The Lasting Effects of Government R&D

Jason Hall*

November 12, 2026

Click here for the most recent version

Abstract

Federal funding for research and development contributes to growth by creating new scientific human capital. To measure the aggregate importance of this channel, I construct a new dataset linking all U.S. patents from 1976 to 2025 to a near-universe of PhD dissertations. Individuals with research doctorates appear on roughly half of patents by U.S. inventors and are especially prevalent in critical technologies such as machine learning, biotechnology, and semiconductors. I show that the institution where individuals pursue their PhD explains surprisingly little of the variation in output, suggesting it is difficult to evaluate research skill ex ante. This implies that the marginal entrant is relatively close to the average, which mutes changes in quality as federal funding expands the pool. Moreover, the quantity of PhDs is highly sensitive to changes in federal funding. I use a new model of endogenous growth to quantify the effects of the roughly 40% reduction in federal support for R&D at universities proposed in the President's Budget Request for Fiscal Year 2026. The long-run effects of these cuts depend crucially on their duration: a 3-year cut generates a peak reduction in growth of roughly half a basis point and welfare costs of roughly 0.1% of permanent consumption, while a 23-year cut leads to a peak reduction in growth of 7 basis points and welfare losses of approximately 1%. Finally, I find that a revenue-neutral reallocation of funding from R&D tax credits to support for universities improves welfare by 0.28%.

*University of Minnesota. hall2160@umn.edu

I thank my advisors Anmol Bhandari, Kyle Herkenhoff, Hannes Malmberg, Todd Schoellman, and Kjetil Storesletten for their guidance on this project. I have also benefited from extensive conversations with Marco Bassetto, Loukas Karabarbounis, and Simon Mongey. I thank Manuel Amador, Lucas Belmudes, Zach Bethune, V.V. Chari, Luigi Falasconi, Andrew Goodman-Bacon, Illenin Kondo, Neil Mehrotra, Angelo Mendes, Shihan Shen, Mitchell VanVuren, Mike Waugh, Jacob Wright, and all seminar participants at the Federal Reserve Bank of Minneapolis, Rice University, University of Iowa, and the Growth/Public and Macro-Labor workshops at the University of Minnesota for their comments and suggestions as well. Any views expressed are those of the author and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3126 (CBDRB-FY26-P3126-R13126).

1 Introduction

The U.S. government has historically devoted nearly a quarter of a percent of GDP to supporting scientific research at universities. To do this research, universities have historically financed large numbers of graduate students. These graduate students then go on to contribute to the creation of new ideas and technologies for a long period of time. While significant efforts have been made to understand the impact of public R&D on innovation through idea spillovers, the effects of government funding on the creation of innovative human capital and the role that human capital plays in shaping aggregate growth have received far less attention. In this paper, I ask and answer two questions. First, what role do STEM PhDs play in aggregate U.S. innovation? Second, what would be the effects of of proposed cuts to federal support for R&D at universities through this human capital channel?

Quantifying the strength of this channel requires measuring three things. First, what percentage of innovation do PhDs account for? Second, how does the quantity of PhDs respond to changes in public R&D? Third, how does the quality of the marginal PhD change in response as the pool of PhD students expands? To do so, I construct a new dataset linking the near universe of U.S. PhD dissertations to the near universe of U.S. patents. This allows me to directly measure, at population scale, the contribution of PhDs to American innovation. In recent years, PhDs appear on roughly 44% of patents by U.S. inventors. Moreover, the number of PhDs is highly sensitive to variation in government support. Understanding how their quality changes empirically is complicated by unobserved selection on ability. To understand how quality responds, I construct a structural model that features selection on two margins. The first is selection into graduate school: individuals who believe themselves to be sufficiently good at research choose to go to graduate school, while others choose to immediately enter the workforce. The second concerns selection into research careers: after graduation, individuals who are less skilled at research exit, while those who are more skilled remain. To recover the degree of uncertainty about ability, I leverage information contained in the assignment of individuals to graduate programs. This, in turn, allows me to identify how quality changes in response to changes in the quantity of PhDs.

I then use the model to quantify the effect of the 41.4% reduction in federal support for university-performed R&D proposed in the President's Budget Request for Fiscal Year 2026. I find that the effect of these cuts depends crucially on their duration: a 3-year reduction leads to a reduction in growth that peaks at half a basis point, while the same cut prolonged an additional 20 years generates a peak reduction in growth of 7 basis points. The large uncertainty the model infers from the sorting patterns in the data also has implications for the design of innovation policy. A revenue-neutral reallocation of funding from R&D tax

credits to support for universities improves welfare by 0.28% of permanent consumption.

I begin my analysis by introducing a new dataset which I have constructed by matching the near universe of PhD dissertations in the United States to patents by U.S. inventors. The expansive nature of the match allows me to measure, at the population level, PhD contributions to innovation. Beyond the aggregate 44% share, PhD inventors are particularly common in frontier areas like medicinal chemistry, semiconductors, and machine learning. For instance, nearly 70% of semiconductor patents have at least one PhD inventor. To measure the responsiveness of PhD completions to government funding, I leverage exogenous variation in funding for the five major R&D accounts of the United States federal government (Fieldhouse and Mertens, 2025). I find that a 1% increase in appropriations relative to GDP leads to a roughly 0.55% increase in the ratio of PhDs in t relative to bachelors degrees in $t-6$, with a peak effect that occurs roughly 6 years after the narrative shock. Finally, I find that in my data, the distribution of patenting output across graduates of programs of different ranks overlap substantially. Where a PhD studied explains a small share, 1.5%, of the variance in patenting outcomes after the PhD.

The empirical patterns I see in the data are potentially polluted by selection into graduate school and into research. To separate the role of selection from underlying heterogeneity in ability, I construct a structural model featuring PhD graduates and endogenous growth. My model builds on the tradition of Schumpeterian growth models as in Grossman and Helpman (1991) and Aghion and Howitt (1992). Relative to the standard framework, my key addition is to explicitly model the process of graduate education. In my model, individuals are born and receive a public signal about their ability. Conditional on this public signal, they choose to apply to graduate programs or enter the labor force. On the other side of the market, not-for-profit graduate programs of different quality decide which applicants to admit, with an objective of admitting as much expected skill as possible. Enrolling in a graduate program entails a period of training, during which individuals produce services sold by the university and receive a stipend. Program size depends on both a university's ability to produce output and a convex cost of admission. Upon graduation, each individual's ability is revealed, and the new graduates decide, conditional on their true ability, whether to become a researcher or exit to the production sector. Increasing government funding for universities expands the quantity of enrolled students, increasing the number of individuals who learn their true research ability.

The graduate program block features a strict hierarchy, with better institutions admitting better graduate students. Nonetheless, in the model - as in the data - ex-post output distributions overlap substantially across program tiers. Through the lens of the model, matching characteristics of the ex-post output distribution conditional on the decile of program rank

allows me to distinguish predictable differences in ability from ex-ante uncertainty. The competitive equilibrium is generically inefficient. Two types of innovation policies - R&D subsidies and subsidies to universities - and an income tax are sufficient to achieve the first best.

I calibrate the full quantitative model to my novel empirical evidence on the output of PhDs. The model fits the data well. To validate the crucial margin for policy – how ability changes in response to changes in the quantity of PhDs – I simulate an increase in the number of PhDs comparable to the increase undertaken by Denmark in the early 2000s. I find that the model generates declines in innovation per PhD comparable in magnitude to those reported in Akcigit et al. (2025). This lends credence to the model's use in counterfactuals.

I then turn to the central policy questions of this paper: what would be the growth effect of the reductions in government support for R&D at universities proposed in the President's Budget Request for Fiscal Year 2026, and how costly would they be in welfare terms? The proposed cuts would have reduced aggregate federal R&D appropriations by $10.7 billion relative to Fiscal Year 2025 totals, a reduction of roughly 6%. This relatively modest top-line cut to federal R&D masked substantial compositional changes. While the Department of Defense (DOD) would have grown by about a quarter, the other major R&D accounts would have seen steep reductions, with the National Science Foundation (NSF) and National Institutes of Health (NIH) each declining by roughly half. Altogether this would have led to a roughly 41.4% reduction in federal support for universities.

I find that even dramatic cuts have substantially muted effects if the cuts are quickly reversed. A 41.4% reduction in federal support for universities would reduce the growth rate by less than a basis point at a welfare cost of 7 basis points in terms of consumption equivalent variation, but only if it were enacted for three years and then immediately reversed. If instead the cut persisted for an additional 20 years, growth would fall at peak by roughly 7 basis points, costing society 1.05% in terms of consumption equivalent variation. Small short-run effects arise because of the stock-flow nature of the talent pool: short term cuts permanently attrite only a small portion of the stock of human capital. Longer cuts attrite more of the human capital stock, and thereby increase the welfare costs of reductions in university support.

Next, I examine the relative efficacy of R&D tax credits and university subsidies. I find that R&D tax credits are less effective than federal support for R&D at universities. This arises for two reasons. The first is that R&D tax credits in the model are poorly targeted instruments. After graduating, researchers in the model know their ability. Many of those individuals earn significantly more than they would in production, even in the absence of the credit. Subsidizing those researchers is the cost of ensuring that researchers at the ability

margin remain involved in research after graduation.

The individuals induced by the university subsidies, on the other hand, do not know their true research ability. The large degree of uncertainty about ability means that, in fact, the marginal entrant is not particularly far from the average. After individuals graduate and realize their ability, only those who are talented at research choose to remain while less able individuals exit. Because many of those who enter will ex-post be far above the research cutoff, the expected skill an additional entrant contributes to research exceeds that of the researcher at the bottom of the distribution, exactly the individual the R&D tax credit pays to retain. The stock of scientific human capital therefore grows substantially. This pushes down the price of ideas, leading individuals who were previously at the bottom of the distribution of ability to exit to production and reinforcing positive selection among the remaining pool of researchers.

1.1 Related Literature

My model builds on endogenous growth models as in Romer (1990) and Aghion and Howitt (1992) and the broader literature examining the link between growth and talent (see e.g. Murphy et al., 1991; Aghion et al., 2017; Bell et al., 2019; Hsieh et al., 2019). Most related to this paper is recent work in Akcigit et al. (2025) which studies the interplay of education policies and research policies in a novel endogenous growth model, calibrated to the universe of Danish PhDs from 2002 to 2013, and finds that a mix of both education policies and R&D subsidies is most effective at increasing growth. Relative to that work, I endogenize the capacity choice of the university and emphasize the role of uncertainty about ability in determining the effects of human capital policy. My model is, to the best of my knowledge, the first to combine endogenous graduate capacity with a model of economic growth, which is crucial given the decentralized structure of the American educational system.

My emphasis on uncertainty also ties this paper to a strand of literature beginning with Manski (1989) that emphasizes the experimentation value of education. Theoretically, Terviö (2009) emphasizes the link between labor market imperfections and inefficient provision of talent. More recently, Hendricks and Leukhina (2018) develops a quantitative model of choice over the college experience that emphasizes the link between learning about one's own ability and dropping out of college. Wu (2026) studies the role of public learning about PhD research quality among computer science PhDs, taking as given their supply. Relative to this literature, I focus on graduate education, and rely on an approach that uses sorting patterns to infer ex-ante uncertainty, crucial given the lack of centralized information about pre-admission characteristics of graduate students.

Finally, I relate to the literature on the returns to government R&D. On the micro-side,

Azoulay et al. (2019) and Myers and Lanahan (2022) exploit windfall designs to recover the returns to NIH and DOE funding respectively. Babina et al. (2023) study the effects of federal funding shocks and find that shocks to government funding are partially offset by private funding, but that this tilts innovation towards patentable outputs. Tham et al. (2024) study the impact of NIH funding delays on graduate education and find that relatively short funding interruptions increase the probability of not working in the United States by 3 percentage points. More directly related to my study are existing studies on the macro impact of government R&D, most notably Dyevre (2023), Fieldhouse and Mertens (2025) and Campbell et al. (2026). Fieldhouse and Mertens (2025) construct a sequence of narrative shocks to the major federal R&D accounts, and estimate returns to government R&D range from 140 to 210 percent. Most closely related to this paper is Campbell et al. (2026), who discuss extensively the human capital channel I am interested in. Relative to them, and the rest of the literature, I focus on how uncertainty in ability magnifies the human capital effects of government R&D support and how this affects future economic growth.

2 Empirical Evidence

In this section, I lay out three pieces of evidence that will be key to measuring the aggregate effects of changes to graduate education. Each of the pieces of evidence addresses one piece of the logical chain discussed in the introduction.

2.1 New Data

The core measurement challenge that I tackle in this paper is that existing data on the educational background of inventors is relatively limited. Datasets such as UMETRICS feature a large quantity of information on the research activities and employees of a limited subset of research universities, while existing matches between the Survey of Earned Doctorates and PatentsView in Census environments are time limited and cover only individuals receiving dissertations from 2000 to 2017. My approach is to directly match two datasets known in the existing innovation literature. The first of these two is ProQuest Dissertation and Theses Global.[1] ProQuest has been the vendor of choice for most American (and to a lesser extent Canadian and British) universities as a repository of dissertations for much of the 20^{th} century.[2] The ProQuest metadata files that I use to construct the mapping provide information on the name of the dissertation author, their advisors, the timing of the dissertation, its

[1] See for instance Shvadron et al. (2026).

[2] In Appendix D I compare the number of dissertations in ProQuest with the number of new PhD awards in a given field over time from the National Science Foundation's Survey of Earned Doctorates (SED). The SED is a census of new PhD recipients that has been run since 1958. The two datasets align well until roughly 2010, after which a number of institutions switched to hosting their own dissertations on local repositories.

underlying content, and the title.

The source of my data on innovation is patents from PatentsView, which is a product of the Office of the Chief Economist at the United States Patent and Trademark Office (USPTO). Patents are a commonly used source of innovation data in the literature.[3] I make use of the disambiguated inventor identities, as well as information on the patent itself, including its title, technological classification, filing and grant dates. While I defer a detailed description of the matching algorithm to Appendix D and validation of its quality, I provide a brief overview here. The algorithm begins by matching candidate matches on name. After a set of candidate matches is generated, the algorithm scores each match based on the timing of patenting relative the dissertation, similarity between the technological content of each patent and the scientific content of the dissertation, and whether or not the advisor appears as a coauthor on granted patents. When restricted to comparable samples, my match produces patent shares that are within a couple of percentage points of the levels reported in the SED-PatentsView match, and when compared to a set of individuals derived from claimed patents on ORCID, the algorithm has a precision of roughly 93%.

2.2 The Contribution of PhDs to Aggregate Innovation

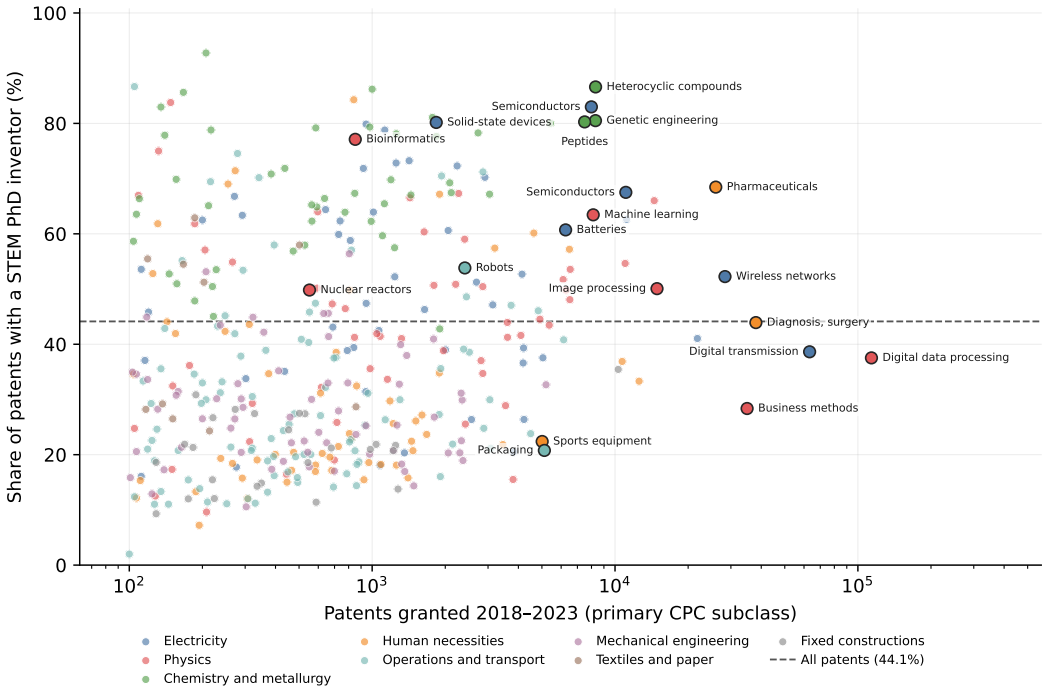
My dataset is, to the best of my knowledge, the most complete match of doctoral inventor education and United States patent data. Its broad coverage allows me to quantify the contributions of PhDs across the full technological spectrum. In aggregate, I find that roughly 44% of patents have an individual with a doctorate on them.

This aggregate number, however, masks substantial heterogeneity across technologies. Figure 1 plots CPC subclasses and compares the number of patents granted from 2018 to 2023 with the share of patents with a PhD inventor in that subclass. Consistent with the close ties between university research and drug discovery, many of the most PhD intensive subclasses are biological in nature (e.g. "Genetic engineering"). Likewise, semiconductor and machine learning patents disproportionately feature PhD inventors. These three technology groupings form three of the ten categories that the National Science Foundation has, under the CHIPS and Science Act, dubbed critical technologies. Across all ten categories, nearly 50% of associated patents have at least one PhD inventor.[4]

[3] See e.g. Akcigit et al. (2025),Akcigit and Kerr (2018),Akcigit and Goldschlag (2023),Babalievsky (2022),Gross and Sampat (2025),Arora et al. (2023), to name a few.

[4] The categories are: AI/ML & Autonomy, High-Performance Computing and Semiconductors, Quantum Information, Robotics and Advanced Manufacturing, Disaster Prevention, Communications & Immersive Technologies, Biotechnology and Medical Technology, Data Storage & Cybersecurity, Energy & batteries, and Advanced Materials. See Appendix E.2 for a complete listing.

Figure 1: The PhD Share Across Technologies



Notes: Each marker corresponds to patents in a primary CPC subclass. The color of each dot corresponds to the associated CPC section. Primary CPC subclass is defined in terms of first inventional CPC assignment on the patent. There are two labeled semiconductor subclasses because of a split in the taxonomy in 2023 and only partial reclassification by USPTO.

2.3 Government Funding and Doctoral Capacity

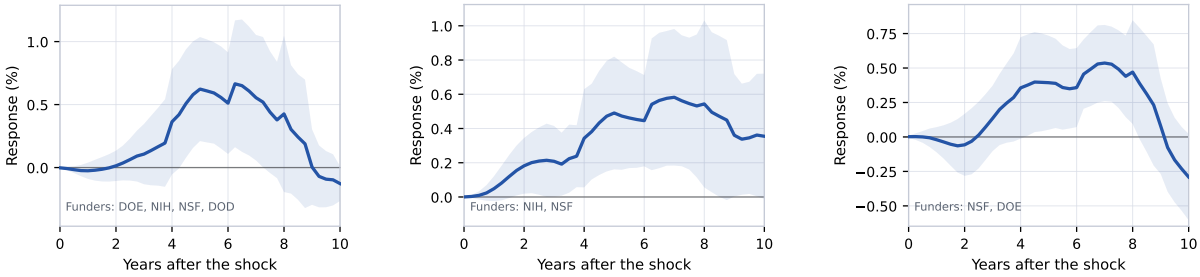
I next turn to measuring the effect of changes in government funding for R&D on the quantity of PhDs introduced. Naive regressions of PhDs awarded on quantity of government funding are subject to concerns of policy endogeneity: both individuals and the government perceive a new technology as a particularly promising approach to solving a problem. I address this concern by using the narrative shocks of Fieldhouse and Mertens (2025). These shocks were constructed directly from the historical context of the R&D appropriations process, and cover the five major R&D budgetary accounts: the National Science Foundation, Department of Health and Human Services, Department of Energy, Department of Defense, and National Aeronautics and Space Administration.

To estimate the elasticity of PhD production to government funding, I project the ratio of PhDs in year t to bachelors in year $t-6$ (a quantity I will refer to as the "PhD yield") at a given horizon h on the narrative shocks, controlling for agency appropriations and a set of standard baseline controls. I map agencies to fields so that e.g. the National Institutes of Health fund biology PhDs and not computer science PhDs, and scale both the narrative shocks and corresponding agency appropriations by prior year agency-linked

university obligations from the Survey of Federal Funds for R&D. Letting f denote a field, t a quarter, and j a lead, I then run a local projection of the form:

$$\overline{\ln \text{PhD}}_{f,t+h} = \alpha + \beta_h \cdot \text{shock}_{f,t} + \sum_{j=1}^{12} \gamma_j \ln A_{f,t-j} + \sum_{j=1}^{12} \delta_j \ln \text{PhD}_{f,t-j} + \sum_{j=1}^{12} \boldsymbol{\phi}_j \mathbf{X}_{t-j} + \varepsilon_{f,t+h} \quad (1)$$

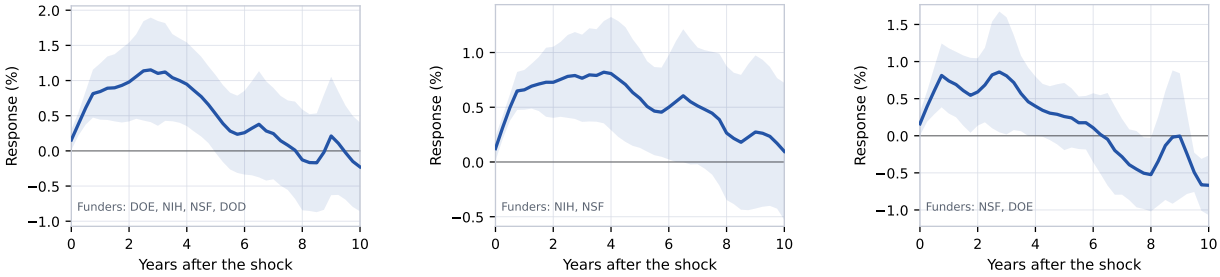
Figure 2: The Response of PhD Yield to Narrative Shocks



(a) Engineering (b) Biology (c) Physics and astronomy

Notes: Response of the field's yield rate – SED doctorates per bachelor's degree awarded six years earlier – to a Fieldhouse-Mertens appropriation shock of 1 percent of the field's university-directed appropriations, horizons 0–10 years. Quarterly local projections with 12 lags of log appropriations, the outcome and the cyclical controls; shaded band is the Lazarus HAR 95 percent interval. Percent per 1 percent shock.

Figure 3: The Response of Research Funding to Narrative Shocks



(a) Engineering (b) Biology (c) Physics and astronomy

Notes: Response of the field's university-directed appropriations as a share of potential GDP to a Fieldhouse-Mertens appropriation shock of 1 percent, horizons 0–10 years. Quarterly local projections with 12 lags of log appropriations, the outcome and the cyclical controls; shaded band is the Lazarus HAR 95 percent interval. Percent per 1 percent shock.

Figure 2 plots the time path of the coefficients of interest β_h for three major fields, while Figure 3 plots the corresponding fiscal impulse.[5] Fiscal impulses peak within the first three years after the shock before decaying downwards. The estimates from the yield rate imply

[5]See Appendix E for the full collections of fields and corresponding fiscal impulses.

that a 1% increase in university directed appropriations relative to potential GDP generates an increase in PhD yield of between 0.5 and 0.6 after 6 years.

2.4 Measuring Quality

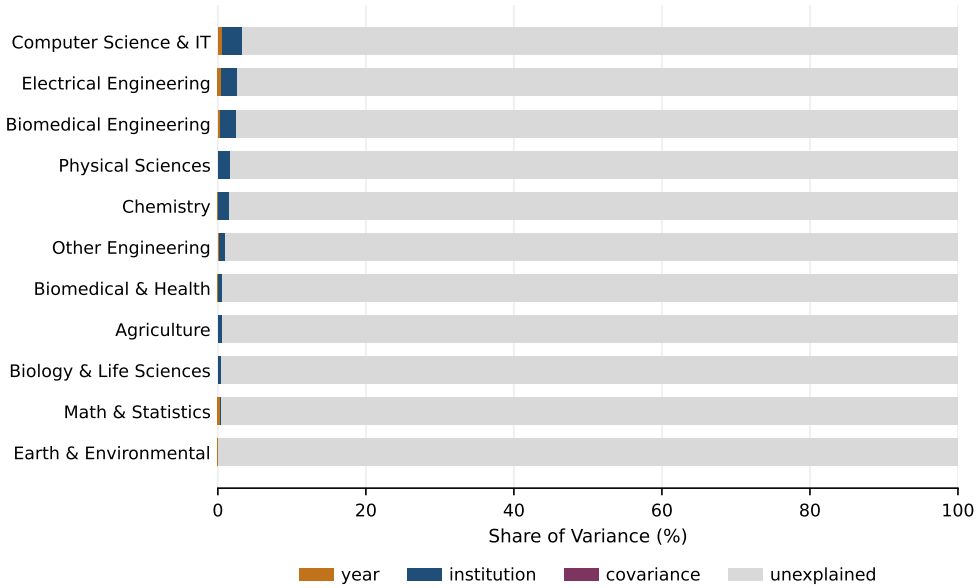
The total effect of changes in federal support depend not just on the number of new PhDs awarded but also how their quality changes as graduate capacity expands. If quality does not move much, then the effect of changes in quantity will be primarily driven by the sheer number of new PhDs. If, on the other hand, quality falls substantially, then effects weighting all the PhDs the same will overestimate the benefit of increasing federal support for universities.

One approach to measuring quality would be to project characteristics of a cohort, such as patents per PhD on the exogenous shock series used to estimate the responsiveness of PhDs to government funding. Appendix E.1 documents the results of this exercise and broadly finds very little statistical evidence of the effects of exogenous increases in government funding on measures of cohort quality. The short time horizon of these shocks and the large surge in patenting during the mid-1990s means that caution should be exercised before inferring that the quality of PhDs does not change in response to changes in quantity.

My preferred approach to measuring quality is instead to measure how well graduate programs perform at screening for research talent. In general, doing so requires the econometrician to measure what is known by the admitting program at time of application. Systematic measures of such information do not currently exist, which necessitates a different approach. My insight is that the initial assignment of students to graduate programs is informative about initial information about graduate quality. To understand the logic, consider a revealed preference argument. There are two programs, one of which is better than the other, and two students, who are likewise ranked. Suppose that better programs prefer better students. In a world of perfect information, the better student will be matched with the better program. In the presence of imperfect information, however, this is no longer always true. The frequency with which these rank reversals occur is informative about uncertainty. More generally, overlap between the ex-post output distributions of graduates from better and worse programs is therefore informative about the degree of uncertainty at the time of application.

I assess the strength of sorting in two complementary ways: first, by asking how much of the variance in ex-post output is explained by the program attended, and second, by comparing the full distribution of graduates' output across programs of different ranks. Let i denote an individual, g the graduate program attended, and t the year the doctorate was

Figure 4: Variance Decomposition by Field



Notes: Output is measured as 10-year post dissertation inventive output. To address finite sample bias in the estimation of fixed effects, I implement the leave-one-out correction of Kline et al. (2020).

awarded. Consider the variance decomposition

$$\mathbb{V}[y_{igt}] = \mathbb{V}[\alpha_t] + \mathbb{V}[\gamma_g] + 2\mathbb{C}[\alpha_t, \gamma_g] + \mathbb{V}_U, \tag{2}$$

Equation 2 splits the variance of patenting output in the first 10 years after graduation into four components: the variance of year effects, the variance of program effects, their covariance, and residual variation. The residual captures both heterogeneity in ability within programs and luck. Because patenting propensity differs substantially across fields, I estimate the decomposition separately by field.

Figure 4 provides the results of that exercise. Across fields, institution explains an average of 1.5% of the variation in inventive output in the 10 years post dissertation. In computer science, I find that doctoral institution explains only 2.75% of variation in ex-post inventive output.[6] This finding of the apparent unimportance of PhD institution is robust to a number of potential concerns. In particular, appendix E.4 shows robustness to alternative horizons of measured output, to quality adjustments of patent output, to decomposing logs instead of counts, to Poisson regression variance decompositions, to winsorizing patent counts at different thresholds, to more granular field taxonomies, and to restricting the sample to individuals with unambiguous names.

These findings suggest that there is significant dispersion in ex-post output that is not

[6]Consistent with my finding for patents, Wu (2026) finds that, among Computer Science PhDs, institution and cohort fixed effects explain only 5.5% of the variation in publications five years after degree.

due to institution. To take the logical step that this is ex-ante unpredictable, would, in general, require additional information on what is known about an individual at time of application. Unlike in the setting of college choice, to the best of my knowledge there is no systematic dataset linking graduate applicants, the programs to which they applied, and their ultimate enrollment decisions.[7] To overcome these data concerns, I rely on an approach that leverages sorting patterns. My approach builds on the intuition that there is a degree of positive assortative matching across programs; better ranked programs match with ex-ante more promising students. To the extent that this is true, the initial assignment of PhDs to programs conveys substantial information about the market belief at admission. My dataset then allows me to directly measure ex post output, and cases where individuals from worse programs outperform those from better programs are informative about the degree of uncertainty at admission time.

To measure the quality of a program, I use the National Research Council 2010 "Assessment of U.S. Doctoral Programs." These rankings cover more than 5000 graduate programs.[8] I restrict consideration to STEM fields. I take the midpoint of the reported program rank and then aggregate programs into 11 broad fields, where broad field rankings are constructed as the mean percentile graduate program ranking. These rankings are not constructed using any direct information about the ex-post patenting behavior of graduates. Nonetheless, they are informative about the three margins of patenting behavior ex-post: graduates of better ranked programs *on average* patent more, are more likely to ever patent, and conditional on patenting are more likely to continue patenting (Appendix E.5).

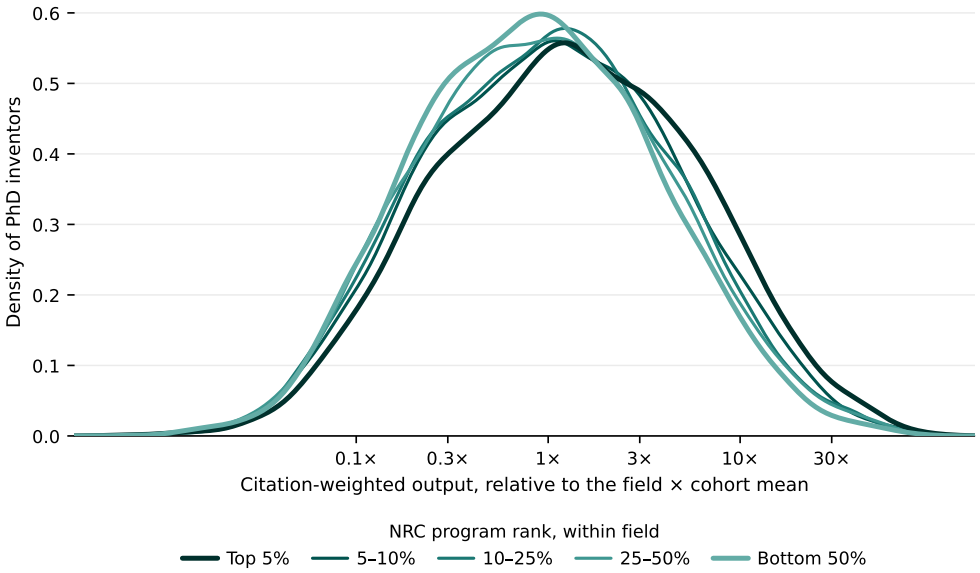
Consistent with Figure 4, cross institution dispersion is dwarfed by within institution dispersion. Figure 5 plots the distribution of quality adjusted patents filed in the 10 years after finishing the PhD, separated by program tiers. To summarize the degree of overlap: an individual from a bottom 50% program will produce more quality adjusted patents than an individual from a top 5% program roughly 42% of the time.[9] Appendix E.6 shows that this result survives rankings estimated from patent output, alternative horizons, winsorizing, finer field definitions, and restricting the sample to individuals with unambiguous names. Taken together, these findings are strongly suggestive of substantial ex-ante uncertainty

[7]Moreover, there does not even exist data on the number of distinct *applicants* across graduate programs.

[8]The rankings were constructed in two different ways. The R rankings were constructed by first surveying faculty, and then projecting faculty rankings on observable characteristics of programs. The second, the S rankings, were constructed by using faculty-provided weights on observable program characteristics. I primarily use the R ranking, but all results are robust to instead using the S.

[9]Despite the difference in number of programs, these two groups account for roughly the same number of patenters. This number does vary across fields, but even in the most separated field, electrical engineering, an individual from a bottom 50% program will produce more patents than one from the top 5% roughly 40% of the time.

Figure 5: Distribution of Quality Adjusted Output by Program Tier



Notes: Output is measured as 10-year post dissertation quality adjusted inventive output. Per patent quality adjustment is performed by dividing 5-year external forward citations by mean of the same within that grant year and primary CPC subclass, while inventor productivity is measured as the sum of quality adjusted output. Densities are constructed using gaussian kernel density estimation on dequantized output residualized by field × cohort fixed effects after dropping zeros.

about the research capability of PhD applicants. This uncertainty would tend to push the quality of the marginal entrant towards that of the average, which would amplify the effects of changes in the quantity of PhDs on aggregate innovation. To quantify precisely how much uncertainty exists, and the implications of that uncertainty for aggregate innovation, I now turn to a model of endogenous growth with heterogeneous graduate programs and educational experimentation.

3 Model

My model has three key features. First, individuals make a choice to apply to graduate school, and, upon graduating, to engage in R&D or exit to the production sector. Second, there is substantial ex-ante uncertainty about the research capability of graduate students. Third, a collection of universities admit applicants optimally using the information available at time of application.

3.1 Physical Environment

Time is continuous and indexed by $t \in \mathbb{R}$. There is a continuum of individuals who discount the future at rate ρ and die at rate δ. An individual, once dead, is replaced by a newborn. Individuals have linear flow utility with respect to consumption, so that $u(c) = c$.

Individuals are characterized by their skill at production z_P and their skill at research z_R. For simplicity, I assume that z_P is common across individuals and will normalize it to 1. There are two important characteristics of research skill. The first is that it is heterogeneous across individuals, and distributed according to some distribution G with positive support and finite mean. The second is that it is observed imperfectly at birth by all agents in the economy. What agents observe is a noisy signal n_R of underlying research ability. I assume that the distribution of z_R conditional on n_R is increasing in n_R in the sense of first-order stochastic dominance: for any $n < n'$, $\mathbb{P}[z_R \geq \bar{z} \mid n_R = n'] \geq \mathbb{P}[z_R \geq \bar{z} \mid n_R = n]$ for every threshold $\bar{z}$, with strict inequality for some $\bar{z}$. An individual's research ability z_R is only revealed after completing graduate education. I will describe this process in more detail after describing the production of physical goods in this economy.

There is a final good in this economy that is produced with a standard unit elastic aggregator across a continuum of intermediate product lines:

$$\ln(Y_t) = \int_0^1 \ln(y_{it}) di \tag{3}$$

Each product line is produced by a frontier, labor-only technology with an endogenous quality q_{it}:

$$y_{it} = q_{it} \ell_{it} \tag{4}$$

Improvements to q_{it} arrive endogenously as the outcome of research undertaken by a match between a researcher and an innovative firm. A researcher of skill z_R generates new ideas at a rate $\varphi_R z_R$ when matched to an innovative firm, and without an innovative firm, a researcher cannot generate new ideas. When an idea arrives, it improves the productivity of a random product line by a quality step of size λ, as in e.g. Grossman and Helpman (1991).

There are J representative graduate programs which enroll students and combine the research ability of those students with their own idiosyncratic productivity α_j. Define the expected research ability of an individual with signal n_R as $m(n_R)$, the stock of enrolled students at university j at time t with signal n_R by $N_{jt}(n_R)$, the flow of newly admitted students by $a_{jt}(n_R)$, and the aggregate productivity frontier by Q_t. Graduate programs produce a flow of university services according to: ‘’

$$A_{jt} = Q_t \int \alpha_j m(n_R) N_{jt}(n_R) dn_R \tag{5}$$

The cost of admission of a mass of $x_{jt} \equiv \int a_{jt}(n_R) dn_R$ students is given by the convex cost function:

$$K(x_{jt}) = Q_t \frac{\kappa x_{jt}^{1+\nu}}{1+\nu} \tag{6}$$

This cost is paid in terms of final goods, and scales with the aggregate productivity level of the economy. University services are perfectly substitutable with the final good.

Graduate Programs. To become a researcher, an individual must enroll in a graduate program. Graduate programs are non-profits that must break even period by period, and take the stipend schedule $Q_t s_t(\cdot)$ as given. A program in tier j chooses an admissions policy $a_{jt}(n) \geq 0$ based on the public signal. Each enrolled student with signal n generates university services $\alpha_j Q_t m(n)$.[10] A student admitted at date τ with signal n is paid $Q_t s_\tau(n)$ at each date $t \geq \tau$ while enrolled. Let $N_{jt}(n)$ denote enrollment of individuals with signal n and let B_{jt} denote the program's productivity-adjusted stipend bill. These evolve according to

$$\dot{N}_{jt}(n) = a_{jt}(n) - (\mu + \delta) N_{jt}(n), \tag{7}$$

$$\dot{B}_{jt} = \int s_t(n)\, a_{jt}(n)\, dn - (\mu + \delta) B_{jt}, \tag{8}$$

Equation (8) reflects that each student's stipend in normalized terms is fixed at admission. The program maximizes quality of its enrolled students, discounted at a rate ρ, subject to the non-profit constraint and the laws of motion for enrollment and stipend bills:

$$\begin{aligned} \max_{\{a_{jt}(\cdot)\geq 0\}_{t\geq 0}} \quad & \int_0^\infty e^{-\rho t} \int m(n)\, N_{jt}(n)\, dn\, dt \\ \text{subject to} \quad & (7),\ (8), \\ & Q_t \Big[\alpha_j \int m(n)\, N_{jt}(n)\, dn - B_{jt}\Big] = K(x_{jt}), \qquad x_{jt} = \int a_{jt}(n)\, dn, \\ & \big(N_{j0}(\cdot),\, B_{j0}\big) \text{ given.} \end{aligned} \tag{9}$$

Individuals. At birth, individuals are endowed with a share in the mutual fund that owns all of the firms in this economy, observe the public signal of their ability n_R, and make a choice as to enter a graduate program or not based on that. If they choose to enter production, they earn a wage $w_t z_P$. If they choose to enter graduate school, they earn a stipend $s_t(n_R)$ and graduate at a rate μ. At graduation, they learn their research skill z_R and choose between receiving $\mathcal{P}_t(z_R)$ in the research sector and $w_t z_P$ in the production sector.

[10] Because universities pick a positive density of individuals of a signal type, a heuristic law of large numbers means this is equivalent to individuals producing university services with their true z_R.

Their life is thus characterized by a set of three Hamilton Jacobi Bellman equations. The first of these is that of an individual working in the production sector:

$$(\rho + \delta)W_t - \dot{W}_t = u(w_t z_P + \Pi_t) \tag{10}$$

The second is that of an individual of skill z_R working in the research sector:

$$(\rho + \delta)\mathcal{R}_t(z) - \dot{\mathcal{R}}_t(z) = u(\mathcal{P}_t^*(z) + \Pi_t)$$

Finally, there is that of an individual in graduate school:

$$(\rho + \delta + \mu)E_t - \dot{E}_t = u(Q_t s_\tau(n_R) + \Pi_t) + \mu \mathbb{E}_{z_R}\left[\max\left\{W_t, \mathcal{R}_t(z)\right\}\middle| n_R\right] \tag{11}$$

Innovative firms. New ideas in this economy improve the productivity of a random product line by a step-size λ. They are created by a combination of a researcher and an innovative firm. This match produces ideas according to the innovation production technology

$$I = \varphi_R z_R \tag{12}$$

To understand the determination of surplus in this economy, it is helpful to define two objects that I will call the pseudo-surplus value of a firm and worker. Specifically, let $\mathcal{P}_t(z_R)$ denote the monetary payment from the firm to the worker at a given point in time. The value of an innovative firm given a sequence of future payments $\{\mathcal{P}_t(z_R)\}$ satisfies the following Hamilton Jacobi Bellman-Equation:

$$(r^* + \delta)\mathcal{V}_t(z_R; \{\mathcal{P}_s(z_R)\}_{s\geq t}) - \dot{\mathcal{V}}_t = V_t \varphi_R z_R - \mathcal{P}_t(z_R) \tag{13}$$

This equation takes the future sequence of payments $\mathcal{P}_{t+\varepsilon}(z_R)$ as given. The analogous expression for the worker is given by:

$$(\rho + \delta)\mathcal{R}_t(z_R; \{\mathcal{P}_s(z_R)\}_{s\geq t}) - \dot{\mathcal{R}}_t = u(\mathcal{P}_t(z_R) + \Pi_t) \tag{14}$$

Let $\theta \in (0, 1)$ denote the bargaining power of a researcher. The researcher's outside option is to exit permanently and enter production, with value W_t, and the firm's outside option is zero. The Nash bargaining problem is given by:

$$\mathcal{P}_t(z_R) \in \arg\max_p \left(\mathcal{R}_t(z_R; p) - W_t\right)^\theta \left(\mathcal{V}_t(z_R; p)\right)^{1-\theta}, \tag{15}$$

where the bargain at date t takes continuation values as given.[11]

Intermediate Goods Producers. Within a product line, goods are perfect substitutes and firms compete on price. Letting π_{it} denote flow profits from product line i, one can write the Hamilton Jacobi Bellman equation of a product line as:

$$(r^* + \iota_t)V_t - \dot{V}_t = \pi_t \tag{16}$$

Resource Feasibility. Resource feasibility in this environment is given by two equations. The first concerns the use of the final good:

$$C_t + \sum_j K(\int a_{jt}(n_R)dn_R) = Y_t + Q_t \sum_j \alpha_j \int \mathbb{E}[z_R|n_R]N_{jt}(n_R)dn_R \tag{17}$$

The second concerns the allocation of labor across the three sectors of the economy. Letting L_t denote the aggregate quantity of production labor, N_t the headcount of graduate students, and R_t the head-count of researchers, we have:

$$L_t + N_t + R_t = 1 \tag{18}$$

This completes the description of the environment. I now define an equilibrium.

