

Early-Career Research Funding and Cumulative Advantage in Scientific Careers: A Mathematical Model and Empirical Analysis

Yuhei Yamaguchi

`yuhei.yamaguchi.t1@dc.tohoku.ac.jp`

Tohoku University

Research Article

Keywords: Matthew effect, Cumulative advantage, Inequality, Early-career research funding, Academic careers, Science of science

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11180729/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Early-Career Research Funding and Cumulative Advantage in Scientific Careers: A Mathematical Model and Empirical Analysis

Yuhei Yamaguchi^{1*}

^{1*}Graduate School of Arts and Letters, Tohoku University, Sendai,
Miyagi, Japan.

Corresponding author(s). E-mail(s):
yuhei.yamaguchi.t1@dc.tohoku.ac.jp;

Abstract

Success breeds success. This simple idea—the Matthew effect, or cumulative advantage—is a classical theme in the science of science. Does obtaining a research grant very early in a researcher’s career, at the very beginning of their academic trajectory, create a performance gap that continues to widen for decades? We approach this question from two sides. First, we build a mathematical model in which talent, publications, research funding, and promotion mutually reinforce one another, and show that even when talent is held constant, the gap in research output between early grant recipients and non-recipients widens over time. We distil this prediction into two hypotheses. Second, we test these hypotheses using a large panel dataset of approximately 40,000 researchers spanning eight broad fields of Japanese academia, drawn from the researchmap database. Using the graduate school of origin as a proxy for talent and prestige, we find that the advantage associated with receiving an early-career fellowship remains positive, statistically significant, and steadily increases over time even after accounting for differences in the institution of origin. In short, even a head start lasting only a few years leaves a lasting mark on a researcher’s career, with its effects continuing to grow more than a decade later. By combining an explicit model of cumulative advantage with large-scale empirical evidence across fields, we quantify how the earliest allocation of research resources shapes the long-term stratification of science.

Keywords: Matthew effect, Cumulative advantage, Inequality, Early-career research funding, Academic careers, Science of science

1 Introduction

1.1 The Matthew effect as a universal inequality process

For to every one who has will more be given, and he will have abundance; but from him who has not, even what he has will be taken away. — Matthew 25:29 (RSV)

Why are some scientists successful while others are not? Or, put differently, why do some researchers steadily accumulate research achievements while others do not? Empirically, inequality in research performance among scientists has been shown to exceed inequality in individual income within developed countries (Petersen & Penner, 2014), and inequality in publication counts among researchers has been shown to widen over time (Allison, Long, & Krauze, 1982). Why, then, do such gaps in performance arise?

The accumulation of research achievement is of course related to innate scientific ability, energy, and motivation (Cole & Cole, 1973). Yet is success fully explained by these factors alone? Are there no factors that shape a career through channels that do not operate via talent or drive?

Merton and Zuckerman proposed that cumulative advantage—the “Matthew effect”—is one such factor promoting success. It is an inequality-generating process found broadly throughout society: wealth begets wealth, attention begets attention, and early success begets later success, so that small initial differences are amplified into large, persistent gaps. Drawing on the verse above, Merton named this principle in the context of science (Merton, 1968, 1988): those who already hold recognition, resources, and visibility more easily attract further recognition and resources. Stated more explicitly in the scientific context, initially advantaged scientists obtain greater opportunities for further achievement and rewards (Zuckerman, 1977).

This self-reinforcing dynamic has been captured under a variety of concepts: cumulative advantage in citation and journal use (de Solla Price, 1976), preferential attachment in networks (Barabási & Albert, 1999), and increasing returns and path-dependent “lock-in” in economics (Arthur, 1989; David, 1985). Such processes are treated as one of the key mechanisms of inequality not only in science but across various domains of society (DiPrete & Eirich, 2006). The phenomenon has been confirmed not only by observational studies but also by experimental designs that permit the estimation of causal effects: Salganik et al.’s artificial cultural-market experiment (Salganik, Dodds, & Watts, 2006)¹ and van de Rijt et al.’s “success-breeds-success” field experiments (van de Rijt, Kang, Restivo, & Patil, 2014) show that small early successes causally raise the probability of subsequent success. The Matthew effect—cumulative advantage—is thus a factor in the generation of inequality that cannot be ignored.

1.2 Early success in academia as a critical juncture

Just as Merton developed the notion of the Matthew effect with scientific research as his subject, academia is a paradigmatic arena in which cumulative advantage can

¹The experiment of Salganik et al. (2006) does not demonstrate the Matthew effect directly; it nevertheless allows us to infer that cumulative advantage was at work as a factor amplifying the inequality that emerged.

emerge. Many studies have shown that the Matthew effect, or cumulative advantage, exists in academia.

Career longevity itself exhibits a Matthew effect (Petersen, Jung, Yang, & Stanley, 2011); citations accumulate through preferential attachment and generate highly skewed, power-law-like distributions (Fortunato et al., 2018); and faculty hiring is strongly stratified by the prestige of researchers’ doctoral institutions (Clauset, Arbesman, & Larremore, 2015). Across eight scientific fields, publishing in prestigious journals early on and being affiliated with a leading institution strongly predict which researchers eventually attain high research impact (Krauss, Danús, & Sales-Pardo, 2023). Early coauthorship with prominent scientists likewise predicts later career success (Li, Aste, Caccioli, & Livan, 2019), and researchers placed in more prestigious research environments publish more and accumulate more citations (Way, Morgan, Larremore, & Clauset, 2019). These empirical results accord with what Cole and Cole suggested roughly half a century ago (Cole & Cole, 1973, pp. 72–75).

Cumulative advantage can also operate at the level of individual papers. Publishing in a high-impact journal is associated with a shorter time to subsequent publications in high-impact journals, suggesting a paper-to-paper feedback process within a researcher’s career (Petersen & Penner, 2014). In the same vein, even after other factors are controlled for, those who have produced many papers in the past go on to produce many papers in the future (Hall & Mairesse, 2024).

Taken together, these findings suggest—as Merton and Zuckerman argued—that *an advantage held at one point in time may generate further advantages, whose effects accumulate and amplify over time*. Such processes raise important questions about the fairness of scientific careers.

In principle, cumulative advantage makes *the first few years of a scientific career a particularly consequential period*. At this early stage, the information available about a researcher’s ability is still limited, so even a single grant or a single paper in a leading journal stands out against an otherwise sparse record. A difference that would be marginal at mid-career can therefore account for *much of what separates contemporaneous researchers* at this stage. The evaluations that follow—for research funding, for positions, and for collaboration—must read these still-sparse records as signals of ability. Early success thus improves one’s standing in the next round of selection, and that standing in turn becomes a source of the next success. Through such a mechanism, a gap that opens early in a career may not be averaged away later, but instead *be carried forward and compound further*.

