorithm [PMRA]

Traditionally, it has been very difficult to assign to individual scientists, or articles, a fixed address in “idea space,” and this data constraint explains in large part why bibliometric analyses typically focus on the determinants of the rate of scientific progress rather than its direction. The empirical exercise in this paper hinges crucially on the ability to relax this constraint in a way that is consistent across death events and also requires little, if any, human judgement.

This challenge is met here by the use of the *PubMed* Related Citations Algorithm [PMRA], a probabilistic, topic-based model for content similarity that underlies the “related articles” search feature in *PubMed*. This database feature is designed to help a typical user search through the literature by presenting a set of records topically related to any article returned by a *PubMed* search query.^{vi} To assess the degree of intellectual similarity between any two *PubMed* records, PMRA relies crucially on MeSH keywords. MeSH is the National Library of Medicine's [NLM] controlled vocabulary thesaurus. It consists of sets of terms arranged in a hierarchical structure that permit searching at various levels of specificity. There are 27,149 descriptors in the 2013 MeSH edition. Almost every publication in *PubMed* is tagged with a set of MeSH terms (between 1 and 103 in the current edition of *PubMed*, with both the mean and median approximately equal to 11). NLM's professional indexers are trained to select indexing terms from MeSH according to a specific protocol, and consider each article in the context of the entire collection (Bachrach and Charen 1978; Névéol et al. 2010). What is key for our purposes is that the subjectivity inherent in any indexing task is confined to the MeSH term assignment process and does not involve the articles' authors.^{vii}

Using the MeSH keywords as input, PMRA essentially defines a distance concept in idea space such that the proximity between a source article and any other *PubMed*-indexed publication can be assessed. The following paragraphs were extracted from a brief description of PMRA:

The neighbors of a document are those documents in the database that are the most similar to it. The similarity between documents is measured by the words they have in common, with some adjustment for document

^v(((("waldman sa"[au] NOT (ether OR anesthesia)) OR ("waldman s"[au] AND (murad OR philadelphia[ad] OR west point[ad] OR wong p[au] OR lasseter kc[au] OR colorectal))) AND 1980:2013[dp])

^{vi}Lin and Wilbur (2007) report that one fifth of “non-trivial” browser sessions in *PubMed* involve at least one invocation of PMRA.

^{vii}This is a slight exaggeration: PMRA also makes use of title and abstract words to determine the proximity of any two pairs of articles in the intellectual space. These inputs are obviously selected by authors, rather than by NLM staff. However, neither the choice of MeSH keywords nor the algorithm depend on cited references contained in publications.

lengths. To carry out such a program, one must first define what a word is. For us, a word is basically an unbroken string of letters and numerals with at least one letter of the alphabet in it. Words end at hyphens, spaces, new lines, and punctuation. A list of 310 common, but uninformative, words (also known as stopwords) are eliminated from processing at this stage. Next, a limited amount of stemming of words is done, but no thesaurus is used in processing. Words from the abstract of a document are classified as text words. Words from titles are also classified as text words, but words from titles are added in a second time to give them a small advantage in the local weighting scheme. MeSH terms are placed in a third category, and a MeSH term with a subheading qualifier is entered twice, once without the qualifier and once with it. If a MeSH term is starred (indicating a major concept in a document), the star is ignored. These three categories of words (or phrases in the case of MeSH) comprise the representation of a document. No other fields, such as Author or Journal, enter into the calculations.

Having obtained the set of terms that represent each document, the next step is to recognize that not all words are of equal value. Each time a word is used, it is assigned a numerical weight. This numerical weight is based on information that the computer can obtain by automatic processing. Automatic processing is important because the number of different terms that have to be assigned weights is close to two million for this system. The weight or value of a term is dependent on three types of information: 1) the number of different documents in the database that contain the term; 2) the number of times the term occurs in a particular document; and 3) the number of term occurrences in the document. The first of these pieces of information is used to produce a number called the global weight of the term. The global weight is used in weighting the term throughout the database. The second and third pieces of information pertain only to a particular document and are used to produce a number called the local weight of the term in that specific document. When a word occurs in two documents, its weight is computed as the product of the global weight times the two local weights (one pertaining to each of the documents).