Definition: A *dynamic competitive equilibrium* given a world interest rate r^* is prices $\{p_{it},\ w_t,\ s_t(n_R),\ \mathcal{P}_t(z_R)\}$, firm values $\{V_t, \mathcal{V}_t(z_R)\}$, quantities of labor $\{L_t,\ N_t,\ R_t\}$,skill $\{H_{Rt}, N_{jt}(n)\ \}$, admitted students $a_{jt}(n_R)$, $\{C_t, Y_t\}$, and $\{g_t, \iota_t\}$ such that the following are true:

1. $\mathcal{P}_t(z_R)$ solves the Nash bargaining problem of a match between a researcher and an innovative firm.
2. The laws of motion for all stocks of labor and skill are satisfied at all points in time.
3. Productivity grows at the rate g_t, and this rate is consistent with the stock of skill in research.
4. $g_t = \iota_t \ln(\lambda)$, where ι_t is the rate of creative destruction.
5. Worker and firm choices are optimal given prices, and their values are consistent with such choices.

[11] See Appendix B.1 for the precise formulation of this problem as the limit of bargaining over a short interval.

6. Graduate programs choose admissions policies optimally given prices and the laws of motion for enrollment and stipend bills.

7. The admissions market and good markets clear.

8. All other labor markets clear.

A balanced growth path is a special case of such an equilibrium in which prices and quantities of output grow at constant rates and the distribution of research skill is stationary.

3.2 Characterization

I now characterize the competitive equilibrium, beginning first with the relatively standard growth block. From now on, I will assume that $r^* = \rho$. The assumption that products within a line are perfectly substitutable and firms compete on price ensures that the optimal behavior is that of limit pricing. After normalizing the price of the final good to 1, standard algebra implies that:

$$\pi_{it} = \left(\frac{\lambda - 1}{\lambda}\right) Y_t \tag{19}$$

In addition, the aggregate production function in this economy is given by:

$$Y_t = Q_t L_t \tag{20}$$

Where Q_t is the standard Cobb-Douglas productivity index. Equilibrium production wages are given by:

$$w_t = \frac{1}{\lambda} Q_t \tag{21}$$

The law of motion for aggregate research skill H_{Rt} is given by the following ODE:

$$\dot{H}_{Rt} = \mu \int_{n_R \geq \underline{n}} \mathbb{E}[z_R \mathbf{1}[z_R \geq \bar{z}_t^*] | n_R] N_t(n_R) dn_R - \delta H_{Rt} - \int z e_t(z) dz \tag{22}$$

The equilibrium growth rate at any instant is then given by:

$$g_t = \ln(\lambda) \iota_t = \ln(\lambda) \varphi_R H_R \tag{23}$$

I now provide some properties of the graduate program block of the model. My first result simplifies the problem. I show that the problem in (9), which requires tracking the history of the full distribution of admissions decisions, is equivalent to a simpler problem that tracks only aggregate payroll obligations. Because individuals graduate at a constant rate,

both the revenue generated by graduate students ($\alpha_j M_{jt}$ below) and payroll expenses (B_{jt}) decline at rates proportional to their current levels. New payroll obligations matter only insofar as they change the aggregate net revenue of the program (P_{jt}). When government subsidies are provided per student, the program must also track the aggregate number of enrolled students.

Lemma 1. *Let $P_{jt} \equiv Q_t\big(\alpha_j M_{jt} - B_{jt}\big)$ and $M_{jt} \equiv \int m(n) N_{jt}(n) dn$. The solution to the university problem* (9) *is equivalent to the solution to the following problem:*

$$\begin{aligned} \max_{\{a_{jt}(\cdot)\geq 0\}_{t\geq 0}} \quad & \int_0^\infty e^{-\rho t} \int m(n)\, a_{jt}(n)\, dn\, dt \\ \textit{subject to} \quad & \dot{P}_{jt} = Q_t \int \big[\alpha_j m(n) - s_t(n)\big]\, a_{jt}(n)\, dn + \Big(\frac{\dot{Q}_t}{Q_t} - \mu - \delta\Big) P_{jt}, \\ & K(x_{jt}) = P_{jt}, \qquad x_{jt} = \int a_{jt}(n)\, dn, \\ & P_{j0} \textit{ given.} \end{aligned} \tag{24}$$

Proof. See Appendix A.1.1. ■

I now prove a useful property of the equilibrium stipend schedule.

Lemma 2. *Order university tiers by productivity so that $\alpha_1 > \alpha_2 > \cdots > \alpha_J$, and suppose that graduate programs take stipends as given. Along the balanced growth path, the equilibrium stipend schedule at any t is piecewise affine in posterior expected productivity, $\mathbb{E}[z_R|n_R]$, continuous, and convex.*

Proof. See Appendix A.1.2. ■

This result is also useful in characterizing the sorting patterns observed along the BGP.

Proposition 1. *Suppose that $K(\cdot)$ is strictly convex and $K(0) = 0$. Then, along the balanced growth path, the equilibrium exhibits a positive ranking on expected research skill, so that the average ability of students at high α_j programs have weakly higher expected research skill than those at lower α_j programs. If the distribution of n_R is atomless, then this ranking is strict.*

Proof. See Appendix A.1.3. ■

The economic content of Proposition 1 is that better programs end up with better students, on average.

3.3 Adding a Government

The government in this economy is endowed with three instruments. The first is a per student subsidy to universities $f_t Q_t$ (so that the total subsidy to a program is $f_t Q_t N_{jt}$). The second is a payroll subsidy to researchers, which I denote by s_R. This treatment is consistent with the structure of the U.S. R&D tax credit, which provides a tax credit on qualified research expenses. The largest component of qualified research expenses is R&D worker wages, which account for roughly 70% of total qualified research expenses.[12] The third instrument is a proportional tax on all income, including dividend income, which I denote by τ_W. I assume that the government balances the budget instant by instant. Let Γ_t^R denote the measure of individuals currently employed in research, with total mass R_t, and let B_{jt} denote program j's aggregate pre-tax stipend bill in units of Q_t. The government's flow budget constraint is thus:

$$\tau_{Wt}\left[w_t z_P L_t + Q_t \sum_{j=1}^{J} B_{jt} + \int_0^\infty \mathcal{P}_t(z_R)\,\mathrm{d}\Gamma_t^R(z_R) + \Pi_t\right] = f_t Q_t \sum_{j=1}^{J} N_{jt} + s_R \int_0^\infty \mathcal{P}_t(z_R)\,\mathrm{d}\Gamma_t^R(z_R). \tag{25}$$

3.4 Efficiency

Before moving on to the quantitative model, I first provide a set of results characterizing the efficiency of the competitive equilibrium in the case where there is a single university, so that $J = 1$. This simplifies the intuition by focusing attention on the distortion between research and production rather than on distortions in the relative sizes of graduate programs. I consider the problem of a planner who values aggregate consumption and can pick the allocation of individuals to research, graduate school, and production, subject to

[12]The R&D Tax Credit is filed as Form 6765, and governed by 26 U.S.C. §41, which specifies the definition of Qualified Research Expenses (QREs). QREs can be divided into four categories: wages, supplies (which specifically excludes anything that must be depreciated), computing resources leased or rented to conduct research, and a portion of contract research expenses paid to non-employees. Of these four categories, the largest by far is research payroll, which accounts for roughly 70% of QREs in the most recent SOI taxstats table for the credit (which comes from 2014). The next largest bucket is supplies at 15.3%, followed by contract research expenses at 14.5%, with the remainder being compute. The credit itself has a complicated structure, with firms being allowed to elect between two schemes, a traditional scheme with a fixed base (the definition of the base itself is quite complicated), and an "alternative simplified credit" which has a moving average basis defined over the prior 3 years. The latter scheme is the larger portion of the credit and distorts the effective value of the credit away from statutory rates. In addition to this, prior to the Protecting Americans from Tax Hikes (PATH) Act of 2015, the credit could only be claimed against income taxes. Because many R&D intensive firms are loss-making in their early years, this substantially limited the effectiveness of the credit. The PATH Act added, among other changes, an option for firms with less than $5 million in gross receipts to apply the credit against employer contributions to social security for a period of at most 5 years, subject to a cap of $250,000, raised to $500,000 during the Inflation Reduction Act. More details on the credit and the tax implications can be found in Pecoraro et al. (2026).

the same feasibility constraints and laws of motion as the competitive equilibrium. Let $k(x_t) \equiv K(x_t)/Q_t$. Omitting initial conditions, the social planner's problem is written in Problem (26).

$$\max_{\{c_t, a_t(n), d_t(z,n), e_t(z)\}} \int_0^\infty e^{-\rho t} Q_t c_t dt$$

Subject to:

$$\begin{aligned}
c_t + k(x_t) &= L_t + \alpha \int m(n) N_t(n) dn \\
\dot{Q}_t &= \varphi_R \ln(\lambda) Q_t \int z R_t(z) dz \\
\dot{N}_t(n) &= a_t(n) - (\mu + \delta) N_t(n) \\
\dot{R}_t(z) &= \mu \int d_t(z,n) h(z|n) N_t(n) dn - \delta R_t(z) - e_t(z) \\
L_t &= 1 - \int N_t(n) dn - \int R_t(z) dz \\
0 &\leq a_t(n) \leq \delta h(n) \\
0 &\leq d_t(z,n) \leq 1 \\
e_t(z) &\geq 0 \,, \quad R_t(z) \geq 0 \,, \quad N_t(n) \geq 0 \,, \quad L_t \geq 0 \,, \quad c_t \geq 0
\end{aligned} \tag{26}$$

Proposition 2. *If an interior balanced growth path exists, then it is generically inefficient.*

Proof. See Appendix A.2.1. ■

The equilibrium is inefficient for seven reasons. The first three are standard quality ladder externalities. The creative destruction externality arises because the value of a product line reflects the fact that it will be displaced by new innovations in the future. The planner discounts at a lower rate because these displacements are pure private transfers. The "shoulders" externality arises because the stock of knowledge is accretive: an innovation today is reflected in the stock of all future knowledge and benefits all future innovators. The consumer surplus externality arises because the innovative firm captures only the discounted stream of monopoly profits, while some consumer surplus generated by new innovations is not captured. In addition to these three standard externalities, a fourth distortion arises from the fact that productivity improvements by research firms also raise the productivity of the university sector, but research firms do not internalize this link. The link exists because the university's revenue and cost must grow at the same rate as the aggregate economy to ensure the existence of the balanced growth path. The ensuing externality will turn out to be quantitatively small. A fifth distortion arises from markups in the production sector,

which lead to inefficiently low production wages and distort individuals' choices away from production.

The last two sources of inefficiency arise from the graduate training block of the model. First, universities are not profit-maximizing. Because universities directly value the quantity of enrolled human capital, they will enroll inefficiently large numbers of graduate students. Second, universities and firms cannot contract on the output of students after they graduate. Universities operate a training technology that will lead to some economic surplus as a result of the match between a researcher and an innovative firm. Only the surplus captured by the researcher can be capitalized into graduate stipends. All else equal, this force pushes the university to admit too few students relative to the social optimum. These two forces push in opposite directions, and so the net effect on the quantity of human capital is ambiguous.

A useful intermediate result in the proof of Proposition 2 is that the equilibrium allocation can be summarized through two cutoffs. The first cutoff pertains to exit from research, and I denote it by z^*. The second cutoff concerns entry into graduate school, and I denote it by n^*. A convenient way to group the externalities is by which cutoff they primarily distort.[13] The quality ladder externalities and the spillover to university productivity can be understood as primarily distorting the research cutoff z^*. To understand these externalities, I compare the competitive research cutoff z_C^* to the planner's research cutoff z_S^*. The planner's cutoff is characterized by the following equation:

$$z_S^* = \frac{1}{\varphi_R v^S}, \quad v^S \equiv \frac{\ln(\lambda)\big(L^S + \alpha M^S - K(x^S)\big)}{\rho - g^S} \tag{27}$$

The competitive cutoff is characterized by the solution to the following equation:

$$z_C^* = \frac{1}{\lambda \varphi_R v^C}, \quad v^C \equiv \left(\frac{\lambda - 1}{\lambda}\right)\frac{L^C}{\rho + \iota^C - g^C} \tag{28}$$

The more interesting externalities arise from the graduate training block of the model. To understand these externalities, I compare the planner's graduate program cutoff n_S^* in Equation 29 to its competitive counterpart n_C^* in Equation 30, where ξ^C is the Lagrange multiplier on the university's non-profit constraint along the balanced growth path.

$$\underbrace{(\rho + \delta + \mu - g^S)K'(x^S)}_{\text{Cost of Admission}} + \underbrace{1}_{\text{Forgone Output}} = \underbrace{\alpha m(n_S^*)}_{\text{University Output}} + \underbrace{\frac{\mu}{\rho + \delta - g^S}\mathbb{E}_{z_R}\left[\left(\varphi_R z_R v^S - 1\right)_+ | n_S^*\right]}_{\text{Experimentation Value}} \tag{29}$$

[13] Markups in the production sector distort both cutoffs by reducing the production wage and hence the outside option.

$$\underbrace{(\rho+\delta+\mu)K'(x^C)}_{\text{Cost of Admission}}+\underbrace{\frac{1}{\lambda}}_{\text{Worker Wage}}=\underbrace{\alpha m(n_C^*)}_{\text{University Output}}+\underbrace{\frac{1}{\xi^C}m(n_C^*)}_{\text{Prestige}}+\underbrace{\frac{\mu\theta}{\rho+\delta-g^C}\mathbb{E}_{z_R}\left[\left(\varphi_R z_R v^C-\frac{1}{\lambda}\right)_+|n_C^*\right]}_{\text{Experimentation Value}} \tag{30}$$

Beginning with the left-hand side of these two expressions, first notice that the annuitized cost of admission differs. In particular, the planner's effective discount rate subtracts the growth rate g^S. In the competitive equilibrium, however, the university cares only about the level of skill, a value that does not depend directly on aggregate productivity.[14] Now compare the second term on the left-hand side of these two expressions. In the planner's expression, the second term is simply the forgone output of an individual who is in graduate school. In the competitive equilibrium, on the other hand, the second term is the wage of a worker in the production sector. Because the allocation of workers across jobs is elastic, markups in the production sector generate inefficiently low quantities of production labor. The next set of terms compares the planner's value of moving the threshold to that of the university. In the planner's expression, the marginal student at n_S^* generates new university output. In the competitive equilibrium, the value of the new student is both that output and the direct contribution of the new student to the university's objective. This force pushes toward over-admission relative to the planner.

The final terms capture the value of experimentation. In the planner's expression, adding a new graduate student at the margin acts like a call option: individuals whose realized research skill will be above the z_S^* threshold become researchers, while those below exit to production. Discounting the value of this call option at the appropriate rate and internalizing that the speed of graduation is governed by μ delivers the last term in Equation 29. The analogous expression in the competitive equilibrium reflects the private value of innovation and compares that value to the wage of a production worker.

The planner's expression reflects the full value of the call option, while the competitive equilibrium incorporates only a fraction θ of the private value. This reflects the fact that firms capture a portion of the surplus as pure rents. These rents arise only because of the admissions decision of the graduate school, but cannot be fully capitalized by the graduate program into lower stipends. This lack of appropriability would, all else equal, result in an inefficiently low level of research skill in the economy.[15] My next result characterizes the attainability of the first best in this economy.

Proposition 3. *The first best is generically implementable through three instruments. Two of the instruments address the externalities: a subsidy to the wages of researchers and a*

[14] If instead a university valued the level of productivity times the stock of enrolled skill, there would be no difference in discount rates.

[15] This is closely related to inefficiencies studied in Acemoglu (1997) and Flinn and Mullins (2015).

subsidy to graduate programs paid per enrolled student. The third instrument is a proportional income tax that allows the government to balance its budget.

Proof. See Appendix A.2.2. ■

The intuition for Proposition 3 is straightforward. Any implementable allocation can be summarized by a feasible allocation of the populace to research, graduate school, and production of the final good. In both the planner's problem and the competitive equilibrium, this allocation of talent is then summarized by two cutoffs: a minimum signal to enter into graduate school n^* and a minimum research skill to remain in the research sector z^*. All of the externalities in the model can be understood as affecting these two cutoffs. As a consequence, any two instruments that can move these two cutoffs to their socially optimal levels are sufficient to correct the underlying distortions. A third, non-distortionary instrument is then required to ensure that the government can balance its budget. Because a proportional income tax leaves relative returns unchanged, it can fill this role.

There are two natural instruments: a subsidy to the wages of researchers and a subsidy to graduate programs paid per enrolled student. The first instrument is sufficient to move the research cutoff to its socially optimal level, while the second instrument is sufficient to move the graduate program cutoff to its socially optimal level. Importantly, neither is generically sufficient to move both cutoffs to their social optimum. In this model, both R&D tax credits and government-funded research at universities are needed. To understand the implications of altering either one, I now turn to a quantitative version of this model.

4 Quantitative Model

This section develops and estimates the quantitative model used for policy analysis in Section 5. I extend the model by allowing universities to borrow, embed a measurement system matching idea production to patents in the data, and calibrate it to match features of the data in Section 2.

4.1 From Model to Data

The quantitative model enriches the baseline model in a number of dimensions. First, I specialize the productivity distribution, and explicitly model graduate school as a series of stages. Second, I modify the innovation section of the model to add an exogenous component of growth and add an explicit measurement system to the model to directly map innovations in the model and patents in the data. Third, I allow universities to borrow to fund PhD admissions rather than paying the full cost of admission using revenue upfront.

Researchers and Workers. Relative to the simple linear utility case discussed in the prior sections, here I assume that flow preferences over consumption are logarithmic, so that $u(c) = \ln(c)$.[16] Because a large proportion of the United States labor force does not have a bachelor's degree, a fact that would typically preclude them from attending graduate school, I further assume that an exogenous fraction π_R of individuals born in each instant are able to choose to attend graduate school, while the remainder directly become production workers. This fraction π_R should be thought of as the fraction of workers with a college degree, which is by and large necessary to attend graduate education. In that vein, I assume that ineligible workers supply $\xi_1 z_P$ efficiency units of production labor. Graduate school provides some benefit to individuals even who do not enter the research sector ex-post. I model this by assuming that an individual with a doctorate supplies $\xi_2 z_P$ efficiency units of production labor if they choose to exit. Conditional on enrolling in graduate school, individuals go through a series of $\mathcal{O}$ stages before graduating. Each of these stages is distributed exponential with rate parameter $\mathcal{O}\mu$, so graduation time follows an Erlang distribution with rate $\mathcal{O}\mu$ and $\mathcal{O}$ stages.[17]

An individual's log research productivity is the sum of two independent draws. The first of these, which I denote by η, is distributed normal with mean 0 and variance σ_R^2. The second of these, which I denote by γ, is distributed exponential with scale parameter τ. Research productivity in levels is therefore Pareto-lognormal with tail index $1/\tau$.

Research productivity is imperfectly observed at time of admission. I assume that the tail component γ is unobserved at time of entry, motivated by the empirical finding that a significant portion of individuals in lower tier programs appear in the upper tail of the patenting distribution.[18] Under the model's sorting mechanism if one assumed that this tail draw was observed, the share of inventors coming from top programs would converge to 1 when considering more productive inventors. Graduates of top tier programs, however are on average more likely to patent and, conditional on patenting, slightly more productive, which suggests that universities have some capacity to screen ex-ante. Accordingly, I assume that universities receive a noisy signal n_R about the normal component of research ability, so that $n_R = \eta + \varepsilon$, where ε is distributed normal with mean 0 and variance σ_ε^2 and independent of η and γ. These assumptions on the distribution of research productivity and what is known about it have a natural check in the data: Breiman's Lemma implies that innovative output has the same tail index across program tiers.[19] Consistent with this, Appendix Figure 33

[16]This is done primarily for computational reasons: it helps ensure existence of equilibria with positive stipend schedules as I calibrate the model.

[17]The mean of this distribution is $1/\mu$ and standard deviation is $1/(\mu\sqrt{\mathcal{O}})$.

[18]See Appendix E.7.

[19]Breiman's Lemma says that, given two independent random variables X and Y with non-negative

shows that the tail index of output is approximately constant across program rank deciles at around 2.1.

Innovation. Because a signficant fraction of innovation comes from non-PhD sources, I assume there is a flow of exogenous ideas χ into the economy. This implies that the aggregate innovation rate can be written as:

$$\iota_t = \chi + \varphi_R H_{Rt} \tag{31}$$

To match the patent data to the model, I assume that an individual of research ability z_R produces patents according to a negative binomial distribution with mean $A\varphi_R z_R$ and dispersion parameter ϱ, which is common across individuals. A pure Poisson formulation implies a tight link between conditional means and variances that is rejected by the data. The negative binomial arises naturally as a result of persistent differences in patenting propensity across individuals with the same ability. See Appendix B.2 for a detailed explanation of how the measurement system and ideas are related. A converts between patents per individual and innovations per product line.

Universities. Universities can borrow from abroad subject to the small open economy interest rate r^* and the restriction that each cohort must pay for itself over the course of their enrollment.[20] In other words, the graduate program must break-even on each cohort in net present value terms.

A tier-j program takes as given the schedule $s_t(n_R)$ and chooses $a_{jt}(n_R)$, the measure of applicants with signal n_R it admits at t. Stipends are contracted at admission and, like the program's service revenue $\alpha_j E[z_R|n_R]$ and the per-student subsidy f_t, grow at the same rate of the economy. I denote the total quantity of students admitted by x_{jt}. This variant of the

support, one of which is regularly varying with tail index α, and the other has a finite moment of order strictly greater than α, the product of the two inherits the tail index α. Informally, the product of a lognormal draw and a Pareto draw inherits the tail index of the Pareto draw.

[20] Without this, the model would be unable to generate the substantial increases in enrollment of graduate students immediately at the onset of the fiscal impulse.

university problem can be written as:

$$
\begin{aligned}
\max_{\{a_{jt}(\cdot)\geq 0\}_{t\geq 0}} \quad & \int_0^\infty e^{-r^* t} \int \mathbb{E}[z_R \mid n_R]\, a_{jt}(n_R)\, dn_R\, dt \\
\text{subject to} \quad & x_{jt} = \int a_{jt}(n_R)\, dn_R, \\
& K(x_{jt}) = \int \Big[A_{\lambda,t}\big(\alpha_j\, \mathbb{E}[z_R \mid n_R] - s_t(n_R)\big) + \mathcal{F}_t \Big]\, a_{jt}(n_R)\, dn_R, \\
& A_{\lambda,t} = \int_0^\infty e^{-\int_0^u (r^*+\delta-g_{t+v})\, dv}\, S(u)\, du, \qquad \mathcal{F}_t = \int_0^\infty e^{-\int_0^u (r^*+\delta-g_{t+v})\, dv}\, S(u)\, f_{t+u}\, du.
\end{aligned}
\tag{32}
$$

The second constraint is the cohort's break-even condition: the cost of admitting the cohort equals the present value of its contracted net revenue over the spell plus the subsidies that the university receives. Because the constraint binds cohort by cohort, the problem separates across cohorts. The enrolled stock, which determines the aggregate stipend bill and the government's outlay $Q_t f_t N_t$, follows the stage counts

$$
\begin{aligned}
\dot{N}_{jt}^{(1)} &= x_{jt} - (\mathcal{O}\mu + \delta)\, N_{jt}^{(1)} \\
\dot{N}_{jt}^{(s)} &= \mathcal{O}\mu\, N_{jt}^{(s-1)} - (\mathcal{O}\mu + \delta)\, N_{jt}^{(s)}, \quad s = 2, \ldots, \mathcal{O} \\
N_{jt} &= \sum_{s=1}^{\mathcal{O}} N_{jt}^{(s)}
\end{aligned}
\tag{33}
$$

with $(N_{j0}^{(s)})$ given. I assume $K(x) = \kappa x^{1+\nu}/(1+\nu)$.

4.2 Identification and Model Fit

I calibrate the model to an average field, constructed as an equal weighted average of the patenting moments by field. Model calibration proceeds in two steps. In the first, I set some parameters externally. In the second step, I internally calibrate the remaining parameters. Section 5 uses the calibrated model to evaluate counterfactuals pertaining to changes in university directed R&D support.

Externally Calibrated Parameters. I begin with parameters fixed from external sources. This set includes the discount, death and world interest rates; the share of a cohort eligible to apply and the two production productivities ξ_1 and ξ_2; the mean and dispersion of time to degree; and the share of growth attributed to sources other than doctoral research, which pins χ given λ. Table 1 lists the values and their sources.

Table 1: Externally calibrated parameters

Parameter	Description	Value	Source
Panel A. Preferences and demography			
ρ	Discount rate	0.03	Standard
δ	Death rate	0.025	Expected working life of 40 years
r^*	World interest rate	0.03	SOE, set equal to ρ
Panel B. Graduate education and the labor market			
π_R	Eligibility	0.46	Bachelor's degree, CPS 25+
ξ_1	Relative productivity of ineligible workers	0.62	BLS median weekly earnings, 2024
ξ_2	Productivity of non-research PhDs	1.198	Master's premium over college in engineering (PSEO)
$1/\mu$	Mean time to degree (years)	5.5	Survey of Earned Doctorates
$\mathcal{O}$	Erlang stages	14	SD of time to degree of 1.5 years
Panel C. Innovation			
$\chi \ln \lambda$	Exogenous growth component	0.0112	$2\% \times (1 -$ patent share$)$
Panel D. Normalizations			
$\mathbb{E}[\eta]$	Mean of the normal component of log research productivity	0	Normalization
z_P	Baseline production productivity	1	Normalization

Macro-Block. Although all parameters are jointly identified, some moments are more informative relative to others. For ease of exposition, I will divide internally calibrated parameters into two groups. Here I discuss the first, which pertains to aggregate quantities and graduate programs. It consists of 17 parameters, 10 of which are the productivity parameters of the graduate programs $\{\alpha_j\}$, and it is disciplined by 18 moments.

Some parameters are identified by standard moments. The innovation step-size λ is identified by the growth rate of 2%. The level of research productivity φ_R is identified by the PhD research payroll as a share of GDP, analogously to the role that the R&D share of GDP typically plays in quality ladder models.[21]

Other parameters target identified estimates in the literature. Researcher bargaining weight θ disciplines the pecuniary compensation of researchers. Kline et al. (2019) find that workers capture roughly 30 cents of every dollar of patent-induced surplus on average

[21] U.S. R&D is approximately 3% of GDP, of which 70% is labor compensation. Doctorate holders are between a sixth and a fifth of the R&D workforce. If they are paid on average twice as much as non-doctoral R&D workers, their share of the R&D wage bill is 29 to 33%, which gives a target of 0.6 to 0.7% of GDP. Since doctorate holders skew toward relatively low-paying academic jobs, I take the bottom of this range.

and about twice that among workers present at the time of the application, so I set θ so that inventors receive 60% of the aggregate revenue from innovation in the form of wages.[22] Increasing θ means that workers capture more of the monetary surplus, while decreasing it does the opposite.

The admissions cost function has both a level shifter κ and a cost elasticity ν. I identify the level κ by the share of individuals with research doctorates in science, engineering and health disciplines. This share is 0.4% in the 2019 Survey of Doctorate Recipients (SDR).[23] The curvature ν is identified by the response of completions to a funding shock, the local-projection impulse response at two horizons, since it governs how strongly admissions respond to a change in per-student resources. An increase in ν means that the same federal funding impulse will result in fewer admissions, flattening the impulse response in the model. Lowering ν pushes in the opposite direction. The relative sizes of the program tiers then identify the profile of program productivity across tiers, and academic R&D expenditure in STEM fields relative to GDP identifies the level. I take the latter from the Higher Education Research and Development (HERD) Survey, which reports that STEM research and development expenses amounted to roughly 0.35% of GDP.

Lastly, there are two federal policy parameters. The first of these is the federal subsidy f. This is identified by the federally financed share of university performed STEM R&D, roughly 0.2% of GDP. The second policy parameter is the R&D credit rate s_R. The structure of the federal R&D tax credit is quite complicated and features very limited takeup (Pecoraro et al., 2026), so instead of setting this to the statutory rate, I set this to match one half of the tax expenditures associated with the credit as reported by the Joint Committee on Taxation. This amounts to 0.06% of GDP.[24] The estimated subsidy rate is around 10%, which lies between the statutory rate of the regular credit (20%) and the effective marginal rate of the Alternative Simplified Credit. The ASC's statutory rate is 14% of QREs in excess of half their trailing three-year average, but because current spending raises the base in subsequent years, its effective marginal rate is closer to 7%. I summarize model fit in this

[22]This assumption is aggressive with respect to the inventor's share. In practice, most patents have multiple inventors, which would have the effect of further reducing the share of pecuniary surplus awarded to any individual researcher. This would exacerbate one important source of inefficiency in the model, which is that graduate programs cannot directly contract with firms. Note also that while Kline et al. (2019) study passthrough of surplus, here I target the level of pecuniary compensation. This is consistent with the standard microfoundation for innovation payments in these models, which specify that researcher wages are a fraction of the total surplus from innovative activities, while implicit credit markets across individuals make the timing of invention irrelevant.

[23]See Table 15-1. This number is the sum of Science, Engineering and Health PhDs less psychology and social science PhDs (639,100), divided by United States employment in 2019 (157,500,000). Of that set of 639,100 PhDs, 68.1% of them report research and development as either their primary or secondary work activity.

[24]The results are not sensitive to how much of the credit I assume is devoted to PhD payroll.

Table 2: Internally calibrated parameters: aggregates and graduate programs

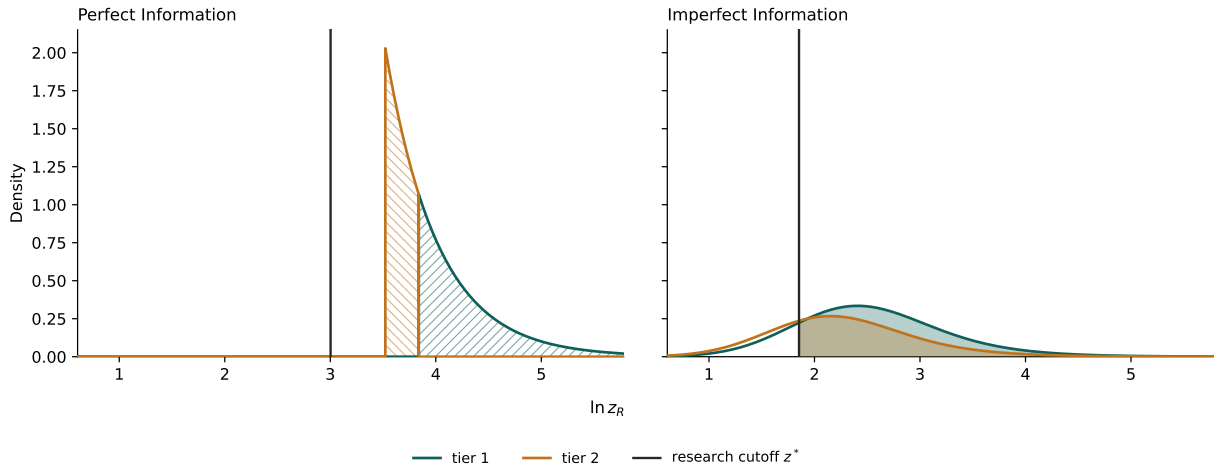
Parameter	Description	Value	Targeted moment	Model	Data
Panel A. Aggregates					
λ	Innovation step size	1.0231	growth rate g	0.0200	0.0200
φ_R	Research productivity	8.154	PhD research payroll / GDP	0.0060	0.0060
θ	Researcher bargaining weight	0.246	researchers' share of innovation value	0.61	0.60
Panel B. Graduate programs					
κ	Admission cost, level	59	PhD holders / labor force (SDR)	0.0041	0.0040
ν	Admission cost, curvature	0.072	completions IRF, years 6 and 7	0.3% / 0.4%	
α_1	Program productivity, top tier	0.386	STEM R&D at universities / GDP	0.0035	0.0035
$\alpha_1, \ldots, \alpha_{10}$	Program productivity profile	0.39–0.30	relative decile sizes	0.4% / 1.4%	
Panel C. Policy instruments					
f	Federal subsidy per student	2.724	federally funded STEM R&D at universities / GDP	0.0020	0.0020
s_R	R&D tax credit rate	0.101	R&D tax credit cost / GDP	0.0006	0.0006

Rows identified by a block of moments report the block's mean and maximum absolute percentage error, model versus data.

collection of moments in Table 2, and Appendix G.1 plots each block of targeted moments against the data.

Ability, Information, and Luck. I next turn to the other set of internally calibrated parameters. This block consists of five parameters and is disciplined by 37 moments. All but one of these moments were constructed from the matched dissertation-inventor data: for each program decile, the share of PhDs with any patent and the mean and the variance of log patent counts among those who patent; pooled across deciles, the mean count among patenters; and the concentration of top-decile graduates among the top 10, 5, 1 and 0.1 percent of inventors, relative to the top decile's share of patenters. Model analogues are computed from the negative binomial count layer over an 11-year window, as in the data. I identify the dispersion of the normal component of log ability σ_R by matching the variance of log output within a decile. I identify the signal noise σ_ε from the gradient of the variance of log output across deciles, as well as the share of PhDs still engaged in research in the SDR. To make the logic clear, Figure 6 shows the distribution of ability among admits in the top two program tiers under perfect (left panel) and imperfect (right panel) information. Imperfect information both generates exit from research after the PhD consistent with the data and flattens the gradient of the within-decile variance of log output.

Figure 6: Distribution of Ability in the Top Two Program Tiers



Notes: Density of log ability among admitted students in the top two program tiers, scaled so that the total area under both curves is equal to 1. The right panel is taken from the calibrated model, while the left panel comes from an alternative perfect information version of the model calibrated only to aggregate moments. Because ability has no natural units, the thresholds should not be directly compared across panels. Shaded areas denote individuals who, along the balanced growth path, remain in research after graduation.

Separating information about the normal component of log ability from the exponential component requires information about the tail of the ability distribution. I identify the scale of the star component τ by matching the share of the top 10%, 5%, 1% and 0.1% of inventors who trained at top-decile programs, relative to that decile's share of PhD inventors, as well as the pooled tail index of patenting output which I compute using Hill's estimator. I identify the dispersion of the negative binomial ϱ from the patenting share in each decile, since given the mean the probability of no patent at all depends on the dispersion, and I identify the nuisance parameter A from the mean count among patenters.[25]

Untargeted Moments. The model fits targeted moments quite well despite being substantially overidentified. Table 4 provides a collection of untargeted moments related to the shape of the underlying distribution. Appendix G.2 plots the full set of untargeted moments by decile. These higher order moments and capture shares by decile contain important information about the shape of the upper tail of the distribution. Given the importance of superstars in science (Lotka, 1926,Azoulay et al., 2010,Agrawal et al., 2017), characteristics of the tail of the distribution are potentially very important for counterfactuals.

[25] A converts between innovations per product line, the unit of the aggregate innovation rate, and patents per person. A researcher of ability z_R has an innovation rate given by $\varphi_R z_R$ per unit of population, so over a window of length T she produces $\varphi_R z_R M_L T/P$ innovations, where M_L is the number of product lines and P the population. If a fraction p of innovations is patented, the expected count is $A\varphi_R z_R$ with $A = p\, M_L T/P$. With the labor force as the population base, one third of innovations patented and period T equal to 11 years, the estimate of A implies approximately $440,000$ product lines in the United States.

Table 3: Internally calibrated parameters: research ability, information and luck

Parameter	Description	Value	Identifying moments	Fit
Panel A. Research ability				
σ_R	Dispersion of the normal component	0.84	within-decile variance of log output	3.7% / 6.5%
τ	Scale of the star component	0.48	pooled Hill tail index, top 5% of patenters	1.99 / 2.06
			elite capture at four depths	4.5% / 6.3%
Panel B. Information at admission				
σ_ε	Signal noise (reliability $\beta = 0.62$)	0.66	research share among PhD holders (SDR)	0.688 / 0.681
			between-decile gradient of the patenting share	4.8% / 17.4%
Panel C. Luck and measurement				
ϱ	Within-person dispersion of output	0.13	patenting share, pooled	0.178 / 0.179
A	Patents per innovation over the window	0.0101	mean patents among patenters	5.41 / 5.60
			mean log count among patenters, by decile	3.5% / 7.4%

Fit is the identifying block's mean and maximum absolute percentage error across program deciles (or depths), model versus data; for scalar targets, model / data.

Table 4: Untargeted moments

	decile 1 (best)		decile 10		abs. % error, 10 deciles	
Untargeted moment	Model	Data	Model	Data	mean	max
Panel A. Count distribution						
share of patenters with exactly one patent	0.323	0.332	0.354	0.409	6.2%	13.3%
share of patenters with exactly two patents	0.163	0.166	0.175	0.163	2.8%	9.4%
share of patents held by the top 1% of PhDs	0.341	0.362	0.400	0.466	9.4%	14.1%
Panel B. Shape of log output among patenters						
kurtosis of $\ln N$ among patenters	3.00	2.56	3.09	2.78	8.8%	17.4%
skewness of $\ln N$ among patenters	0.73	0.47	0.79	0.66	23.1%	54.9%

Mean and maximum absolute percentage errors are across the ten program deciles. Kurtosis and skewness are of log output centred on the field × cohort cell's geometric mean and pooled by decile. For consistency, I simulate a sample of comparable size to the data from the estimated parameters.

In general, the model performs well at matching most of the untargeted moments, although the cross tier gradient in variables is slightly more muted than in the data. Where the model misses worst is on skewness: among patenters, log patent counts are more right-skewed in the model than in the data. Direct evidence on the responsiveness of the ability of

the marginal entrant to changes in entrant quantities is quite rare. What evidence that the literature does have mostly comes from policy experiments run by Denmark which doubled the number of PhD slots in the early 2000s. Appendix G.3 shows that, when I simulate a doubling of the PhDs from the steady-state distribution, the model generates a proportional reduction in patents per PhD that is quite similar to what was reported in Akcigit et al. (2025).

5 Counterfactuals

In this section I use the calibrated model to quantify the effect of counterfactual policies. I consider two sets of exercises. The first of these two implements proposed cuts to federal support for funding in R&D, and evaluates the long-run impact on the researcher talent distribution and thereby growth. The second exercise examines the relative efficacy of federal funding for universities vis-à-vis R&D tax credits.

5.1 Changing Federal R&D Expenditures

In the summer of 2025, the second Trump administration proposed a budget for the United States Federal government covering Fiscal Year 2026 outlays. The proposed budget would have dramatically reduced appropriations to the civilian R&D accounts. Specifically, these cuts would have reduced NIH funding by 43.5% in nominal terms, NSF by 57%, NASA by 24%, and DOE by 16%.[26] I estimate that, had the budget been enacted, Federal government support for university based R&D would have fallen by 41.4%.[27] Congress did not ultimately enact them, opting to instead keep funding at roughly fiscal year 2025 baseline levels.

Baseline Cuts. To understand the impact of these cuts, I implement a policy experiment where, at time 0, there is a 3-year, 41.4% reduction in federal research expenditures at the university as a percentage of GDP. Universities correctly perceive that these cuts will end at the end of those three years.[28] Figure 7 plots the impulse responses of funding, admissions, researcher ability, and growth over the forty years after the initial cut.[29] There are three key

[26]The NIH and NASA figures are from AAAS, NSF's is from its Fiscal Year 2026 budget justification, and DOE's is its R&D from CRS R48694.