Among these forms of early success, we focus on one that is *concrete, measurable, and policy-relevant: competitive research funding obtained early in a career*. Research funding is a key instrument supporting scientific productivity and creativity, although its effects vary with the size of the award and the conditions attached to it (Azoulay, Graff Zivin, & Manso, 2011; Ebadi & Schiffauerova, 2016; Gush, Jaffe, Larsen, & Laws, 2018; Hottenrott & Lawson, 2017; Jacob & Lefgren, 2011).²

Our aim is to examine the possibility that *obtaining research funding early in a career generates subsequent advantage in a Matthew-effect fashion*. Research output

²Some studies point out that research funding does not necessarily increase the number of publications (Ding & Moreira, 2025; Lawson, Geuna, & Finardi, 2021).

produced under a grant becomes part of the visible track record that is evaluated in later funding decisions. A self-reinforcing cycle may therefore arise in which an early award raises the probability of obtaining additional funding.

That early-career funding is positively linked to later career performance has already been shown in several studies. Examining a large Dutch funding program, Bol et al. compare *applicants who fell just above the funding threshold with those who fell just below* (Bol, de Vaan, & van de Rijt, 2018). Over the following eight years, researchers who narrowly won funding obtained *more than twice as much subsequent funding* as nearly identical applicants who narrowly lost, a pattern the authors call a “Matthew effect in science funding.” Importantly, this difference cannot be explained by the research output produced under the initial grant alone, which suggests that *the first award itself becomes an asset in later funding decisions*.

Another study shows that recipients of early-career grants subsequently publish more and attain higher academic positions than applicants who were rejected but were otherwise comparable (van den Besselaar & Sandström, 2015).

There are, to be sure, exceptions to this pattern. Among researchers who continue their research activity after narrowly missing an early grant, later research output can even surpass that of narrow winners (Wang, Jones, & Wang, 2019). Nevertheless, *surveying the existing literature as a whole, the evidence predominantly indicates that early-career success in academia generates further advantage and can widen the gap between researchers who succeed early and those who do not*.

1.3 Limitations of prior work and the present study

This body of work, however, leaves several things unattempted and several questions unexamined.

First, these studies do not formalize the mechanism that generates inequality in academic research. A narrative account of how the gap is amplified is available, and it has been shown that early advantages accumulate and compound within academia. In most cases, however, it is left unstated *which* factors, through *what* dynamics, actually produce the disparity. Writing this mechanism down explicitly, as a mathematical model, brings substantial benefits. First, it converts a verbal account into a rigorously defined, analyzable structure. Hypotheses can then be *derived* from the model’s behavior rather than merely posited, which places the empirical analysis that follows on firmer ground. It also makes it possible to draw out policy implications. Once the process is expressed through parameters, one can ask *which* parameter an intervention should change in order to shrink the gap. This turns the search for remedies from an intuitive exercise into an evidence-based examination of possible interventions.

Second, few studies focus on early-career success and then track subsequent research performance and academic standing over the long run. Some of the work described above—the Dutch funding program, the follow-up of early-career grant applicants, and the study of coauthorship with prominent scientists—does analyze the relationship between an early grant or an early collaboration and later career success (Bol et al., 2018; Li et al., 2019; van den Besselaar & Sandström, 2015). Such studies nonetheless remain few in number, and the horizon over which they follow careers is usually short. As a consequence, several outcomes have yet to be examined adequately

in relation to early-career success. How early funding relates to the number of grants subsequently won, to publication counts, to the h -index, and to citations has been analyzed (Bol et al., 2018). How it relates to the ease of *promotion*, by contrast, has been addressed only on a small scale and within a single national system (van den Besselaar & Sandström, 2015), so that this relationship remains far from clear.

Third, a number of studies rest on a single field or cohort, or on a small sample (Bol et al., 2018; Petersen & Penner, 2014; van den Besselaar & Sandström, 2015). Even work that covers several fields at once sometimes restricts attention to a selected group of researchers who are already prominent (Krauss et al., 2023). The careers that never reached prominence—the very cases in which the consequences of a missed early advantage should appear—thus drop out of the analysis. Few studies analyze across multiple fields and cohorts, and fewer still interpret their findings in light of the disciplinary background and historical context that produced them. It therefore remains largely unknown whether the advantage conferred by early research funding is specific to particular fields or is instead a universal phenomenon, and how the magnitude of the effect has shifted across scientific generations. This matters all the more because the relationship between research funding and research productivity can be context-dependent, varying in sign and in magnitude by field, by period, and by metric.

By addressing all three of these issues, the present study aims to refine what is known about the effect that early-career disparities in success exert on the long-run course of a career.

A further consideration is that the findings of existing work may depend on the institutional and social context of the country in which they were obtained. Carrying out the same kind of analysis in other countries and regions makes it possible to test how far those findings generalize, and opens the way to international comparison.

To overcome these limitations, the present study proceeds in two steps. We first construct an *explicit mathematical model* of the mechanism by which early-career research funding widens the gap. The model represents a positive feedback among a researcher’s ability, papers, research funding, and promotion. From it we derive not merely that “a gap exists,” but sharp, testable predictions about the *entire time path* the gap traces over a researcher’s career.

We then collect long-horizon, cross-field performance data from Japan’s largest researcher database, split researchers by enrollment cohort and by research field, and test the model’s predictions within each subgroup. In this way we can trace how the shape of the disparity *changes across eras* and *varies across disciplines*, instead of collapsing the evidence into a single pooled estimate over all researchers. Furthermore, by avoiding a sample that is disproportionately composed of a subset of researchers and instead encompassing a broader range of researchers, we expect to obtain more reliable results. In addition, by putting Japanese academia—where cumulative advantage has not previously been tested—on the examining table, we aim to establish where our results agree with and where they diverge from those of existing studies, and so to deepen what is known about cumulative advantage.

The remainder of the paper is organized as follows. Section 2 formalizes the mathematical model of how early-career funding generates a widening performance gap,

reports the model’s numerical analysis, and derives from it the hypotheses to be tested. Section 3 describes the data and empirical methods; Section 4 presents the empirical results; and Section 5 concludes with the study’s novelty, positioning relative to prior work, and contributions.