The global weight of a term is greater for the less frequent terms. This is reasonable because the presence of a term that occurred in most of the documents would really tell one very little about a document. On the other hand, a term that occurred in only 100 documents of one million would be very helpful in limiting the set of documents of interest. A word that occurred in only 10 documents is likely to be even more informative and will receive an even higher weight.

The local weight of a term is the measure of its importance in a particular document. Generally, the more frequent a term is within a document, the more important it is in representing the content of that document. However, this relationship is saturating, i.e., as the frequency continues to go up, the importance of the word increases less rapidly and finally comes to a finite limit. In addition, we do not want a longer document to be considered more important just because it is longer; therefore, a length correction is applied.

The similarity between two documents is computed by adding up the weights of all of the terms the two documents have in common. Once the similarity score of a document in relation to each of the other documents in the database has been computed, that document's neighbors are identified as the most similar (highest scoring) documents found. These closely related documents are pre-computed for each document in PubMed so that when one selects Related Articles, the system has only to retrieve this list. This enables a fast response time for such queries.^{viii}

The algorithm uses a cut-off rule to determine the number of related citations associated with a given source article. First, the 100 most related records by similarity score are returned. Second, a reciprocity rule is applied to this list of 100 records: if Publication A is related to Publication B, Publication B must also be related to publication A. As a result, the set of related citations for a given source article may contain many more than 100 publications.^{ix}

Given our set of source articles, we delineate the scientific fields to which they belong by focusing on the set of articles returned by PMRA that satisfy three additional constraints: (i) they are original articles (as opposed to editorials, comments, reviews, etc.); (ii) they were published in or before 2006 (the end of our observation period); and (iii) they appear in journals indexed by the *Web of Science* (so that follow-on citation information can be collected).

^{viii} Available at <http://ii.nlm.nih.gov/MTI/related.shtml>

^{ix} The effective number of related articles returned by PMRA varies between 58 and 2,097 in the sample of 3,074 source articles published by the 452 star scientists in the five years preceding their death. The mean is 185 related articles, and the median 141.

To summarize, PMRA is a modern implementation of *co-word analysis*, a content analysis technique that uses patterns of co-occurrence of pairs of items (i.e., title words or phrases, or keywords) in a corpus of texts to identify the relationships between ideas within the subject areas presented in these texts (Callon et al. 1989; He 1999). One long-standing concern among practitioners of this technique has been the “indexer effect” (Whittaker 1989). Clustering algorithm such as PMRA assume that the scientific corpus has been correctly indexed. But what if the indexers who chose the keywords brought their own “conceptual baggage” to the indexing task, so that the pictures that emerge from this process are more akin to their conceptualization than to those of the scientists whose work it was intended to study?

Indexer effects could manifest themselves in three distinct ways. First, indexers may have available a lexicon of permitted keywords which is itself out of date. Second, there is an inevitable delay between the publication of an article and the appearance of an entry in *PubMed*. Third, indexers, in their efforts to be helpful to users of the database, may use combinations of keywords which reflect the conventional views of the field. The first two concerns are legitimate, but probably have only a limited impact on the accuracy of the relationships between articles which PMRA deems related. This is because the NLM continually revises and updates the MeSH vocabulary, precisely in an attempt to neutralize keyword vintage effects. Moreover, the time elapsed between an article’s publication and the indexing task has shrunk dramatically, though time lag issues might have been a first-order challenge when MeSH was created, back in 1963. The last concern strikes us as being potentially more serious; a few studies have asked authors to validate ex post the quality of the keywords selected by independent indexers, with generally encouraging results (Law and Whittaker 1992). Inter-indexer reliability is also very high (Wilbur 1998).