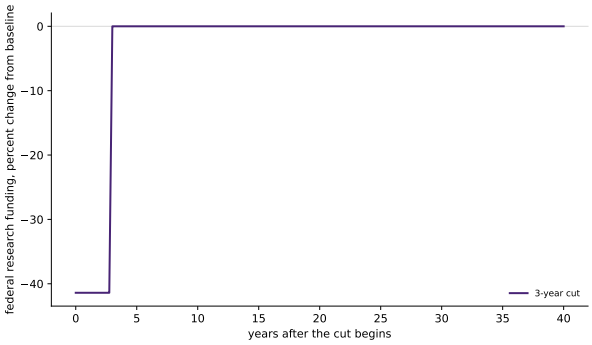
[27]I estimated this number using both proposed reductions in appropriations in the FY2026 President's budget and the composition, by performer, of R&D obligations in the Survey of Federal Funds for R&D. Appendix F explicitly details the calculations.

[28]There is one technical nuisance in this exercise, which is that cohorts initially enrolled prior to the unexpected shock will no longer be self-financing. In lieu of directly modeling the savings decision of the university, I assume that universities can raise additional funds to cover the funding shortfall for incumbent cohorts.

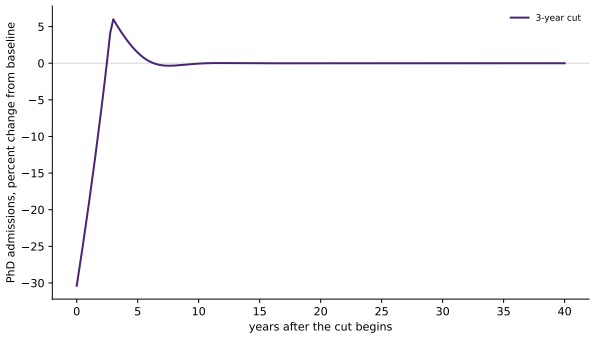
[29]Although I show the impulse response here only for the first 40 years, I solve the model until the transition is sufficiently close to the terminal balanced growth path.

takeaways from this exercise. First, drastic changes even in the short-run can move aggregate admissions very substantially, with the 41.4% cut leading to a roughly 30% reduction in admissions at onset. The second finding pertains to the ability of lost researchers. Figure 7c plots the decline in share of admitted students at different percentiles of the balanced growth path distribution of ex-post research ability. Although the marginal student is, on average, worse, there is still a roughly 18% reduction in the enrollment of individuals who would have been ex post in the top 1% of the balanced growth path research ability distribution. This inability to perfectly screen ex-post ability raises the cost of cuts. Despite that, the third implication of this exercise is that the growth response itself is relatively modest. This is because growth effects are dominated by the stock of graduated research ability. Because individuals in the model have a roughly 40 year long life, and spend only 5.5 of those years on average in graduate school, the size of any individual cohort is small relative to the stock of research skill in the economy. In all, growth moves by less than a basis point, consistent with short-run cuts seeming disastrous for future scientific talent even while having relatively modest macroeconomic consequences in the long-run. Nonetheless, the policy is relatively costly to society, with the reduction in federal funding leading to welfare costs of around 0.07% of permanent consumption. Contrast this with the fiscal gain of the policy. Federal R&D support to universities in the model is targeted to hit 0.20% of GDP. Reducing that number by 41.4% saves roughly 0.08% of GDP for three years. The foregone growth in the new equilibrium is worth roughly 0.9% of consumption in those same three years, which implies that the benefit to cost ratio of the cuts is roughly 0.1. In other words, the cuts cost about 10 times more than they save.

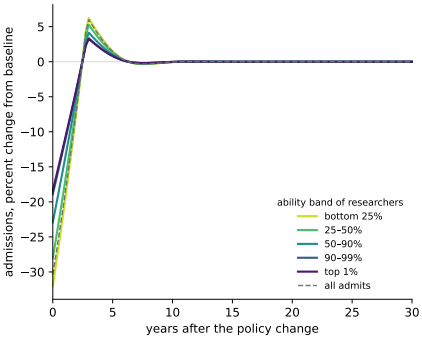
More Persistent Cuts. The prior policy exercise assumes that funding returns to the baseline level after three years. Given how persistent appropriations levels are, this seems unlikely. To understand how the length of these cuts interacts with long-run growth effects, I vary the duration of the 41.4% cut, starting from the 3 year baseline before moving to 15 years and then to 23 years. Figure 8 plots the growth effects at these three-horizons on the left panel, and the welfare consequences of these cuts at a more granular horizon. Longer duration cuts naturally degrade the stock of research skill more. This, in turn, leads to larger reductions in the growth rate, and correspondingly larger welfare costs. A 41.4% cut that persists for an additional 20 years would reduce welfare by roughly 1.05% of permanent consumption. The stock-flow nature of the pool here interacts with discounting to mute the welfare effects of these cuts, as the gains from increased consumption today are discounted substantially less than the losses from reduced consumption in the future.



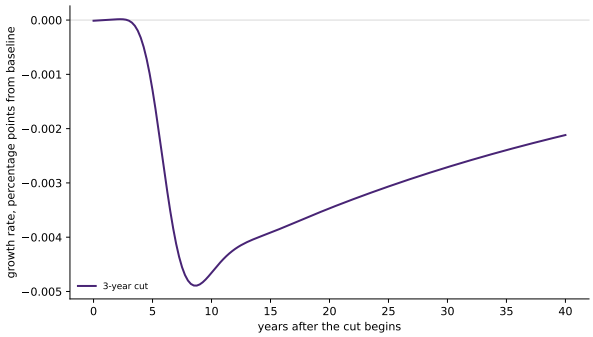
(a) Funding



(b) Admissions



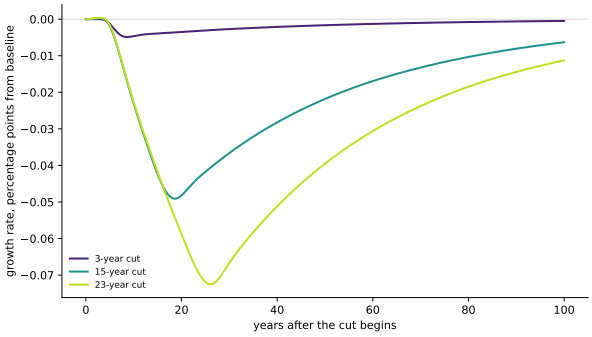
(c) Researcher ability



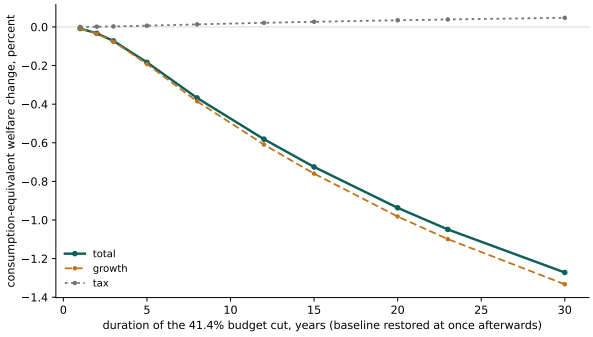
(d) Growth

Figure 7: Impulse responses to a temporary cut in federal research funding

Figure 8: Growth and welfare effects of 41.4% cuts of different durations



(a) Growth

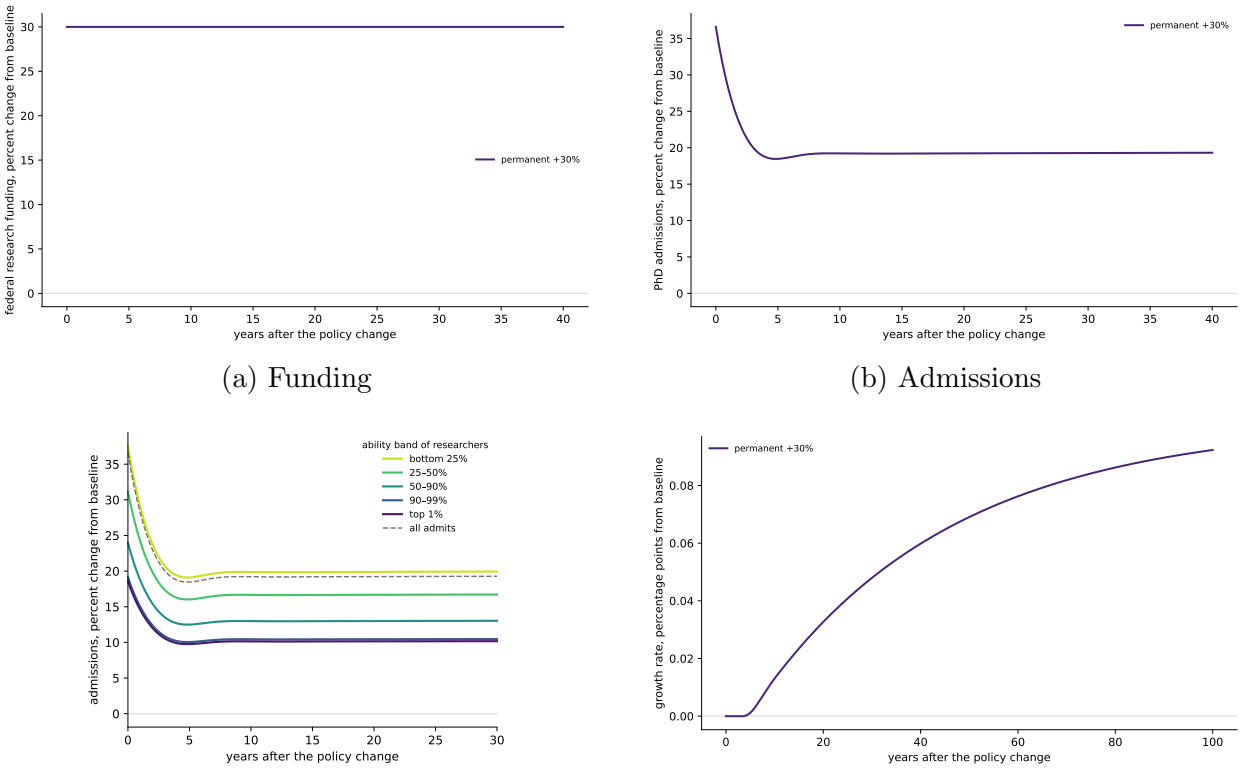


(b) Welfare

Notes: Plots show welfare and growth effects of 41.4% cuts of different durations.

An Increase in Funding. Next I consider the growth consequences of a permanent, 30% increase in federal funding for universities. Figure 9 plots the result.

Figure 9: Impulse responses to a permanent 30% increase in federal research funding



(a) Funding

(b) Admissions

(c) Researcher ability

(d) Growth

Notes: Plots show changes in responses to a permanent 30% increase in federal research funding.

The key takeaway from this exercise is that growth effects - even for a substantial and permanent increase in funding - arise gradually. The growth rate ultimately rises by about 10 basis points, but it takes nearly 200 years for the growth rate to converge to its new level. An observer in 2045 would have difficulty differentiating a true 3 basis point increase in the growth rate from a temporary shock. These policies nonetheless substantially increase welfare, with a permanent 30% increase in federal funding for universities leading to a roughly 1.28% increase in welfare.

5.2 Tax Credits versus University Research

As discussed in Section 3.4, this model features two sets of inefficiencies. The first set, common to all quality ladder models, are the three competing growth externalities. Creative destruction forces mean that innovative firms in general discount the value of their innovation too fast. "Shoulders" externalities mean that innovative firms do not internalize that other firms will later build on their advancements. Both of these forces tend to push towards

under-investment in R&D, while business-stealing, the fact that new firms receive the full flow of profits, not just the incremental, tends to push towards too much R&D.

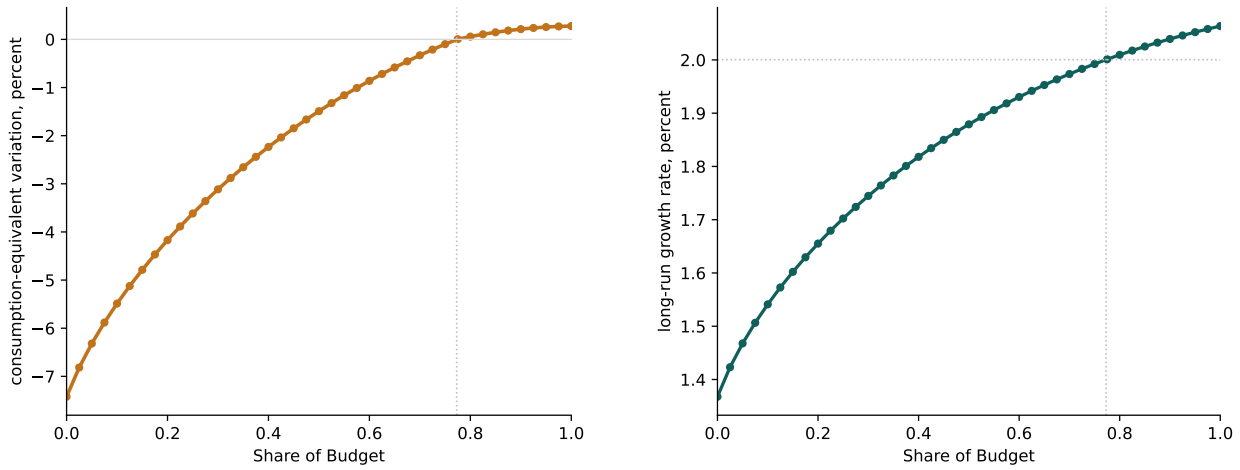
The second set of inefficiencies in this model are related to the training of graduate students. Universities in this model maximize the total ability of enrolled students. This force tends to push towards the over-provision of graduate students, or excessive human capital. On the other hand, the inability of universities to contract with innovative firms ex-post means that they are unable to fully capitalize the private value of their graduates' output. This force, closely related to the inefficiency studied in Acemoglu (1997), tends to push towards the underprovision of human capital.

Both sets of inefficiencies in this model are addressed by different instruments. The classical growth externalities in this model can be addressed by a proportional subsidy to the payroll of innovative firms, analogous to a R&D tax credit. Although these have general equilibrium effects on the choice to enter research, the primary effect that these subsidies to the payroll have is to manipulate the entry margin: a higher research subsidy increases payments to researchers ex post and thereby induces relatively worse individuals to remain in the research sector. The training externalities in the model are instead addressed by instruments that instead directly change the quantity of admitted students, as in the case of a subsidy per enrolled student. These subsidies provide resources to universities that are then used to admit additional students.

How should a government with limited fiscal capacity choose between the two instruments? To understand the relative efficacy of these two sets of instruments, I consider the problem of a government that holds fixed current expenditure and varies only its composition across the two instruments. For each of the combinations, I solve for the welfare of the new policy, inclusive of a transition where the composition immediately jumps to the proposed value. Across all scenarios, I hold fixed the size of support relative to output, which necessitates time-varying rates. Figure 10 below plots the set of feasible combinations of instruments, and the consumption equivalent variation associated with them. My main finding is that, R&D tax credits perform substantially worse than direct support to universities. The constrained planner would want to dramatically increase the amount of support for universities, even if it means removing support for R&D tax credits ex post. This complete reallocation would increase welfare by roughly 0.28% of permanent consumption, while raising the long-run growth rate by roughly 6.3 basis points. Computing the transition is crucial here. The removal of the credit reduces the quantity of people involved in research immediately, which leads to a reduction in growth in the short run. Because the stocks of talent are relatively long-lived, recovering from this immediate shock takes a period of time. A naive comparison of balanced growth paths would generate a welfare effect of 2.1% of

permanent consumption.

Figure 10: Welfare and Growth When Varying the Composition of Federal Innovation Policy



(a) Welfare (b) Long-run growth

Notes: The horizontal axis is the share of the federal innovation budget (the university appropriation plus the R&D credit, $\approx 0.26\%$ of GDP) spent on graduate training. The dotted vertical line marks today's split, and the dotted horizontal line in panel (b) the baseline growth rate.

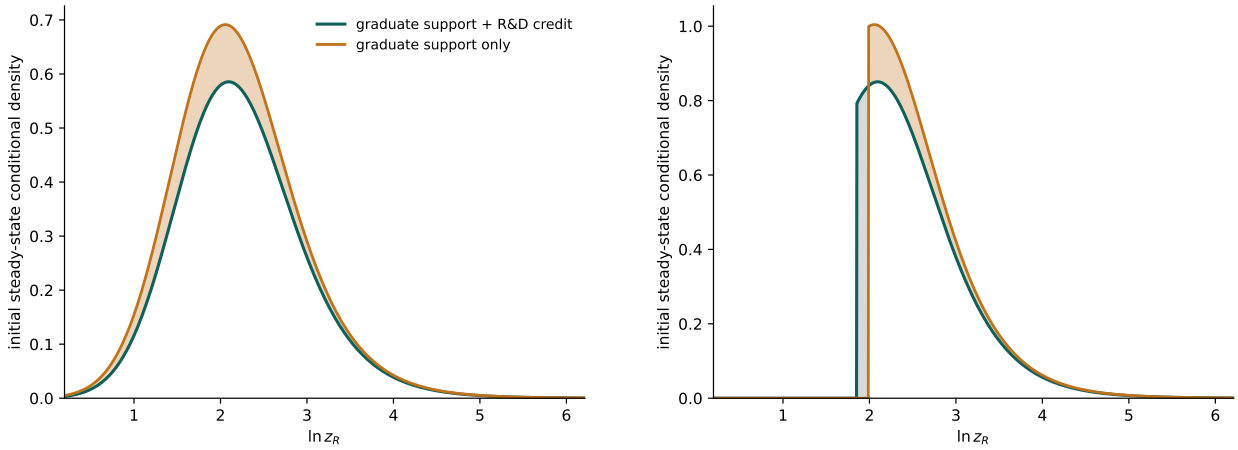
To understand this finding, it is useful to decompose the effects of these subsidies into two pieces. The first pertains to the direct quantity response as a result of changes in government spending: how many more PhDs does a sustained increase in each of the instruments buy? Figure 11a plots the quantity of first year graduate students at the two different equilibria, where, I have scaled the density of the new distribution by the total quantity of students in the baseline. The main takeaway is that graduate support generates substantially more entrants than R&D tax credits. This arises for two reasons. The first is that R&D tax credits in the model are poorly targeted instruments: most of their benefits accrue to inframarginal researchers, consistent with the evidence in Goolsbee (1998). The second reason is that the estimated admissions cost function is almost linear. Subsidies to graduate programs thus generate substantial increases in enrollment. In a model with perfect information, this could potentially lead to a dramatic reduction in the quality of researchers.

In this model, however, the substantial uncertainty about ability mitigates this quality dilution channel. To understand the differing ways in which the two instruments affect selection, Equation 34 rewrites the researcher entry condition.

$$\ln z^* = \ln \xi_2 - \ln \lambda - \ln \varphi_R + \ln(1 - s_R) - \ln V, \qquad V = \frac{(1 - 1/\lambda)\, L}{r^* + \iota - g}. \tag{34}$$

A reduction in the R&D tax credit has a mechanical effect: it reduces the private gains from trade between a firm and a researcher, which raises the threshold for entry. But it also

Figure 11: Change in Skill of the Marginal Entrant under Different Policies



(a) PhD entrants (b) Researchers

Notes: These plots depict the changes in the distribution of ability across different groups when I vary policy between the baseline and pure training policy economies. I scale the counterfactual economy densities relative to the baseline. Panel (a) shows the distribution of ability for PhD entrants, while Panel(b) shows the distribution of ability for researchers.

generates general equilibrium effects by changing the value of a product line. This general equilibrium effect does not, in general, offset the mechanical effect, and so the threshold for research entry increases.

Subsidies to universities behave differently. They increase the quantity of students admitted, which increases the number of researchers in the economy. This, in turn, reduces the value of a product line, which raises the threshold for research entry. The net effect is that the marginal researcher becomes somewhat better. Figure 11b plots the distribution of ability for researchers under these two policies. These differing effects arise because of the timing of selection and information revelation. Consider first the R&D tax credit. These gains from trade arise when research productivity is fully known. As a consequence, the matches that arise are less productive.[30] The other instrument induces universities to expand their graduate programs. The marginal student here has a worse public signal, so they are, in expectation, worse at research than the existing pool. Because there is substantial noise about ability at time of admission, however, they are not much worse. Many of the new researchers drawn in will be highly productive after graduation. The influx of skill increases the rate of creative destruction, reducing the value of a product line. This, in turn, reduces the incremental revenue of an innovative firm, pushing out individuals who are less productive from the pool of researchers, providing additional GE feedback that improves the overall quality of the research pool. Quantitatively, the ultimate effect is that individuals who were

[30]It is not crucial that productivity is fully known when researcher and firm meet. What is crucial is that the R&D tax credit selects on a more precise signal of ability than the subsidy to graduate programs.

Table 5: Research Entry Decomposition

	$\Delta \ln z^*$	$\Delta \ln(1 - s_R)$	$\Delta \ln(r^* + \iota - g)$
Baseline to corner	0.137	0.107	0.030
Credit removed first	0.088	0.107	−0.018
then appropriation raised	0.049	0	0.048
Appropriation raised first	0.052	0	0.051
then credit removed	0.085	0.107	−0.021

Log points. Each row compares across balanced growth paths. The scale term, $-\Delta \ln L$, is below 0.001 in every row and is omitted.

ex-ante at the bottom of the realized productivity distribution exit to production. To quantify the relative importance of these two channels, one can directly decompose the change in the cutoff in Equation 34. Because the equation holds exactly on any balanced growth path, the change in $\ln z^*$ between two of them is the sum of a credit term, $\Delta \ln(1 - s_R)$, a creative destruction term, $\Delta \ln(r^* + \iota - g)$, and a scale term, $-\Delta \ln L$. Table 5 decomposes the split between the status-quo and the no tax-credit equilibrium using this identity.[31] In general, the general equilibrium effects of the training policy account for about a quarter of the change in the cutoff.

6 Conclusion

In this paper, I have argued that federal funding plays an important role in the development of scientific human capital in the United States. Empirically, my argument proceeded in three stages. I began by showing that individuals with doctorates account for a substantial portion of American innovation, and that these individuals are particularly profound in critical technologies like semiconductors, machine learning, and biotechnology. In the second stage, I used a causal strategy based off of narrative shocks and showed that federal funding substantially affects the quantity of admitted graduate students, with a 1% increase in relative appropriations leading to a 0.55% increase in the ratio of doctorates awarded relative to bachelors 6 years prior.

Finally, I argued that there is very limited information about ex-post research productivity available at the time of admission. I did this by showing that, across all fields, the dispersion within programs of a given rank dwarfs the dispersion across programs. But that dispersion, although suggestive, is not sufficient to quantify the extent of uncertainty. To overcome that challenge, I constructed a model of endogenous growth that featured public learning and endogenous university size, and calibrated the model to match the new empirical evidence I have introduced. Through the lens of the model, I was able to quantify how

[31] Because the decomposition is sensitive to the order in which the changes are applied, I report the results for both orders.

little is known by graduate programs ex-ante.

I then used the model as a laboratory to evaluate the effect of university directed government R&D on growth and found that a 3-year, 41.4% reduction in federal funding would have a peak effect of reducing growth from 2% to roughly 1.995%, and impose a welfare cost of roughly 7 basis points in consumption equivalent terms. Under alternative assumptions about the length of the proposed cuts, I showed that longer cuts become progressively more costly because they attrite more of the pool of employed research skill. If the cuts were to persist for an additional 20 years, they would lead to a reduction in welfare of about 1%.

I then turned to the design of policy. I found that reallocating funding from R&D tax credits towards universities could lead to improvements in welfare of 0.28% of permanent consumption without increasing fiscal outlays. In the model, this is because the two instruments have differing effects on selection in the model. Subsidies to graduate programs induce individuals to enter graduate school. Because there is substantial uncertainty about one's ability, many of those individuals will ex-post be quite productive researchers. This, in turn, reduces the price of an idea, inducing individuals who were before at the margin to exit, leading to positive selection among continuing researchers ex-post. R&D tax credits on the other hand, are poorly targeted in two respects. First, they subsidize inframarginal researchers, raising their fiscal costs. Secondly, they incentivize individuals who are relatively less productive to remain in research, an effect that further reduces their efficacy.

In this paper, I have inferred uncertainty from the distribution of output conditional on ranks. My strategy relied on an assumption that individuals either do not know or cannot signal their ability to graduate programs. But individuals do know something about themselves. Measuring precisely what individuals believe about themselves, and how they came to that view, is an important avenue for future work. Secondly, I have abstracted from the presence of value added across universities. Universities may differ in their ability to train students, and this may, in turn lead to misallocation of students across programs. Finally, I have abstracted from the role of PhDs in innovation in other dimensions, for instance entrepreneurship and (especially) scientific publications. These are all important avenues for future work.

References

Acemoglu, Daron (1997) "Training and innovation in an imperfect labour market," *The Review of Economic Studies*, 64 (3), 445–464.

Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen (2017) "The social origins of inventors,"Technical report, National Bureau of Economic Research.

Aghion, Philippe and Peter Howitt (1992) "A MODEL OF GROWTH THROUGH CREATIVE DESTRUCTION," *Econometrica*, 60 (2), 323–351.

Agrawal, Ajay, John McHale, and Alexander Oettl (2017) "How stars matter: Recruiting and peer effects in evolutionary biology," *Research Policy*, 46 (4), 853–867.

Akcigit, Ufuk and Nathan Goldschlag (2023) "Measuring the characteristics and employment dynamics of US inventors,"Technical report, National Bureau of Economic Research.

Akcigit, Ufuk and William R Kerr (2018) "Growth through heterogeneous innovations," *Journal of Political Economy*, 126 (4), 1374–1443.

Akcigit, Ufuk, Jeremy Pearce, and Marta Prato (2025) "Tapping into talent: Coupling education and innovation policies for economic growth," *Review of Economic Studies*, 92 (2), 696–736.

Arora, Ashish, Sharon Belenzon, Larisa C Cioaca, Lia Sheer, and Hansen Zhang (2023) "The effect of public science on corporate R&D,"Technical report, National Bureau of Economic Research.

Azoulay, Pierre, Joshua S Graff Zivin, Danielle Li, and Bhaven N Sampat (2019) "Public R&D investments and private-sector patenting: evidence from NIH funding rules," *The Review of economic studies*, 86 (1), 117–152.

Azoulay, Pierre, Joshua S Graff Zivin, and Jialan Wang (2010) "Superstar extinction," *The Quarterly Journal of Economics*, 125 (2), 549–589.

Babalievsky, Fil (2022) "Misallocation in the Market for Inventors."

Babina, Tania, Alex Xi He, Sabrina T Howell, Elisabeth Ruth Perlman, and Joseph Staudt (2023) "Cutting the innovation engine: how federal funding shocks affect university patenting, entrepreneurship, and publications," *The Quarterly Journal of Economics*, 138 (2), 895–954.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen (2019) "Who becomes an inventor in America? The importance of exposure to innovation," *The Quarterly Journal of Economics*, 134 (2), 647–713.

Calvo, Guillermo A and Maurice Obstfeld (1988) "Optimal time-consistent fiscal policy with finite lifetimes," *Econometrica: Journal of the Econometric Society*, 411–432.

Campbell, Sheila, Jaeger Nelson, Eli A Schrag, Heidi L Williams, and Caleb K Wroblewski (2026) "Estimating the Economic Effects of Federally Funded R&D,"Technical report, National Bureau of Economic Research.

Cohan, Arman, Sergey Feldman, Iz Beltagy, Doug Downey, and Daniel S. Weld (2020) "SPECTER: Document-level Representation Learning using Citation-informed Transformers," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

Dyevre, Arnaud (2023) "Public R&D spillovers and productivity growth," Working paper.

Fieldhouse, Andrew and Karel Mertens (2025) "The returns to government R&D: Evidence from US appropriations shocks."

Flinn, Christopher and Joseph Mullins (2015) "Labor market search and schooling investment," *International Economic Review*, 56 (2), 359–398.

de Fraisse, Antoine Hubert, Peter Maxted, Kunal Sangani, and David Sraer (2026) "Pre-Trends Ahead of Monetary Policy Surprises: Recovering the Response of Corporate Investment."

Goolsbee, Austan (1998) "Does government R&D policy mainly benefit scientists and engineers?."

Gross, Daniel P and Bhaven N Sampat (2025) "The Government Patent Register: A new resource for measuring US government-funded patenting," *Research Policy*, 54 (1), 105142.

Grossman, Gene M and Elhanan Helpman (1991) "Quality ladders in the theory of growth," *The review of economic studies*, 58 (1), 43–61.

Hendricks, Lutz and Oksana Leukhina (2018) "The return to college: selection and dropout risk," *International Economic Review*, 59 (3), 1077–1102.

Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow (2019) "The allocation of talent and us economic growth," *Econometrica*, 87 (5), 1439–1474.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy (2021) "Measuring technological innovation over the long run," *American Economic Review: Insights*, 3 (3), 303–320.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar (2019) "Who profits from patents? rent-sharing at innovative firms," *The quarterly journal of economics*, 134 (3), 1343–1404.

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten (2020) "Leave-out estimation of variance components," *Econometrica*, 88 (5), 1859–1898.

Lazarus, Eben, Daniel J Lewis, James H Stock, and Mark W Watson (2018) "HAR inference: Recommendations for practice," *Journal of Business & Economic Statistics*, 36 (4), 541–559.

Lotka, Alfred J (1926) "The frequency distribution of scientific productivity," *Journal of the Washington academy of sciences*, 16 (12), 317–323.

Manski, Charles F (1989) "Schooling as experimentation: a reappraisal of the postsecondary dropout phenomenon," *Economics of Education review*, 8 (4), 305–312.

Murphy, Kevin M, Andrei Shleifer, and Robert W Vishny (1991) "The allocation of talent: Implications for growth," *The quarterly journal of economics*, 106 (2), 503–530.

Myers, Kyle R and Lauren Lanahan (2022) "Estimating spillovers from publicly funded R&D: Evidence from the US Department of Energy," *American Economic Review*, 112 (7), 2393–2423.

Pecoraro, Brandon, Nicholas Hoffman, Martin Lopez-Daneri, Elena Derby, Rachel Moore, and Shannon Sledz (2026) "Estimating the Present Value of R&D Tax Benefits in the United States," in *Entrepreneurship and Innovation Policy and the Economy, volume 6*: National Bureau of Economic Research, Inc, `https://ideas.repec.org/h/nbr/nberch/15434.html`.

Romer, Paul M (1990) "Endogenous technological change," *Journal of political Economy*, 98 (5, Part 2), S71–S102.

Shvadron, Dror, Hansen Zhang, Lee Fleming, and Daniel P Gross (2026) "Measuring the production of scientific human capital: New data, methods, and evidence on the US scientific training ecosystem," *Research Policy*, 55 (10), 105605.

Singh, Amanpreet, Mike D'Arcy, Arman Cohan, Doug Downey, and Sergey Feldman (2023) "SciRepEval: A Multi-Format Benchmark for Scientific Document Representations," in *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.

Terviö, Marko (2009) "Superstars and mediocrities: Market failure in the discovery of talent," *The Review of Economic Studies*, 76 (2), 829–850.

Tham, Wei Yang, Joseph Staudt, Elisabeth Ruth Perlman, and Stephanie D Cheng (2024) "Scientific talent leaks out of funding gaps," *arXiv preprint arXiv:2402.07235*.

Wu, Alice (2026) "Employer Learning, Sorting, and Productivity: Evidence from Computer Scientists,"Technical report, ROCKWOOL Foundation Berlin (RFBerlin).

A Simple Model

In this section, I solve a simple version of the model to illustrate the tradeoffs a planner faces in the presence of uncertainty about scientific talent and characterize the sorting properties of the equilibrium.

A.1 The Competitive Allocation

I first characterize the competitive allocation of the model. I begin with product lines before moving to characterize the equilibrium payoffs of production workers. Next I derive equilibrium payoffs for innovative firms and research workers. I then use these payoffs to characterize the sorting properties of the equilibrium in the simple case.

Product Lines. The value of a product line is characterized by the following Hamilton Jacobi Bellman equation:

$$(r^* + \iota_t)V_t - \dot{V}_t = \pi_t \tag{35}$$

Begin by normalizing the value function V_t so that $V_t = v_t Q_t$, where Q_t is frontier productivity at time t. It follows that

$$\begin{aligned}\dot{V}_t &= \dot{v}_t Q_t + v_t \dot{Q}_t \\ \frac{\dot{V}_t}{Q_t} &= \dot{v}_t + v_t \frac{\dot{Q}_t}{Q_t}\end{aligned}$$

Hence:

$$\begin{aligned}(r^* + \iota_t)V_t - \dot{V}_t &= \left(\frac{\lambda - 1}{\lambda}\right) Q_t L_t \\ (r^* + \iota_t)v_t - \frac{\dot{V}_t}{Q_t} &= \left(\frac{\lambda - 1}{\lambda}\right) L_t \\ (r^* + \iota_t)v_t - \dot{v}_t - g_t v_t &= \left(\frac{\lambda - 1}{\lambda}\right) L_t \\ (r^* + \iota_t - g_t)v_t - \dot{v}_t &= \left(\frac{\lambda - 1}{\lambda}\right) L_t\end{aligned}$$

On a balanced growth path, $\dot{v}_t = 0$ and the aggregate quantity of labor is constant,

$L_t = L$. Hence, the value of a product line along the balanced growth path is given by:

$$v = \left(\frac{\lambda - 1}{\lambda}\right)\frac{L}{r^* + \iota - g} \tag{36}$$

Production Workers. Production workers in this economy make no choices. They are hand to mouth individuals who consume however much they receive in income, and have linear utility with respect to that consumption.

Standard results imply that the equilibrium wage in efficiency units is given by:

$$w_t = \frac{1}{\lambda}Q_t$$

Using the fact that profits are also proportional to aggregate productivity, we can write the Hamilton Jacobi Bellman equation of a worker as follows:

$$(\rho + \delta)\mathcal{W}_t - \dot{\mathcal{W}}_t = \frac{1}{\lambda}Q_t + \hat{\Pi}_t Q_t$$

To normalize this, we can write $\mathcal{W}_t = \hat{\mathcal{W}}_t Q_t$. Then we have that:

$$\begin{aligned}
(\rho + \delta)\frac{\mathcal{W}_t}{Q_t} - \frac{\dot{\mathcal{W}}_t}{Q_t} &= \frac{1}{\lambda} + \hat{\Pi}_t \\
(\rho + \delta)\hat{\mathcal{W}}_t - \frac{\dot{\mathcal{W}}_t}{Q_t} &= \frac{1}{\lambda} + \hat{\Pi}_t \\
(\rho + \delta)\hat{\mathcal{W}}_t - \dot{\hat{\mathcal{W}}}_t - g_t\hat{\mathcal{W}}_t &= \frac{1}{\lambda} + \hat{\Pi}_t \\
(\rho + \delta - g_t)\hat{\mathcal{W}}_t - \dot{\hat{\mathcal{W}}}_t &= \frac{1}{\lambda} + \hat{\Pi}_t
\end{aligned}$$

Along the balanced growth path, $\dot{\hat{\mathcal{W}}}_t = 0$ and $\hat{\Pi}_t = \hat{\Pi}$. Hence, the value of a production worker along the balanced growth path is given by:

$$\hat{\mathcal{W}} = \frac{1}{\rho + \delta - g}\left(\frac{1}{\lambda} + \hat{\Pi}\right)$$

Note that $\hat{\Pi}$ is not the flow profits from intermediates in the economy. It is the flow profits from the mutual fund, which is the flow profits from product lines, net of research and development expenses. I next characterize the determination of these research and development expenses.

The Research Sector. The equilibrium expenses for research and development depend on the outcome of Nash bargain between an innovative firm and a research worker who is named by their productivity z_R. The match between the two generates a flow of new ideas at a rate given by $\varphi_R z_R$. When an idea arrives, the innovative firm auctions the ideas to the highest bidder. Because there is a unit mass of intermediate firms, this implies that the equilibrium price of an idea is bid up to the value of an individual product line, which is given by V_t. Hence, the flow of profits from the match between an innovative firm and a research worker is given by $\varphi_R z_R V_t$, net of a payment $\mathcal{P}_t(z_R)$. Given this payment, we can thus write the Hamilton Jacobi Bellman equation of an innovative firm as:

$$(r^* + \delta)\mathcal{V}_t - \dot{\mathcal{V}}_t = \varphi_R z_R V_t - \mathcal{P}_t(z_R)$$

On the other side, research workers at every instant have a choice of whether to continue working for their matched innovative firm or exit to the production sector. Exit is irreversible. Hence their problem is given by the following Hamilton Jacobi Bellman Variational Inequality:

$$\min\left\{(\rho + \delta)\mathcal{R}_t(z_R) - \dot{\mathcal{R}}_t(z_R) - \mathcal{P}_t(z_R) - \Pi_t, \mathcal{R}_t(z_R) - \mathcal{W}_t\right\} = 0$$

Because researchers in this economy have a fixed productivity and there is no aggregate uncertainty, this HJBVI can be directly expressed as a HJB in the region where the research worker continues to work for the innovative firm and a boundary condition that governs exit. Equivalently, the solution requires only value matching without imposing any condition on smooth pasting.

To characterize the solution to this problem, first normalize the problem. As with the prior problems, guess that the normalized value of a research worker and innovative firm are given by $\mathcal{R}_t(z_R) = \hat{\mathcal{R}}_t(z_R)Q_t$ and $\mathcal{V}_t(z_R) = \hat{\mathcal{V}}_t(z_R)Q_t$, respectively. Following the same process as above then implies that we can write the normalized value of a research worker as:

$$(\rho + \delta - g_t)\hat{\mathcal{R}}_t(z_R) - \dot{\hat{\mathcal{R}}}_t(z_R) = \hat{\mathcal{P}}_t(z_R) + \hat{\Pi}_t$$

The equivalent expression for a research firm is given by:

$$(r^* + \delta - g_t)\hat{\mathcal{V}}_t(z_R) - \dot{\hat{\mathcal{V}}}_t(z_R) = \varphi_R z_R v_t - \hat{\mathcal{P}}_t(z_R)$$

I characterize the solution to this problem along the balanced growth path.

$$\max_{\hat{\mathcal{P}}_t(z_R)} \left(\hat{\mathcal{R}}_t(z_R;\hat{\mathcal{P}}_t(z_R)) - \hat{W}_t\right)^{\theta} \left(\hat{\mathcal{V}}_t(z_R;\hat{\mathcal{P}}_t(z_R))\right)^{1-\theta}$$

Along the balanced growth path, an individual researcher chooses to work for the innovative firm only if their flow payoff from working for the firm is weakly greater than their flow payoff from working in production. Because firms that do not match die and thus receive no value, they are weakly better off when providing the full flow value of the match to the worker. As a consequence, when this flow value is higher than the wage of production, there are gains from trade and a match is formed:

$$\varphi_R z_R v \geq \frac{1}{\lambda}$$

When this condition is satisfied, the match is formed. Instantaneous renegotiation along the balanced growth implies that we must hold fixed future payoffs when considering a deviation today. To that end, define the difference in payoffs $\mathcal{D} = \hat{\mathcal{R}} - \hat{\mathcal{W}}$ and $\mathcal{V} = \hat{\mathcal{V}}$. Write that:

$$\begin{aligned}
\hat{\mathcal{D}}_t^{\Delta} &= (\hat{P}(z_R) - \hat{w})\Delta + e^{-(\rho+\delta-g)\Delta}\hat{\mathcal{D}}_{t+\Delta} + o(\Delta) \\
\hat{\mathcal{V}}_t^{\Delta} &= (\varphi_R z_R v - \hat{P}(z_R))\Delta + e^{-(r^*+\delta-g)\Delta}\hat{\mathcal{V}}_{t+\Delta} + o(\Delta)
\end{aligned}$$

Write the bargaining problem as follows:

$$\max_{\hat{\mathcal{P}}(z_R)} \left((\hat{\mathcal{P}}(z_R) - \hat{w})\Delta + e^{-(\rho+\delta-g)\Delta}\hat{\mathcal{D}}_{t+\Delta}\right)^{\theta} \left((\varphi_R z_R v - \hat{P}(z_R))\Delta + e^{-(r^*+\delta-g)\Delta}\hat{\mathcal{V}}_{t+\Delta}\right)^{1-\theta}$$

Log differentiating with respect to $\hat{\mathcal{P}}(z_R)$ delivers:

$$
\begin{aligned}
0 &= \theta \frac{\Delta}{\hat{D}_t^{\Delta}} - (1-\theta)\frac{\Delta}{\hat{V}_t^{\Delta}} \\
\theta \hat{V}_t^{\Delta} &= (1-\theta)\hat{D}_t^{\Delta} \\
\theta \hat{V} &= (1-\theta)\hat{D}
\end{aligned}
$$

Where in the last step I have taken the limit as $\Delta \to 0$. Substituting in for the balanced growth path values of $\hat{V}^{\Delta}$ and $\hat{D}^{\Delta}$ then delivers the following expressions:

$$
\hat{D} = \frac{\hat{\mathcal{P}}(z_R) - \frac{1}{\lambda}}{\rho + \delta - g} \text{ and } \hat{V} = \frac{\varphi_R z_R v - \hat{\mathcal{P}}(z_R)}{r^* + \delta - g}
$$

Substituting these expressions into the Nash bargaining condition then delivers the following expression for the equilibrium payment schedule:

$$
\hat{\mathcal{P}}(z_R) = \frac{1}{\lambda} + \theta \left(\frac{\rho + \delta - g}{\theta(\rho + \delta - g) + (1-\theta)(r^* + \delta - g)} \right) \left(\varphi_R z_R v - \frac{1}{\lambda} \right)
$$

In the case where $r^* = \rho$, this simplifies to:

$$
\hat{\mathcal{P}}(z_R) = (1-\theta)\frac{1}{\lambda} + \theta \varphi_R z_R v
$$

Graduate Programs. Graduate programs in this economy are non-profit institutions that maximize the expected ability of their students. They decide which students to admit, conditional on a convex cost of admission levied on bodies admitted to the program. I now characterize properties of their choices, both on and off the balanced growth path. Note that I restrict consideration to paths where $s_t(n_R) \geq -\Pi_t$ for all n_R and t.