2 Mathematical Model

We express the ‘path dependent’ influence of early-career funding on performance as a discrete-time stochastic dynamic model with a positive feedback among talent, papers, funding, and promotion: funding raises paper output, cumulative papers raise the chance of the next grant, and cumulative papers also raise the chance of promotion. Papers are the visible, objective performance, corresponding to Mertonian “recognition.” Throughout, uppercase letters denote random variables and lowercase their realizations, and $\text{logit}^{-1}(x) = 1/(1 + e^{-x})$.

2.1 The basic idea

Before the formal specification, we state the idea of the model in plain terms. A researcher’s career is described as the repetition of a simple cycle in which resources and output feed each other:

- I. Talent determines the first grant.** Each researcher is endowed with a latent talent, and whether the *first* research grant of the career is obtained depends on that talent: the more talented a researcher is, the more likely they are to be funded at the outset. This is the “selection” step, and it is the only point at which talent alone decides an allocation.
- II. Funding and talent determine paper output.** In each period, the number of papers a researcher produces is determined jointly by their talent and by whether they currently hold a grant. Holding a grant raises the expected output, so two researchers of identical talent diverge in output as soon as their funding status differs.
- III. Papers determine the next grant.** Whether the *next* grant is obtained depends on the papers accumulated so far. Past output has become a visible credential, so the researcher who was funded—and therefore published more—is now more likely to be funded again.
- IV. Steps II and III repeat.** The cycle “funding → papers → funding” is iterated over the career. Because each pass through the loop converts an advantage in resources into an advantage in output and back again, a difference created once at step I is not merely preserved but amplified at every subsequent period.

The essential point is that, from step II onward, *talent no longer enters the allocation rule*: grants are awarded on accumulated papers, not on ability. Any gap that keeps growing after the first period is therefore attributable to the feedback loop itself rather than to the underlying difference in talent. We additionally attach promotion to the same accumulated output, so that the career consequences of the loop can be traced in a third, institutionally meaningful outcome. The remainder of this section formalizes steps I–IV and examines their consequences numerically.

2.2 Setup and dynamics

Individuals, talent, and relative rank.

Consider a set of n individuals $N = \{1, \dots, n\}$ over discrete time $t = 1, \dots, T$. Individual i has talent $E_i \sim \text{Beta}(a, b)$ (realization e_i), with percentile rank

$$F(e_i) = \frac{1}{n} \sum_{j \in N: e_j \leq e_i} 1. \quad (1)$$

First grant (talent-based selection).

The first-grant (early-fellowship) outcome depends on the talent rank:

$$D_{1i} \sim \text{Bernoulli}(\text{logit}^{-1}(\alpha + \beta F(e_i))), \quad \beta > 0, \quad (2)$$

so higher-talent individuals are more likely to be funded (selection on ability).

Paper production.

Given the funding indicator d_{ti} , the number of papers is Poisson with a talent-scaled rate that is higher when funded ($\lambda_1 > \lambda_2$); cumulative papers are G_{ti} :

$$A_{ti} \sim d_{ti} \text{Poisson}(\lambda_1 e_i) + (1 - d_{ti}) \text{Poisson}(\lambda_2 e_i), \quad G_{ti} = \sum_{\tau=1}^t a_{\tau i}. \quad (3)$$

Subsequent grants (papers beget funding).

The next-period grant depends on cumulative papers:

$$D_{(t+1)i} \sim \text{Bernoulli}(\text{logit}^{-1}(\gamma + \delta G_{ti})), \quad \delta > 0. \quad (4)$$

Equations (3)–(4) form the positive feedback “funding $\rightarrow$ papers $\rightarrow$ funding,” the source of the Matthew effect (rich-get-richer).

Promotion (papers beget promotion).

Promotion to associate professor or above is an absorbing state. With $p_{1i} = 0$ for all i ,

$$\pi_{(t+1)i} = p_{ti} + (1 - p_{ti}) \text{logit}^{-1}(\epsilon + \zeta G_{ti}), \quad (5)$$

$$P_{(t+1)i} \sim \text{Bernoulli}(\pi_{(t+1)i}), \quad \zeta > 0, \quad (6)$$

so cumulative papers also raise the promotion hazard.

2.3 Numerical analysis and results

We simulate the model with $n = 1000$, $T = 15$, $E_i \sim \text{Beta}(5, 5)$, $\lambda_1 = 5$, $\lambda_2 = 2$, and the logit coefficients in (2), (4), and (5). To evaluate the effect net of talent, we run a

Monte Carlo of $S = 100$ replications and, at each period t , estimate the least-squares regression

$$G_{ti} = c_0 + c_1 e_i + \beta_t d_{1i} + u_{ti}, \quad (7)$$

recovering the first-grant coefficient β_t conditional on talent e_i (and, replacing G_{ti} with cumulative grants and with promotion, the analogous talent-controlled coefficients for those outcomes).

In simulation, the first-grant group exceeds the no-grant group on all three outcomes—cumulative papers, cumulative competitive grants, and promotion to associate professor—and the gaps widen over time (Fig. 1). In the Monte Carlo regressions, the talent-controlled first-grant coefficient increases with t for cumulative papers and, in the same way, for cumulative grants, while for promotion it stays near zero over roughly the first half of the horizon and then climbs steeply (Fig. 2). Thus the initial award compounds through the feedback (3)–(4), producing a *pure cumulative advantage independent of selection*.

[Figure 1 about here]

[Figure 2 about here]

2.4 Hypotheses

The numerical analysis tells us what the feedback loop of steps I–IV implies: an advantage created once, at the very first allocation, does not fade as the grant expires but is converted into an ever-larger lead in cumulative output, in subsequent funding, and eventually in promotion—and it does so *even between researchers of identical talent*. This is a statement about a mechanism, and a mechanism is only as good as the predictions it makes about the world. Read as a prediction, the model says something concrete and checkable about actual research careers: if we take real researchers who did and did not win an early-career fellowship and follow them forward, their trajectories should separate in a particular way—not by a constant amount, but by an amount that *grows with the years elapsed*, and it should keep growing even once we hold constant the ability and environment on which the fellowship was awarded.

The empirical part of this paper asks whether that prediction is correct. Our hypotheses are as follows:

- H1.** The descriptive performance gap between early-funding recipients and non-recipients widens with elapsed time.
- H2.** Even at equal research talent, the gap between recipients and non-recipients widens with elapsed time.