In Table C1, we illustrate the use of PMRA with an example taken from our sample. Ira Herskowitz is a faculty member in our sample who died in 2003. “*The transcriptional program of sporulation in budding yeast*” (*PubMed* ID #9784122) is a publication from his lab which appeared in the October 23rd 1998 issue of the journal *Science* and lists 15 MeSH terms and 5 substances. *PubMed* ID #12242283 is its most related paper according to the PMRA algorithm; it appeared in *Molecular and Cell Biology* in October of 2002 and has 24 MeSH terms (resp. 11 substances). The keywords that overlap exactly have been highlighted in dark blue; those whose close ancestors in the MeSH keyword hierarchical tree overlap have been highlighted in light blue. These terms include common terms such as **Saccharomyces cerevisiae** and **Transcription Factors** as well as more specific keywords including **NDT80 protein**, **S cerevisiae** and **Gene Expression Regulation, Fungal**.

Table C1: PMRA and MeSH Term Overlap—An Example

Source Article	PMRA-Linked Article
<p>Chu et al., “The transcriptional program of sporulation in budding yeast.” <i>Science</i>, 1998.</p> <p style="text-align: center;">PMID #9784122</p>	<p>Sopko et al. “Phosphorylation and maximal activity of <i>Saccharomyces cerevisiae</i> meiosis-specific transcription factor Ndt80 is dependent on Ime2.” <i>MCB</i>, 2002.</p> <p style="text-align: center;">PMID #12242283</p>
MeSH Terms	MeSH Terms
Animals	Active Transport, Cell Nucleus
Chromosomes, Fungal	Binding Sites
DNA-Binding Proteins*	Cell Cycle Proteins*
Fungal Proteins	Cell Nucleus
Gene Expression Regulation, Fungal*	DNA-Binding Proteins*
Genes, Fungal	Fungal Proteins*
Genome, Fungal	Gene Expression Regulation, Fungal*
Humans	Genes, Fungal
Meiosis	Intracellular Signaling Peptides and Proteins
Morphogenesis	Meiosis*
Organelles	Phosphorylation
<i>Saccharomyces cerevisiae</i> *	Promoter Regions, Genetic
Spores, Fungal	Protein Kinases*
Transcription Factors	Protein-Serine-Threonine Kinases
Transcription, Genetic*	Recombinant Fusion Proteins
	<i>Saccharomyces cerevisiae</i>
	<i>Saccharomyces cerevisiae</i> Proteins*
	Spores, Fungal
	Substrate Specificity
	Transcription Factors*
	Transcriptional Activation
Substances	Substances
DNA-Binding Proteins	Cell Cycle Proteins
Fungal Proteins	DNA-Binding Proteins
NDT80 protein, <i>S cerevisiae</i>	Fungal Proteins
<i>Saccharomyces cerevisiae</i> Proteins	Intracellular Signaling Peptides and Proteins
Transcription Factors	NDT80 protein, <i>S cerevisiae</i>
	Recombinant Fusion Proteins
	<i>Saccharomyces cerevisiae</i> Proteins
	Transcription Factors
	Protein Kinases
	IME2 protein, <i>S cerevisiae</i>
	Protein-Serine-Threonine Kinases

Appendix D: Construction of the Control Group

We detail the procedure implemented to identify the control subfields that help pin down the life-cycle and secular time effects in our difference-in-differences (DD) specification. Happenstance might yield a sample of stars clustered in decaying scientific fields. More plausibly, activity in the typical subfield might be subject to idiosyncratic life-cycle patterns, with their productive potential first increasing over time, eventually peaking, and thereafter slowly declining. Relying solely on subfields treated earlier or later as an implicit control group raises the worry that these time-varying omitted variables will not be fully captured by subfield age controls, particularly since dating the birth of a subfield is a process fraught with hazards.