A.1.1 Proof of Lemma 1

Conceptually, this lemma leverages the fact that universities care only about the contributions of individual students insofar as they contribute to the flow objective and the budget constraint. The flow objective is linear in the stock of enrolled students, and the budget constraint is linear in the stock of stipends. Hence, the history of admissions only matters insofar as it affects these two stocks.

Proof. The state of Problem (9) is the full distribution of enrollment $N_{jt}(\cdot)$ together with the stipend bill B_{jt}. Integrating Equation (7) against $m(n)$ gives $\dot{M}_{jt} = \int m(n)\, a_{jt}(n)\, dn -$

$(\mu+\delta)M_{jt}$. Combining this with Equation (8),

$$\dot{P}_{jt} = \frac{\dot{Q}_t}{Q_t}P_{jt} + Q_t\big(\alpha_j\dot{M}_{jt} - \dot{B}_{jt}\big) = Q_t\int\left[\alpha_j m(n) - s_t(n)\right]a_{jt}(n)\,dn + \Big(\frac{\dot{Q}_t}{Q_t} - \mu - \delta\Big)P_{jt}.$$

For the objective, note that M_{jt} is bounded: the inflow of students with signal n is at most $\delta h(n)$, so M_{jt} is bounded by a multiple of $\mathbb{E}[z_R] < \infty$ plus a term from the initial stock that decays at rate $\mu+\delta$. Hence $e^{-\rho t}M_{jt} \to 0$. Integrating $e^{-\rho t}\dot{M}_{jt}$ by parts and substituting the law of motion for M_{jt} gives

$$\int_0^\infty e^{-\rho t}M_{jt}\,dt = \frac{1}{\rho+\mu+\delta}\Big[M_{j0} + \int_0^\infty e^{-\rho t}\int m(n)\,a_{jt}(n)\,dn\,dt\Big].$$

Since M_{j0} is given, the two objectives differ by a positive affine transformation. The budget constraint depends on the state only through P_{jt}, with $P_{j0} = Q_0(\alpha_j M_{j0} - B_{j0})$ given. Hence the two problems are equivalent. ■

A.1.2 Proof of Lemma 2

I will prove the result for a case where there is a government subsidy $f_t Q_t$ in levels per body. The case in the maintext just sets $f_t = 0$. Recall the claim is that:

Order university tiers by productivity so that $\alpha_1 > \alpha_2 > \cdots \geq \alpha_J$. Along the balanced growth path, the equilibrium stipend schedule is piecewise affine in posterior belief about productivity $\mathbb{E}[z_R|n_R]$.

Proof. Begin by observing that the sequence formulation is given in (9). After detrending the problem, one can construct the current value Hamiltonian as:

$$\begin{aligned}\mathcal{H}_{jt} = Z_{jt} &+ v_{jt}(x_{jt} - (\mu+\delta)N_{jt}) \\ &+ \eta_{jt}\left[\int \mathbb{E}[z_R \mid n_R]a_{jt}(n_R)\,dn_R - (\mu+\delta)Z_{jt}\right] \\ &- p_{jt}\left[\int s_t(n_R)a_{jt}(n_R)\,dn_R - (\mu+\delta)B_{jt}\right] \\ &+ \xi_{jt}\left[\alpha_j Z_{jt} + f_t N_{jt} - B_{jt} - K(x_{jt})\right].\end{aligned}$$

Where v_{jt} denotes the costate on the stock of enrolled bodies, η_{jt} denotes the costate on Z_{jt} which is the stock of enrolled skill, $-p_{jt}$ denotes the costate on the normalized stipend, ξ_{jt} denotes the multiplier on the period by period non-profit constraint, and $x_{jt} \equiv \int a_{jt}(n_R)dn_R$.

Maximizing $\mathcal{H}_{jt}$ via the admissions policy, we recover the following set of optimality conditions:

$$v_{jt} + \eta_{jt}\mathbb{E}[z_R|n_R] \leq p_{jt}s_t(n_R) + \xi_{jt}K'(x^*_{jt})$$

With equality if $a_{jt}(n_R) > 0$. The economic intuition of this expression is straightforward. The university considering admitting an additional student trades off marginal benefit and marginal cost. The marginal benefit consists of a shadow value of an additional student v_{jt}, which arises here because of the government subsidy and the shadow value of that particular student through both the contribution to the flow objective and through an increase in university revenue, captured in the expression prefaced by η_{jt}. The marginal cost consists of the shadow value of stipends they must pay, captured by p_{jt} and the shadow value of the non-profit condition ξ_{jt}.

Next, we must show that $p_{jt} > 0$. To do this, observe that maximizing the Hamiltonian implies that:

$$\dot{p}_{jt} = (\rho + \mu + \delta)p_{jt} - \xi_{jt}$$

With terminal condition that:

$$\lim_{T\to\infty} e^{-(\rho+\delta+\mu)(T-u)}p_{jT} = 0$$

Integrating forward, one recovers that:

$$p_{jt} = \int_t^\infty e^{-(\rho+\delta+\mu)(u-t)}\xi_{ju}du$$

Next, I claim that p_{jt} is strictly positive along the balanced growth path. To do so, I will first show that $\xi_{ju} > 0$ along the balanced growth path. To that end, note that

$$\dot{v}_{jt} = (\rho + \mu + \delta)v_{jt} - f_t\xi_{jt}$$

$$\dot{\eta}_{jt} = (\rho + \delta + \mu)\eta_{jt} - 1 - \alpha_j\xi_{jt}$$

Along the balanced growth path, $\dot{v}_{jt} = \dot{\eta}_{jt} = \dot{p_{jt}} = 0$. Hence, we can write:

$$
\begin{aligned}
f\xi_j &= (\rho + \mu + \delta)v_j \\
1 + \alpha_j \xi_j &= (\rho + \delta + \mu)\eta_j \\
p_j &= \frac{\xi_j}{\rho + \delta + \mu}
\end{aligned}
$$

Along the balanced growth path, stocks are stationary, so that we have:

$$
\begin{aligned}
Z_j &= \frac{1}{\mu + \delta} \int \mathbb{E}[z_R | n_R] a_j(n_R) dn_R \\
B_j &= \frac{1}{\mu + \delta} \int s(n_R) a_j(n_R) dn_R \\
N_j &= \frac{1}{\mu + \delta} x_j
\end{aligned}
$$

From the non-profit constraint, we have that:

$$
\begin{aligned}
K(x_j) &= \alpha_j Z_j + f N_j - B_j \\
(\mu + \delta) K(x_j) &= f x_j + \alpha_j \int \mathbb{E}[z_R | n_R] a_j(n_R) dn_R - \int s(n_R) a_j(n_R) dn_R
\end{aligned}
$$

Now take the optimality conditions, which hold on any n_R where $a_j(n_R) > 0$, and integrate over the admissions distribution. This yields:

$$
\begin{aligned}
0 &= \int \big(v_j + \eta_j \mathbb{E}[z_R|n_R] - p_j s(n_R) - \xi K'(x_j)\big) a_j(n_R) dn_R \\
0 &= x_j v_j + \eta_j \int \mathbb{E}[z_R|n_R] a_j(n_R) dn_R - p_j \int s(n_R) a_j(n_R) dn_R - x_j \xi K'(x_j) \\
0 &= \frac{f x_j \xi_j}{\rho + \mu + \delta} + \frac{1 + \alpha_j \xi_j}{\rho + \mu + \delta} \int \mathbb{E}[z_R|n_R] a_j(n_R) dn_R \\
&\quad - \frac{\xi_j}{\rho + \mu + \delta} \int s(n_R) a_j(n_R) dn_R - x_j \xi_j K'(x_j) \\
0 &= f x_j \xi_j + (1 + \alpha_j \xi_j) \int \mathbb{E}[z_R|n_R] a_j(n_R) dn_R \\
&\quad - \xi_j \int s(n_R) a_j(n_R) dn_R - (\rho + \mu + \delta) x_j \xi_j K'(x_j) \\
0 &= \int \mathbb{E}[z_R|n_R] a_j(n_R) dn_R + \xi_j \left[f x_j + \int \big(\alpha_j \mathbb{E}[z_R|n_R] - s(n_R)\big) a_j(n_R) dn_R \right] \\
&\quad - (\rho + \mu + \delta) \xi_j x_j K'(x_j) \\
0 &= \int \mathbb{E}[z_R|n_R] a_j(n_R) dn_R + \xi_j \Big[(\mu + \delta) K(x_j) - (\rho + \mu + \delta) x_j K'(x_j) \Big] \\
\int \mathbb{E}[z_R|n_R] a_j(n_R) dn_R &= \xi_j \Big[(\rho + \mu + \delta) x_j K'(x_j) - (\mu + \delta) K(x_j) \Big]
\end{aligned}
$$

Note that the left hand side of this expression is strictly positive because productivity has positive support. Hence ξ_j shares the same sign as $(\rho + \mu + \delta) x_j K'(x_j) - (\mu + \delta) K(x_j)$. I claim that the latter expression is strictly positive.

Recall that $K(0) = 0$ and $K(\cdot)$ is convex. By standard properties of convex functions, we have that

$$K'(x_j) x_j \geq K(x_j)$$

Since $\rho > 0$, we then have that:

$$
\begin{aligned}
(\rho + \mu + \delta) x_j K'(x_j) &\geq (\rho + \mu + \delta) K(x_j) \\
&> (\mu + \delta) K(x_j)
\end{aligned}
$$

Hence it must be that $\xi_j > 0$. Since $\xi_j > 0$, it follows that $p_j > 0$. As a consequence, divide through by p_j in the optimality condition to recover the following expression for the stipend schedule:

$$s_t(n_R) = \underbrace{\frac{\eta_{jt}}{p_{jt}}}_{\equiv\beta_{jt}} \mathbb{E}[z_R|n_R] + \underbrace{\frac{v_{jt} - \xi_{jt}K'(x^*_{jt})}{p_{jt}}}_{\equiv\gamma_{jt}}$$

The consequence of this is that stipend schedule is given by:

$$s_t(n_R) = \max_{j\in J} \left\{\beta_{jt}\mathbb{E}[z_R|n_R] + \gamma_{jt}\right\}$$

This is the upper envelope of affine functions, and is thus piecewise affine, continuous and convex.

■

A.1.3 Proof of Proposition 1

Here I show that the balanced growth path features sorting on means, so that higher α_j universities match with individuals who on average have higher n_R.

Proof. Again let us suppose that we have ordered $\alpha_1 > \alpha_2 > \cdots > \alpha_J$. Lemma 2 says $s_j(n_R) = \max_{j\in J}\left\{\beta_j\mathbb{E}[z_R|n_R] + \gamma_j\right\}$. We wish to show that $\beta_i > \beta_k$ for all $i < k$, and that $\gamma_i < \gamma_k$ for all $i < k$.

Suppose that $K(\cdot)$ is strictly convex and $K(0) = 0$. Let

$$x_j \equiv \int a_j(n_R)dn_R \text{ and } \bar{m}_j \equiv \frac{\int \mathbb{E}[z_R|n_R]a_j(n_R)dn_R}{x_j}$$

We wish to show that $\alpha_j > \alpha'_j \implies \bar{m}_j \geq \bar{m}_{j'}$.

From Lemma 2, we have that $\xi_j > 0$. Define $m(n_R) \equiv \mathbb{E}[z_R|n_R]$. Then we have that:

$$(\mu+\delta)K(x_j) + \int s(n_R)a_j(n_R)dn_R = \alpha_j x_j \bar{m}_j + f x_j$$

Substitute using the expression for $s(n_R)$ from Lemma 2 to recover:

$$(\mu+\delta)K(x_j) + \beta_j\bar{m}_j x_j + \gamma_j x_j = \alpha_j x_j \bar{m}_j + f x_j$$

Rearranging, we have that:

$$(\mu+\delta)\frac{K(x_j)}{x_j} + (\gamma_j - f) = (\alpha_j - \beta_j)\bar{m}_j$$

$$\gamma_j = f + (\alpha_j - \beta_j)\bar{m}_j - (\mu+\delta)\frac{K(x_j)}{x_j}$$

Now fix a type n_R admitted by j. Then the stipend expression holds with equality and we have that:

$$s_t(n_R) = \beta_j m(n_R) + \gamma_j$$

Comparing with k, we have that:

$$s_t(n_R) \geq \beta_k m(n_R) + \gamma_k$$

Hence we have that:

$$\begin{aligned} \beta_j m(n_R) + \gamma_j &\geq \beta_k m(n_R) + \gamma_k \\ (\beta_j - \beta_k) m(n_R) &\geq (\gamma_k - \gamma_j) \end{aligned}$$

Integrating over the admissions policies of j, we have that:

$$(\beta_j - \beta_k)\bar{m}_j \geq \gamma_k - \gamma_j$$

Doing the same thing for types admitted to k and comparing the two implies that:

$$(\beta_j - \beta_k)\bar{m}_j \geq (\gamma_k - \gamma_j) \geq (\beta_j - \beta_k)\bar{m}_k$$

Note that this as well implies that:

$$(\beta_j - \beta_k)(\bar{m}_j - \bar{m}_k) \geq 0$$

Now suppose $\bar{m}_j < \bar{m}_k$. Then we have that:

$$\beta_j \leq \beta_k$$

Substituting for γ_j and γ_k, we have that:

$$f - (\rho + \mu + \delta)K'(x_j) \geq f - (\rho + \mu + \delta)K'(x_k)$$

By strict convexity we then have that $x_j \leq x_k$. Strict convexity then further implies that $K(x_j)/x_j \leq K(x_k)/x_k$. Substituting into the expression for γ_j and γ_k, we have that:

$$\begin{aligned}
\big(\beta_k - \alpha_k\big)\bar{m}_k - (\beta_j - \alpha_j)\bar{m}_j + (\mu+\delta)\left[\frac{K(x_k)}{x_k} - \frac{K(x_j)}{x_j}\right] &\leq \big(\beta_k - \beta_j\big)\bar{m}_k \\
(\beta_j - \alpha_k)\bar{m}_k - (\beta_j - \alpha_j)\bar{m}_j + (\mu+\delta)\left[\frac{K(x_k)}{x_k} - \frac{K(x_j)}{x_j}\right] &\leq 0 \\
(\alpha_j - \alpha_k)\bar{m}_k + (\beta_j - \alpha_j)(\bar{m}_k - \bar{m}_j) + (\mu+\delta)\left[\frac{K(x_k)}{x_k} - \frac{K(x_j)}{x_j}\right] &\leq 0
\end{aligned}$$

This is a contradiction because the first two terms are strictly positive and the third is non-negative. Hence it must be that $\bar{m}_j \geq \bar{m}_k$ for all $j < k$. This establishes weak sorting in mean posterior productivity.

To strengthen this to a strict inequality, we need only specify that the distributions of n_R be atomless. ■

Graduate Students. Graduate students in this economy are hand to mouth and are named by their signal n_R. Their HJB is given by:

$$(\rho + \delta + \mu - g_t)\hat{E}_t - \dot{\hat{E}}_t = s_t(n_R) + \hat{\Pi}_t + \mu \mathbb{E}_{z_R}\left[\max\left\{\hat{\mathcal{W}}_t, \hat{\mathcal{R}}_t(z_R)\right\} \mid n_R\right]$$

We wish to characterize the equilibrium threshold at which people enter. Maintaining the assumption that $r^* = \rho$, we can simplify the HJB as follows:

$$\begin{aligned}
\hat{\mathcal{R}}(z_R) &= \frac{\hat{\mathcal{P}}(z_R) + \hat{\Pi}}{\rho + \delta - g} \\
&= \frac{(1-\theta)\frac{1}{\lambda} + \theta\varphi_R z_R v + \hat{\Pi}}{\rho + \delta - g} \\
&= \frac{\frac{1}{\lambda} + \hat{\Pi}}{\rho + \delta - g} + \theta\frac{\varphi_R z_R v - \frac{1}{\lambda}}{\rho + \delta - g} \\
\hat{\mathcal{R}}(z_R) - \hat{\mathcal{W}} &= \theta\frac{\varphi_R z_R v - \frac{1}{\lambda}}{\rho + \delta - g} \\
&= \theta\frac{\varphi_R z_R\big(\frac{\lambda-1}{\lambda}\big)\frac{L}{\rho+\iota-g} - \frac{1}{\lambda}}{\rho + \delta - g}
\end{aligned}$$

Substituting this into the HJB for graduate students and subtracting from a suitably transformed version of the worker HJB, we have that:

$$(\rho + \delta + \mu - g)\big[\hat{E}(n_R) - \mathcal{W}\big] = s(n_R) - \frac{1}{\lambda} + \frac{\mu\theta}{\rho + \delta - g}\mathbb{E}_{z_R}\left[\left(\varphi_R z_R v - \frac{1}{\lambda}\right)_+ \middle| n_R\right]$$

At an interior cutoff n_R^*, it must be that $\hat{E}(n_R) = \mathcal{W}$, so that we have that:

$$s(n_R^*) = \frac{1}{\lambda} - \frac{\mu\theta}{\rho + \delta - g}\mathbb{E}_{z_R}\left[\left(\varphi_R z_R v - \frac{1}{\lambda}\right)_+ \middle| n_R^*\right]$$

This expression says that the stipend at the cutoff is equal to the flow value of production minus an option value research term. Note that this option value is scaled by θ, the bargaining power of the researcher. If $\theta = 1$, the stipend at the cutoff captures the full continuation value, but for any $\theta < 1$, there is some private surplus that the university is not able to appropriate. This is one key source of inefficiency in the model.

To characterize the total quantity of graduate students in the model, observe that for $n_R > n_R^*$, $\hat{E}(n_R) > \mathcal{W}$, so that all students with $n_R > n_R^*$ enter graduate school. The total quantity of graduate students admitted at any point along the balanced growth path is then given by:

$$\sum_j x_j = \delta \bar{F}(n_R^*)$$

Where $\bar{F}$ denotes the complementary CDF of the signal distribution.

Stocks. Next we characterize the flow stocks of workers and skill in this economy. Recall that the flow of workers enrolled in graduate school is given by $\delta\bar{F}(n_R^*)$. The flow of students graduating or dying is given by $(\mu + \delta)N$. Hence the stock of students in graduate school is given by:

$$N = \frac{\delta\bar{F}(n_R^*)}{\mu + \delta}$$

With the stock of graduate students stationary, the stock of researchers in this economy is characterized by the flow equation:

$$\dot{R} = \frac{\mu\delta}{\mu + \delta}\int_{z_R^*}^{\infty} \mathbb{P}[n_R \geq n_R^* | z_R] dG(z_R) - \delta R$$

Hence along the balanced growth path, we have that:

$$R = \frac{\mu}{\mu + \delta}\int_{z_R^*}^{\infty} \mathbb{P}[n_R \geq n_R^* | z_R] dG(z_R)$$

Efficiency units of research skill H_R are then given by the expression:

$$H_R = \frac{\mu}{\mu+\delta}\int_{z_R^*}^{\infty} z_R \mathbb{P}[n_R \geq n_R^* | z_R] dG(z_R)$$

Finally population feasibility requires that:

$$L = 1 - N - R$$

Innovation and Growth. Now we move towards characterizing the growth rate of the economy. Recall that the growth rate of the economy is given by:

$$g = \ln(\lambda)\iota = \ln(\lambda)\varphi_R H_R$$

Recall that an individual researcher enters research if and only if $z_R \geq z_R^*$, where z_R^* is given by:

$$z_R^* = \frac{1}{\lambda \varphi_R v} = \frac{r^* + \varphi_R H_R - g}{(\lambda - 1) L \varphi_R}$$

Now observe that we can write:

$$\begin{aligned}\varphi_R z_R v - \frac{1}{\lambda} &= \frac{\varphi_R z_R v z_R^*}{z_R^*} - \frac{1}{\lambda z_R^*} z_R^* \\ &= \frac{z_R}{\lambda z_R^*} - \frac{1}{\lambda z_R^*} z_R^*\end{aligned}$$

This implies that we can write the stipend at the cutoff as:

$$s(n_R^*) = \frac{1}{\lambda} - \frac{\mu\theta}{\lambda z_R^*(\rho + \delta - g)} \mathbb{E}_{z_R}\left[\left(z_R - z_R^*\right)_+ \middle| n_R^*\right]$$

We can directly characterize total R&D spending in this economy as:

$$B_R = \frac{\mu}{\mu+\delta}\int_{z_R^*}^{\infty} \hat{\mathcal{P}}(z_R)\mathbb{P}[n_R \geq n_R^* | z_R] dG(z_R) = \frac{R}{\lambda} + \theta\left(\varphi_R H_R v - \frac{R}{\lambda}\right)$$

The balance sheet of the mutual fund records sales of patents as both revenue and cost, and so the net dividends dispersed at each instant are given by:

$$\hat{\Pi} = \left(\frac{\lambda - 1}{\lambda}\right) L - B_R$$

A.2 The Planner Problem

I now turn to characterizing a special case of the social planner problem with a single university. This assumption that $J = 1$ substantially simplifies the analysis, but is not crucial. In addition, since the planner can frictionlessly reallocate consumption across individuals, I write the problem only in terms of aggregate consumption. Let n denote the signal of a graduate student and let z denote the realized research ability of a graduate.

It is convenient to define two shorthand stocks:

$$Z_t \equiv \int m(n)N_t(n)dn \text{ and } x_t \equiv \int a_t(n)dn$$

The planner problem can be written as:

$$\begin{aligned}
&\max_{\{c_t, a_t(n), d_t(z,n), e_t(z)\}} \int_0^\infty e^{-\rho t} Q_t c_t \\
&\quad \text{Subject to:} \\
&\quad c_t + K(x_t) = L_t + \alpha Z_t \\
&\quad \dot{Q}_t = \varphi_R \ln(\lambda) Q_t \int z \mathcal{R}_t(z) dz \\
&\quad \dot{E}_t(n) = a_t(n) - (\mu + \delta) E_t(n) \\
&\quad \dot{R}_t(z) = \mu \int d_t(z,n) h(z|n) E_t(n) dn - \delta R_t(z) - e_t(z) \\
&\quad L_t = 1 - \int E_t(n) dn - \int R_t(z) dz \\
&\quad 0 \leq a_t(n) \leq \delta f(n) \\
&\quad 0 \leq d_t(z,n) \leq 1 \\
&\quad e_t(z) \geq 0 \ , \ R_t(z) \geq 0 \ , \ N_t(n) \geq 0 \ , \ L_t \geq 0 \ , \ c_t \geq 0
\end{aligned} \tag{37}$$

To solve this, write the current value Hamiltonian as:

$$
\begin{aligned}
\mathcal{H}_t = {} & Q_t c_t + \\
& \Psi_t \varphi_R \ln(\lambda) Q_t \int z \mathcal{R}_t(z) dz + \\
& \int p_t^N(n) \big[a_t(n) - (\mu + \delta) E_t(n) \big] dn + \\
& \int p_t^R(z) \big[\mu \int d_t(z,n) h(z|n) E_t(n) dn - \delta R_t(z) - e_t(z) \big] dz + \\
& \nu_t \big[L_t + \alpha Z_t - K(x_t) - c_t \big] \\
& \gamma_t \big[1 - L_t - \int E_t(n) dn - \int R_t(z) dz \big]
\end{aligned}
\tag{38}
$$

Note that this is subject to same flow and state controls as the primal problem. Using the fact that $Q_t = \nu_t$, we can substitute for c_t to remove one of the constraints. The Hamiltonian can then be written as:

$$
\begin{aligned}
\mathcal{H}_t = {} & Q_t \big[L_t + \alpha Z_t - K(x_t) \big] + \\
& \Psi_t \varphi_R \ln(\lambda) Q_t \int z \mathcal{R}_t(z) dz + \\
& \int p_t^N(n) \big[a_t(n) - (\mu + \delta) E_t(n) \big] dn + \\
& \int p_t^R(z) \big[\mu \int d_t(z,n) h(z|n) E_t(n) dn - \delta R_t(z) - e_t(z) \big] dz + \\
& \gamma_t \big[1 - L_t - \int E_t(n) dn - \int R_t(z) dz \big]
\end{aligned}
$$

The costate equation for Ψ_t is given by:

$$
\dot{\Psi}_t = \rho \Psi_t - \frac{\partial \mathcal{H}_t}{\partial \mathcal{Q}_t} = (\rho - g_t) \Psi_t - c_t
$$

Along a balanced growth path, $\dot{\Psi}_t = 0$ and $g_t = g$. Hence we have that:

$$
\Psi_t = \frac{c_t}{\rho - g}
$$

Define the social value of an innovation by: $v^S = \Psi_t \ln(\lambda)$. Then we have that:

$$
v^S = \frac{c \ln(\lambda)}{\rho - g}
$$

With the social value of an innovation in hand, we can now analyze the allocation rule

for a researcher of type z. The costate equation for $p_t^R(z)$ is given by:

$$\dot{p}_t^R(z) = \rho p_t^R(z) - \frac{\partial \mathcal{H}_t}{\partial R_t(z)} = (\rho + \delta) p_t^R(z) - Q_t\big[\Psi_t \varphi_R \ln(\lambda) z - 1\big]$$

Where I have used the fact that $Q_t = \gamma_t$. In normalized terms, we can write this as:

$$(\rho + \delta - g_t)\hat{p}_t^R(z) - \dot{\hat{p}}_t^R(z) = v^S \varphi_R z - 1$$

Along the balanced growth path, $\dot{\hat{p}}_t^R(z) = 0$ and $g_t = g$. Hence we have that:

$$\hat{p}^R(z) = \frac{v^S \varphi_R z - 1}{\rho + \delta - g}$$

Exiting research is optimal whenever naively $\hat{p}^R(z) < 0$, which occurs whenever $z < z_S^*$ where z_S^* is defined by:

$$z_S^* = \frac{1}{\varphi_R v^S}$$

This characterizes the research choice functional $d(z, n)$. In particular, it takes the form of a cutoff rule.

$$d(z, n) = \mathbb{1}\{z \geq z_S^*\}$$

Incorporating exit then gives the social value of a researcher of type z as:

$$\boxed{\hat{p}^R(z) = \frac{\big(v^S \varphi_R z - 1\big)_+}{\rho + \delta - g}}$$

Next, I characterize the optimal allocation of graduate students. The costate equation for $p_t^N(n)$ is given by:

$$\begin{aligned}
\dot{p}_t^N(n) &= \rho p_t^N(n) - \frac{\partial \mathcal{H}_t}{\partial N_t(n)} \\
&= (\rho + \mu + \delta) p_t^N(n) - Q_t\big[\alpha m(n) - 1\big] \\
&\quad - \mu \int p_t^R(z) d_t(z, n) h(z \mid n)\, dz.
\end{aligned}$$

The optimal graduation rule implies that $d_t(z, n)\hat{p}_t^R(z) = \hat{p}_t^R(z)$. Normalizing the costates by Q_t therefore gives:

$$\dot{\hat{p}}_t^N(n) = (\rho + \mu + \delta - g_t)\hat{p}_t^N(n) - \big[\alpha m(n) - 1\big] - \mu \mathbb{E}\big[\hat{p}_t^R(z) \mid n\big].$$

Along the balanced growth path, the normalized costate is constant. Hence:

$$\hat{p}^N(n) = \frac{\alpha m(n) - 1 + \mu \mathbb{E}\big[\hat{p}^R(z) \mid n\big]}{\rho + \mu + \delta - g}.$$

From admissions optimality, we have that:

$$\frac{\delta \mathcal{H}_t}{\delta a_t(n)} = p_t^N(n) - Q_t K'(x_t) = Q_t \big[\hat{p}_t^N(n) - K'(x_t)\big].$$

The optimal policy is a cutoff. This cutoff n_S^* satisfies:

$$\boxed{K'(x) = \frac{\alpha m(n_S^*) - 1 + \mu \mathbb{E}\big[\hat{p}^R(z) \mid n_S^*\big]}{\rho + \mu + \delta - g}}$$

This expression is important because, coupled with the density of n, it characterizes the efficient quantity of graduate students. Distortions in the competitive equilibria away from this quantity are the novel source of inefficiency in the model.

Having now characterized the decision rules for graduate students and researchers, what remains is to characterize the way that those decision rules translate into aggregate quantities along the balanced growth path.

The quantity of admitted students at any instant is given by:

$$x = \int a(n) dn = \delta \bar{F}(n_S^*)$$

Imposing that the stock is stationary delivers:

$$N = \frac{\delta \bar{F}(n_S^*)}{\mu + \delta}$$

The quantity of researchers of type z is given by:

$$R(z) = \frac{\mu}{\mu + \delta} \mathbb{1}\{z \geq z_S^*\} \int_{n_S^*}^{\infty} h(z|n) f(n) dn$$

The total quantity is given by:

$$R = \frac{\mu}{\mu + \delta} \int_{z_S^*}^{\infty} \int_{n_S^*}^{\infty} h(z|n) f(n) dn dz$$

Given this, we have that:

$$L = 1 - N - R$$

Finally, the aggregate quantity of research skill in this economy is given by:

$$H_R = \frac{\mu}{\mu + \delta} \int_{z_S^*}^{\infty} z \int_{n^*}^{\infty} h(z|n) f(n) dn dz$$

This characterizes an interior solution to the planner problem. The next step is to characterize the efficiency properties of the competitive equilibria. In particular, I will show that the competitive equilibrium is generically inefficient. The proof strategy is to fix a candidate set of thresholds (n^*, z^*) and then show that efficiency implies three conditions on the underlying parameters. This immediately implies that the locus of parameters under which the competitive equilibrium is efficient is measure zero.

A.2.1 Proof of Proposition 2

Proof. As with before, fix a single university type so that $J = 1$. Let n_C^* and z_C^* denote the thresholds in the competitive equilibrium. It is easy to see the thresholds for z_C^* and n_C^* are sufficient for equilibrium quantities of research skill, graduate students, researchers, and hence production labor. Call the equilibrium quantities induced by a tuple (n_C^*, z_C^*) under a given parameter vector (H_R^C, N^C, R^C, L^C). The competitive equilibrium is efficient if and only if $(H_R^C, N^C, R^C, L^C) = (H_R^S, N^S, R^S, L^S)$, where again the superscript S denotes the social planner allocation. This is equivalent to requiring that $n_C^* = n_S^*$ and $z_C^* = z_S^*$. Hence we can focus on the conditions under which the competitive equilibrium thresholds coincide with the social planner thresholds.

Let n_S^* and z_S^* denote the thresholds in the social planner problem. Suppose that $n_C^* = n_S^*$ and $z_C^* = z_S^*$. As was shown in the prior section, the thresholds for z_S^* and z_C^* are pinned down by the following two equations:

$$z_S^* = \frac{1}{\varphi_R v^S} \text{ and } z_C^* = \frac{1}{\lambda \varphi_R v^C}$$

Where $v^S = \frac{c^S \ln(\lambda)}{\rho - g}$ and $v^C = \frac{\left(\frac{\lambda - 1}{\lambda}\right) L^C}{\rho + \iota - g}$. Suppose that $z_S^* = z_C^*$. Then we have that:

$$\frac{c^S \ln(\lambda)}{\rho - \varphi_R \ln(\lambda) H_R^C} = \frac{\big(\lambda - 1\big) L^C}{\rho + \varphi_R H_R^C - \ln(\lambda) \varphi_R H_R^C}$$

Algebraic manipulations then deliver the following restriction on aggregates:

$$\big(\lambda - 1\big) L = c^S \ln(\lambda) + \frac{H_R(n^*, z^*)}{z^*}$$

The first claim is that there is exactly one value of λ such that this is true. To see this, define the function $F(\lambda) = \frac{c^S \ln(\lambda) + \frac{H_R(n^*, z^*)}{z^*})}{L(\lambda - 1)}$.

This function $F(\lambda)$ is continuous and strictly decreasing in λ for $\lambda > 1$. In addition, we have that $\lim_{\lambda \to 1} F(\lambda) = \infty$ and $\lim_{\lambda \to \infty} F(\lambda) = 0$. Hence there is exactly one value of λ such that $F(\lambda) = 1$. This establishes the first claim.

The next claim is to show that this, in turn, pins down a unique value of φ_R. To see this, note that:

$$\rho = \varphi_R \ln(\lambda)\big(c^S(n^*, z^*)z_S^* + H_R(n^*, z^*)\big)$$

Hence we have that:

$$\varphi_R = \frac{\rho}{\ln(\lambda)\big(c^S(n^*, z^*)z_S^* + H_R(n^*, z^*)\big)}$$

This implies we can write the growth rate as:

$$g = \ln(\lambda)\varphi_R H_R(n^*, z^*) = \frac{\rho H_R(n^*, z^*)}{c^S(n^*, z^*)z_S^* + H_R(n^*, z^*)} < \rho$$

Note that we can write the private surplus owed to a researcher worker as a function of (n^*, z^*) as well. This implies that we can write the equilibrium stipend schedule at the cutoff as:

$$\begin{aligned} s(n^*) &= \frac{1}{\lambda} - \frac{\mu\theta}{\rho + \delta - g}\mathbb{E}_{z_R}\left[\left(\varphi_R z_R v - \frac{1}{\lambda}\right)_+ \middle| n^*\right] \\ &= \frac{1}{\lambda} - \frac{\mu\theta}{\rho + \delta - g}\mathcal{S}(n^*, z^*) \end{aligned} \tag{39}$$

Note also that we have:

$$s(n^*) = \left(\alpha + \frac{1}{\xi(n^*)}\right)m(n^*) - (\rho + \delta + \mu)K'(x(n^*))$$

The latter expression makes explicit that $s(n^*)$ is pinned. Now call the θ that makes 39 true $\theta(n^*, z^*)$, and restrict attention to (n^*, z^*) such that $\theta(n^*, z^*) \in (0, 1)$. This establishes the second claim. Hence we have that for any (n^*, z^*) there is at most one tuple $(\theta, \varphi_R, \lambda)$ that can support a competitive equilibrium that is efficient. It follows that the locus of parameters under which the competitive equilibrium is efficient is measure zero, given that the underlying map from (n^*, z^*) to $(\theta, \varphi_R, \lambda)$ is continuously differentiable. This establishes Proposition 2.

■

The next step having proven inefficiency is to show that the first best is attainable with

three instruments: a research subsidy τ_R, a graduate student subsidy f, and a proportional income tax. I now do so.

A.2.2 Proof of Proposition 3

Formally, the claim is that a research subsidy τ_R and a graduate student subsidy $f_t Q_t$ are necessary and sufficient to implement the first best when paired with a proportional income tax. The proof proceeds in two steps, beginning with sufficiency.

Sufficiency.

Proof. Fix a vector of parameters that induces an interior optimal balanced growth path. As in the prior proof, it is convenient to work directly with n_S^* and z_S^*, the thresholds for graduate students and researchers, respectively. The first step is to show that there exists a unique candidate tuple (τ_R, f) that implements the first best. Begin by defining the competitive value of an innovation at the planner's allocation:

$$v_C^S = \frac{\frac{\lambda-1}{\lambda} L^S}{\rho + \iota - g}$$

The social value of an innovation is given by:

$$v^S = \frac{c^S \ln(\lambda)}{\rho - g}$$

For the competitive research threshold to coincide with the social planner threshold, it must be that:

$$z_S^* = \frac{1}{\varphi_R v^S} = \frac{1}{\lambda \varphi_R v_C^S (1 + \tau_R)}$$

Hence we must have that:

$$1 + \tau_R = \frac{1}{\lambda} \left[\frac{v^S}{v_C^S} \right]$$

From the threshold stipend condition for the CE evaluated at the planner's allocation, we have that:

$$s(n_S^*) = \frac{1}{\lambda} - \frac{\mu\theta}{\lambda(\rho + \delta - g)} \mathbb{E}_{z_R} \left[\left(\frac{z_R}{z_S^*} - 1 \right)_+ \middle| \, n_S^* \right]$$

Next, define the average posterior productivity of a graduate student at the university as:

$$\bar{m}^S = \frac{\int \mathbb{E}[z_R \mid n_R] a_j(n_R)\, dn_R}{x^S}$$

Where E^S is the total number of graduate students and $x^S = (\mu + \delta)E^S$. From the CE equations, we have that:

$$s(n) = \left(\alpha + \frac{1}{\xi^S}\right) m(n) + f - (\rho + \delta + \mu) K'(x^S)$$

Now recall that we have that:

$$\frac{1}{\xi^S} = \frac{K'(x^S)(\rho + \mu + \delta) - (\mu + \delta)\frac{K(x^S)}{x^S}}{\bar{m}^S}$$

Notice that f raises the stipend schedule uniformly. Allowing signed policy rates, there is therefore a unique f that equates the university's offer at n_S^* to the required stipend. By the university characterization and monotonicity of the entry surplus, this implements the first best, subject again to strictly positive individual consumption in the implementing equilibrium. This establishes sufficiency. ■

Necessity.

Proof. To prove generic necessity among interior implementations, I approach the problem with a similar strategy used to prove generic inefficiency. In particular, fix all parameters but $(\lambda, \varphi_R, \theta)$ and consider the planner's optimal cutoffs (n_S^*, z_S^*).

By the prior proof, we have that:

$$g(n, z) = \frac{\rho H_R(n, z)}{c^S(n, z) z + H_R(n, z)}$$

Substituting this expression into the planner's admission condition then delivers:

$$\begin{aligned} A(n, z) \equiv (\rho + \delta + \mu - g(n, z)) K'(x(n)) - \alpha m(n) + 1 \\ - \frac{\mu}{\rho + \delta - g(n, z)} \mathbb{E}\left[\left(\frac{z_R}{z} - 1\right)_+ \middle| n\right] = 0. \end{aligned}$$

This is continuously differentiable under the maintained smoothness assumptions. Holding n fixed, we have that:

$$(c^S)_z = -R_z \geq 0, \qquad (H_R)_z = -z(c^S)_z, \qquad \frac{\partial}{\partial z}(c^S z + H_R) = c^S > 0.$$

Hence $g_z < 0$. Since $K'(x) > 0$ and the expected research option is nonnegative and weakly decreasing in z, it follows that $A_z > 0$. By the implicit function theorem, we can write the set of candidate cutoffs locally as $z = \zeta(n)$.

Now fix $\tau_R = 0$. By the prior proof, matching research cutoffs requires:

$$(\lambda - 1)L(n, z) = c^S(n, z)\ln(\lambda) + \frac{H_R(n, z)}{z}.$$

Along $z = \zeta(n)$, this uniquely pins down λ given n. The planner's research condition then pins down φ_R, leaving (n, θ) free.

Now instead fix $f = 0$. Given (n, λ) and $z = \zeta(n)$, the planner's research condition pins down φ_R, while the entry condition pins down θ whenever the marginal student's expected research option is positive. If that option is zero, entry instead pins down λ, leaving (n, θ) free.

In both cases, the exceptional parameter sets are covered by countably many continuously differentiable images of two-dimensional parameter domains in the three-dimensional space $(\lambda, \varphi_R, \theta)$, and hence have Lebesgue measure zero. Thus, generically, neither instrument alone implements an interior planner balanced growth path.

To ensure budget balance, the government can set a uniform tax on all income (including dividends).

■

B Mathematical Appendix

This section provides additional mathematical characterizations useful for results in the main text.

B.1 Nash Bargaining

Although agents in the model are hand to mouth, they remain forward looking. This forward-looking behavior matters for separation decisions. In particular, a researcher compares not only the current payment from remaining in research with current consumption in production, but also the continuation value associated with preserving the research match.

Let $\mathcal{P}_t(z_R)$ denote the flow payment from an innovative firm to a researcher of type z_R. Conditional on a fixed future sequence of payments $\{\mathcal{P}_s(z_R)\}_{s \geq t}$, the value of the firm satisfies

$$(r^* + \delta)\mathcal{V}_t\left(z_R; \{\mathcal{P}_s(z_R)\}_{s \geq t}\right) - \dot{\mathcal{V}}_t = V_t \varphi_R z_R - \mathcal{P}_t(z_R), \tag{40}$$

while the value of the researcher satisfies

$$(\rho+\delta)\mathcal{R}_t\left(z_R;\{\mathcal{P}_s(z_R)\}_{s\geq t}\right)-\dot{\mathcal{R}}_t=u(\mathcal{P}_t(z_R)+\Pi_t)\,. \tag{41}$$

Let W_t denote the continuation value associated with permanently exiting research and entering production. The researcher's continuation surplus is therefore

$$\mathcal{D}_t(z_R)\equiv\mathcal{R}_t(z_R)-W_t. \tag{42}$$

Because changing the payment at a single instant has measure zero in continuous time, it is useful to define the bargaining problem over a short interval of length $\Delta>0$ and then take the limit as $\Delta\downarrow 0$. Fix a date t. Conditional on the match continuing, suppose that the worker and firm bargain over a constant flow payment p on $[t,t+\Delta)$ and take the equilibrium continuation allocation from $t+\Delta$ onward as given.