H1 concerns the descriptive gap in research output: visualizing the actual data, we ask whether the career trajectories of researchers who did and did not obtain funding early on reproduce the pattern simulated in Fig. 1. H2 asks whether, between researchers of comparable talent, a performance gap nevertheless opens up depending on whether the early grant was obtained—that is, whether the result of Fig. 2 is also recovered in an

analysis of the actual data. If, as the model assumes, the first grant is allocated to the more talented researchers, then a widening gap between recipients and non-recipients is in a sense self-evident. In other words, talent may act as a confounder between early receipt and research output, so testing whether early funding itself produces a long-run performance disparity requires holding talent constant. Sections 3–4 test both hypotheses against the careers of about forty thousand Japanese researchers, using an early-career fellowship as the empirical counterpart of the model’s first grant and the graduate school of origin as the empirical counterpart of the model’s talent control.

3 Data and Methods

3.1 The empirical case: an early-career fellowship (JSPS DC)

As a concrete instance of “early-career research funding,” we use the *Research Fellowship for Young Scientists (DC)* of the Japan Society for the Promotion of Science (JSPS) ([Japan Society for the Promotion of Science, 2026](#)). DC is the most prestigious research-support program for doctoral students in Japan and the *first competitive funding in academia* to which a large fraction of graduate students apply.

DC is a scheme that supports doctoral students in a Japanese graduate school with both *living expenses* and *research funds*. Students selected for the fellowship receive from JSPS a stipend of about 200,000 JPY per month, together with research funds of up to 1.5 million JPY per year. Support runs from the time of appointment until the end of the standard doctoral programme, that is, for one to three years depending on the year of study in which the student is selected.

The success rate was originally around 20%, but with the recent growth in the number of applicants it has been trending downward; for fellows appointed from fiscal year 2026 it stands at roughly 11%. The fellowship is therefore a sharply competitive filter applied at the very entrance to a research career, and one that has become markedly more so over the period our data cover—a point we return to when we examine heterogeneity across enrollment cohorts.

We operationalize the model’s first-grant indicator D_{1i} (realization d_{1i}) as *DC receipt*. Because DC is competitively selected, recipients may have higher talent; separating this selection from cumulative advantage is the empirical focus (H2).

3.2 Data collection

From the Japanese researcher database *researchmap*, we collected researcher records (publications, professional history, research projects, education, profile) for eight broad research fields defined by the research-field code table ([researchmap, 2026](#)), de-duplicating across fields by researcher id (we exclude two small fields—Energy and Aerospace & Marine—which have few researchers). Of about 272,000 unique researchers, we retained those with at least one publication (about 182,000; 66.7%) and, among them, those whose doctoral-enrollment year could be identified (Table 1). *researchmap* users with no publications are likely to have dropped out of research activity. To identify active researchers for the analysis, we applied this filtering criterion. The analysis sample is about 40,000 researchers (14.7% of the collected set).

DC receipt was determined by normalizing professional-history text and detecting co-occurrence of “JSPS” and “research fellow,” classified into DC1/DC2/DC (number unknown)/unspecified (other fellowships by JSPS such as PD, SPD, RPD, and overseas fellows excluded). In addition, we use the competitive-funding records: because the JSPS Fellow research allowance (*Grant-in-Aid for JSPS Fellows*) is awarded only to special research fellows, a researcher who received it *during the doctoral period* is also counted as DC. This funding channel recovers DC status that the career text alone would miss, while the timing restriction (receipt while enrolled in the doctorate) separates DC recipients from post-doctoral PD fellows, who receive the same allowance only after graduation. Detection relies on self-reported histories. It should be noted, however, that precisely because these histories are self-reported, our procedure does not necessarily identify DC recipients with 100% accuracy.

[Table 1 about here]

3.3 Variables

Elapsed year is $t = (\text{calendar year} - \text{doctoral-enrollment year})$, $t = 0, \dots, 15$, with enrollment $+t > 2025$ treated as right-censored. The main outcomes are (i) cumulative papers, (ii) cumulative competitive grants (excluding the DC research allowance), and (iii) reaching associate professor or above (0/1). For each field and cohort we plot the mean $\pm 95\%$ confidence interval for recipients and non-recipients by elapsed year.

3.4 Analysis

Because talent is unobserved, we use the *graduate school of origin* as a proxy: students at the same school are assumed to have comparable talent and to have satisfied the same degree requirements, and prestigious schools tend to have more DC recipients, so educational background is also a confounder affecting both DC receipt and later performance.

For each elapsed year t we fit

$$y_{i,t} = \alpha_t + \beta_{1,t} \text{School}_i + \beta_{2,t} D_i + \varepsilon_{i,t}, \quad (8)$$

implemented with *graduate-school fixed effects* (the finest control) and standard errors clustered by school. Here $\beta_{2,t}$ is the effect of DC receipt on performance net of educational background; $\beta_{2,t} > 0$ means DC amplifies performance, and a $\beta_{2,t}$ that grows with elapsed time ($\beta_{2,t} \leq \beta_{2,t+1}$) indicates cumulative advantage (H1, H2). We analyze each field separately and by five-year enrollment cohort.

Use of AI tools.

The collection of the researcher records from *researchmap* and the statistical analysis reported below were carried out with the assistance of a generative AI coding tool (Claude Code, Anthropic; Claude Opus 4.8). The author specified the analyses,

inspected what the generated code actually did, and validated its output. The mathematical model of Section 2, including the ideas behind it and the code for its numerical analysis, was developed by the author.

4 Results

4.1 Descriptive statistics: the research landscape

Before isolating the fellowship’s effect, we describe the raw terrain on which it acts. Figures 3 and 4 give two views of one map: each plots mean cumulative papers (horizontal axis) against mean cumulative competitive grants (vertical axis) at ten years since enrollment, with the *marker size* scaled to the promotion rate (the share reaching associate professor or above) and the *marker shape* denoting the graduate-school prestige tier (T1: the former imperial universities plus Tokyo Tech and Hitotsubashi; T2: other selective national and top private universities; T3: all others). The two differ only in how the points are grouped and coloured—by five-year enrollment cohort (Fig. 3) and by field (Fig. 4). Three patterns stand out.

[Figure 3 about here]

[Figure 4 about here]

Successive cohorts write more but wait longer. Moving from the oldest cohort (dark) to the newest (bright) in Figure 3, the markers march *rightward and upward*—cumulative output at year ten climbs from about 13 papers in the early-1980s cohort to roughly 21 in the 2010s, and the grants held by year ten rise with it, from about two to four in the earliest cohorts to roughly three and a half to five in the most recent—while at the same time they *shrink*: the share reaching associate professor within ten years falls from 40% in the oldest cohort to about 25% in the recent ones.