To address this concern, we create an additional level of difference by selecting control subfields. Recall that selecting a subfield in our framework is akin to first selecting a source article and then using PMRA to harvest all the related articles to this source in intellectual space. Since the second step is fully automated, only the first step is really of concern. Practically, we will recruit control source articles from the set of articles authored by star scientists who do not die prematurely. But what makes a satisfactory control group? It is important to distinguish between *ex ante* vs. *ex post* criteria. *Ex ante*, one would like control source articles to have the following properties:

1. to be published contemporaneously with the source article for the treated subfield;
2. to be unrelated in both an intellectual and a social sense, to the source article for the treated subfield;
3. to be of similar expected impact and fruitfulness, relative to the source article for the treated subfield;
4. to have a similar number of authors as the source article for the treated subfield;
5. to have a superstar author in the same authorship position and of approximately the same age as that occupied by the deceased superstar on the authorship roster of the source article for the treated subfield.

Ex post, it will be important for the control subfields to satisfy an additional condition: the treated and control subfields should exhibit very similar trends in publication activity and funding flows up to the year of treatment (i.e., the year of death for the treated superstar).

Coarsened Exact Matching. To meet these goals, we implement a “Coarsened Exact Matching” (CEM) procedure (Blackwell et al. 2009). The first step is to select a relatively small set of covariates on which we need to guarantee balance *ex ante*. This choice entails judgement, but is strongly guided by the set of criteria listed above. The second step is to create a large number of strata to cover the entire support of the joint distribution of the covariates selected in the previous step. In a third step, each observation is allocated to a unique strata, and for each observation in the treated group, control observations are selected from the same strata.

The procedure is coarse because we do not attempt to precisely match on covariate values; rather, we coarsen the support of the joint distribution of the covariates into a finite number of strata, and we match a treated observation if and only if a control observation can be recruited from this strata. An important advantage of CEM is that the analyst can guarantee the degree of covariate balance *ex ante*, but this comes at a cost: the more fine-grained the partition of the support for the joint distribution (i.e., the higher the number of strata), the larger the number of unmatched treated observations.

Implementation. We identify controls based on the following set of covariates (t denotes the year of death): star scientist career age, citations received by the article up to year t , number of authors; position of the star

author on the authorship roster (only first or last authorship positions are considered); journal; and year of publication. The first three covariates only need to match within relatively coarse bins. For instance, we create nine career age categories: less than 10 years; between 10 and 20 years; between 20 and 25 years; between 25 and 30 years; between 30 and 35 years; between 35 and 40 years; between 40 and 45 years; between 45 and 50 years, over 50 years of career age. Similarly, we coarsen the distribution of citations at baseline into five mutually exclusive bins: zero citations; between one and 10 citations; between 10 and 50 citations; between 50 and 120 citations; and more than 120 citations. In contrast, we impose an exact match on journal, publication year, and the star’s authorship position.

We match approximately 75% of the treated source articles in this way. Unfortunately, some further trimming of the control articles is needed. First, we eliminate any control that shares any author with the treated source. Second, we eliminate any control article with a dead star scientist on its authorship roster, even if s/he appears in an intermediate position in the authorship list. Third, we drop every control that also happens to be related intellectually to its source as per PMRA. Finally, we drop from the data any source article that finds itself an orphan (i.e., not paired with any control) at the conclusion of this process. Figure III provides an illustrative example.

The final sample has 3,074 treated source articles and 31,142 control source articles. As can be seen in Figure IV, the distribution of activity levels, measured by cumulative publications up to the baseline year, is very similar between treated and control subfields. As well, there is no evidence of preexisting trends in activity, as demonstrated by the coefficient estimates graphed in Figure V. In Table II, treated and control subfields are very well-balanced on the covariates that formed the basis of the CEM matching procedure. This is true almost by construction. What is more surprising (and also welcome) is that the procedure balances a number of covariates that were not used as inputs for matching, such as various metrics of star eminence. For other covariates, we can detect statistically significant mean differences, though they do not appear to be substantively meaningful (e.g., 6.7% of control stars vs. 9.9% of treated stars are female).

Sensitivity Analyses. Human judgement matters for the outcome of the CEM procedure insofar as one must draw a list of “reasonable” covariates to match on, as well as decide on the degree of coarsening to impose. We have verified that slight variations in the implementation (e.g., varying slightly the number of cutoff points for the stock of baseline citations for the source; focusing on birth age as opposed to career age for the stars) have little impact on the main results.