Let

$$\mathcal{D}_{t+\Delta}(z_R)\equiv\mathcal{R}_{t+\Delta}(z_R)-W_{t+\Delta}$$

denote the worker's continuation surplus conditional on remaining in research at $t+\Delta$, and let $\mathcal{V}_{t+\Delta}(z_R)$ denote the corresponding continuation value of the firm. For a candidate payment p, the worker's gain from agreement at date t is

$$\begin{aligned}\mathcal{D}_t^{\Delta}(p;z_R)=&\int_t^{t+\Delta}e^{-(\rho+\delta)(s-t)}\left[u(p+\Pi_s)-u(c_s^P)\right]ds\\&+e^{-(\rho+\delta)\Delta}\mathcal{D}_{t+\Delta}(z_R),\end{aligned} \tag{43}$$

where c_s^P denotes consumption following permanent exit to production.

The firm's gain from agreement is

$$\begin{aligned}\mathcal{V}_t^{\Delta}(p;z_R)=&\int_t^{t+\Delta}\exp\left\{-\int_t^s(r^*+\delta)d\tau\right\}\left[V_s\varphi_R z_R-p\right]ds\\&+\exp\left\{-\int_t^{t+\Delta}(r^*+\delta)d\tau\right\}\mathcal{V}_{t+\Delta}(z_R).\end{aligned} \tag{44}$$

An agreement is feasible whenever there exists a payment p such that

$$\mathcal{D}_t^{\Delta}(p;z_R)\geq 0\qquad\text{and}\qquad\mathcal{V}_t^{\Delta}(p;z_R)\geq 0. \tag{45}$$

If no such payment exists, the match separates and the worker permanently exits to produc-

tion. Conditional on agreement being feasible, the payment solves

$$\mathcal{P}_t^{\Delta}(z_R) \in \arg\max_p \left[\mathcal{D}_t^{\Delta}(p; z_R)\right]^{\theta} \left[\mathcal{V}_t^{\Delta}(p; z_R)\right]^{1-\theta}, \tag{46}$$

where $\theta \in (0,1)$ is the researcher's Nash bargaining weight.

For an interior agreement, the first-order condition associated with (46) can be written as

$$\theta \frac{\int_t^{t+\Delta} e^{-(\rho+\delta)(s-t)} u'(\mathcal{P}_t^{\Delta}(z_R) + \Pi_s)\, ds}{\mathcal{D}_t^{\Delta}(\mathcal{P}_t^{\Delta}(z_R); z_R)} = (1-\theta) \frac{\int_t^{t+\Delta} \exp\left\{-\int_t^s (r^* + \delta) d\tau\right\} ds}{\mathcal{V}_t^{\Delta}(\mathcal{P}_t^{\Delta}(z_R); z_R)}. \tag{47}$$

Next I claim that $\mathcal{P}_t^{\Delta}(z_R) \to \mathcal{P}_t(z_R)$ as $\Delta \downarrow 0$.

Lemma 3. *Suppose that u is strictly increasing, strictly concave, and twice continuously differentiable with $u'' < 0$. Consider a point (t, z_R) in the interior of the continuation region, so that*

$$\mathcal{D}_t(z_R) > 0, \qquad \mathcal{V}_t(z_R) > 0.$$

Suppose further that the limiting first-order condition admits an interior solution $\mathcal{P}_t(z_R)$. Then, for all sufficiently small $\Delta > 0$, the finite-Δ Nash problem admits a unique interior solution $\mathcal{P}_t^{\Delta}(z_R)$, and

$$P_t^{\Delta}(z_R) \longrightarrow P_t(z_R) \qquad as \qquad \Delta \downarrow 0.$$

Proof. For fixed $\Delta > 0$, define the log Nash objective

$$\mathcal{N}_t^{\Delta}(p) = \theta \log \mathcal{D}_t^{\Delta}(p; z_R) + (1-\theta) \log \mathcal{V}_t^{\Delta}(p; z_R).$$

The worker surplus $\mathcal{D}_t^{\Delta}(p; z_R)$ is increasing and concave in p, while the firm surplus $\mathcal{V}_t^{\Delta}(p; z_R)$ is decreasing and affine in p. Hence

$$\log \mathcal{D}_t^{\Delta}(p; z_R)$$

is concave and

$$\log \mathcal{V}_t^{\Delta}(p; z_R)$$

is strictly concave on the interior of the feasible set. Therefore $\mathcal{N}_t^{\Delta}(p)$ is strictly concave, so any interior first-order condition characterizes the unique maximizer.

The first-order condition is

$$\theta \frac{\partial_p \mathcal{D}_t^{\Delta}(p; z_R)}{\mathcal{D}_t^{\Delta}(p; z_R)} + (1-\theta) \frac{\partial_p \mathcal{V}_t^{\Delta}(p; z_R)}{\mathcal{V}_t^{\Delta}(p; z_R)} = 0.$$

Using

$$\partial_p \mathcal{D}_t^\Delta(p; z_R) = \Delta u'(p + \Pi_t) + o(\Delta)$$

and

$$\partial_p \mathcal{V}_t^\Delta(p; z_R) = -\Delta + o(\Delta),$$

together with

$$\mathcal{D}_t^\Delta(p; z_R) = \mathcal{D}_t(z_R) + o(1), \qquad \mathcal{V}_t^\Delta(p; z_R) = \mathcal{V}_t(z_R) + o(1),$$

uniformly on a neighborhood of the limiting solution, the first-order condition can be written

$$u'(p + \Pi_t) = \frac{1-\theta}{\theta} \frac{\mathcal{D}_t(z_R)}{\mathcal{V}_t(z_R)} + o(1).$$

The limiting equation

$$u'(P_t(z_R) + \Pi_t) = \frac{1-\theta}{\theta} \frac{\mathcal{D}_t(z_R)}{\mathcal{V}_t(z_R)}$$

has a unique interior solution because u' is strictly decreasing. Since the derivative of the limiting first-order condition with respect to p is

$$u''(P_t(z_R) + \Pi_t) < 0,$$

the implicit function theorem implies that, for sufficiently small Δ, there exists a unique solution $P_t^\Delta(z_R)$ in a neighborhood of $P_t(z_R)$ and

$$P_t^\Delta(z_R) \to P_t(z_R).$$

Strict concavity implies that this solution is the unique finite-Δ Nash maximizer. ■

Now suppose that the equilibrium lies in the interior of the continuation region. By continuity,

$$\int_t^{t+\Delta} e^{-(\rho+\delta)(s-t)} u'(\mathcal{P}_t^\Delta(z_R) + \Pi_s)\, ds = \Delta u'(\mathcal{P}_t(z_R) + \Pi_t) + o(\Delta), \tag{48}$$

$$\int_t^{t+\Delta} \exp\left\{ -\int_t^s (r^* + \delta) d\tau \right\} ds = \Delta + o(\Delta). \tag{49}$$

Moreover,

$$\mathcal{D}_t^\Delta(\mathcal{P}_t^\Delta(z_R); z_R) \to \mathcal{D}_t(z_R) = \mathcal{R}_t(z_R) - W_t,$$

and

$$\mathcal{V}_t^{\Delta}(\mathcal{P}_t^{\Delta}(z_R); z_R) \to \mathcal{V}_t(z_R).$$

Dividing (47) by Δ and taking the limit yields the continuous-time Nash bargaining condition

$$\theta \frac{u'(\mathcal{P}_t(z_R) + \Pi_t)}{\mathcal{R}_t(z_R) - W_t} = (1-\theta)\frac{1}{\mathcal{V}_t(z_R)}. \tag{50}$$

Or equivalently,

$$\theta \mathcal{V}_t(z_R) = (1-\theta)\frac{\mathcal{R}_t(z_R) - W_t}{u'(\mathcal{P}_t(z_R) + \Pi_t)}. \tag{51}$$

Equation (51) motivates defining the worker's continuation surplus in pecuniary terms as

$$\mathcal{M}_t(z_R) \equiv \frac{\mathcal{R}_t(z_R) - W_t}{u'(\mathcal{P}_t(z_R) + \Pi_t)}, \tag{52}$$

and adjusted match surplus as

$$\mathcal{S}_t(z_R) \equiv \mathcal{V}_t(z_R) + \mathcal{M}_t(z_R). \tag{53}$$

The Nash condition can then be written in the familiar form

$$\mathcal{M}_t(z_R) = \theta \mathcal{S}_t(z_R), \qquad \mathcal{V}_t(z_R) = (1-\theta)\mathcal{S}_t(z_R). \tag{54}$$

Under log utility,

$$u(c) = \log c,$$

the monetized worker surplus becomes

$$\mathcal{M}_t(z_R) = \left(\mathcal{P}_t(z_R) + \Pi_t\right)\left(\mathcal{R}_t(z_R) - W_t\right). \tag{55}$$

The continuous-time first-order condition is necessary and sufficient only in the interior of the continuation region. The separation rule is instead the limit of the finite-Δ feasibility condition in (45).

B.1.1 Normalizations

Having established that the bargaining problem is well formed, I now normalize the value functions so that the bargaining problem is stationary along a balanced growth path.

Let Q_t denote aggregate productivity, with

$$g_t \equiv \frac{\dot{Q}_t}{Q_t}.$$

All pecuniary flows scale with Q_t. Define normalized wages, transfers, research payments, and innovation values by

$$w_t = Q_t \hat{w}_t, \qquad \Pi_t = Q_t \hat{\Pi}_t, \qquad \mathcal{P}_t(z_R) = Q_t \hat{\mathcal{P}}_t(z_R), \qquad V_t = Q_t \hat{V}_t. \tag{56}$$

Worker values. Begin with the value of permanently entering production, which in levels satisfies

$$(\rho + \delta) W_t - \dot{W}_t = \ln\left((1 - \tau_W)(w_t z_P + \Pi_t)\right). \tag{57}$$

Using (56),

$$(\rho + \delta) W_t - \dot{W}_t = \ln Q_t + \ln\Big((1 - \tau_W)(\hat{w}_t z_P + \hat{\Pi}_t)\Big). \tag{58}$$

Define the common growth component A_t as the solution to

$$(\rho + \delta) A_t - \dot{A}_t = \ln Q_t, \tag{59}$$

and define the normalized production value

$$\hat{W}_t \equiv W_t - A_t. \tag{60}$$

It follows that

$$(\rho + \delta) \hat{W}_t - \dot{\hat{W}}_t = \ln\Big((1 - \tau_W)(\hat{w}_t z_P + \hat{\Pi}_t)\Big). \tag{61}$$

Analogously, define

$$\hat{\mathcal{R}}_t(z_R) \equiv \mathcal{R}_t(z_R) - A_t. \tag{62}$$

In the continuation region, the normalized researcher HJB is

$$(\rho + \delta) \hat{\mathcal{R}}_t(z_R) - \dot{\hat{\mathcal{R}}}_t(z_R) = \ln\Big((1 - \tau_W)[\hat{\mathcal{P}}_t(z_R) + \hat{\Pi}_t]\Big). \tag{63}$$

Because the same additive growth component is removed from both value functions, the researcher's continuation surplus is invariant to this normalization:

$$\mathcal{D}_t(z_R) \equiv \mathcal{R}_t(z_R) - W_t = \hat{\mathcal{R}}_t(z_R) - \hat{W}_t. \tag{64}$$

Firm values. The value of a research match is a pecuniary object and therefore scales multiplicatively with Q_t. Define

$$\hat{\mathcal{V}}_t(z_R) \equiv \frac{\mathcal{V}_t(z_R)}{Q_t}. \tag{65}$$

The level HJB,

$$(r^* + \delta)\mathcal{V}_t(z_R) - \dot{\mathcal{V}}_t(z_R) = V_t \varphi_R z_R - \mathcal{P}_t(z_R), \tag{66}$$

then becomes

$$(r^* + \delta - g_t)\hat{\mathcal{V}}_t(z_R) - \dot{\hat{\mathcal{V}}}_t(z_R) = \hat{V}_t \varphi_R z_R - \hat{\mathcal{P}}_t(z_R). \tag{67}$$

Nash payment. From (50), the interior Nash condition is

$$\theta \frac{(1-\tau_W)u'(c_t^R(z_R))}{\mathcal{D}_t(z_R)} = (1-\theta)\frac{1}{\mathcal{V}_t(z_R)}, \tag{68}$$

where

$$c_t^R(z_R) = (1-\tau_W)\left[\mathcal{P}_t(z_R) + \Pi_t\right].$$

Equivalently,

$$u'(c_t^R(z_R)) = \frac{1-\theta}{\theta}\frac{\mathcal{D}_t(z_R)}{\mathcal{V}_t(z_R)}\left(\frac{1}{1-\tau_W}\right). \tag{69}$$

Specializing to log utility, $u'(c) = 1/c$, gives

$$c_t^R(z_R) = (1-\tau_W)\frac{\theta}{1-\theta}\frac{\mathcal{V}_t(z_R)}{\mathcal{D}_t(z_R)}. \tag{70}$$

Using the normalizations above,

$$\mathcal{V}_t(z_R) = Q_t \hat{\mathcal{V}}_t(z_R),$$

and

$$c_t^R(z_R) = (1-\tau_W)Q_t\left[\hat{\mathcal{P}}_t(z_R) + \hat{\Pi}_t\right].$$

Cancelling Q_t therefore yields

$$\left[\hat{\mathcal{P}}_t(z_R) + \hat{\Pi}_t\right] = \frac{\theta}{1-\theta}\frac{\hat{\mathcal{V}}_t(z_R)}{\mathcal{D}_t(z_R)}. \tag{71}$$

Hence the normalized research payment is

$$\boxed{\hat{\mathcal{P}}_t(z_R) = \frac{\theta}{1-\theta}\frac{\hat{\mathcal{V}}_t(z_R)}{\mathcal{D}_t(z_R)} - \hat{\Pi}_t.} \tag{72}$$

Equations (61), (63), (67), and (72) characterize the bargaining problem in the interior of the continuation region, which is exactly the region where $\hat{\mathcal{R}}_t(z_R) > \hat{W}_t$.

B.2 Deriving the Measurement Block

In the main text, I assume that ideas arrive at a Poisson rate $\varphi_R z$ while patents are derived from a negative binomial distribution with parameters $\tilde{A}\varphi_R z$ and ϱ. In this section, I make explicit the connection between the two.

Suppose that an innovator generates ideas at a rate $\varphi_R z$. In the data, the variance of patents is considerably higher than the variance implied by a Poisson process. This overdispersion arises naturally in the case of persistent inventor heterogeneity in how frequently one of their ideas is patented. This heterogeneity might arise from legal conditions about patentability, employer behavior, or many other sources of variation in propensity to patent that do not reflect variation in innovation. I model this heterogeneity as an additional draw γ from a patenting propensity distribution that I assume to be independent of z and distributed $\Gamma(\varrho, \varrho)$, so that it has mean 1. Hence, patents arrive at a Poisson rate $\gamma\tilde{A}\varphi_R z_R$. Notice that while conditional on (z, γ), this is a Poisson process, from the perspective of an econometrician observing only z, this process is itself negative binomial. Importantly, this heterogeneity in patenting propensity is not priced. Individuals in my model are paid for ideas, not for patents.

The cross-sectional distribution of patents is then given by:

$$\begin{aligned}
\mathbb{P}[n = k \mid z] &= \int_0^\infty f(\gamma)\, e^{-\gamma\tilde{A}\varphi_R z}\, \frac{(\gamma\tilde{A}\varphi_R z)^k}{k!}\, d\gamma \\
&= \frac{(\tilde{A}\varphi_R z)^k\, \varrho^\varrho}{k!\,\Gamma(\varrho)} \int_0^\infty \gamma^{k+\varrho-1} e^{-(\tilde{A}\varphi_R z+\varrho)\gamma}\, d\gamma \\
&= \frac{\Gamma(k+\varrho)}{\Gamma(\varrho)\, k!} \left(\frac{\varrho}{\varrho + \tilde{A}\varphi_R z}\right)^\varrho \left(\frac{\tilde{A}\varphi_R z}{\varrho + \tilde{A}\varphi_R z}\right)^k, \\
\mathbb{P}[n = k] &= \int_z g_R(z)\, \mathbb{P}[n = k \mid z]\, dz.
\end{aligned}$$

Now I claim that the expected number of patents and the expected number of ideas are proportional. To see this, observe that the expected rate of idea arrival is given by:

$$\mathbb{E}[\varphi_R z_R] = \varphi_R \mathbb{E}[z_R] = \varphi_R H_R$$

The expected production of patents is given by:

$$\begin{aligned}\mathbb{E}[N] &= \mathbb{E}\big[\mathbb{E}[N|z,\gamma]\big] \\ &= \mathbb{E}[\varphi_R z \tilde{A}\gamma] \\ &= \varphi_R \tilde{A}\mathbb{E}[z]\mathbb{E}[\gamma] \\ &= \varphi_R \tilde{A} H_R\end{aligned}$$

C Computation

This appendix describes how I compute the competitive equilibrium of the quantitative model in Section 5. I work directly with the detrended model. I write $m(n) \equiv \mathbb{E}[z_R \mid n_R = n]$ for expected research ability given the signal and order program tiers so that $\alpha_1 > \alpha_2 > \cdots > \alpha_J$, with $J = 10$.

C.1 Discretization

Log research ability $\ln z_R = \eta + \gamma$ is the sum of a normal and an exponential draw, so its density is exponentially modified Gaussian and available in closed form. I place log ability on an evenly spaced grid of 400 points. The grid runs from five standard deviations of η below its mean to an upper edge chosen so that the ability-weighted mass beyond it is of order e^{-12}. The signal has a normal distribution, since $n_R = \eta + \varepsilon$ does not load on the tail component. Conditional on the signal, ability is the sum of a normal posterior for η and the unobserved exponential component. I compute the posterior objects the model requires, namely $m(n)$, $\mathbb{P}[z_R \geq \bar{z} \mid n]$, $\mathbb{E}[z_R \mathbf{1}\{z_R \geq \bar{z}\} \mid n]$ and the full posterior density on the ability grid, by 24-node Gauss–Laguerre quadrature over the exponential component, with closed-form normal integrals inside each node. Integrals over signal bands, such as the mean ability $\bar{m}_j$ of the students admitted to tier j, use 40-node Gauss–Legendre quadrature, and signal quantiles are computed by Newton's method on the signal CDF.

C.2 The Balanced Growth Path

On the balanced growth path the growth rate $g = \iota \ln \lambda$ and every detrended object are constant, and the equilibrium reduces to a square system of $2J + 3$ nonlinear equations. The unknowns are the innovation rate ι, the detrended value of a product line V, admissions per program in each tier $x_1, \ldots, x_J$, the lower signal cutoff of each tier $n_1 > n_2 > \cdots > n_J$, and the per-capita dividend d paid by the mutual fund.[32] Given a guess for these unknowns, the residuals are built in three blocks.

[32] Strictly speaking, given the mass x_j, one can infer the cutoffs and vice versa. I keep both because empirically it seems to improve the convergence properties of the algorithm.

Individuals and bargaining. Given (g, V, d), the research cutoff z^* solves $\varphi_R z^* V/(1 - s_R) = \xi_2 w$, where $w = 1/\lambda$ is the detrended production wage, $\xi_2 w$ is the outside option of a graduate, and $1/(1 - s_R)$ is the markup on the firm's payoff implied by the credit at rate s_R on the research payroll: a dollar paid to a researcher costs the firm $1 - s_R$, so a match forms whenever $\varphi_R z V \geq (1 - s_R)\xi_2 w$. Above the cutoff, the bargained payment $\mathcal{P}(z)$ solves the balanced growth path version of the Nash sharing rule under continuous renegotiation,

$$\theta \, \frac{\rho + \delta}{r^* + \delta - g} \Big(\frac{\varphi_R z V}{1 - s_R} - \mathcal{P} \Big) = (1 - \theta) \big(\mathcal{P} + d \big) \ln \Big(\frac{\mathcal{P} + d}{\xi_2 w + d} \Big),$$

which I solve pointwise on the ability grid by bisection followed by Newton polishing. The researcher's surplus over permanent exit is $\mathcal{D}(z) = \ln \big((\mathcal{P}(z) + d)/(\xi_2 w + d) \big)/(\rho + \delta)$. The reservation stipend of an applicant with signal n then solves

$$\mathcal{A} \ln \big(s^{\mathrm{res}}(n) + d \big) + q^{\mathcal{O}} \Big(W^g + \mathbb{E} \big[\mathcal{D}(z_R) \mid n \big] \Big) = W,$$

where $q = \mathcal{O}\mu/(\rho + \delta + \mathcal{O}\mu)$ is the discounted probability of clearing one Erlang stage, $\mathcal{A} = (1 - q^{\mathcal{O}})/(\rho + \delta)$ is the utility annuity of a graduate-school spell, and W and W^g are the values of production work at w and $\xi_2 w$. The expectation is taken against the posterior density on the ability grid.

Graduate programs. Because the break-even constraint binds cohort by cohort, each tier's problem is static given the financial annuity of a spell, $A_\lambda = (1 - q_\lambda^{\mathcal{O}})/(r^* + \delta - g)$ with $q_\lambda = \mathcal{O}\mu/(r^* + \delta - g + \mathcal{O}\mu)$. The first-order conditions of problem (32) imply that the stipend schedule of tier j is affine in expected ability, $s_j(n) = \sigma_j m(n) + c_j$. The slope and intercept are

$$\sigma_j = \alpha_j + \frac{1}{\Lambda_j}, \qquad \Lambda_j = \frac{\bar{m}_j A_\lambda (1 + \nu)}{\nu K'(x_j)}, \qquad c_j = f - \frac{K(x_j)}{x_j A_\lambda} - \frac{\bar{m}_j}{\Lambda_j},$$

where Λ_j is the present value of the multiplier on the cohort's budget. The band means $\bar{m}_j$ are computed from the guessed cutoffs. Schedules that satisfy these conditions satisfy each tier's break-even constraint by construction.

Equilibrium conditions. The $2J + 3$ residuals are:

1. consistency of the innovation rate, $\iota = \chi + \varphi_R H_R$, where H_R integrates $\mathbb{E}[z_R \mathbf{1}\{z_R \geq z^*\} \mid n]$ over the admitted signal distribution and the Erlang survival probability to graduation;

2. the product-line value, $V = (1-1/\lambda)L/(r^*+\iota-g)$, with production labor L computed from the eligible and ineligible populations, the enrolled stock, researchers, and PhDs working in production at ξ_2;

3. participation at the bottom of the bottom tier, $s_J(n_J) = s^{\text{res}}(n_J)$;

4. indifference at each interior cutoff, $s_j(n_j) = s_{j+1}(n_j)$ for $j = 1, \ldots, J-1$, which ensures that the equilibrium stipend schedule is continuous as required by Lemma 2;

5. clearing in each tier, so that admissions $M_j x_j$ equal the flow $\delta\pi_R$ of eligible newborns times the mass of the signal band $[n_j, n_{j-1})$;

6. consistency of the dividend, $d = (1 - 1/\lambda)L - \text{payroll} + s_R\,\text{payroll}$, the fund's profits net of the research payroll plus the R&D credit it receives on that payroll.

The proportional income tax that finances f and s_R does not appear in any residual. The proportional income tax in this environment is non-distortionary and can therefore be inferred as a residual.

Solution algorithm. I solve the system by damped Newton's method, with the Jacobian computed by forward-mode automatic differentiation. Each step is scaled so that no unknown other than the dividend changes by more than 20% of its current value, and the dividend step is capped at 20% of a production wage. If Newton's method fails to reach a sup-norm residual of 10^{-7}, typically from a cold start far from the root, I restart from the same guess with a trust-region Levenberg–Marquardt solver on the same residual and Jacobian, and then polish with Newton. I accept a root only if its cutoffs are strictly ordered, its admission masses are positive, and its schedule slopes σ_j are decreasing in j. Within the calibration, I warmstart every new evaluation from the prior one. All derivatives are computed via the implicit function theorem applied directly to the underlying equilibrium system.

C.3 Transition Paths

A counterfactual specifies an exogenous path for the federal appropriation to universities, and in the instrument-frontier exercise also for the cost of the R&D credit. The economy starts from the stocks of an initial balanced growth path and converges to a terminal balanced growth path at the new long-run policy. For temporary changes, the terminal path equals the initial one. For permanent changes, I first solve the terminal path by walking the policy parameter from its old value to its new one with warm starts. When the policy is stated as an aggregate budget, I also solve by the secant method for the per-student rate f at which $fN(f)$ equals the new appropriation.

Unknowns. Time is discretized at quarterly nodes, $\Delta t = 0.25$, over a horizon of 120 years (480 nodes) for temporary changes in policy and 200 years (800 nodes) for permanent ones. At each node the unknowns are the innovation rate ι_t, production labor L_t, each tier's multiplier on its cohort budget $\lambda_{1t}, \ldots, \lambda_{Jt}$, the per-student subsidy f_t (endogenous because a fixed appropriation is spread over endogenous enrollment), and the dividend d_t. This gives 14 unknowns per node, and 6,720 or 11,200 in total. Given a guess for the stacked path U, the residual $\mathcal{E}(U)$ is computed in two sweeps.

Backward sweep. Starting from the terminal balanced growth path's values and working backward, I integrate every value function in the model. These are the product-line value V_t, the values of production work W_t and W_t^g, the firm's value of a match $\mathcal{V}_t(z)$, the researcher's surplus $\mathcal{D}_t(z)$, the students' stage-by-stage values, and the programs' stage-by-stage annuities $A_{\lambda,t}$ and subsidy present values $\mathcal{F}_t$. Every one of these solves a linear ODE of the form $a_t y_t - \dot{y}_t = \mathrm{src}_t$ given the guessed aggregates, and the Erlang stage values solve a triangular system of such ODEs. I integrate them exactly over each interval under the assumption that the source is linear in time between nodes, which yields

$$y_t = e^{-a\Delta t} y_{t+\Delta t} + \omega_0 \,\mathrm{src}_t + \omega_1 \,\mathrm{src}_{t+\Delta t},$$

where the weights ω_0 and ω_1 depend only on $a\Delta t$. The bargaining problem is solved node by node. Given the continuation values at $t+\Delta t$, which are zero where the match separates, the payment solves $\theta\,\mathcal{V}_t(z) = (1-\theta)(\mathcal{P}_t + d_t)\,\mathcal{D}_t(z)$, where $\mathcal{V}_t$ and $\mathcal{D}_t$ are themselves one-step functions of $\mathcal{P}_t$. I solve this pointwise by Newton's method in $\ln(\mathcal{P}_t + d_t)$ from the next node's payment. The research cutoff z_t^* is the zero of the feasibility gap, the ability at which no payment leaves both sides with non-negative surplus. Below the next node's cutoff this zero is available in closed form, and above it I locate via linear interpolation on the grid.

Forward sweep. Starting from the initial stocks and working forward, I integrate the stock variables. These are the enrolled population by Erlang stage and ability, the stock of living graduates by ability, researchers by ability, and, for each tier, enrollment, the stock of enrolled expected ability, and the stocks of contracted net revenue, subsidy and stipend obligations. Given their inflows, all of these obey linear ODEs with a constant generator G, the Erlang chain of stages with death at rate δ. I integrate them exactly over each interval, with inflows linear in time between nodes, using the matrix exponential of an augmented generator (Van Loan's method). The inflows at $t + \Delta t$ depend on the policies at $t + \Delta t$, which in turn depend on the stocks at $t + \Delta t$, so the step is implicit. I resolve it with a zero-order-hold predictor followed by three fixed-point corrections.

Because exit from research is permanent, researchers cannot be recovered from the graduate stock and the current cutoff alone, and hence one must track the mass of researchers in each grid cell. Between nodes, the fraction of a cell's occupied mass that falls newly below the cutoff exits for good, and new graduates enter each cell only above the cutoff. Within-cell masses are computed from a cubic Hermite reconstruction of the graduate density, so that the sharp edge at the cutoff is integrated exactly and the residual is continuous in the cutoff.

At each node the policy block computes, given the stocks and the values from the backward sweep:

1. the admissions schedule, which solves the fixed point $x_{jt} = \left(\bar{m}_{jt} A_{\lambda,t}(1+\nu)/(\nu\kappa\Lambda_{jt})\right)^{1/\nu}$, where the band means $\bar{m}_{jt}$ depend on x_t through the cutoffs, by damped Newton's method in $\ln x_t$ from the balanced growth path's schedule.

2. the band edges, including the entry cutoff n_{Jt}, from market clearing given x_t via inversion on the signal CDF

3. the stipend intercepts: c_{Jt} from the participation condition at n_{Jt}, and the remaining intercepts by recursion at the tier intersections

4. the multipliers λ_{jt} implied by each cohort's break-even constraint, with the subsidy valued at its present value $\mathcal{F}_t$ over the spell

5. the aggregates implied by the stocks: $\iota_t^{\text{impl}} = \chi + \varphi_R \sum_z z\, R_t(z)$, L_t^{impl}, the dividend, and aggregate subsidy outlays $f_t N_t$.

The residual stacks the gaps between the guessed and implied innovation rate, labor, multipliers and dividend at every node, together with the gap between $f_t N_t$ and the exogenous appropriation. When computing the efficacy of R&D tax credits vs. university spending, the R&D credit rate is an additional unknown at each node.

Solution algorithm. I solve $\mathcal{E}(U) = 0$ with a hybrid chord–Newton method. First, I compute the full Jacobian of the stacked system once, at the terminal balanced growth path tiled across all nodes, from Jacobian–vector products evaluated in batches of up to 1,024 by forward-mode automatic differentiation. I factor it by LU decomposition and reuse the factorization across every scenario that shares a terminal balanced growth path. Starting from the terminal path, I take chord steps $U \leftarrow U - J_0^{-1}\mathcal{E}(U)$. These are cheap and converge quickly for moderate shocks. If a chord step fails to reduce the sup-norm residual by at least a factor of five, I switch to inexact Newton steps. The Newton system is solved by restarted GMRES with restart length 8, using exact matrix-free Jacobian–vector products

of $\mathcal{E}$ at the current iterate and the stored LU factorization as a preconditioner. I tighten the inner tolerance as the outer residual falls, using an Eisenstat–Walker forcing term. Every step is damped so that no unknown moves by more than half its own magnitude. I iterate to a residual of 10^{-9}. For long or large scenarios that stall from the balanced growth path start, I walk across nearby scenarios.

Accuracy. I measure every counterfactual relative to a transition computed with the same code under unchanged policy. On a flat path the growth rate agrees with the balanced growth path's to within 2×10^{-4} basis points. Second, I check sensitivity to the grids. I re-solve the balanced growth path on ability grids of 200, 400 and 800 points, and I recompute the welfare results on a 200-point grid. Temporary changes in policy leave the economy close to the balanced growth path, so I solve them with a 120-year horizon. Permanent changes converge more slowly, so I solve them over a 200-year horizon.

Initial cohorts. Cohorts enrolled before an unanticipated change in policy were admitted under the old subsidy and are no longer self-financing afterward. I hold their stipend at its booked value, and the government continues to pay these legacy cohorts the original per-student subsidy f_0. The resulting top-up, $(f_0 - f_t)N_t^{\text{leg}}$, where N_t^{leg} is the date-0 enrolled stock aged without inflow, is financed through an additional proportional income tax.

C.4 Welfare

Following Calvo and Obstfeld (1988), the planner discounts the future at a rate ρ. Given this, I measure welfare as the utilitarian, ρ-discounted sum of the flow utility of everyone alive at a given instant.

With log utility, the flow utility of an individual with detrended pre-tax income y_{it} is $\ln Q_t + \ln(1 - \tau_t) + \ln(y_{it} + d_t)$. The welfare change between a counterfactual and the flat baseline therefore decomposes into four pieces: a growth term $\Delta \ln Q_t$, a tax term $\Delta \ln(1 - \tau_t)$, a dividend term, and a composition term from the reallocation of individuals across occupations and incomes. The composition term is computed exactly from the transition's sums of $\ln(y_{it} + d_t)$ over students, researchers and graduates. I integrate each piece over the horizon at the nodes. Beyond the horizon T I append the continuation value implied by the terminal balanced growth path, discounted by $e^{-\rho T}$: the terminal level gap divided by ρ, plus, for the growth term, the terminal growth gap divided by ρ^2. The consumption-equivalent variation reported in Section 5 is $\exp(\rho\, \Delta\mathcal{W}) - 1$. This is the permanent proportional change in consumption in the baseline that delivers the same welfare as the counterfactual.

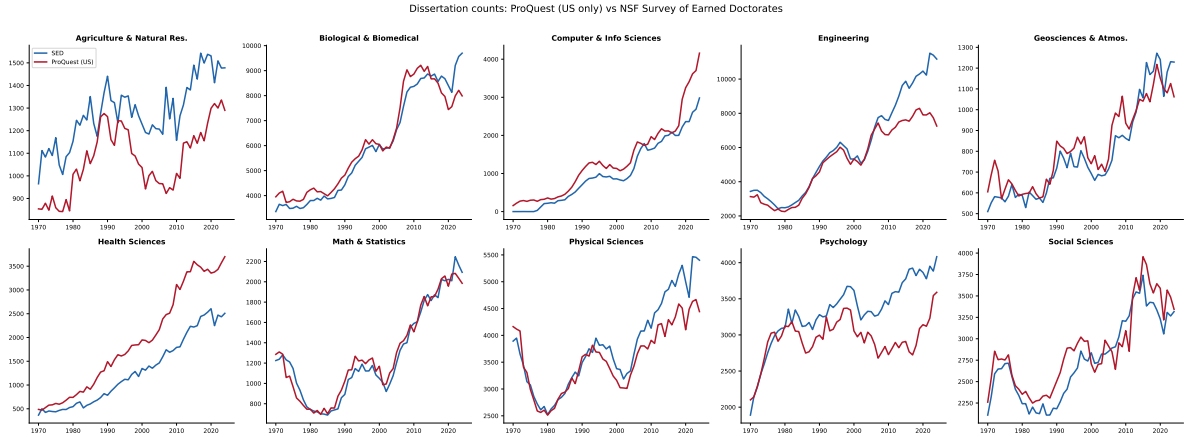
D Data

This section describes the dataset used in the paper, including the construction of the patenting outcomes and the coverage of the ProQuest Dissertation and Thesis Database relative to the Survey of Earned Doctorates.

D.1 Coverage

For idiosyncratic reasons, ProQuest was the de-facto primary repository for dissertations in the United States for more than half a century. Figure 12 compares the number of dissertations in ProQuest to the number of doctorates reported in the Survey of Earned Doctorates (SED) by field and year. Coverage is generally quite good until circa 2010, when institutional repositories begin to supplant ProQuest at some institutions. As a result, most analyses in the baseline restrict the sample to dissertations between either 1980 or 1990 and 2010.

Figure 12: ProQuest Coverage Relative to the Survey of Earned Doctorates



D.2 Data Construction

This subsection describes the algorithm that links dissertation authors in ProQuest to inventors in PatentsView. The rough structure of the algorithm is as follows:

1. Block the data by name to generate candidate pairs of patents and dissertations.

2. Score each patent against every dissertation in its block using a combination of text similarity, name similarity, temporal plausibility and advisor co-inventorship.

3. Assign inventors to dissertations one-to-one by a greedy algorithm that prioritizes the earliest matched patent.

4. Keep only inventor–dissertation pairs that a large share of the inventor's patents corroborate, and that have a high mean text similarity across corroborating patents.

The dissertation side consists of ProQuest records for doctoral dissertations in STEM fields. Each record reports the author's first, middle and last names, the names of the dissertation advisors, the degree date, the title, and the subject terms, classification terms and keywords attached to the dissertation. I drop dissertations from institutions whose name places them in East, South or Southeast Asia due to well known issues with matching Asian names.

The patent side consists of the 3,832,919 granted patents in PatentsView with at least one US-based inventor and an assignee. For each patent I use the disambiguated inventor identities, the title, the abstract and the grant date. Before matching, I merge 4,424 PatentsView inventor identifiers into 2,182, consolidating identifiers that the disambiguation splits for a single inventor. The identifiers in each merged group share a last name and at least one assignee, and their first names are identical or variants of one another.[33] All names are lower-cased and stripped of punctuation, and a two-part first name with no recorded middle name (for instance, Chun-Chen) is split into a first and a "middle" name.

The first step of the matching process is to block the data. Specifically, I compare a patent and dissertation only if the names of the inventor and the author agree. I construct blocks from tightest to loosest in four stages:

1. Exact match on last name, full first name and middle initial.

2. Exact match on last name and full first name.

3. Exact match on last name, first initial and middle initial.

4. Exact match on last name and first initial.

Inventor records can be matched in at most one stage. As is well-known, pure name-based blocking generally performs horribly. Fortunately, both patents and dissertations contain substantial additional information that can be used to improve matching. I next turn to describing how I use this information to score candidate patent–dissertation pairs within each block.

I represent each dissertation by the concatenation of its title, subject terms, classification terms and keywords, and each patent by its title and abstract. I then embed the concatenated terms using SPECTER2 (Singh et al., 2023), a transformer model trained on the scientific literature to place documents that cite one another close together (Cohan et al., 2020), and

[33]For example, Abhijit Abhyankar and Abhijit M. Abhyankar are consolidated.

normalize every embedding to unit length. As is standard, I define text similarity $\mathrm{sim}(p, d)$ of a patent p and a dissertation d as the cosine similarity between their embeddings.

Within a block, I then score every patent–inventor record p against every dissertation d by

$$S(p, d) = 0.5\,\mathrm{sim}(p, d) + 0.25\,N(p, d) + 0.25\,T(p, d) + 0.15\,A(p, d). \tag{73}$$

$N(p, d) \in [0, 1]$ measures agreement of first and middle names beyond what the block requires. It rewards exact matches and matches on the first three letters of full first and middle names and penalizes conflicts. $T(p, d)$ measures temporal plausibility as a function of the gap g between the patent's grant year and the year of the degree:

$$T(g) = \begin{cases} \exp\!\left(-\frac{\ln 20}{5}\,(-3 - g)\right) & g < -3, \\ 1 & -3 \le g \le 5, \\ \exp\!\left(-\frac{\ln 20}{20}\,(g - 5)\right) & g > 5, \end{cases} \tag{74}$$

so a patent granted between three years before and five years after the degree receives full weight, and the weight falls to 0.05 eight years before the degree and twenty-five years after it. The weights were manually tuned by hand-reviewing the results of the matching process. $A(p, d)$ equals one if the surname of one of the dissertation's advisors appears among the inventor's co-inventors on any patent. To avoid typical concerns about agglomeration due to common surnames, I limit consideration to surnames shared by fewer than 50 inventor identifiers. Each patent–inventor record keeps its highest-scoring dissertation, and the match is retained if $S(p, d) \ge 0.50$ in stages 1 and 2 or $S(p, d) \ge 0.55$ in stages 3 and 4.

With the patent-level matches in hand, I then aggregate patent-level matches to inventor–dissertation pairs. First, I discard patent-level matches in which the full first names of the inventor and the author have a Jaro–Winkler similarity below 0.85. This removes matches admitted by the initial-only blocks despite conflicting first names. For each remaining inventor–dissertation pair, I record the number of the inventor's patents that independently selected that dissertation ("corroborating patents"), and their mean text similarity. Pairs are assigned one-to-one by a greedy algorithm, which first sorts pairs by the score of the inventor's earliest matched patent. This yields 488,775 candidate inventor–dissertation pairs.

The next step is to prune false candidate matches. To understand how I do this, I introduce notation here. Let P_i denote the total number of patents of inventor i and C_i the number of corroborating patents for the dissertation assigned to i. I keep a pair if the following three things are true:

1. $C_i/P_i \ge 0.8$ when $P_i \le 10$, or $C_i \ge 8$ when $P_i > 10$.

2. The mean text similarity across corroborating patents is at least 0.80.

3. The pair was first matched in stages 1 to 3.

After completing this pruning step, I drop dissertations whose primary subject term lies outside STEM. The resulting analysis sample includes 264,996 matched PhDs.

Before using the sample for analysis, I first address a recurring issue with the PatentsView disambiguation. Occasionally, PatentsView assigns the patents of several distinct people to a single inventor identifier. After extensive manual review, I constructed an algorithm designed to detect this agglomeration ("globbing"). Specifically, I flag an identifier as pooled if its patents spread over at least three secondary assignee clusters and at least three secondary CPC subclasses, each with at least five patents, and at least two of the secondary clusters lie in a different CPC section from the inventor's primary one. The last condition separates pooled identifiers from single people whose patents span several assignees. Hand review discovered that most identifiers meeting only the first two conditions belong to one person.[34] The rule flags 340 of the 264,996 matched inventors, who I then drop.

D.2.1 Validation

As with any matching exercise, there are two primary concerns when evaluating the quality: precision (the frequency with which claimed matches are true) and recall (the proportion of true matches that are correctly identified).