Fields differ enormously, and along different axes (Fig. 4). Output is stratified nearly threefold: a decade in, nanotechnology and materials (about 32 papers) and information and communication (about 31) roughly triple the humanities and social sciences (about 12), with the remaining fields strung out in between. Funding follows a *different* ranking—nanotechnology again leads (about eight grants), the environmental and agricultural sciences punch well above their paper count (about five), and the natural sciences, despite healthy output, raise the fewest (about three). Promotion reshuffles the order once more: by year ten the applied, engineering, and humanities fields have sent roughly a third of their people to associate professor or above, whereas the life sciences—with their long postdoctoral queues—have promoted barely one in six (0.16, the smallest markers on the map). The fields do fan out along a gentle positive diagonal, then, but a paper, a grant, and a professorship are plainly earned on very different terms from one field to the next.

Prestige buys funding more than papers. In every cohort and in every field the circles (T1) sit higher on the vertical axis than the squares and triangles (T2 and T3): the hierarchy of graduate-school prestige may act most strongly on the acquisition

of research *resources*. Papers are far less orderly—in information and communication and in manufacturing it is the T2 graduates, not the T1 graduates, who accumulate more—and promotion shows almost no such gradient at all, running at roughly 0.25–0.30 in all three tiers. That output and funding nevertheless vary appreciably with the school of origin is precisely why we use it, throughout the analysis, as the variable that absorbs in a single bundle what we cannot observe directly—a researcher’s talent and the resources available to them.

Every DC effect estimated below must be read against these baselines: “one paper,” “one grant,” or “one professorship” means very different things across these strata, and it is precisely to strip them out that our regressions absorb the graduate school of origin.

4.2 The DC gap by enrollment cohort

Pooling the eight fields, we follow DC and non-DC researchers within five-year enrollment cohorts (1980–2019). The picture is a pair of scissors that opens with time: in every cohort the two groups set out together and then pull apart. Figures 5–7 show cumulative papers, competitive grants and promotion. For the two output measures the separation is unmistakable and grows year on year, precisely the widening the model predicts (Fig. 1); for promotion the gap is smaller and, when it appears at all, arrives late.

[Figure 5 about here]

[Figure 6 about here]

[Figure 7 about here]

4.3 The DC gap by field

Broken out by field, the widening output gap proves to be no quirk of a single discipline but a near-universal law (Figs. 8–10). For papers and grants the DC curve lies above the non-DC curve in most fields except the life sciences, and the separation is largest exactly where output is highest. By year ten the raw paper gap reaches about 17 papers in information and communication and 16 in manufacturing, roughly 12 in infrastructure and seven to ten in nanotechnology and the natural sciences, but only a couple of papers in the humanities and essentially none in the life sciences; and where the gap is large it keeps growing—by year fifteen the information and manufacturing gaps swell to roughly 25–35 papers. Funding tells the same story in miniature: recipients lead in every field, by roughly one to four grants, most clearly in nanotechnology, manufacturing and information. Promotion arrives on a different schedule, but in the end it points the same way. In every one of the eight fields the two curves cross at around ten years since enrollment (roughly years nine to twelve, depending on the

field), and from that point onward the share of DC recipients who have reached associate professor or above exceeds that of non-recipients (Fig. 10). Within the first ten years, by contrast, it is the non-recipients who are ahead in most fields. One likely reason is the presence of researchers who travel the reverse route—those who are already serving as associate professors on the strength of a master’s degree and only afterwards enroll in a doctoral programme to obtain the doctorate.

[Figure 8 about here]

[Figure 9 about here]

[Figure 10 about here]

4.4 The education-adjusted effect on papers

Do these gaps survive once talent are partly held fixed? Re-estimating β_2 on cumulative papers with graduate-school fixed effects (8)—separately within each cohort and, in turn, within each field—yields what is arguably the paper’s central empirical result: *the education-adjusted DC coefficient is positive, statistically significant, and rising with elapsed time.*

The cohort view (Fig. 11) shows that this is no fluke of one lucky generation. A few cohorts are exceptions—the estimate hugs zero in the 1980s cohorts, and in the 1990–94 panel it even drifts below zero in the final years—but in the majority of cohorts β_2 is positive and grows in magnitude as the years pass: about +3.8 papers at year ten for the 1995–99 cohort, +4.8 for 2000–04 and +5.7 for 2005–09, each comfortably significant, with the same upward slope visible in the 2010s cohorts over the span that can be observed. The exceptional cohorts are precisely the small ones that predate the fellowship’s expansion, and their confidence bands are wide enough to accommodate almost any story; the youngest cohorts, for their part, are censored past about ten years. Wherever the question can be posed cleanly, then, the answer repeats, generation after generation.

Broken out by field, the picture is equally uniform: *in all eight fields the education-adjusted coefficient is positive, statistically significant, and rising with elapsed time* (Fig. 12). Its size tracks the field baselines. The gap is largest in the high-output engineering and information fields—about +16 papers by year ten in information and communication and +14 in manufacturing—next roughly +8 to +11 in infrastructure, the natural sciences, the environmental sciences, and nanotechnology, and smallest, though still clearly positive, in the humanities and social sciences (+2.8) and the life sciences (+1.8). These estimates do not establish causality, but the association they describe is a striking one: DC is awarded for only one to three years and ends the day the doctorate is conferred, yet a decade later receipt of it is still associated with a disparity in published output. And the gap between those who obtained the fellowship

and those who did not does not settle: it continues to widen year after year, which is just the shape a Matthew effect would take.

[Figure 11 about here]

[Figure 12 about here]

4.5 The education-adjusted effect on competitive grants

We now apply exactly the design of Figures 11 and 12—the within-cohort and within-field trajectories of the graduate-school-fixed-effects coefficient β_2 (8)—to the second outcome, cumulative competitive grants. Grants behave almost exactly like papers. The cohort view (Fig. 13) tells the familiar story: apart from the small, noisy 1980s cohorts, the grant advantage is positive in every generation and widens as the career unfolds. Broken out by field, the pattern is just as uniform: in every one of the eight fields the coefficient departs from zero and climbs steadily with elapsed time (Fig. 14); by year ten it reaches roughly +2 to +3 grants in the most active fields (nanotechnology, manufacturing, information and communication) and stays clearly positive even in the lowest, the natural sciences (+0.9). Amplified disparities explained by mathematical models are thus visible in the data—a single early grant is followed by a growing, education-adjusted lead in the money a researcher goes on to raise.