Precision. To evaluate the precision of the algorithm, I first construct a sample of clean matches. To do this, I scrape profile records on ORCID. ORCID is a non-profit organization that provides unique identifiers for researchers. Importantly for me, individuals can both list their PhD institution and claim patents on their ORCID profile.

Recall. To evaluate recall, I use confidential microdata from the United States Census Bureau and the National Center for Science and Engineering Statistics (NCSES). The Census Bureau maintains a confidential version of the Survey of Earned Doctorates (SED) that includes Census Protected Identification Keys ("PIKs") for individuals who received a doctorate from a U.S. institution between 2000 and 2017. Using PIKs and the crosswalk in Akcigit and Goldschlag (2023), I am able to construct an analogous sample of matched inventors and PhDs. I then compare the number of patents in the SED-PatentsView sample

[34] Typically, this was an inventor whose employer changed its name, whose patents were reassigned to licensees, who founded a series of companies, or was associated with an Intellectual Ventures "invention session."

Table 6: Precision of the Match Against ORCID Profiles

Stage	Name block	Linked	Correct	Incorrect	Precision
1	Last name, full first name, middle initial	323	311	12	96.3%
2	Last name, full first name	366	328	38	89.6%
3	Last name, first initial, middle initial	7	6	1	85.7%
	All stages	696	645	51	92.7%

Notes: ORCID profiles come from a search of the ORCID public API for "patent" (77,408 profiles); 51,576 list at least one patent among their works and 21,181 of those list a doctorate. A profile is linked to the analysis sample when it claims a patent in the PatentsView universe on which an inventor has the profile's last name and first name or first initial (1,157 profiles), and that inventor is in the analysis sample (696). Each profile is linked at most once. A link is correct when the doctoral institution on the ORCID profile matches the institution of the matched dissertation, after harmonizing abbreviations and university-system and campus names. Stage is the name block at which the inventor was first matched; stage 4 matches are not in the analysis sample.

with at least one PhD inventor to the analogous sample of inventors in my data, separately by the largest CPC classes in my field. The results of this exercise are shown in Table 7.

Table 7: PhD Shares of Inventors and Patents: This Match and the SED–PatentsView Linkage

		Inventor share (%)			Patent share (%)		
CPC	Technology	SED–PV	This paper	Difference	SED–PV	This paper	Difference
H10	Semiconductors	29.5	23.1	−6.4	61.1	60.7	−0.4
H01	Basic electronics	17.6	14.9	−2.7	38.3	36.6	−1.7
G01	Instruments	16.7	14.7	−2.0	32.2	32.8	+0.6
A61	Medical/pharma	14.4	12.3	−2.1	25.2	26.9	+1.7
H02	Power generation	10.0	9.7	−0.3	22.6	23.2	+0.6
A01	Agriculture	9.1	9.7	+0.6	20.0	24.1	+4.1
G06	Computing	9.0	7.6	−1.4	22.5	23.4	+0.9
H04	Telecom	8.0	7.6	−0.4	25.9	27.6	+1.7
B60	Vehicles	5.7	5.4	−0.3	12.5	14.9	+2.4
Mean absolute difference				1.8			1.6
Rank correlation (Spearman)				0.98			0.93

Notes: SED–PatentsView links U.S. doctorates awarded in 2000–2017 in the confidential Survey of Earned Doctorates to PatentsView inventors through Census PIKs and the crosswalk of Akcigit and Goldschlag (2023), for STEM PhDs and patents granted in 2018–2023. "This paper" restricts the analysis sample to the same population: PhDs from US institutions with degrees in 2000–2017, and patents granted in 2018–2023, each classified by its primary CPC class. I define patent share by the fraction of patents with at least one matched PhD inventor.

In general, I tend to undermatch individual inventors but overmatch patents by a bit relative to the SED-PatentsView linkage. This pattern likely reflects imperfection in the matching process but also legitimate matches. PhDs are highly assortative on patents, and the at least one share I use to define patenting shares is robust to this.

Additional Validation. Although I use a substantial amount of information directly in the construction of the match, I also examine the extent to which the match is consistent with well known university specializations. Consider, for instance, the University of Akron, which is primarily known for its world-class polymer science program. Consistent with that, I find that graduates of the University of Akron disproportionately patent in polymer related CPC subclasses. Table 8 provides the results of this and analogous exercises across a number of well-known programs.

Table 8: Patenting of Matched PhDs from Universities with Known Specializations

	Patents		
University	In class	All	Share (%)
C08: Polymers			
SUNY College of Environmental Science and Forestry	309	1,047	29.5
University of Southern Mississippi	1,062	2,512	42.3
University of Akron	3,034	7,195	42.2
Average university			2.9
G02: Optics			
Kent State University	675	2,108	32.0
University of Arizona	2,807	19,159	14.7
University of Central Florida	482	4,081	11.8
University of Rochester	2,460	18,992	13.0
Average university			2.1
A01: Agriculture			
North Dakota State University	708	2,544	27.8
University of Nebraska–Lincoln	1,079	5,788	18.6
Iowa State University	3,392	22,157	15.3
Average university			2.9
B60: Vehicles			
University of Detroit Mercy	170	603	28.2
Wayne State University	785	8,949	8.8
University of Michigan	2,106	47,514	4.4
Average university			2.3

Notes: For each university, the share of the patents of its matched PhDs that carry the CPC class. Patents in multiple CPC classes are counted separately in each class, but each patent can only be assigned to a university once. Average university is the mean of the same share across every university in the analysis sample. Universities are chosen for well-known specializations in given technologies.

I also examine the extent to which dissertations and patents match to fields that make sense. Grouping dissertations into fields and comparing them to the CPC subclasses of their matched patents, I find that the fields of the dissertation and the patents broadly align with what one would expect. Table 9 provides the results of this exercise.

Table 9: Dissertation Fields of Matched PhD Inventors by CPC Subclass (%)

CPC subclass		CS & IT	Chem.	Elec. eng.	Biology	Biomed.	Physics	All other
G06F	Digital data processing	**40**	3	25	4	4	5	19
C07D	Heterocyclic compounds	0	**77**	0	8	10	1	3
A61K	Medical preparations	1	**42**	1	23	22	2	8
H04L	Digital transmission	27	2	**49**	3	3	3	14
C07K	Peptides	1	31	1	**43**	17	1	7
H04W	Wireless networks	15	1	**67**	2	2	2	11
G01N	Analysing materials	2	**30**	8	20	11	13	17
H01L	Semiconductor devices	2	12	32	1	1	**32**	19
C12N	Genetic engineering	1	25	1	**49**	12	2	11
A61B	Diagnosis and surgery	6	9	14	9	**26**	9	27

Notes: Each row is a CPC subclass. Columns give the percentage of matched PhD–patent pairs in the subclass by the field of the PhD's dissertation. A patent carrying several subclasses counts in each. All other is the sum over residual fields. Bold indicates largest field share.

E Additional Empirical Results

This appendix presents additional empirical results.

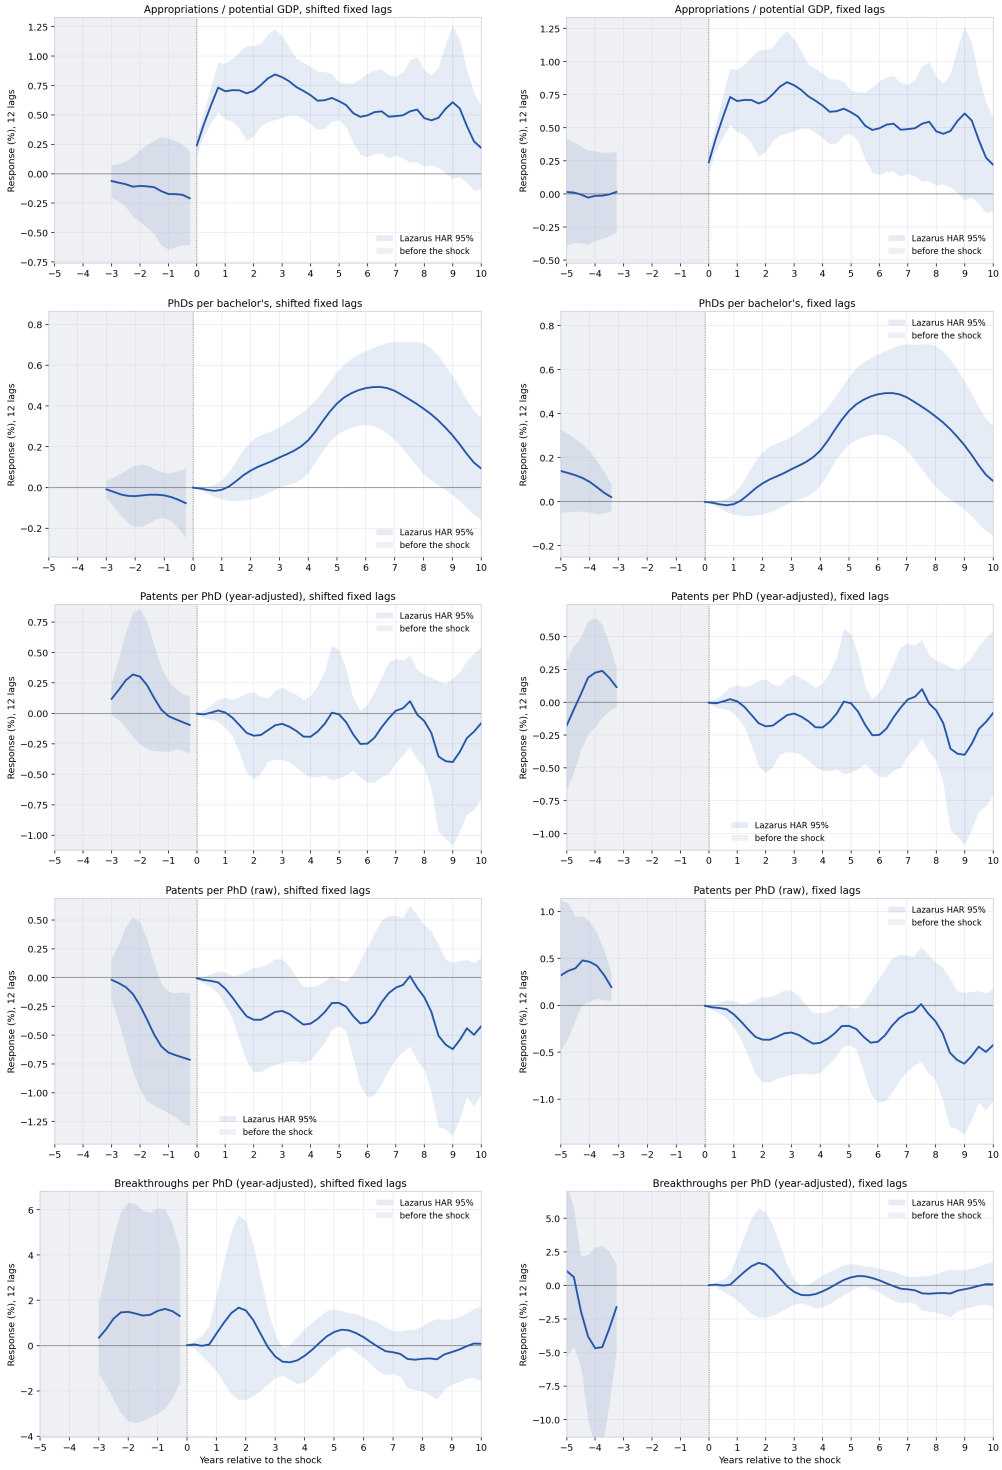
E.1 Quality Local Projections

A natural approach to examine the effect of changes in government spending on the quality of researchers is to project measures of cohort quality on a funding shock. Unfortunately, my measures of patenting output are only available for cohorts that span at least 10 years in my patenting data. This means that I am only able to estimate the shock's effect on cohorts that run from 1980 to 2010. Because this is a small sample, I implement a design that stacks the local projections of each field and restricts the coefficients on shocks, lags, and controls to be identical across fields, while allowing intercepts to vary via field fixed effects. This design, which is listed in Equation 75, improves power by pooling restrictions by requiring that all fields have identical responses to a given shock.

$$\overline{\ln Y}_{f,t+h} = \alpha_f + b_h \, \varepsilon_{f,t} + \sum_{k=1}^{12} \beta_{h,k} \ln Y_{f,t-k} + \sum_{k=1}^{12} \gamma_{h,k} \ln G_{f,t-k} + \sum_{k=1}^{12} \delta'_{h,k} X_{t-k} + u_{f,t+h} \quad (75)$$

One other complication in estimating the effect of funding shocks in such a short sample is the proliferation of patenting in the mid-1990s to early 2000s. This proliferation is unlikely to be related to short-run funding shocks but coincides with a significant expansion of federal support for R&D. Because patenting outcomes must be cumulated over a longer horizon, this proliferation can potentially bias estimates of changes in patents per proquest dissertation downwards in finite sample. To address this concern, I construct a time series of cohort effects by estimating a Poisson regression of total patents by a given cohort on cohort and year fixed effects, and then extract the estimated cohort effects. I then use the estimated cohort effects as one outcome, which I denote by "patents per PhD (year-adjusted)" in Figure 13. I also present results for patents per PhD without this adjustment, which I denote by "patents per PhD (raw)." Breakthroughs are taken from Kelly et al. (2021) and normalized per PhD analogously to the year adjusted patents per PhD measure. Each cohort's breakthroughs are counted over its first eight filing years, which the KPST scores cover only in part for the latest cohorts.

Figure 13: The Response of Cohort Quality to Narrative Shocks



Notes: Stacked local projections of each outcome on the field's Fieldhouse-Mertens appropriation shock, pooling the nine fields with field fixed effects, common lag coefficients and 12 lags of log appropriations, the outcome and the cyclical controls, from 5 years before to 10 years after the shock. Each row is one outcome, and the panels are, from top to bottom: funding as a share of potential GDP, PhDs per bachelor's degree, patents per PhD (year-adjusted), patents per PhD (raw) and breakthroughs per PhD (each cohort's first eight filing years). Left column: in the three years before the shock ($h = -12, \ldots, -1$), the controls are held fixed at $t-27, \ldots, t-16$, clear of every pre-shock outcome. Right column follows de Fraisse et al. (2026), and holds controls fixed in pre-period $t-1, \ldots, t-12$; horizons inside the lag window cannot be estimated and are left blank. Standard errors Lazarus et al. (2018). Outcomes are percent per 1% shock.

The broad takeaway from Figure 13 is that there is little evidence that the quality of PhD cohorts, whether measured by patents per PhD or breakthroughs per PhD, decays substantially as cohorts grow in response to funding shocks. Caution should be exercised in interpreting these results given the weak power of the design. What is quite robust, however, is the relationship between federal funding shocks and PhD yield.

E.2 Critical Technologies

Table 10 reports the PhD share of patents in each of the ten critical technology areas of the CHIPS and Science Act, and Table 11 lists the CPC codes that define each area.

Table 10: PhD Share of Critical and Emerging Technology Patents

	Technology Area	Patents	All PhD	US-Trained PhD
1	AI/ML & Autonomy	31,731	52.5%	49.5%
2	HPC & Semiconductors	33,678	69.9%	67.6%
3	Quantum Information	1,119	76.1%	73.2%
4	Robotics & Adv. Manufacturing	15,077	42.8%	40.2%
5	Disaster Prevention	3,808	30.7%	27.2%
6	Communications & Immersive	148,660	44.2%	41.3%
7	Biotech & Medical	124,735	55.4%	52.0%
8	Data Storage & Cybersecurity	63,858	33.2%	30.6%
9	Energy & Batteries	19,179	51.8%	49.2%
10	Advanced Materials	15,740	69.8%	66.7%
	All Critical Tech	443,339	49.3%	46.3%
	All Patents	891,018	44.1%	41.1%

Notes: Technology areas from the CHIPS & Science Act §10331, mapped to the CPC codes in Table 11. Patents granted 2018–2023, classified by their first-listed CPC code. The PhD share is the fraction of patents with at least one matched STEM PhD inventor; the US-trained column counts only PhDs from US institutions. A patent can count in more than one area, so area totals overlap; All Critical Tech counts each patent once. PhDs are identified from ProQuest dissertations, which are mostly US, so foreign-trained PhD inventors are largely missed and the gap between the two columns understates them.

Table 11: CPC Codes by Critical Technology Area

CPC	Description	Patents	All PhD	US-Trained PhD
(1) AI/ML & Autonomy				
G06N	ML/Neural Networks	8,132	63.4%	60.4%
G01S	Radio Navigation/Radar	6,540	53.6%	50.3%
G06V	Image Analysis	6,114	51.7%	48.6%
B60W	Vehicle Control Systems	3,734	47.1%	44.0%
G01C	Navigation/Surveying	3,609	41.3%	38.8%
G05D	Control Systems	3,602	43.9%	41.1%
	Total	31,731	52.5%	49.5%
(2) HPC & Semiconductors				
H01L	Semiconductor Devices	11,078	67.5%	65.0%

Table 11, continued

CPC	Description	Patents	All PhD	US-Trained PhD
H10D	Semiconductors (reclassified)	7,991	83.0%	80.9%
G11C	Static Stores/Memories	6,289	49.8%	47.5%
H10F	Photovoltaic Devices	2,238	72.3%	70.2%
H10N	Electric Solid-State Devices	1,836	80.2%	77.9%
H10P	Semiconductor Mfg. Processes	1,255	72.8%	70.8%
H10H	Semiconductor Light Sources	1,127	78.8%	76.1%
H10K	Organic Semiconductors	944	79.9%	78.7%
H10W	Semiconductor Packaging	920	71.8%	69.2%
	Total	33,678	69.9%	67.6%
(3) Quantum Information				
G06N10	Quantum Computing	523	79.3%	73.8%
H10N60	Superconducting Devices	249	81.5%	80.7%
H04L9/0852-58	Quantum Key Distribution	109	49.5%	49.5%
H10D48	Quantum Effect Devices	98	68.4%	68.4%
B82Y10	Nano Info Processing	77	87.0%	84.4%
H01F6	Superconducting Magnets	40	60.0%	60.0%
H10N69	Superconducting Integrated	23	95.7%	95.7%
	Total	1,119	76.1%	73.2%
(4) Robotics & Adv. Manufacturing				
B29C	Shaping of Plastics	6,144	40.8%	38.3%
G05B	Control/Regulating Systems	4,096	41.6%	37.9%
B25J	Manipulators/Robots	2,408	53.8%	51.9%
B23K	Soldering/Welding	2,380	39.1%	36.9%
B33Y	Additive Manufacturing	49	42.9%	40.8%
	Total	15,077	42.8%	40.2%
(5) Disaster Prevention				
G08B	Signalling/Alarm Systems	3,542	28.9%	25.6%
G01W	Meteorology	266	54.9%	48.9%
	Total	3,808	30.7%	27.2%
(6) Communications & Immersive				
H04L	Digital Transmission	63,269	38.6%	35.6%
H04W	Wireless Networks	28,363	52.3%	49.1%
H04N	Pictorial Communication	21,835	41.0%	38.6%
G06T	Image Processing	14,889	50.1%	47.6%
H04B	Transmission Systems	11,162	62.6%	59.0%
H04M	Telephonic Communication	4,219	26.4%	23.8%
H04R	Loudspeakers/Microphones	4,185	36.6%	32.1%
H04J	Multiplex Communication	738	62.3%	60.0%
	Total	148,660	44.2%	41.3%
(7) Biotech & Medical				
A61B	Diagnosis/Surgery	38,018	43.9%	40.7%
A61K	Pharmaceuticals	25,990	68.5%	64.8%
A61F	Prostheses/Implants	12,564	33.3%	30.2%
A61M	Devices for Body Cavities	10,701	36.9%	33.3%
C12N	Genetic Engineering	8,308	80.5%	77.2%
C07K	Peptides	7,497	80.3%	76.9%
A61N	Electrotherapy/Radiation	6,491	57.2%	54.0%
C12Q	Measuring/Testing (Bio)	5,424	80.0%	75.2%

Table 11, continued

CPC	Description	Patents	All PhD	US-Trained PhD
G16H	Health Informatics	5,359	43.5%	39.6%
C12P	Fermentation Processes	1,257	78.1%	75.7%
G16B	Bioinformatics	852	77.1%	74.8%
C12M	Bioreactors	848	67.3%	65.4%
A61P	Therapeutic Activity	840	84.3%	81.0%
C07H	Sugars/Nucleosides	586	79.2%	73.4%
	Total	124,735	55.4%	52.0%
(8) Data Storage & Cybersecurity				
G06Q	Data Processing/Business	34,979	28.4%	25.6%
G06F21	Security Arrangements	11,718	33.8%	30.8%
G11C	Static Stores/Memories	6,289	49.8%	47.5%
H04L9	Cryptography/Security Protocols	5,042	33.7%	32.1%
G11B	Recording/Reproducing	2,855	50.4%	47.8%
G06K	Graphical Data Reading	2,855	34.7%	32.3%
H04K	Secret Communication	120	45.8%	40.8%
	Total	63,858	33.2%	30.6%
(9) Energy & Batteries				
H01M	Batteries/Fuel Cells	6,259	60.7%	58.1%
H02J	Power Distribution	5,047	37.5%	34.8%
H02K	Electric Machines	2,003	39.0%	36.7%
H10N	Electric Solid-State Devices	1,836	80.2%	77.9%
H02S	Solar Electric Generators	840	39.4%	37.0%
H01G	Capacitors	728	59.9%	58.4%
F03D	Wind Motors	662	41.4%	39.1%
G21C	Nuclear Reactors	552	49.8%	43.3%
C10L	Fuels	473	56.9%	53.7%
F24S	Solar Heat Collectors	237	43.5%	43.0%
G21F	Radiation Protection	177	36.2%	34.5%
G21K	Nuclear Techniques	109	67.0%	63.3%
G21G	Nuclear Transmutation	87	57.5%	56.3%
G21B	Fusion Reactors	80	58.8%	58.8%
G21D	Nuclear Power Plant	54	59.3%	55.6%
G21H	Nuclear Energy Sources	35	71.4%	71.4%
	Total	19,179	51.8%	49.2%
(10) Advanced Materials				
C09D	Coating Compositions	2,117	67.5%	64.0%
C23C	Coating Metallic Material	2,087	69.2%	66.6%
C08G	Polycondensation Polymers	1,839	77.7%	75.8%
C08F	Addition Polymers	1,771	81.1%	77.3%
C08L	Polymer Compositions	1,444	67.0%	62.9%
C04B	Lime/Ceramics	1,237	57.5%	54.6%
C01B	Non-Metallic Elements/Compounds	1,195	69.8%	66.9%
C08J	Polymer Processing	1,123	65.4%	63.2%
C08K	Polymer Additives	637	66.4%	64.1%
C22C	Alloys	564	65.2%	62.4%
C30	Crystal Growth	384	70.8%	69.3%
B81B	Microstructural Devices	342	70.2%	61.1%
B81C	Microstructural Fabrication	279	74.6%	72.4%
C22B	Production/Refining of Metals	228	53.5%	49.1%

Table 11, continued

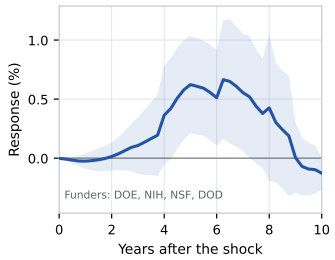
CPC	Description	Patents	All PhD	US-Trained PhD
C22F	Working Metallic Powder	212	65.1%	61.3%
C01G	Metallic Compounds	140	77.9%	75.0%
B82Y	Nanotechnology Applications	105	86.7%	83.8%
B82B	Nanostructures	36	97.2%	97.2%
	Total	15,740	69.8%	66.7%

Notes: Codes at every level of the CPC hierarchy; a patent counts under each code its first-listed CPC code falls in, so rows overlap within and across areas. Totals count each patent once within its area. Grey: fewer than 100 patents.

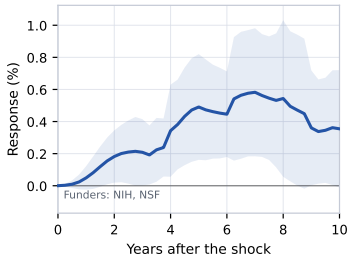
E.3 Local Projections for All Fields

Figures 14 and 15 repeat Figures 2 and 3 for every field.

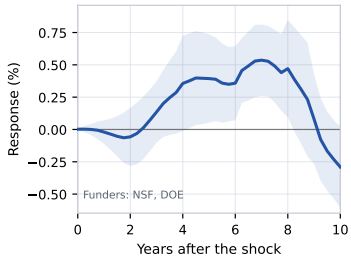
Figure 14: The Response of PhD Yield to Narrative Shocks, All Fields



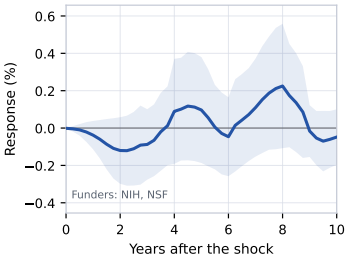
(a) Engineering



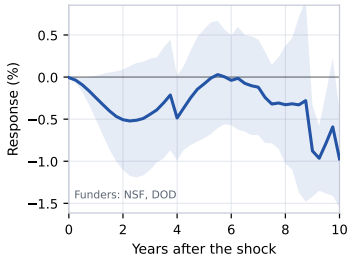
(b) Biology



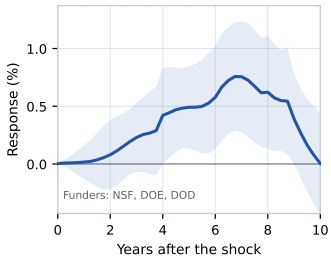
(c) Physics and astronomy



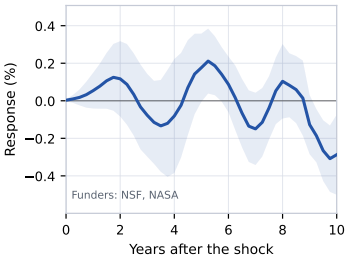
(d) Chemistry



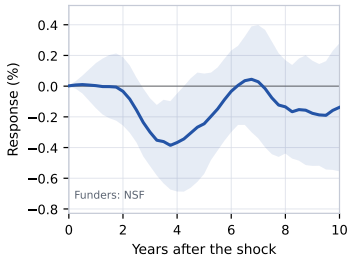
(e) Computer science



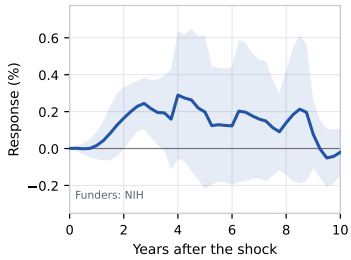
(f) Mathematics and statistics



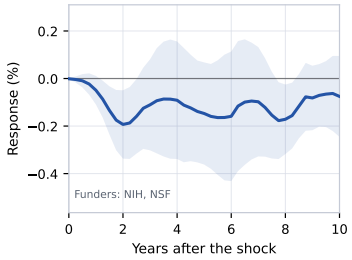
(g) Geosciences



(h) Agriculture



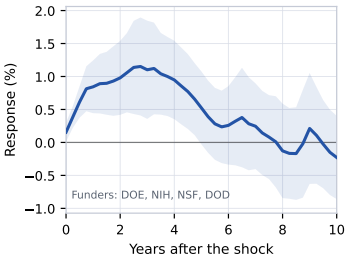
(i) Health sciences



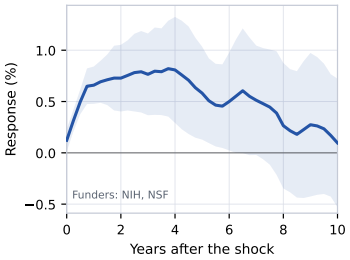
(j) Psychology

Notes: Response of the field's yield rate – SED doctorates per bachelor's degree awarded six years earlier – to a Fieldhouse-Mertens appropriation shock of 1 percent of the field's university-directed appropriations, horizons 0–10 years. Quarterly local projections with 12 lags of log appropriations, the outcome and the cyclical controls; shaded band is the Lazarus HAR 95 percent interval. Percent per 1 percent shock.

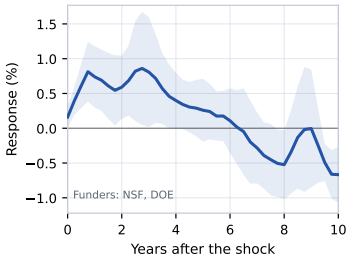
Figure 15: The Response of Research Funding to Narrative Shocks, All Fields



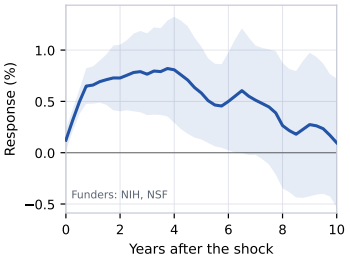
(a) Engineering



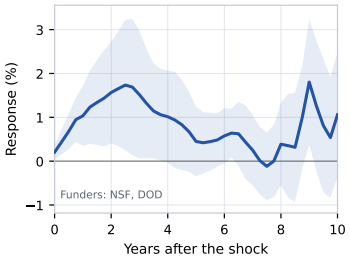
(b) Biology



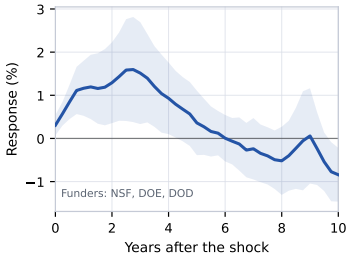
(c) Physics and astronomy



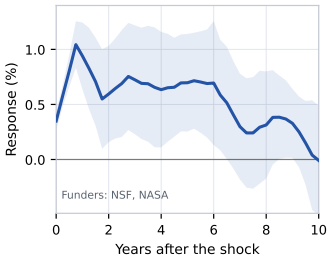
(d) Chemistry



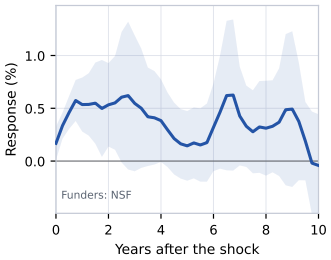
(e) Computer science



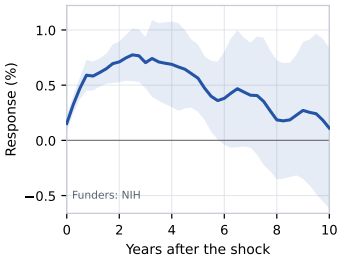
(f) Mathematics and statistics



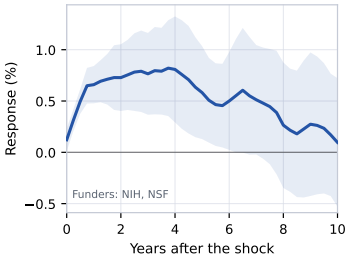
(g) Geosciences



(h) Agriculture



(i) Health sciences



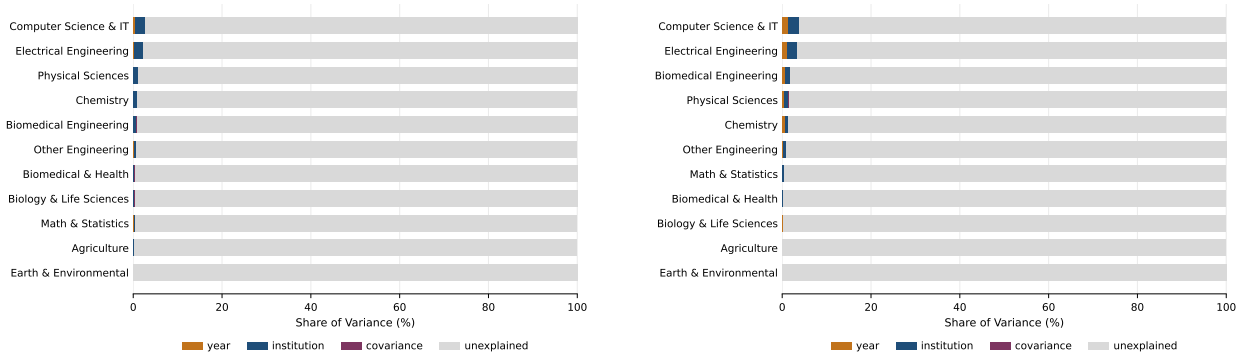
(j) Psychology

Notes: Response of the field's university-directed appropriations as a share of potential GDP to a Fieldhouse-Mertens appropriation shock of 1 percent, horizons 0–10 years. Quarterly local projections with 12 lags of log appropriations, the outcome and the cyclical controls; shaded band is the Lazarus HAR 95 percent interval. Percent per 1 percent shock.

E.4 Robustness of the Variance Decomposition

Figure 16 repeats the variance decomposition of Figure 4 for four alternative outcomes. Tables 12, 13 and 14 repeat it on patent counts over alternative horizons and for individuals with unambiguous names, under alternative winsorizing caps and Poisson regression, and with finer field definitions, and Table 15 repeats each check for each outcome. For patents, quality-adjusted patents, citations as well as log patents, doctoral institution accounts for less than 5% of the variance regardless of field or specification. The largest shares are for the extensive margin in computer science, which rise to at most 8%.

Figure 16: Variance Decomposition by Field: Alternative Outcome Measures



(a) Quality-adjusted patents

(b) Forward citations

(c) Ever patent

(d) Log patents, patenters only

Notes: Each panel repeats the decomposition in Figure 4 on a different outcome: (a) quality-adjusted patents, (b) five-year external forward citations, (c) extensive margin patenting, and (d) log patents among individuals with at least one patent. All specifications use use the leave-out finite sample correction of Kline et al. (2020).

Table 12: Variance Decomposition by Field: Horizon and Name Ambiguity

		Cohorts 1990–2005			
Field	Baseline	5 years	10 years	15 years	Unambiguous
Computer Science & IT	2.79	2.76	3.12	3.46	3.26
Electrical Engineering	2.31	2.53	2.74	2.89	2.62
Biomedical Engineering	2.29	2.28	2.39	2.06	2.74
Physical Sciences	1.64	1.88	1.78	1.74	2.08
Chemistry	1.55	1.48	1.73	1.77	2.02
Other Engineering	0.86	0.73	0.97	1.08	1.10
Biomedical & Health	0.56	0.26	0.31	0.29	0.69
Agriculture	0.50	0.24	0.44	0.80	0.67
Biology & Life Sciences	0.47	0.31	0.42	0.56	0.66
Math & Statistics	0.27	0.22	0.33	0.56	0.38
Earth & Environmental	0.07	0.04	0.13	0.22	0.21
Mean across fields	1.21	1.16	1.31	1.40	1.49

Notes: Each cell is the share of the variance of patents filed within the window that is attributable to the doctoral institution, in percent, from a decomposition run separately within each field with institution and degree-year effects and the leave-out correction of Kline et al. (2020). Baseline is Figure 4: ten years, counts winsorized at 100. Energy is omitted, as in the figure. Columns two to four count patents filed within 5, 10 and 15 years of the degree for the same individuals, cohorts 1990 to 2005, so that every window is complete. The last column keeps individuals whose surname and first initial are shared by at most ten PatentsView inventors.

Table 13: Variance Decomposition by Field: Winsorizing and Poisson Regression

		Winsorizing cap				
Field	Baseline	None	200	50	20	Poisson
Computer Science & IT	2.79	2.47	2.50	3.37	4.62	2.91
Electrical Engineering	2.31	1.22	1.85	2.94	3.86	1.48
Biomedical Engineering	2.29	2.20	2.20	2.37	2.33	2.70
Physical Sciences	1.64	1.41	1.52	1.81	2.11	1.77
Chemistry	1.55	1.42	1.45	1.73	2.06	1.70
Other Engineering	0.86	0.77	0.78	1.00	1.31	0.89
Biomedical & Health	0.56	0.43	0.43	0.65	0.73	0.66
Agriculture	0.50	0.50	0.50	0.54	0.58	0.69
Biology & Life Sciences	0.47	0.24	0.40	0.54	0.70	0.46
Math & Statistics	0.27	0.27	0.27	0.28	0.35	0.89
Earth & Environmental	0.07	0.05	0.05	0.12	0.20	0.37
Mean across fields	1.21	1.00	1.09	1.40	1.71	1.32

Notes: Each cell is the share of the variance of patents filed within the window that is attributable to the doctoral institution, in percent, from a decomposition run separately within each field with institution and degree-year effects and the leave-out correction of Kline et al. (2020). Baseline is Figure 4: ten years, counts winsorized at 100. Energy is omitted, as in the figure. Columns two to five winsorize patent counts at the stated cap. The Poisson column estimates institution and degree-year effects by Poisson pseudo-maximum likelihood on uncapped counts; the institution share of the variance of the fitted log index is scaled by the share of total variance explained by the fitted mean and corrected for estimation error using the correlation of institution effects estimated on odd and even cohorts.

Table 14: Variance Decomposition by Field: Finer Field Definitions

		Discipline × year		
Field	Baseline	≥ 1,000	≥ 250	By discipline
Computer Science & IT	2.79	2.71	2.66	2.78
Electrical Engineering	2.31	2.31	2.31	2.31
Biomedical Engineering	2.29	2.25	2.25	2.39
Physical Sciences	1.64	1.20	1.21	1.96
Chemistry	1.55	1.33	1.34	1.96
Other Engineering	0.86	0.66	0.67	1.37
Biomedical & Health	0.56	0.44	0.44	0.58
Agriculture	0.50	0.43	0.42	0.90
Biology & Life Sciences	0.47	0.36	0.37	0.46
Math & Statistics	0.27	0.30	0.30	0.43
Earth & Environmental	0.07	0.04	0.03	0.09
Mean across fields	1.21	1.09	1.09	1.39

Notes: Each cell is the share of the variance of patents filed within the window that is attributable to the doctoral institution, in percent, from a decomposition run separately within each field with institution and degree-year effects and the leave-out correction of Kline et al. (2020). Baseline is Figure 4: ten years, counts winsorized at 100. Energy is omitted, as in the figure. Disciplines are ProQuest subject classifications (the leading class term of each dissertation). Columns two and three replace degree-year effects with discipline-by-degree-year effects within each field, pooling disciplines with fewer than the stated number of PhDs into a residual category. The last column runs the decomposition separately within each discipline with at least 1,500 PhDs, so that institution effects are specific to an institution and discipline, and averages the institution variance across disciplines weighted by PhDs.

Table 15: Variance Decomposition by Field: Robustness Across Outcome Measures

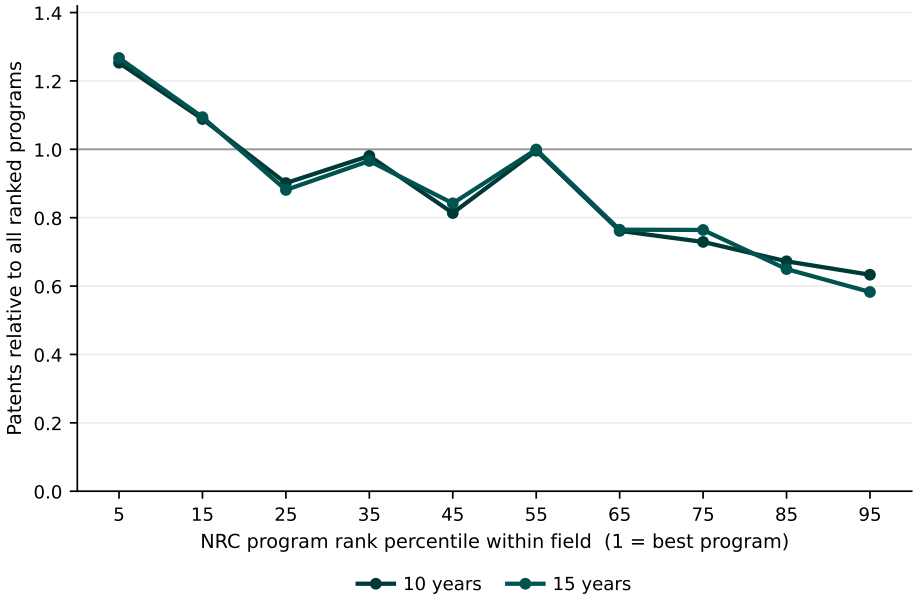
	Patents	Quality-adj.	Citations	Ever patent	Log patents
Baseline	1.21 (2.79)	0.80 (2.25)	0.84 (2.66)	2.10 (6.36)	1.38 (2.85)
Horizon, cohorts 1990–2005					
5 years	1.16 (2.76)	0.81 (2.28)	0.77 (2.35)	1.82 (5.66)	0.75 (2.78)
10 years	1.31 (3.12)	1.00 (2.96)	0.99 (3.13)	2.18 (6.60)	1.43 (3.10)
15 years	1.40 (3.46)	1.08 (3.27)	1.03 (3.49)	2.37 (7.04)	1.47 (3.78)
Sample					
Unambiguous names	1.49 (3.26)	0.97 (2.64)	0.97 (2.76)	2.65 (7.69)	1.91 (2.72)
Patent cap (baseline 100)					
None	1.00 (2.47)	0.54 (1.63)	0.44 (1.53)	–	1.39 (2.84)
200	1.09 (2.50)	0.65 (1.85)	0.68 (2.26)	–	1.39 (2.84)
50	1.40 (3.37)	1.03 (2.90)	1.03 (3.28)	–	1.36 (2.84)
20	1.71 (4.62)	1.44 (4.36)	1.34 (4.43)	–	1.27 (2.72)
Estimator					
Poisson	1.32 (2.91)	0.97 (2.13)	0.85 (2.01)	–	–
Finer field definitions					
Discipline × year, ≥ 1,000	1.09 (2.71)	0.75 (2.21)	0.81 (2.61)	1.70 (6.10)	1.43 (2.85)
Discipline × year, ≥ 250	1.09 (2.66)	0.75 (2.14)	0.81 (2.59)	1.69 (5.98)	1.39 (2.85)
By discipline	1.39 (2.78)	0.91 (2.33)	0.96 (2.68)	2.14 (6.22)	1.20 (3.45)

Notes: Each cell is the institution share of the variance of the outcome, in percent, averaged across the fields of Figure 4, with the largest field in parentheses. Outcomes are those of Figure 16: patents, quality-adjusted patents, five-year external forward citations, an indicator for ever patenting, and log patents among individuals with at least one patent; the last conditions on patenting and is descriptive. Each specification is otherwise as in Tables 12 to 14. Horizon rows measure every outcome over the stated window. Patent caps winsorize patent counts; quality-adjusted patents are capped at the same percentile and citations at ten times the cap (baseline 1,000). An indicator for ever patenting does not change with the cap, and Poisson regression applies only to the three count measures; those cells are marked with a dash. Field-level estimates of a component near zero can be negative; they enter the means as estimated.