[Figure 13 about here]

[Figure 14 about here]

4.6 The education-adjusted effect on promotion

The same design applied to promotion—the probability of having reached associate professor or above—uncovered a strikingly different, and more interesting, shape: a coefficient that is *J-shaped*. For the first several years β_2 is flat or even slightly negative—just what one expects if DC recipients spend a little longer in research or postdoctoral positions before taking up a faculty post—and then, from around year ten, it turns up sharply. The cohort view (Fig. 15) finds this shape in every generation: the coefficient sits at or just below zero for the first five to nine years—dipping to about -0.04 to -0.05 in the 1990–94 and 2010–14 cohorts—before turning up to roughly $+0.07$ to $+0.18$ by year fifteen, most strongly in the 1995–99 cohort; only the youngest cohort (2015–19) is censored too early for the turn to be completed.

Broken out by field (Fig. 16), the same J appears in almost every field and is clearly positive by year fifteen (about $+0.07$ in the humanities and social sciences, $+0.09$ in nanotechnology, $+0.13$ in information and communication, and $+0.13$ in the natural sciences). The one laggard is the environmental and agricultural sciences, where the

early dip is deepest (-0.08 at year ten) before a partial recovery. The lesson is that the *status* advantage of an early fellowship is real but *delayed*: it trails the advantages in output and funding by roughly a decade, surfacing only once the accumulated record of papers and grants has had time to pay off. Whatever else one makes of it, the delayed arrival of the promotion gap — exactly as the mathematical model predicted (Fig. 2) — is borne out by the actual data.

[Figure 15 about here]

[Figure 16 about here]

5 Conclusion and Discussion

5.1 Overview

Using a mathematical model together with large-scale empirical evidence, we have shown that research funding won at the very start of a career can open a cumulative gap in research performance. The model, built around a positive feedback among talent, papers, funding, and promotion, predicts that *even at equal talent the gap between those who win early funding and those who do not widens over time*. Following about 40,000 researchers registered in researchmap, spanning eight broad fields and four five-year enrollment cohorts, we find exactly what the model predicts: for cumulative papers and for competitive grants alike, both the raw gap between early recipients and non-recipients and the education-adjusted coefficient β_2 grow with elapsed time, in essentially every field and every generation. Promotion to associate professor obeys the same logic on a delayed schedule—the education-adjusted advantage is flat or slightly negative for roughly a decade and then turns sharply positive—so a fellowship that lasts one to three years and ends the day the doctorate is conferred leaves an imprint that reaches output first, resources next, and career status last, and that is still widening more than a decade later.

Does it follow that those who fail to win funding at the outset must resign themselves to giving up an academic career? That conclusion would be premature. It is true that this study, like much work before it, shows through a mathematical model and empirical analysis alike that those who secure funding first go on to accumulate more research output. Yet a study of decisions on U.S. National Institutes of Health (NIH) grants finds that applicants who fell just below the funding threshold left science at a higher rate, and that those who stayed went on, over the long run, to *systematically outperform* researchers who scored just above it. That result accords with the idea that “what does not kill me makes me stronger,” and the authors suggest that an early-career setback can lead to improved performance among those who persevere (Wang et al., 2019). It has likewise been shown that the point in a career at which a scientist produces their highest-impact work is essentially random (Sinatra, Wang, Deville, Song, & Barabási, 2016). There are also dimensions of career success that counts of

papers, of grants won, or of promotions cannot measure. Failing to accumulate output at the start does not mean that a career has been thrown away.

5.2 Contribution

Three features distinguish this account of the Matthew effect and of cumulative advantage in academia. First, it *joins theory to evidence*. An explicit generative model of inequality predicts that the gap grows even when talent is held constant. We then test whether this prediction of the mathematical model is replicated in an analysis using empirical data. In this way, the theoretical model and its empirical test are directly connected, without a conceptual leap between them.

Second, it *catches the process at its origin*. DC is the first competitive funding most Japanese doctoral students apply for, not a mid- or late-career grant awarded to those already established as researchers, which makes it unusually well-suited material for demonstrating that early success amplifies later success and widens inequality.

Third, it *tests the theory against big data* drawn from a large-scale researcher database, which allows us to expose how the widening of the gap differs across fields and across cohorts. For the science of science, this shows that cumulative advantage—long established as a qualitative principle (Merton, 1968), given an early stochastic formalization and cohort-level test (Allison et al., 1982), and, in economics, developed as a theory of increasing returns and path dependence (Arthur, 1989; David, 1985)—operates measurably at the level of individual careers, with a magnitude that can be tracked year by year and compared across fields and generations. Carrying out the analysis condition by condition in this way makes it possible to identify the fields in which the gap is largest and the generations that have been treated least favorably. Findings of this kind are useful when choosing where to aim, should measures to redress inequality be put into practice.

5.3 Policy Implication

These results carry several implications for efforts to redress inequality in scientific output. The long-run gap that may be opened by early success—in the present study, by winning research funding at the start of a career—is different in kind from a gap that arises out of differences in individual ability such as talent or research productivity. The inequality we have focused on, amplified in the manner of the Matthew effect, is structural: it is produced by the institution itself. Any prospect of redressing such a gap therefore argues for a reexamination of how the funding-allocation system is designed. Fortunately, this study formalizes the mechanism of the inequality process in a mathematical model. On the assumption that the model is sound, designers of the system can use it to debate what interventions would serve to redress inequality: that is, they can weigh, on the evidence the model provides, which of its parameters to adjust in order to reduce the gaps in papers, in grants, and in promotion that follow from winning funding early in a career. If the sole objective were to make inequality smaller, funders should avoid allocating grants preferentially to the most talented or most promising applicants. (Although we do not develop the point explicitly here, numerical analysis of the theoretical model under alternative rules for allocating the

initial success confirms that such a policy is effective in reducing inequality.) Doing so, however, comes at a price: while the gap generated by early success does shrink, the total number of papers produced by the system as a whole may fall. Aggregate research productivity and the degree of inequality therefore stand in a trade-off, and funding agencies should choose an allocation rule with that trade-off in view. Moreover, if success or failure at the very beginning determines later career success, those who fail at that first step may lose motivation. Avoiding this calls for institutional designs in which a single selection outcome does not lock in the rest of a career. Beyond this, enlarging the number of awards at the early-career stage, and preparing alternative funding or “second-chance” schemes for those who do not obtain one, are likely to be both effective and important. Finally, to arrest the widening of the performance gap, it may also be important that decisions on mid- and late-career funding be made without counting the size of the past track record (recognition) among the criteria. Past record is a standard item on grant applications, yet it favors precisely those who won funding early in their careers and were boosted by it. Excluding, from subsequent funding decisions, those elements whose advantage or disadvantage derives from the presence or absence of early funding could partly halt the growth of cumulative inequality. Which intervention to adopt in practice depends on what a given funding agency is trying to achieve, but the findings of this study should be of use in thinking the question through.