E.5 Patenting by Program Rank

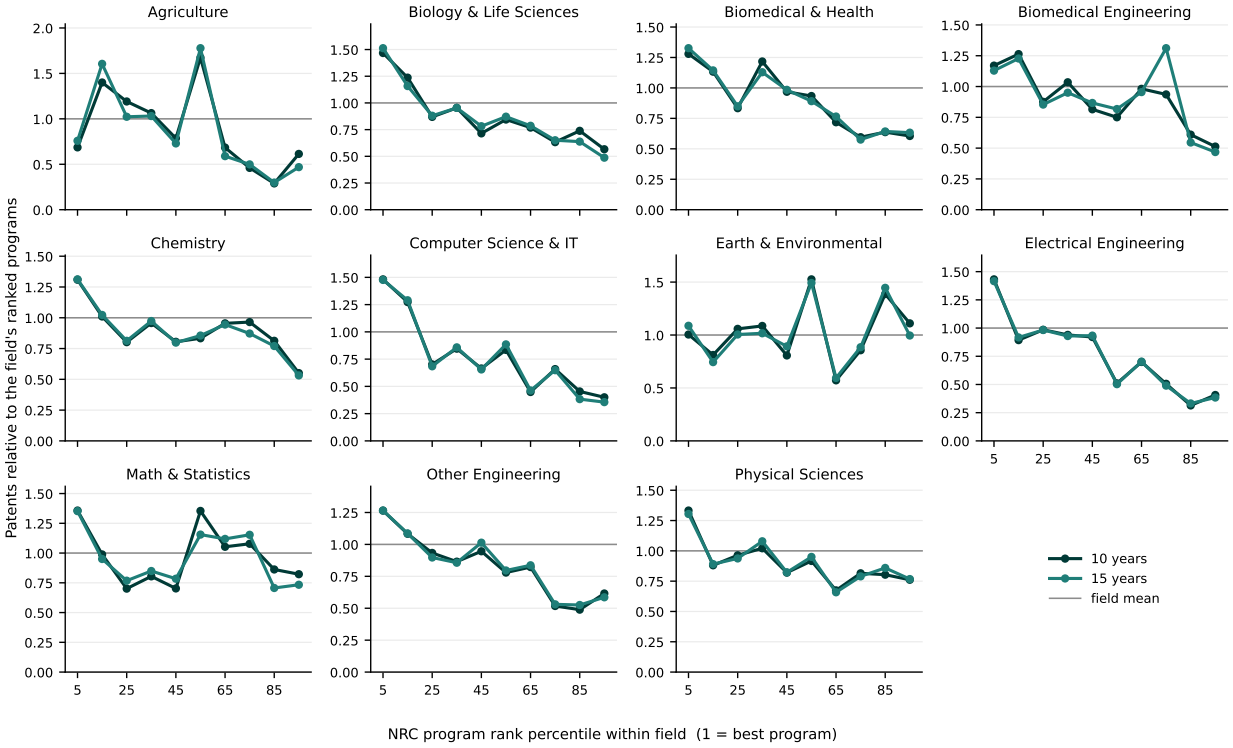
Figures 17, 19 and 21 plot three margins of patenting against the National Research Council rank of the doctoral program: mean patents, the share of graduates who ever patent, and, conditional on patenting, what percentage continue to patent. I construct a synthetic field via equal weighted averaging across fields.

Figure 17: Mean Patents by Program Rank Decile



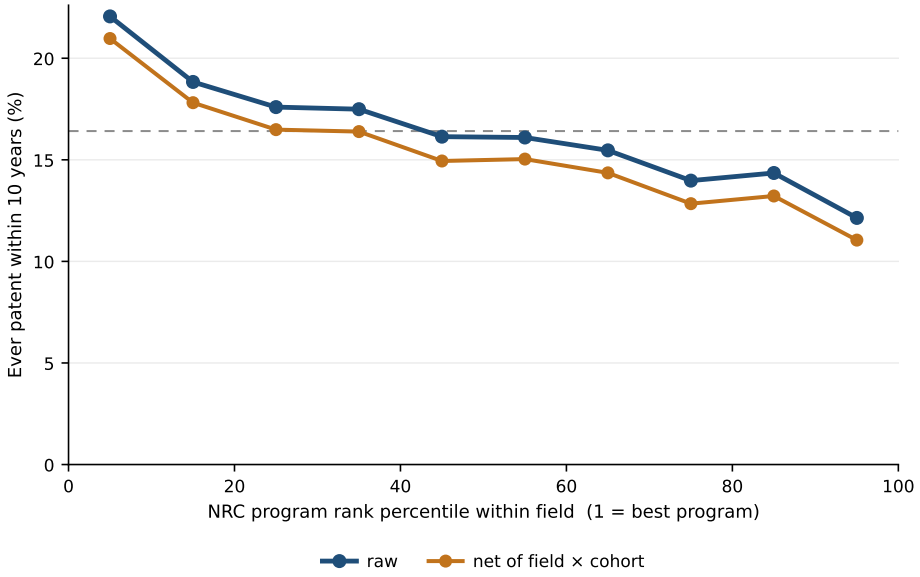
Notes: Programs are ranked within field by the NRC ranking and grouped into deciles of programs. Each point is the mean number of patents filed within either 10 or 15 years of the degree by graduates of that decile's programs, relative to the mean across all ranked programs in that field. I then average across fields with equal weights. Cohorts run from 1990 to 2005 to ensure comparability across cohorts.

Figure 18: Mean Patents by Program Rank Decile, by Field



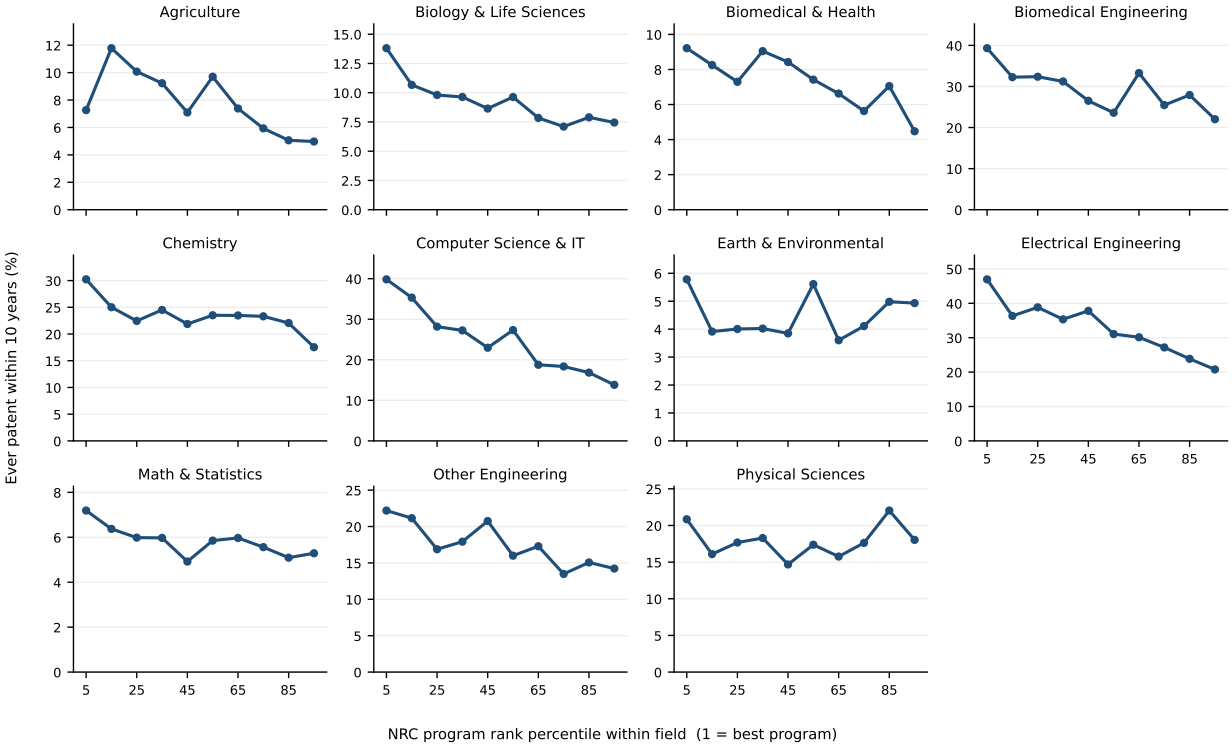
Notes: Each panel is one field of Figure 17: mean patents filed within either 10 or 15 years of the degree by graduates of each decile of the field's NRC-ranked programs, relative to the mean across all of the field's ranked programs (the horizontal line). Deciles with fewer than 40 PhDs are omitted. Cohorts run from 1990 to 2005. Vertical axes differ across panels due to differences in patenting propensity across fields.

Figure 19: Share Ever Patenting by Program Rank Decile



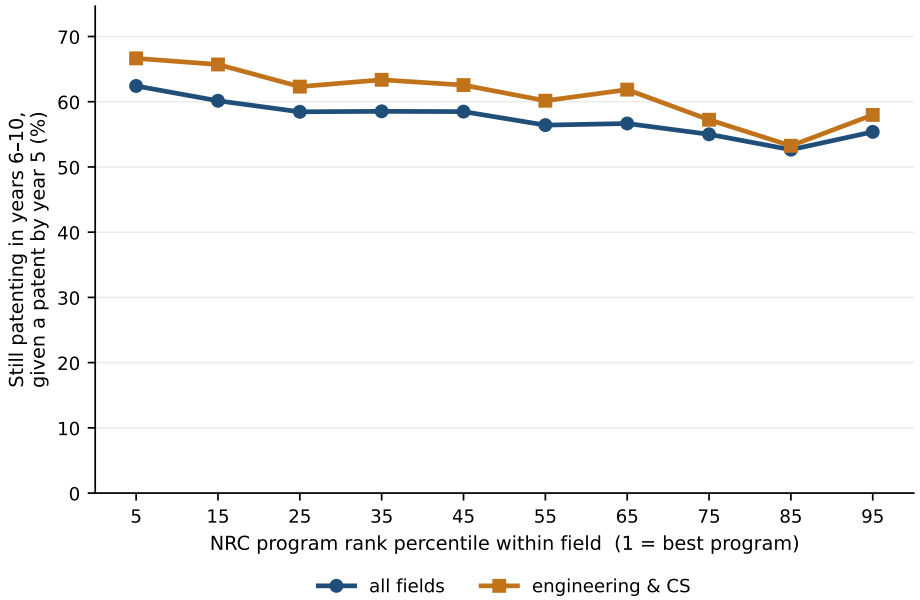
Notes: Share of PhDs with at least one patent filed within ten years of the degree, by decile of NRC-ranked programs within field. The series is generated by averaging across fields with equal weights. The net series removes field-by-cohort means before averagingm, while the dashed line is the mean across deciles of the raw series. Cohorts run from 1990 to 2010.

Figure 20: Share Ever Patenting by Program Rank Decile, by Field



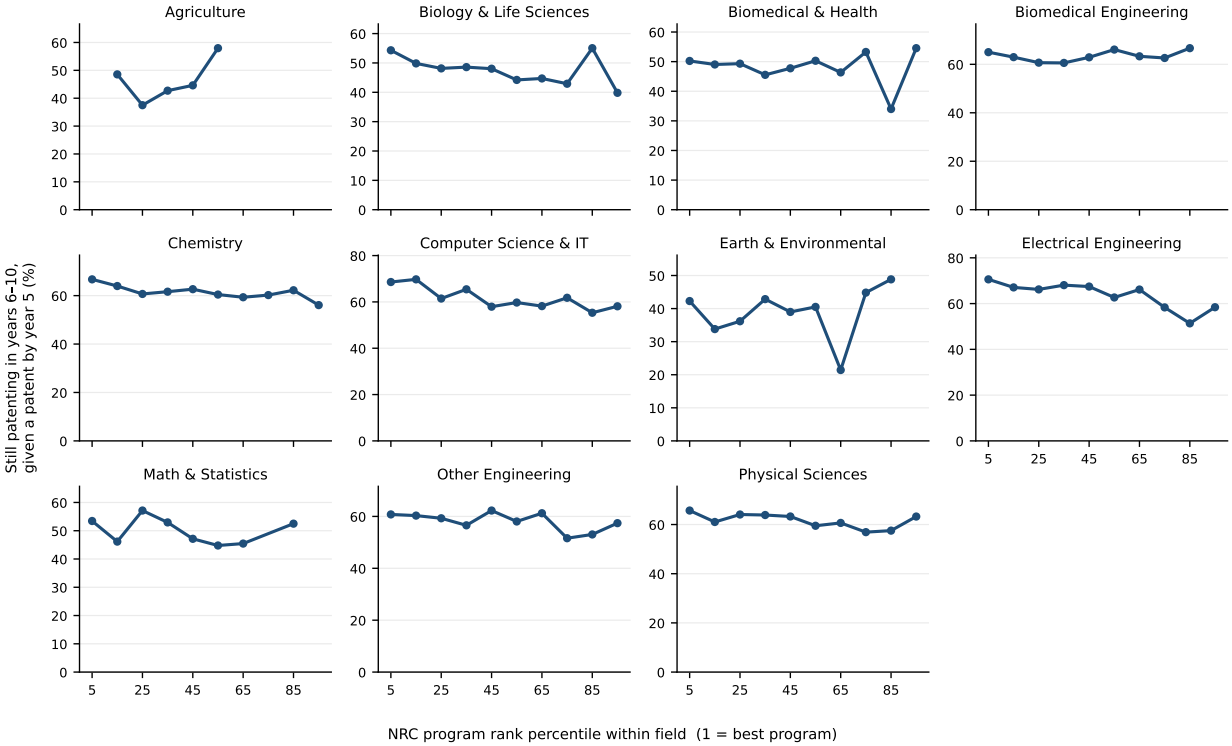
Notes: Each panel contains one field of Figure 19: the share of PhDs with at least one patent filed within ten years of the degree, again sorted by decile of the given field's NRC-ranked programs. Cohorts run from 1990 to 2010. Vertical axes differ across panels due to differences in patenting propensity across fields.

Figure 21: Continued Patenting by Program Rank Decile



Notes: This panel contains share of PhDs from a given decile, who, conditional on patenting within five years of recieviing their degree, were also active in the following five years. I then average with equal weights across the seven fields with at least 40 PhDs inventors in every decile. These fields are biology, biomedical and health sciences, chemistry, computer science, electrical engineering, other engineering and physical sciences. The engineering and computer science series averages computer science, electrical engineering and other engineering. Cohorts run from 1990 to 2010.

Figure 22: Continued Patenting by Program Rank Decile, by Field

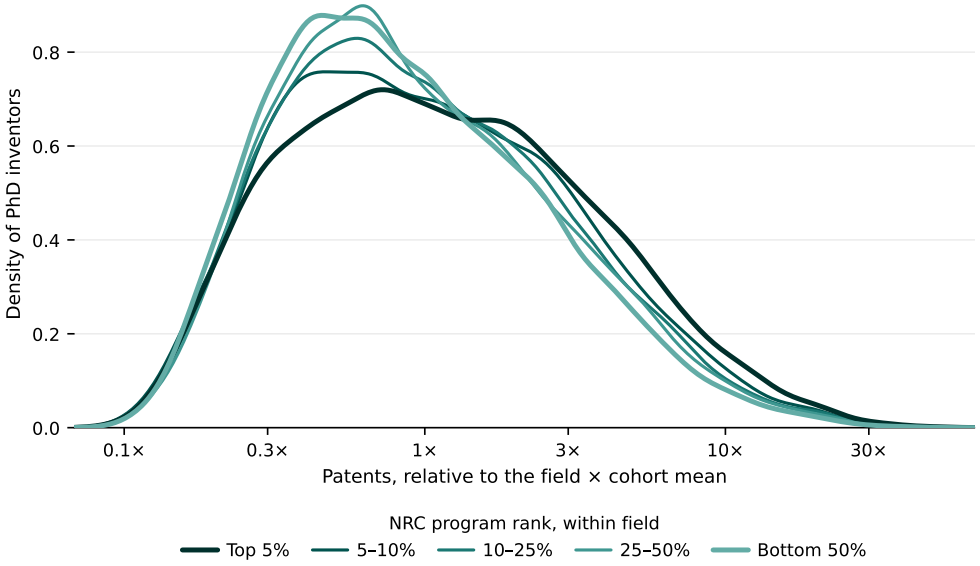


Notes: Each panel contains the share of PhDs with a patent filed within five years of the degree, who also have a patent filed six to ten years after it, by decile of that field's NRC-ranked programs. Deciles with fewer than 40 such PhDs are omitted. Cohorts run from 1990 to 2010. Vertical axes differ across panels due to differences in patenting propensity across fields.

E.6 Robustness for Conditional Output Densities

Figure 23 repeats Figure 5 for raw patent counts, Figures 24 and 25 show both outcomes separately by field, and Figure 26 compares within-program dispersion across program ranks. Table 16 reports the probability that a graduate of a bottom-50% program out-produces one from a top-5% program under alternative program rankings, including rankings estimated from patent output itself, for individuals with unambiguous names, over alternative horizons and winsorizing caps, and with output measured relative to finer discipline-by-cohort cells. Tables 17 and 18 test that probability against one half, for quality-adjusted patents and for patents, pooled and by field, using both standard Mann–Whitney and clustered Mann–Whitney tests.[35] Table 19 repeats both tests on all PhDs, with those who never patent at zero, so that the comparison combines the participation margin with the inventor margin. Figure 27 and Table 20 conduct quantile comparisons.

Figure 23: Distribution of Raw Patent Output by Program Tier



[35]The clustered Mann–Whitney test is Van Elteren's test.

Figure 24: Distribution of Quality-Adjusted Output by Program Tier and Field

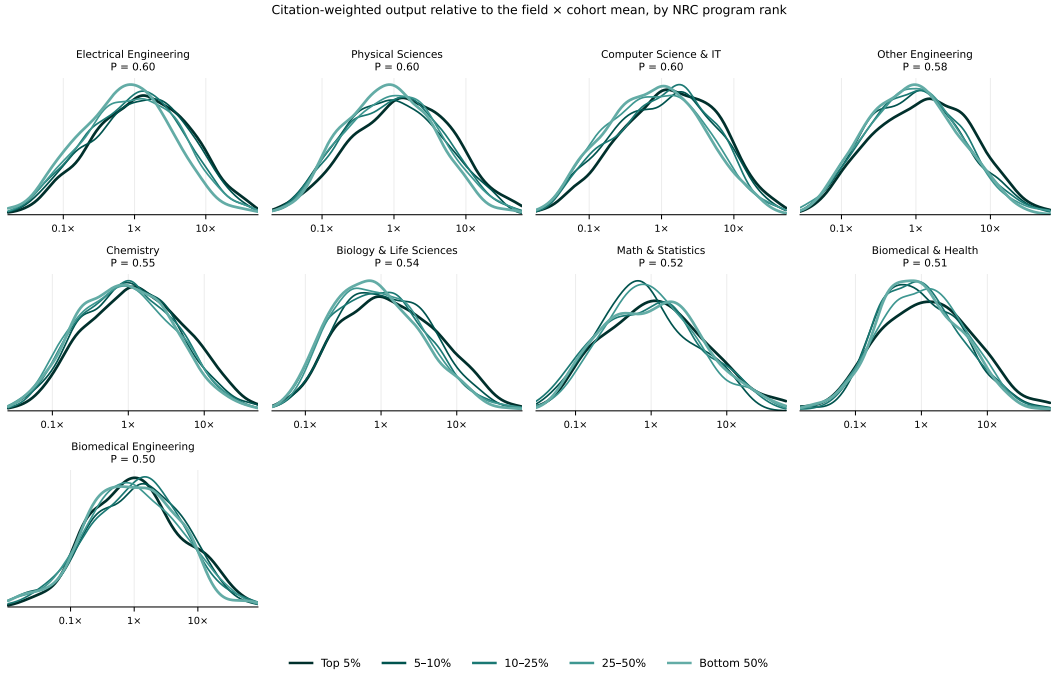


Figure 25: Distribution of Raw Patent Output by Program Tier and Field

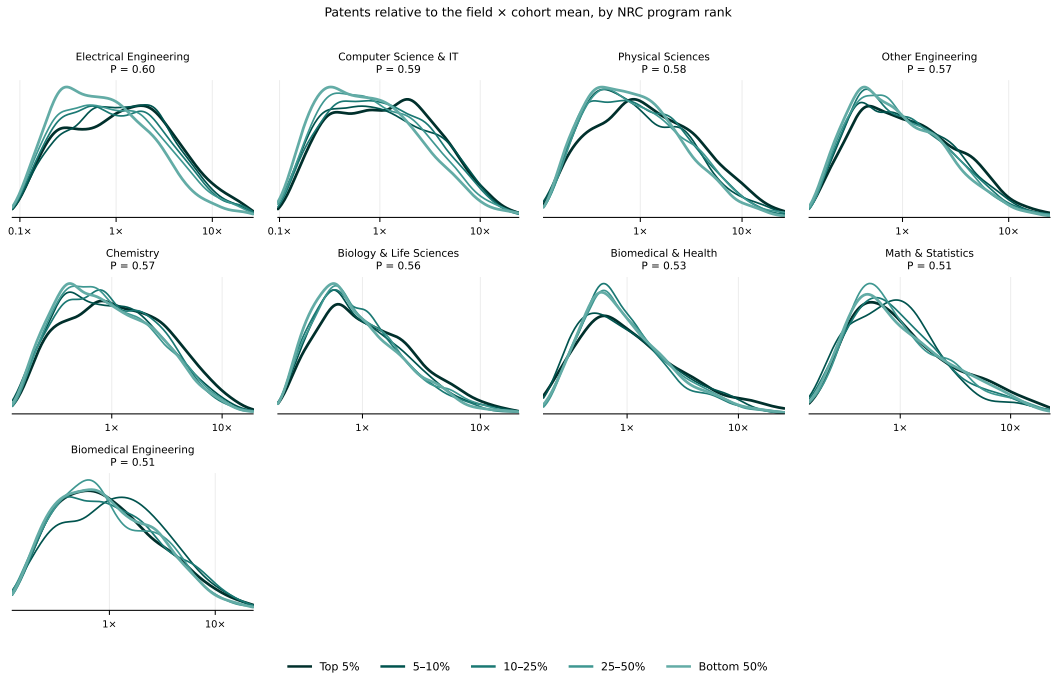


Figure 26: Within-Program Dispersion is not Concentrated at Top Programs

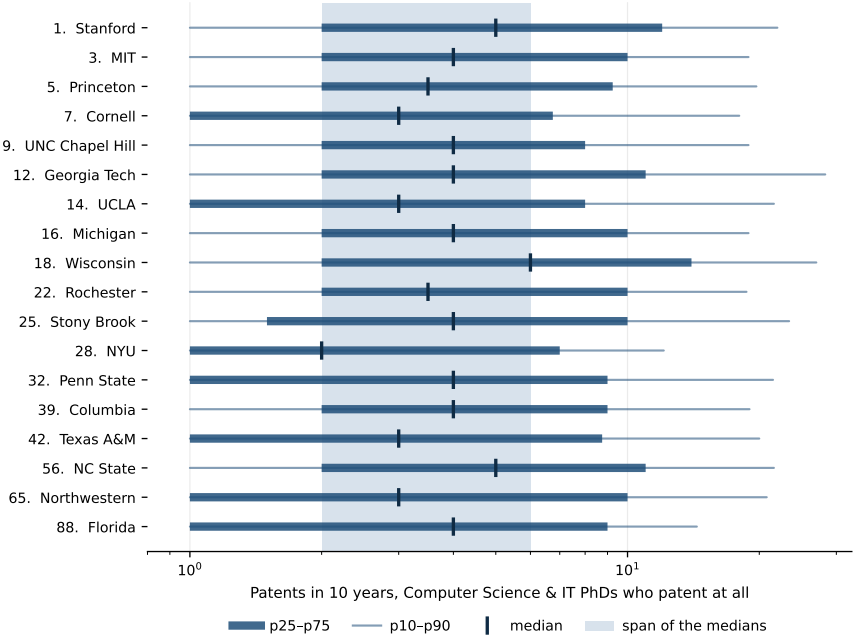
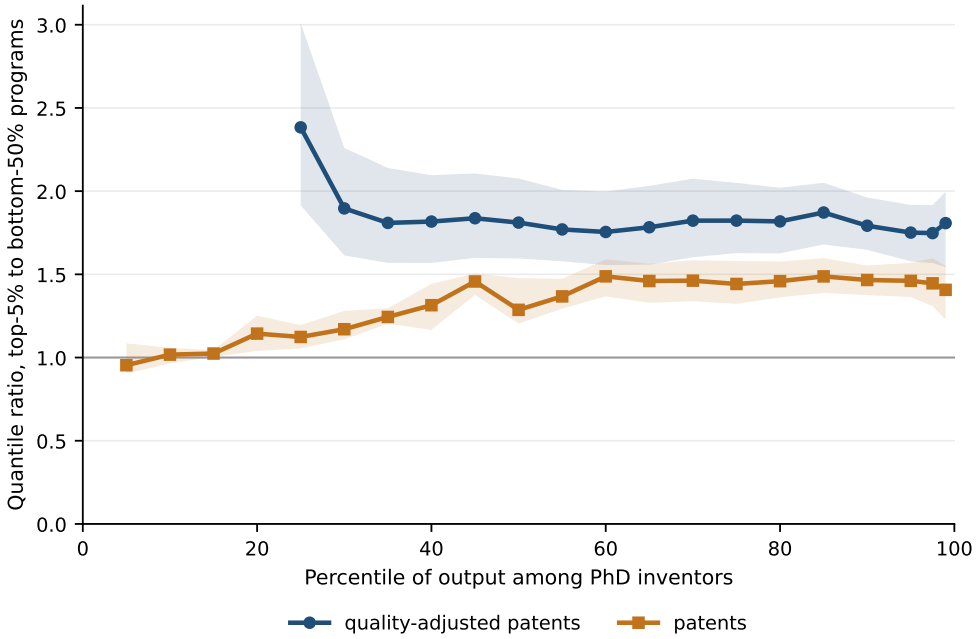


Figure 27: Quantile Ratios of Output: Top-5% Against Bottom-50% Programs



Notes: At each percentile of output among PhD inventors, the top-5% group's quantile divided by the bottom-50% group's. Programs are ranked within field using NRC rankings. Output is measured relative to the geometric mean of positive output in the same field and cohort. Every inventor has (by definition) at least one patent. Shaded areas denote 95% confidence intervals constructed via resampling at program (institution by field) level within field and group.

Table 16: Overlap of Output Distributions: Robustness

	P(bottom beats top)			Most separated	
	Quality-adj.	Patents	Overlap	field	Inventors
Baseline	0.420	0.435	0.883	Elec. eng. (0.396)	67,611
Program ranking					
NRC S ranking	0.425	0.436	0.886	Elec. eng. (0.388)	67,611
Patent FE, leave-out	0.404	0.410	0.856	Elec. eng. (0.363)	36,666
Patent FE, in-sample	0.366	0.371	0.803	CS & IT (0.343)	71,428
Sample					
Unambiguous names	0.415	0.432	0.879	Elec. eng. (0.397)	43,173
Horizon, cohorts 1990–2005					
5 years	0.418	0.445	0.887	Physical sci. (0.403)	34,189
10 years	0.419	0.440	0.885	Elec. eng. (0.390)	47,633
15 years	0.416	0.438	0.881	Elec. eng. (0.377)	55,123
Winsorizing cap (baseline 100)					
None	0.420	0.435	0.883	Elec. eng. (0.396)	67,611
200	0.420	0.435	0.883	Elec. eng. (0.396)	67,611
50	0.420	0.434	0.883	Elec. eng. (0.396)	67,611
20	0.420	0.434	0.882	Elec. eng. (0.396)	67,611
Discipline × cohort cells					
$\geq$ 1,000 PhDs	0.421	0.434	0.884	Physical sci. (0.395)	67,611
$\geq$ 250 PhDs	0.422	0.433	0.884	Physical sci. (0.395)	67,611

Notes: The first two columns give the probability that a randomly drawn PhD inventor from a program in the bottom half of its field's ranking produces more than one drawn from the top 5%. Ties are counted as half and outcomes are quality-adjusted patents and raw patent counts. I normalized output by measuring relative to the geometric mean of inventors in the same field and cohort among PhDs with at least one patent filed within ten years of degree. Overlap is the overlap coefficient of the two quality-adjusted densities. The most separated field is the field with the lowest probability, with that probability in parentheses. The NRC S ranking replaces the regression-based R ranking with the survey-based one. Patent FE rankings order programs by their shrunk Poisson fixed effect on patents. Leave-out version is estimated on odd cohorts and evaluated on even ones. The in-sample version estimates on the outcome of interest (and hence bounds the performance of any ranking). Unambiguous names keeps individuals whose surname and first initial are shared by at most ten PatentsView inventors. Horizon rows count patents filed within 5, 10 and 15 years of the degree for the same cohorts. Discipline rows measure output relative to discipline-by-degree-year cells, pooling disciplines below the mentioned size.

Table 17: Bottom-50% Against Top-5% Programs: Mann–Whitney and van Elteren Tests, Quality-Adjusted Patents

	Inventors		Mann–Whitney			van Elteren		
Field	Bottom 50%	Top 5%	P(beats)	95% CI	p	P(beats)	95% CI	p
All fields	13,680	11,430	0.420	[0.402, 0.437]	<0.001	0.426	[0.413, 0.440]	<0.001
Biology	1,752	1,128	0.458	[0.426, 0.484]	0.005	0.454	[0.421, 0.484]	0.004
Biomed. & health	792	365	0.486	[0.417, 0.544]	0.661	0.467	[0.392, 0.534]	0.367
Biomed. eng.	592	323	0.497	[0.443, 0.561]	0.918	0.500	[0.445, 0.566]	0.991
Chemistry	2,924	2,326	0.447	[0.424, 0.470]	<0.001	0.435	[0.412, 0.460]	<0.001
CS & IT	1,127	1,520	0.401	[0.369, 0.431]	<0.001	0.404	[0.374, 0.431]	<0.001
Elec. eng.	1,657	2,201	0.396	[0.345, 0.448]	<0.001	0.395	[0.346, 0.447]	<0.001
Math & stat.	328	217	0.477	[0.429, 0.534]	0.386	0.477	[0.427, 0.528]	0.364
Other eng.	1,979	1,930	0.417	[0.395, 0.454]	<0.001	0.422	[0.401, 0.459]	<0.001
Physical sci.	1,811	1,338	0.401	[0.369, 0.431]	<0.001	0.400	[0.367, 0.431]	<0.001

Notes: P(beats) is the probability that a randomly drawn PhD inventor from a program in the bottom half of its field's NRC ranking has more quality-adjusted patents than one drawn from a program in the top 5%, with ties (3% of sample) counted as one half. Output is measured relative to the geometric mean of positive output in the same field and degree year. The Mann–Whitney columns compare every inventor in the bottom group with every inventor in the top group. The van Elteren columns compare inventors only within the same field and degree year. Confidence intervals and p-values, which test P(beats) = 0.5, resample whole programs within field and group 9,999 times. Agriculture and Earth & Environmental are omitted because they have fewer than 100 inventors.

Table 18: Bottom-50% Against Top-5% Programs: Mann–Whitney and van Elteren Tests, Patents

	Inventors		Mann–Whitney			van Elteren		
Field	Bottom 50%	Top 5%	P(beats)	95% CI	p	P(beats)	95% CI	p
All fields	13,680	11,430	0.435	[0.421, 0.449]	<0.001	0.432	[0.419, 0.444]	<0.001
Biology	1,752	1,128	0.436	[0.407, 0.467]	<0.001	0.437	[0.409, 0.469]	<0.001
Biomed. & health	792	365	0.467	[0.408, 0.518]	0.247	0.461	[0.406, 0.511]	0.149
Biomed. eng.	592	323	0.495	[0.441, 0.560]	0.875	0.502	[0.448, 0.569]	0.958
Chemistry	2,924	2,326	0.433	[0.406, 0.459]	<0.001	0.434	[0.408, 0.460]	<0.001
CS & IT	1,127	1,520	0.413	[0.384, 0.437]	<0.001	0.417	[0.390, 0.439]	<0.001
Elec. eng.	1,657	2,201	0.401	[0.354, 0.449]	<0.001	0.399	[0.355, 0.445]	<0.001
Math & stat.	328	217	0.486	[0.446, 0.530]	0.516	0.479	[0.432, 0.524]	0.353
Other eng.	1,979	1,930	0.432	[0.409, 0.463]	<0.001	0.437	[0.415, 0.468]	<0.001
Physical sci.	1,811	1,338	0.423	[0.387, 0.457]	<0.001	0.425	[0.388, 0.458]	<0.001

Notes: P(beats) is the probability that a randomly drawn PhD inventor from a program in the bottom half of its field's NRC ranking has more patents within the 10 years post-dissertation than one drawn from a program in the top 5%, counting ties (under 0.1% of sample) as one half. Output is measured relative to the geometric mean of positive output in the same field and degree year. The Mann–Whitney columns compare every inventor in the bottom group with every inventor in the top group. The van Elteren columns compare inventors only within the same field and degree year. Confidence intervals and p-values, which test P(beats) = 0.5, resample whole programs within field and group 9,999 times. Agriculture and Earth & Environmental are omitted because they have fewer than 100 inventors.

Table 19: Bottom-50% Against Top-5% Programs: All PhDs, Including Those Without Patents

	PhDs		Mann–Whitney			van Elteren		
Field	Bottom 50%	Top 5%	P(beats)	95% CI	p	P(beats)	95% CI	p
Panel A: Quality-adjusted patents								
All fields	102,675	44,195	0.441	[0.429, 0.453]	<0.001	0.455	[0.447, 0.462]	<0.001
Agriculture	4,897	544	0.500	[0.494, 0.529]	0.974	0.500	[0.493, 0.529]	0.968
Biology	21,679	7,321	0.475	[0.464, 0.482]	<0.001	0.475	[0.463, 0.482]	<0.001
Biomed. & health	12,378	3,087	0.483	[0.464, 0.493]	0.023	0.482	[0.463, 0.493]	0.022
Biomed. eng.	2,167	821	0.455	[0.438, 0.481]	<0.001	0.455	[0.439, 0.481]	<0.001
Chemistry	12,997	6,977	0.452	[0.434, 0.469]	<0.001	0.451	[0.433, 0.468]	<0.001
CS & IT	5,684	3,587	0.383	[0.359, 0.409]	<0.001	0.383	[0.358, 0.408]	<0.001
Earth & env.	8,141	1,041	0.504	[0.496, 0.510]	0.327	0.503	[0.496, 0.510]	0.329
Elec. eng.	5,942	4,669	0.393	[0.350, 0.431]	<0.001	0.392	[0.349, 0.430]	<0.001
Math & stat.	5,828	2,529	0.487	[0.476, 0.496]	0.009	0.487	[0.476, 0.495]	0.007
Other eng.	12,756	7,670	0.451	[0.438, 0.470]	<0.001	0.451	[0.438, 0.470]	<0.001
Physical sci.	10,206	5,949	0.470	[0.446, 0.496]	0.018	0.470	[0.446, 0.496]	0.019
Panel B: Patents								
All fields	102,675	44,195	0.435	[0.423, 0.446]	<0.001	0.449	[0.441, 0.457]	<0.001
Agriculture	4,897	544	0.497	[0.490, 0.521]	0.764	0.498	[0.490, 0.520]	0.794
Biology	21,679	7,321	0.463	[0.449, 0.473]	<0.001	0.462	[0.449, 0.473]	<0.001
Biomed. & health	12,378	3,087	0.473	[0.449, 0.486]	0.004	0.473	[0.449, 0.486]	0.004
Biomed. eng.	2,167	821	0.439	[0.422, 0.461]	<0.001	0.440	[0.424, 0.461]	<0.001
Chemistry	12,997	6,977	0.441	[0.419, 0.462]	<0.001	0.441	[0.419, 0.462]	<0.001
CS & IT	5,684	3,587	0.380	[0.356, 0.405]	<0.001	0.379	[0.355, 0.405]	<0.001
Earth & env.	8,141	1,041	0.504	[0.495, 0.510]	0.343	0.503	[0.496, 0.510]	0.360
Elec. eng.	5,942	4,669	0.391	[0.348, 0.429]	<0.001	0.389	[0.347, 0.427]	<0.001
Math & stat.	5,828	2,529	0.485	[0.475, 0.493]	0.001	0.485	[0.475, 0.493]	<0.001
Other eng.	12,756	7,670	0.449	[0.435, 0.470]	<0.001	0.448	[0.434, 0.469]	<0.001
Physical sci.	10,206	5,949	0.473	[0.446, 0.503]	0.062	0.473	[0.446, 0.503]	0.060

Notes: P(beats) is the probability that a randomly drawn PhD from a program in the bottom half of its field's NRC ranking has more output within 10 years of the dissertation than one drawn from a program in the top 5%, with ties counted as one half. Output is quality-adjusted patents in Panel A and patents in Panel B, both filed within ten years of the degree. Every PhD at a ranked program is included, with those who never patent at zero. In the all-fields rows 70% (Panel A) and 64% (Panel B) of pairs are ties. Output is measured relative to the geometric mean of positive output in the same field and degree year. The Mann–Whitney columns compare every PhD in the bottom group with every PhD in the top group. The van Elteren columns compare PhDs only within the same field and degree year. Confidence intervals and p-values, which test P(beats) = 0.5, resample whole programs within field and group 9,999 times.

Table 20: Quantiles of Output: Top-5% Against Bottom-50% Programs

	Quality-adjusted patents			Patents		
Percentile	Bottom 50%	Top 5%	Ratio [95% CI]	Bottom 50%	Top 5%	Ratio [95% CI]
10	0.00	0.00	–	0.30	0.30	1.02 [0.96, 1.06]
25	0.08	0.18	2.38 [1.91, 3.01]	0.41	0.46	1.12 [1.05, 1.20]
50	0.49	0.88	1.81 [1.60, 2.08]	0.79	1.02	1.29 [1.20, 1.48]
75	1.73	3.15	1.82 [1.63, 2.05]	1.79	2.58	1.44 [1.32, 1.58]
90	4.66	8.35	1.79 [1.65, 1.96]	3.68	5.39	1.47 [1.37, 1.55]
95	7.96	13.93	1.75 [1.58, 1.92]	5.59	8.17	1.46 [1.36, 1.57]
99	19.66	35.53	1.81 [1.54, 2.00]	12.36	17.38	1.41 [1.23, 1.55]
Mean	1.86	3.21	1.73 [1.58, 1.88]	1.60	2.24	1.40 [1.32, 1.49]

Notes: Quantiles of output among PhD inventors (PhDs with at least one patent filed within ten years of the degree) from programs in the bottom half and the top 5% of their field's NRC ranking, 13,680 and 11,430 inventors. Output is measured relative to the geometric mean of positive output in the same field and degree year, as in Figure 5. Ratio divides the top-5% quantile by the bottom-50% quantile. Dashes marks quantiles at which both groups have zero quality-adjusted output. The last row compares means. Intervals resample whole programs within field and group, 9,999 times.