5.4 Limitation & Future work

Several limitations qualify these conclusions. Our estimates are not necessarily causal and indicate association only. The mathematical model shows that the presence or absence of early funding amplifies career inequality even once talent is held fixed, but talent is difficult to observe in real data. Among the available variables we take the graduate school of origin as the most plausible proxy for talent, yet it remains an imperfect one, because differences in ability and in research motivation persist among researchers within the same graduate school. Both DC status and the research outcomes also rest on self-reported records. A paper that exists in reality but has not been registered in the research database does not enter the data, and this concern about the accuracy of the database must be kept in mind. A further avenue for development is to build a theoretical model that can be fitted to real data by statistical modeling. Fitting the model to observed careers would make it possible to estimate the parameters of the real system, and hence to discuss which components are weak and which of them, if intervened upon, would redress inequality. We hope that future work will build on this study by overcoming the challenges set out above.

Acknowledgements. Not applicable.

Statements and Declarations

Funding. The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Competing interests. The author has no competing interests to declare that are relevant to the content of this article.

Ethics approval. Not applicable. This study analyses information that researchers have themselves made public in the *researchmap* database; it involves no intervention with, and no collection of private information from, human participants.

Consent to participate and consent for publication. Not applicable.

Data availability. The data analysed in this study were collected from the publicly accessible researcher database *researchmap* (<https://researchmap.jp/>), and are subject to that database’s terms of use.

Author contributions. The author confirms sole responsibility for the study conception and design, data collection, analysis, and manuscript preparation.

Use of AI tools. The mathematical model presented in this article—both the ideas behind it and the code used for its numerical analysis—was developed by the author. For the empirical part, the collection of the researcher data and the statistical analysis were carried out with the assistance of a generative AI coding tool (Claude Code, Anthropic; Claude Opus 4.8); the author validated what the AI-generated code did and checked its output. Generative AI was also used for formatting the manuscript and for English-language editing. The author takes full responsibility for the content of this article.

References

- Allison, P.D., Long, J.S., Krauze, T.K. (1982). Cumulative advantage and inequality in science. *American Sociological Review*, 47(5), 615–625, <https://doi.org/10.2307/2095162>
- Arthur, W.B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116–131, <https://doi.org/10.2307/2234208>
- Azoulay, P., Graff Zivin, J.S., Manso, G. (2011). Incentives and creativity: evidence from the academic life sciences. *RAND Journal of Economics*, 42(3), 527–554, <https://doi.org/10.1111/j.1756-2171.2011.00140.x>
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512, <https://doi.org/10.1126/science.286.5439.509>
- Bol, T., de Vaan, M., van de Rijt, A. (2018). The Matthew effect in science funding. *Proceedings of the National Academy of Sciences*, 115(19), 4887–4890, <https://doi.org/10.1073/pnas.1719557115>

- Clauset, A., Arbesman, S., Larremore, D.B. (2015). Systematic inequality and hierarchy in faculty hiring networks. *Science Advances*, 1(1), e1400005, <https://doi.org/10.1126/sciadv.1400005>
- Cole, J.R., & Cole, S. (1973). *Social stratification in science*. University of Chicago Press.
- David, P.A. (1985). Clio and the economics of QWERTY. *The American Economic Review*, 75(2), 332–337,
- de Solla Price, D. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27(5), 292–306, <https://doi.org/10.1002/asi.4630270505>
- Ding, Y., & Moreira, F. (2025, March). Funding and productivity: Does winning grants affect the scientific productivity of recipients? evidence from the social sciences and economics. *Scientometrics*, 130(3), 1831–1870, <https://doi.org/10.1007/s11192-025-05277-3>
- DiPrete, T.A., & Eirich, G.M. (2006). Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments. *Annual Review of Sociology*, 32, 271–297, <https://doi.org/10.1146/annurev.soc.32.061604.123127>
- Ebadi, A., & Schiffauerova, A. (2016). How to boost scientific production? a statistical analysis of research funding and other influencing factors. *Scientometrics*, 106, 1093–1116, <https://doi.org/10.1007/s11192-015-1825-x>
- Fortunato, S., Bergstrom, C.T., Börner, K., Evans, J.A., Helbing, D., Milojević, S., ... Barabási, A.-L. (2018). Science of science. *Science*, 359(6379), eaao0185, <https://doi.org/10.1126/science.aao0185>
- Gush, J., Jaffe, A.B., Larsen, V., Laws, A. (2018). The effect of public funding on research output: the New Zealand Marsden Fund. *New Zealand Economic Papers*, 52(2), 227–248, <https://doi.org/10.1080/00779954.2017.1325921>

- Hall, B., & Mairesse, J. (2024, March). *Explorations of cumulative advantage using data on french physicists* (NBER Working Paper No. 32285). National Bureau of Economic Research.
- Hottenrott, H., & Lawson, C. (2017). Fishing for complementarities: Research grants and research productivity. *International Journal of Industrial Organization*, 51, 1–38, <https://doi.org/10.1016/j.ijindorg.2016.12.004>
- Jacob, B.A., & Lefgren, L. (2011). The impact of research grant funding on scientific productivity. *Journal of Public Economics*, 95(9–10), 1168–1177, <https://doi.org/10.1016/j.jpubeco.2011.05.005>
- Japan Society for the Promotion of Science (2026). *Research fellowships for young scientists (DC)*. <https://www.jsps.go.jp/english/e-pd/>. (Accessed 2026)
- Krauss, A., Danús, L., Sales-Pardo, M. (2023, November). Early-career factors largely determine the future impact of prominent researchers: evidence across eight scientific fields. *Scientific Reports*, 13(1), , <https://doi.org/10.1038/s41598-023-46050-x>
- Lawson, C., Geuna, A., Finardi, U. (2021, April). The funding-productivity-gender nexus in science, a multistage analysis. *Res. Policy*, 50(3), 104182, <https://doi.org/10.1016/j.respol.2020.104182>
- Li, W., Aste, T., Caccioli, F., Livan, G. (2019). Early coauthorship with top scientists predicts success in academic careers. *Nature Communications*, 10, 5170, <https://doi.org/10.1038/s41467-019-13130-4>
- Merton, R.K. (1968). The Matthew effect in science. *Science*, 159(3810), 56–63, <https://doi.org/10.1126/science.159.3810.56>
- Merton, R.K. (1988). The Matthew effect in science, II: Cumulative advantage and the symbolism of intellectual property. *ISIS*, 79(4), 606–623, <https://doi.org/10.1086/354848>
- Petersen, A.M., Jung, W.-S., Yang, J.-S., Stanley, H.E. (2011). Quantitative and empirical demonstration of the Matthew effect in a study of career longevity. *Proceedings of the National Academy of Sciences*, 108(1), 18–23, <https://doi.org/10.1073/pnas.1016733108>

- Petersen, A.M., & Penner, O. (2014, October). Inequality and cumulative advantage in science careers: a case study of high-impact journals. *EPJ Data Science*, 3(1), , <https://doi.org/10.1140/epjds/s13688-014-0024-y>
- researchmap (2026). *Research-field code table (rmapv2, based on e-rad)*. https://researchmap.jp/outline/rmapv2/area/rmapV2_ResearchArea.pdf. (Accessed 2026)
- Salganik, M.J., Dodds, P.S., Watts, D.J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854–856, <https://doi.org/10.1126/science.1121066>
- Sinatra, R., Wang, D., Deville, P., Song, C., Barabási, A.-L. (2016). Quantifying the evolution of individual scientific impact. *Science*, 354(6312), aaf5239, <https://doi.org/10.1126/science.aaf5239>
- van den Besselaar, P., & Sandström, U. (2015, October). Early career grants, performance, and careers: A study on predictive validity of grant decisions. *Journal of Informetrics*, 9(4), 826–838, <https://doi.org/10.1016/j.joi.2015.07.011>
- van de Rijt, A., Kang, S.M., Restivo, M., Patil, A. (2014). Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences*, 111(19), 6934–6939, <https://doi.org/10.1073/pnas.1316836111>
- Wang, Y., Jones, B.F., Wang, D. (2019). Early-career setback and future career impact. *Nature Communications*, 10, 4331, <https://doi.org/10.1038/s41467-019-12189-3>
- Way, S.F., Morgan, A.C., Larremore, D.B., Clauset, A. (2019, April). Productivity, prominence, and the effects of academic environment. *Proceedings of the National Academy of Sciences*, 116(22), 10729–10733, <https://doi.org/10.1073/pnas.1817431116>
- Zuckerman, H. (1977). *Scientific elite: Nobel laureates in the united states*. Free Press.

Figure




Fig. 1 Simulated model trajectories (mean $\pm 95\%$ CI): cumulative papers, cumulative competitive grants, and promotion to associate professor, for first-grant recipients vs. non-recipients. Parameters: $n = 1000$, $T = 15$, $a = b = 5$, $\alpha = -5$, $\beta = 5$, $\gamma = -2$, $\delta = -10$, $\zeta = 0.25$, $\lambda_1 = 5$, $\lambda_2 = 2$.




Fig. 2 Monte Carlo first-grant coefficient, controlling for talent, for the three outcomes ($S = 100$ replications; thin lines are individual runs, the thick line their mean with a 95% CI band). The talent-controlled effect on cumulative papers (purple) and on cumulative grants (green) is positive and grows throughout, whereas the effect on promotion (orange) is negligible over the early periods and then climbs steeply, echoing the delayed emergence found in the data.




Fig. 3 Descriptive map grouped by *five-year enrollment cohort* (colour) and graduate-school prestige tier (marker shape: T1 circle, T2 square, T3 triangle). Axes are mean cumulative papers and grants at ten years since enrollment; *marker size is proportional to the promotion rate at year ten*. Whiskers are 95% confidence intervals; thin lines connect the three tiers within a cohort. All legends are placed to the right of the panel.




Fig. 4 The same map grouped by *field* (colour) and graduate-school prestige tier (marker shape: T1 circle, T2 square, T3 triangle). Axes are mean cumulative papers and grants at ten years since enrollment; *marker size is proportional to the promotion rate* at year ten. Whiskers are 95% confidence intervals; thin lines connect the three tiers within a field. All legends are placed to the right of the panel.




Fig. 5 Cumulative papers by DC status, by enrollment cohort (eight fields pooled; mean $\pm$ 95% CI).




Fig. 6 Cumulative competitive grants by DC status, by enrollment cohort (eight fields pooled; mean $\pm$ 95% CI).




Fig. 7 Reaching associate professor or above (rate) by DC status, by enrollment cohort (eight fields pooled).




Fig. 8 Cumulative papers by DC status, by field (mean $\pm$ 95% CI).




Fig. 9 Cumulative competitive grants by DC status, by field (mean $\pm$ 95% CI).




Fig. 10 Reaching associate professor or above (rate) by DC status, by field.




Fig. 11 First-grant (DC) coefficient β_2 on cumulative papers by elapsed year, estimated separately within each five-year enrollment cohort (eight fields pooled; graduate-school fixed effects, cluster-robust 95% CI). Within most cohorts the effect grows with elapsed time; the oldest cohorts are noisy and the youngest is right-censored.




Fig. 12 First-grant (DC) coefficient β_2 on cumulative papers by elapsed year, by field (graduate-school fixed effects; shaded band is the 95% CI). The effect grows over time in every field.




Fig. 13 Education-adjusted DC coefficient β_2 on *cumulative competitive grants* by elapsed year, estimated within each five-year enrollment cohort (eight fields pooled; graduate-school fixed effects, cluster-robust 95% CI). Same design as Fig. 11.




Fig. 14 Education-adjusted DC coefficient β_2 on *cumulative competitive grants* by elapsed year, estimated within each field (graduate-school fixed effects; shaded band is the 95% CI). Same design as Fig. 12. The effect grows over time in every field.




Fig. 15 Education-adjusted DC coefficient β_2 on *promotion* (reaching associate professor or above) by elapsed year, estimated within each five-year enrollment cohort (graduate-school fixed effects, cluster-robust 95% CI). Same design as Fig. 11.




Fig. 16 Education-adjusted DC coefficient β_2 on *promotion* (reaching associate professor or above) by elapsed year, estimated within each field (graduate-school fixed effects; shaded band is the 95% CI). Same design as Fig. 12. The coefficient is J-shaped: flat-to-negative early, then rising to clearly positive by year fifteen.

Table

Table 1 Sample funnel (unique researchers).

Stage	Count	Share
Collected (researcher_detail)	272,122	100%
+ at least one publication	181,641	66.7%
+ identifiable doctoral-enrollment year (analysis sample)	44,809	16.5%