E.7 Capture Curves and Tail Index

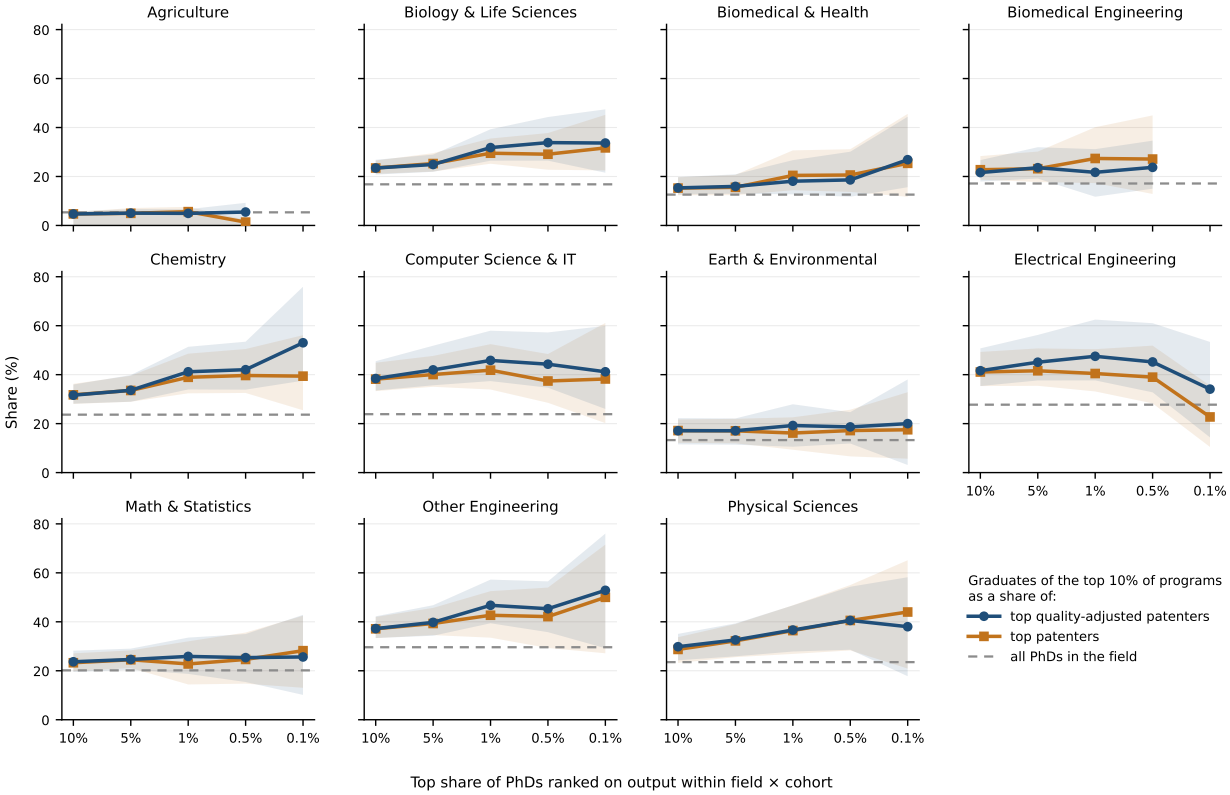
Figure 28 plots capture curves for quality-adjusted output, and Figures 30 and 31 report the tail index of output by band of program rank, normalized by field and cohort and on raw patent counts. Figures 32 and 33 report it by program decile. Figure 29 and Table 21 report the same capture shares separately for each field, with intervals for the ratio of each share to the field's share of PhDs.

Figure 28: Most of the Tail Comes from Outside the Top Programs



Notes: Each solid line is the share of top producers who trained at a program in the top 5, 10, 25 or 50% of their field's NRC ranking. Top producers are PhDs with at least one patent who rank, within a field and degree year, in the top 10, 5, 1, 0.5 or 0.1 percent of all PhDs on quality-adjusted patents within the 10 years post-dissertation. Because most PhDs never patent, these percentiles combine extensive and intensive margin effects. The dashed lines are the same programs' share of all PhDs. Fields are weighted equally, and fields with fewer than 25 producers in the top 0.1% are excluded. Cohorts run from 1990 to 2010.

Figure 29: Capture of Top Producers by Top-Ranked Programs, by Field



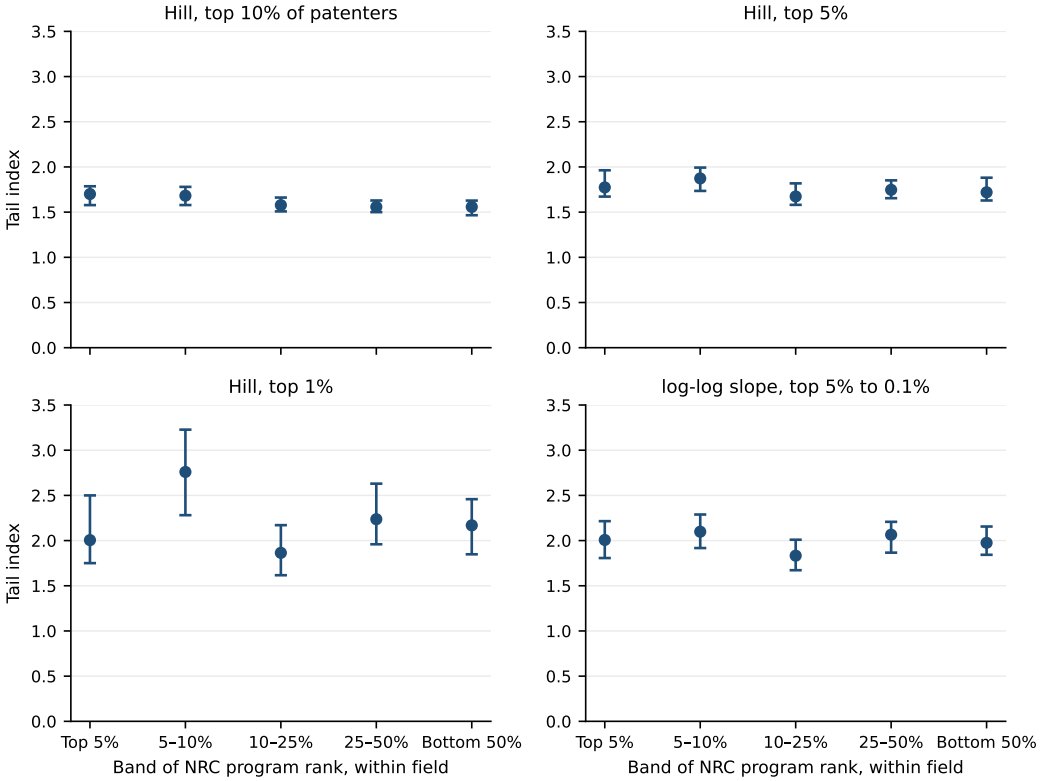
Notes: Each panel plots the share of the field's top producers who trained at a program in the top 10% of the field's NRC ranking, for the top 10, 5, 1, 0.5 and 0.1 percent of PhDs (within field and degree year) and separately for quality-adjusted and non quality-adjusted patents. Dashed line is that programs' share of all PhDs in the field. Shaded areas denote 95% confidence intervals, constructed via resampling at program level, and then scaled by the share of all PhDs. Depths with fewer than 25 patenters are not shown.

Table 21: Capture of Top Producers by Top-Ranked Programs, by Field

Field	Share of PhDs	Share of the top producers			Ratio, top 1% [95% CI]	Top 1% PhDs
		Top 10%	Top 5%	Top 1%		
Agriculture	5.4	4.7	5.0	4.9	0.92 [0.00, 1.24]	142
Biology & Life Sciences	16.8	23.5	24.9	31.8	1.89 [1.57, 2.34]	915
Biomedical & Health	12.6	15.3	15.9	18.1	1.43 [1.01, 2.12]	592
Biomedical Engineering	17.1	21.6	23.6	21.7	1.27 [0.68, 1.82]	106
Chemistry	23.7	31.6	33.6	41.2	1.74 [1.43, 2.17]	578
Computer Science & IT	23.9	38.4	42.0	45.8	1.92 [1.57, 2.43]	251
Earth & Environmental	13.3	17.1	17.1	19.2	1.45 [0.80, 2.10]	260
Electrical Engineering	27.7	41.6	45.1	47.5	1.71 [1.36, 2.25]	341
Math & Statistics	20.2	23.7	24.6	25.9	1.28 [0.93, 1.66]	259
Other Engineering	29.6	37.2	39.8	46.7	1.58 [1.33, 1.93]	614
Physical Sciences	23.5	29.8	32.6	36.6	1.56 [1.18, 1.98]	423
Mean across fields	19.4	25.9	27.7	30.9	1.59	

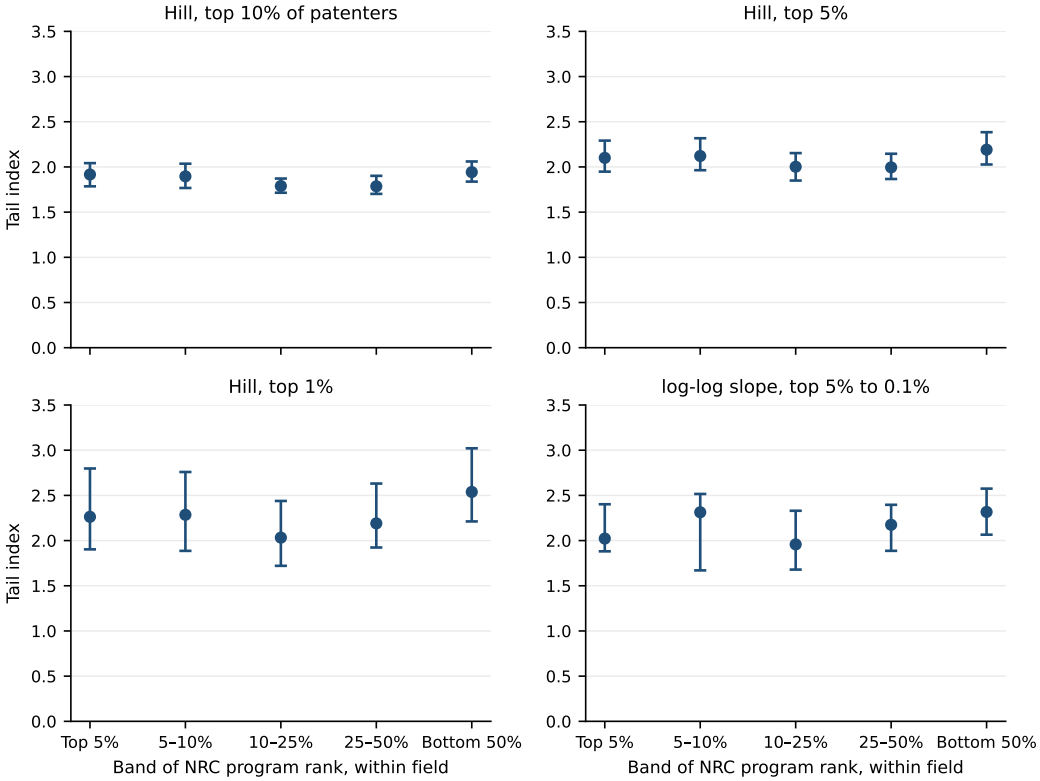
Notes: Top producers are PhDs with at least one patent who rank, within a field and degree year, in the top 10, 5 and 1 percent of all PhDs on quality-adjusted patents within the 10 years post-dissertation. Share of PhDs is the percentage of all PhDs in the field who trained at a program in the top 10% of the field's NRC ranking. The next three columns give the same percentage among each group of top producers. Ratio divides the top-1% share by the share of PhDs, so a ratio of one means programs in that bucket capture exactly their overall share in the tail. Confidence intervals resample whole institutions within field 9,999 times. The last column gives the number of top-1% producers. The mean row weights fields equally.

Figure 30: Tail Index by Program Rank Band, Normalized by Field and Cohort



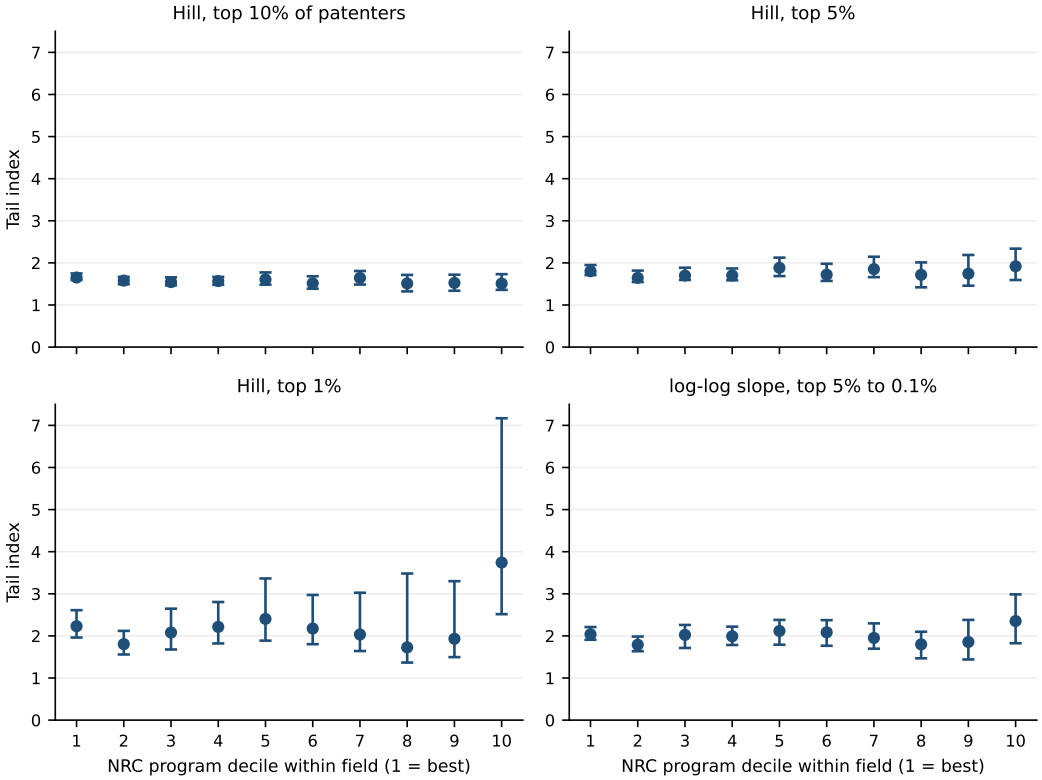
Notes: Tail index of patents filed within ten years of the degree among PhDs with at least one patent, by band of each field's NRC ranking. This is pooled across fields. Cohorts run from 1990 to 2010. Each PhD's patents are divided by the mean of all PhDs in the same field and degree year, so that every band is measured on a common scale. The panels give Hill estimates at the top 10, 5 and 1 percent of inventors and the slope of the log survival function between the top 5 and the top 0.1 percent. Bars are 95% confidence intervals from resampling whole programs (institution by field) within field.

Figure 31: Tail Index by Program Rank Band, Raw Patent Counts



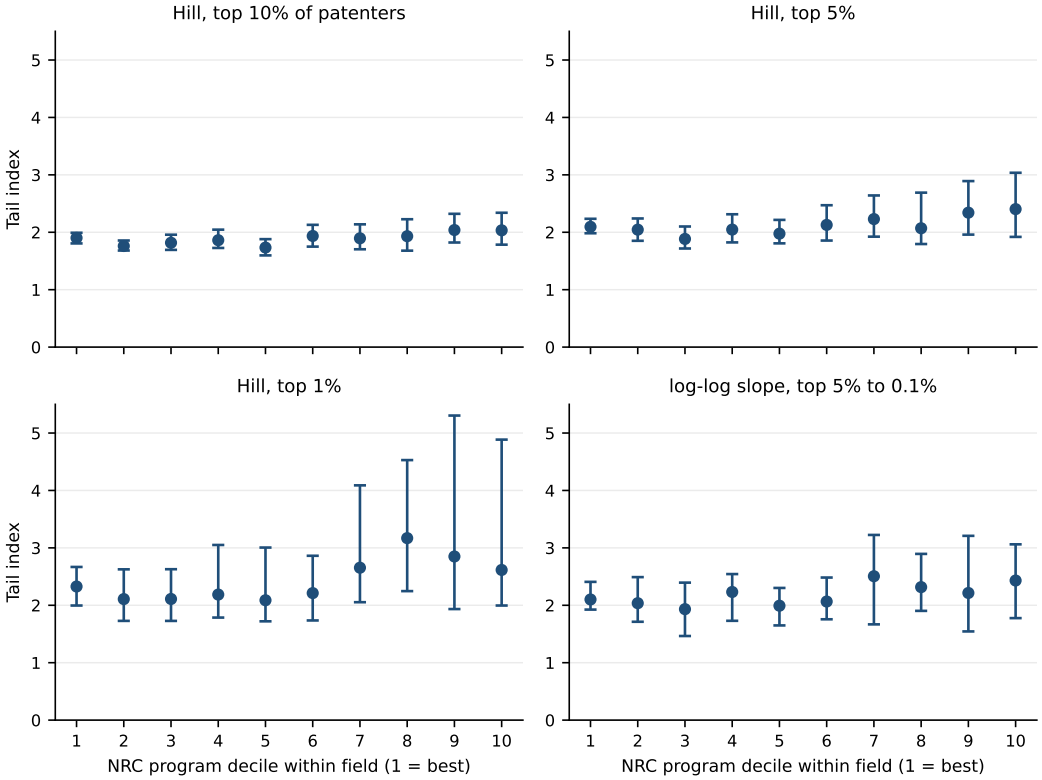
Notes: As Figure 30, on raw patent counts. A common tail index across bands is not rejected at the 5 percent level for any of the four estimators.

Figure 32: Tail Index by Program Decile, Normalized by Field and Cohort



Notes: Tail index of patents filed within ten years of the degree among PhDs patenters. This is constructed separately across within-field deciles and then pooled across fields. Cohorts run from 1990 to 2010. Each PhD's patents are divided by the mean of all PhDs in the same field and degree year. Bars are 95% confidence intervals from resampling graduate programs within field.

Figure 33: Tail Index by Program Decile, Raw Patent Counts



Notes: Same as in Figure 32, but instead on raw patent counts.

E.8 Quality Adjustment with Ten-Year Citations

Quality-adjusted output weighs each patent by external forward citations within five years of disclosure, but normalized relative to patents granted in the same year in the same CPC subclass to handle truncation and changes in the meaning of citations over time . Table 22 repeats every quality-adjusted result using 10-year citations instead.

Table 22: Quality Adjustment with Ten-Year Citations

	Citation window	
	5 years (baseline)	10 years
Rank correlation with 5-year measure	–	0.95
Inventors with no cited output (%)	19.7	14.1
P(bottom 50% beats top 5%)		
Mann–Whitney	0.420	0.420
	[0.402, 0.437]	[0.404, 0.436]
van Elteren	0.426	0.424
	[0.413, 0.440]	[0.410, 0.438]
Quantile ratio, top 5% to bottom 50%		
Median	1.81	1.77
	[1.60, 2.08]	[1.58, 2.00]
90th percentile	1.79	1.77
	[1.65, 1.96]	[1.62, 1.93]
99th percentile	1.81	1.73
	[1.54, 2.00]	[1.46, 1.97]
Mean	1.73	1.71
	[1.58, 1.88]	[1.57, 1.85]
Top 1% of producers, top 10% of programs		
Share (%)	30.9	30.9
Ratio to share of PhDs	1.59	1.59
Institution share of variance (%)	0.80	0.81

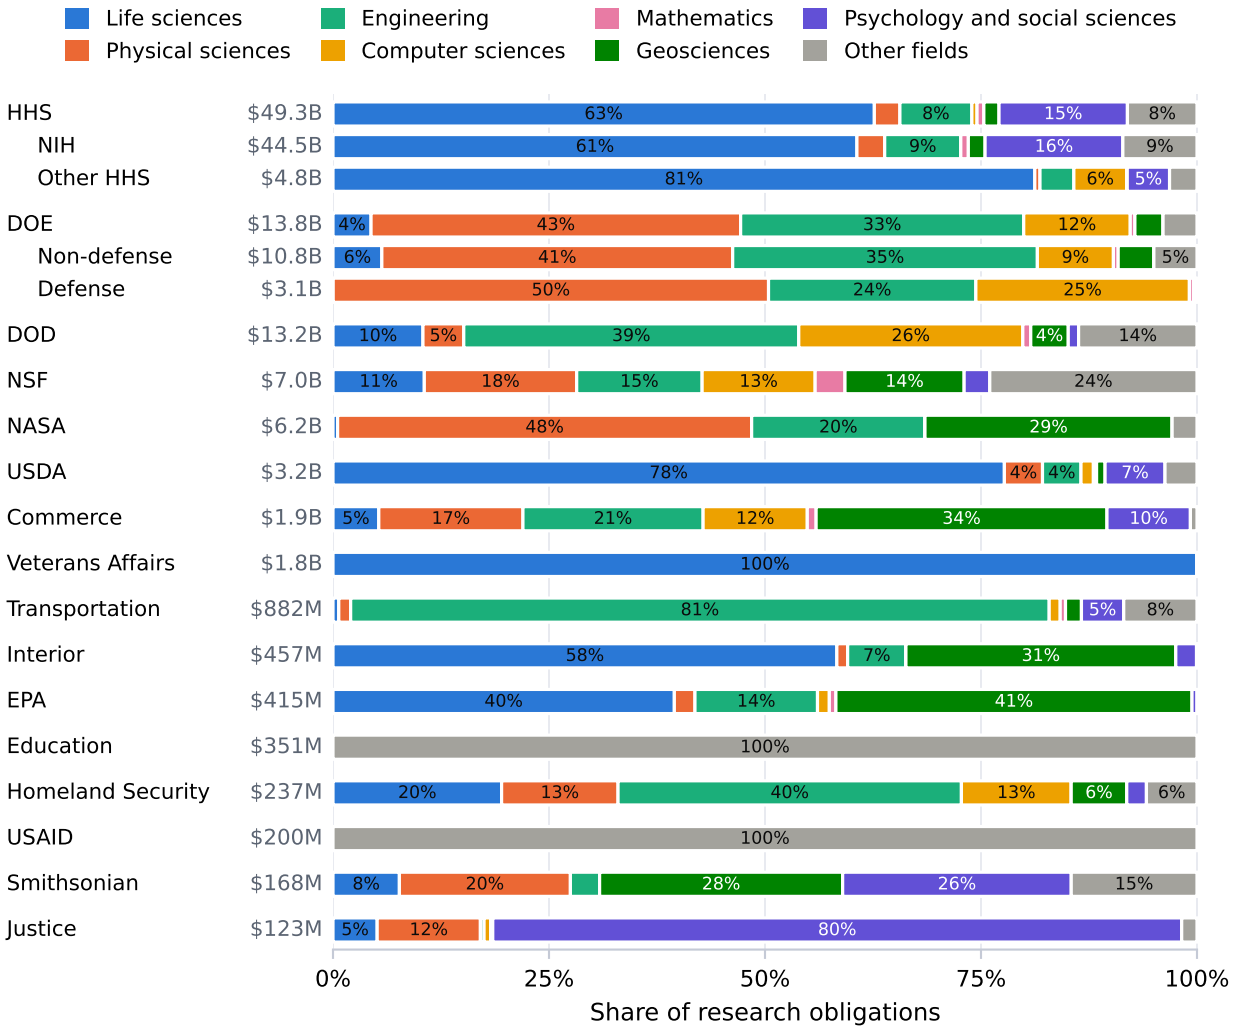
Notes: Quality-adjusted output divides each patent's external forward citations by the mean for patents granted in the same year in the same CPC subclass and sums over patents filed within ten years of the degree. The baseline counts citations within five years of disclosure; the alternative within ten. Because each patent is compared only with patents granted in the same year, the incomplete ten-year windows of recent grants affect every patent in a cell alike. Rows other than the last are among PhD inventors. P(beats) is as in Table 17, quantile ratios as in Table 20, and capture as in Table 21, averaged across fields; bracketed 95% intervals resample whole programs. The last row is the institution share of the variance of quality-adjusted output among all PhDs, from the by-field decomposition of Figure 16, averaged across fields.

F Proposed Reductions in University Support

The United States federal government is one of the world's largest funders of research and development. The bulk of this spend is concentrated in a subset of five federal agencies: the National Science Foundation (NSF), the National Institutes of Health (NIH), the Department of Energy (DOE), the National Aeronautics and Space Administration (NASA), and the Department of Defense (DOD).

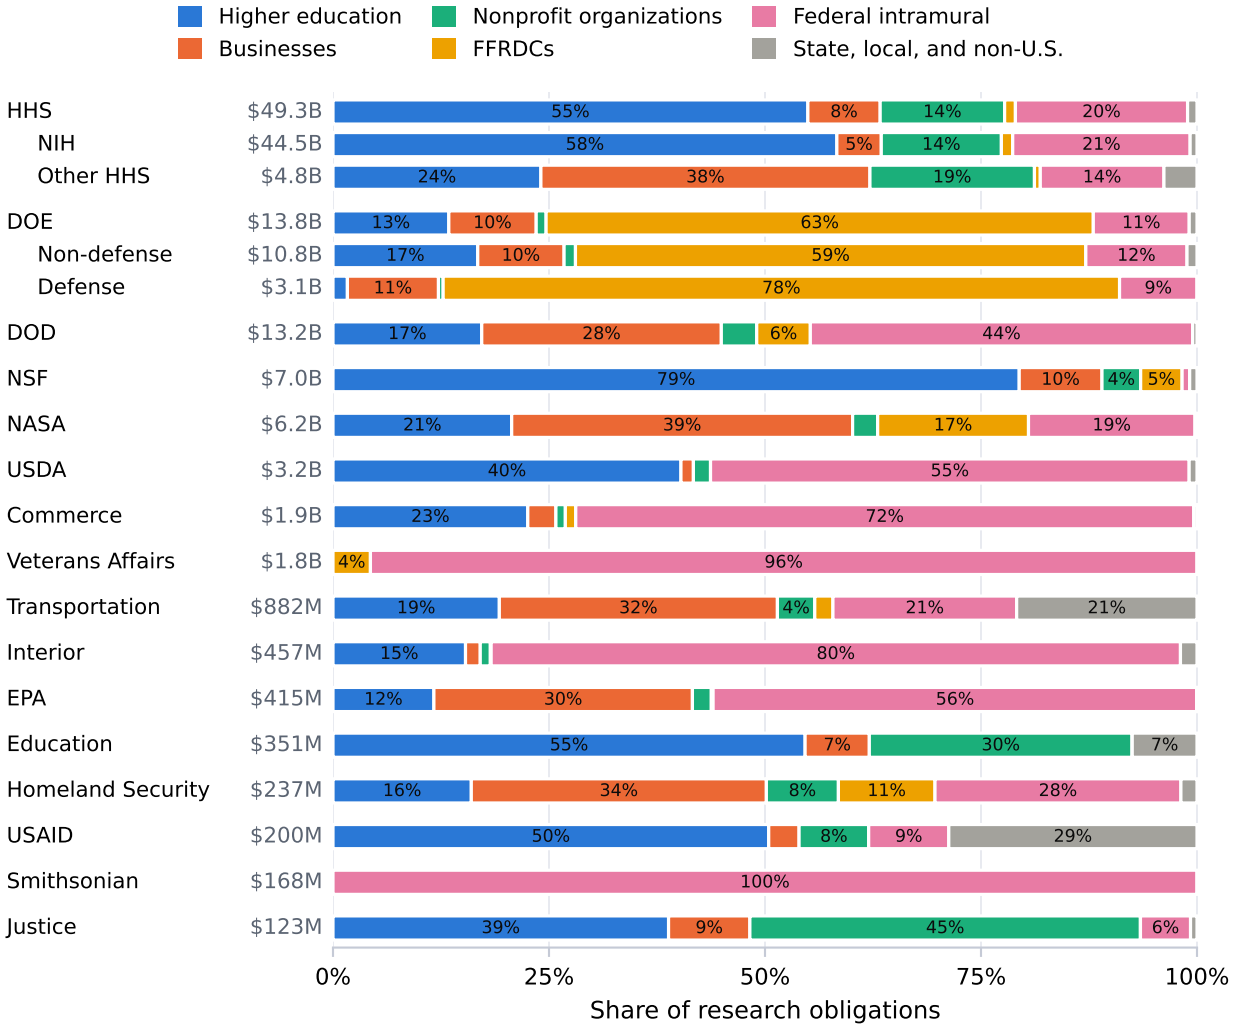
The R&D accounts within each of these agencies vary enormously in both the fields that they support, and the composition of support across different entities. Figure 34 shows the distribution of research *obligations* (not appropriations) by field. Two things are important to keep in mind. First, HHS provides dramatically more funding for research than any other Federal agency. Second, the split across performers varies substantially when comparing research and experimental development (Figures 35 and 37). Research activities are substantially more likely to be performed at universities, while development activities are more often performed at private firms: universities perform 41% of federal research obligations but only 4% of experimental development obligations, while businesses perform 13% of research and 60% of development.

Figure 34: Federal Research Obligations by Field, All Agencies



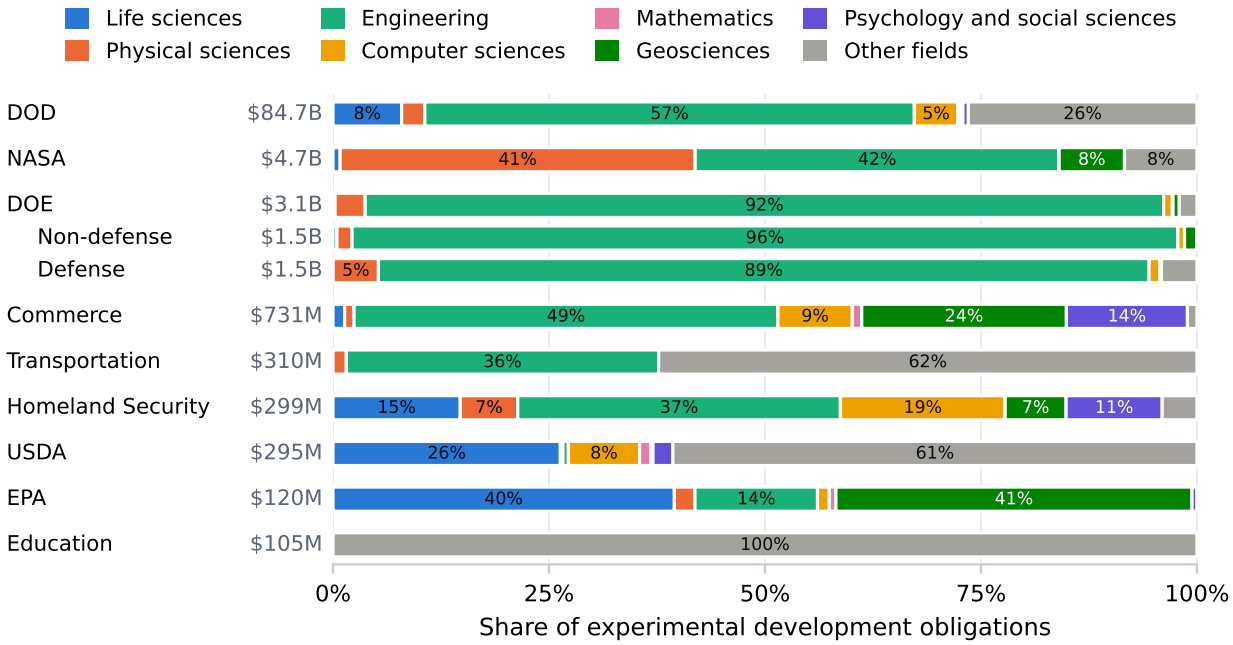
Notes: Includes every department and independent agency with at least $100 million of research obligations (99.8% of federal research obligations). DOE non-defense accounts include Office of Science, Energy Efficiency and Renewable Energy, Fossil Energy, Nuclear Energy, ARPA-E, Electricity, and Environmental Management, while DOE defense accounts include NNSA and Cybersecurity, Energy Security, and Emergency Response. HHS includes the Patient-Centered Outcomes Research Trust Fund. Dollar amounts beside each entity are total obligations. Source: NCSES, Survey of Federal Funds for Research and Development, FY 2024 (NSF 26-316, table 24).

Figure 35: Federal Research Obligations by Performer, All Agencies



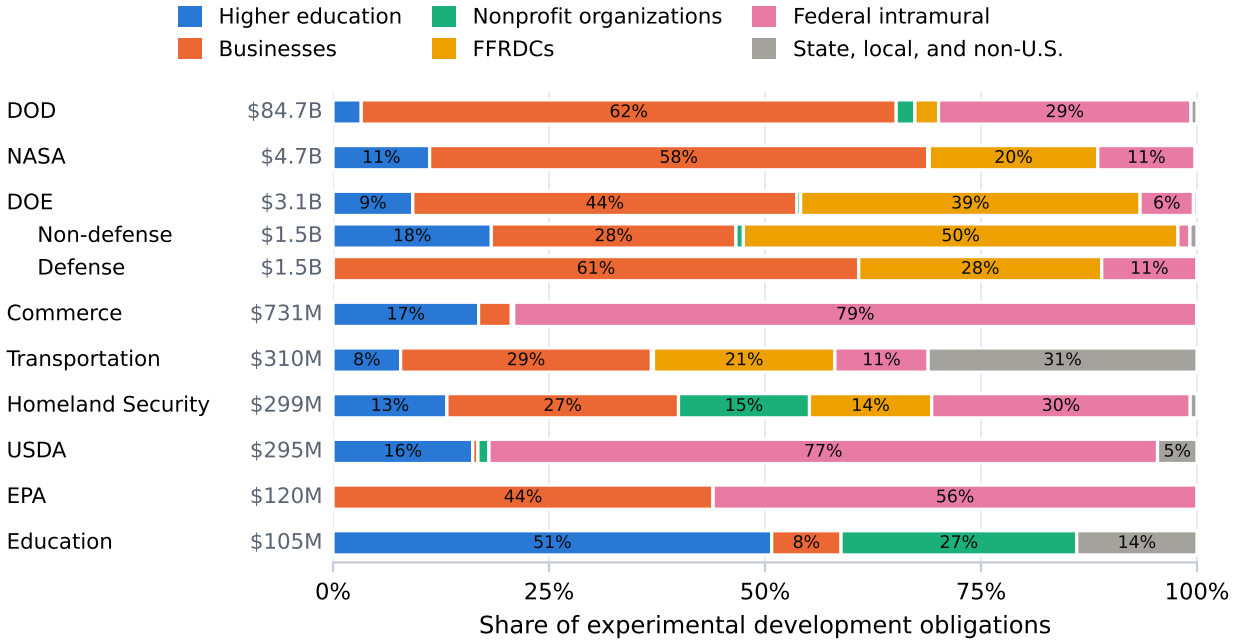
Notes: Includes every department and independent agency with at least $100 million of research obligations (99.8% of federal research obligations). DOE non-defense accounts include Office of Science, Energy Efficiency and Renewable Energy, Fossil Energy, Nuclear Energy, ARPA-E, Electricity, and Environmental Management, while DOE defense accounts include NNSA and Cybersecurity, Energy Security, and Emergency Response. HHS includes the Patient-Centered Outcomes Research Trust Fund. Dollar amounts beside each entity are total obligations. State and local governments and non-U.S. performers are combined. FFRDCs are federally funded research and development centers. Source: NCSES, Survey of Federal Funds for Research and Development, FY 2024 (NSF 26-316, table 21).

Figure 36: Federal Experimental Development Obligations by Field, All Agencies



Notes: Includes every department and independent agency with at least $100 million of experimental development obligations (99.6% of federal experimental development obligations). DOE non-defense accounts include Energy Efficiency and Renewable Energy, Nuclear Energy, ARPA-E, Electricity, and Environmental Management, while DOE defense accounts include NNSA (CESER has no experimental development obligations). Dollar amounts beside each entity are total obligations. Source: NCSES, Survey of Federal Funds for Research and Development, FY 2024 (NSF 26-316, table 39).

Figure 37: Federal Experimental Development Obligations by Performer, All Agencies



Notes: Includes every department and independent agency with at least $100 million of experimental development obligations (99.6% of federal experimental development obligations). DOE non-defense accounts include Energy Efficiency and Renewable Energy, Nuclear Energy, ARPA-E, Electricity, and Environmental Management, while DOE defense accounts include NNSA (CESER has no experimental development obligations). Dollar amounts beside each entity are total obligations. State and local governments and non-U.S. performers are combined. FFRDCs are federally funded research and development centers. Source: NCSES, Survey of Federal Funds for Research and Development, FY 2024 (NSF 26-316, table 36).

Table 23: Proposed Changes to Federal R&D Accounts and the Implied Cut to University Support

	Budget authority ($M)			Universities ($M)	
Account	FY 2025	FY 2026 request	Change	FY 2024 obligations	Implied change
Department of Health and Human Services					
National Institutes of Health‡	–	–	−43.5%	32,142	−13,982
Cancer	7,224	4,531	−37%		
Body Systems	7,229	4,152	−43%		
Neuroscience and Brain Research	4,107	2,445	−40%		
Allergy and Infectious Diseases	6,562	4,175	−36%		
General Medicine	5,774	3,677	−36%		
Disability Related Research	2,293	1,414	−38%		
Aging	4,425	2,687	−39%		
Behavioral Health	4,532	2,785	−39%		
Eliminated (NINR, NCCIH, NIMHD, Fogarty)	997	0	−100%		
Environmental Health Sciences (leaves NIH)	997	728	−27%		
ARPA-H	1,500	945	−37%	555	−205
Other HHS R&D†			−37.8%	619	−234
National Science Foundation	9,060	3,903	−56.9%		
Research and Related Activities	7,177	2,988	−58.4%	5,556	−3,245
STEM Education	1,172	288	−75.4%		
Major Research Equipment and Facilities	234	251	+7.3%		
Agency Operations and Award Management	448	355	−20.8%		
Inspector General and Science Board	30	21	−28.8%		
Department of Defense					
Science and technology (6.1–6.3)					
Army	3,874	2,587	−33%	315	−104
Navy	3,177	2,635	−17%	938	−160
Air Force	3,521	2,538	−28%	765	−214
Space Force	1,042	904	−13%	8	−1
Defense-wide	8,443	9,195	+9%	426	+38
RDT&E, all budget activities					
Army	14,363	14,549	+1%	48	0
Navy	25,967	25,708	−1%	381	−4
Air Force	46,881	52,017	+11%	190	+21
Space Force	18,553	15,486	−17%	73	−12
Defense-wide	35,239	33,922	−4%	326	−13
DARPA	4,146	4,566	+10%	653	+65
Defense Health Program R&D	1,695	1,047	−38%	911	−346

continued on the next page

Table 23, continued

Account	Budget authority ($M)			Universities ($M)	
	FY 2025	FY 2026 request	Change	FY 2024 obligations	Implied change
Department of Energy					
Office of Science	8,240	7,092	−14%	1,191	−167
Energy Efficiency and Renewable Energy	3,460	888	−74%	383	−284
Fossil Energy	–	–	−31%	275	−85
Nuclear Energy	–	–	−19%	57	−11
ARPA-E	–	–	−57%	174	−99
NNSA Weapons Activities	–	–	+29%	8	+2
NNSA Defense Nuclear Nonproliferation	–	–	−5%	13	−1
Other DOE R&D†			−16.2%	37	−6
NASA	24,838	18,800	−24%		
Science Mission Directorate	7,334	3,908	−47%	1,805	−848
Department of Agriculture					
National Institute of Food and Agriculture	1,679	1,049	−38%	999	−379
Agricultural Research Service	1,831	1,700	−7%	253	−18
Other USDA R&D†			−17.3%	84	−15
Department of Commerce					
NIST Scientific and Technical Research	857	707	−18%	176	−32
NOAA Oceanic and Atmospheric Research	761	0	−100%	367	−367
Other Commerce R&D†			−9.5%	17	−2
Department of Education†			−75.6%	245	−186
Department of Transportation					
FAA Research, Engineering and Development	223	165	−26%	63	−16
Other Transportation R&D†			−2.7%	131	−4
Department of the Interior					
USGS Surveys, Investigations, and Research	1,450	892	−38%	52	−20
Other Interior R&D†			−43.9%	23	−10
EPA Science and Technology	756	501	−34%	48	−16
All other departments and agencies†			−4.7%	284	−13
All accounts			−41.4%	50,593	−20,971

Notes: Budget authority is from AAAS. FY 2025 is the full-year continuing resolution, and rows without dollar amounts are reported by AAAS only as a percentage change. NSF is from NSF's FY 2026 budget justification: the request moves STEM Education into Research and Related Activities, so the R&RA request here excludes STEM Education's $288 million. University obligations are FY 2024 federal obligations to higher education for research and experimental development (NCSES, Survey of Federal Funds for Research and Development, NSF 26-316, tables 21 and 36); for NIH (‡) they also include $6,186 million to nonprofit performers. Each row's obligations take the change in its row; the implied change is the product, and the total row's change is the obligation-weighted average. I apply the R&RA change to all of NSF's university

obligations, and, analogously for NASA, I use the Science Mission Directorate's. Following AAAS, the NIH institute rows group the current institutes into the eight the request proposes, the four it would eliminate, and Environmental Health Sciences, which it moves out of NIH. This grouping and the resulting cuts are taken from AAAS's institute table and do not add up to the NIH-wide cut because of the treatment of ARPA-H and Environmental Health Sciences. All together, the budget would have reduced NIH appropriations by 43.5%. Department of Defense science and technology (S&T) is budget activities 6.1–6.3 of each service's RDT&E.

†Accounts without a row of their own. These are typically smaller AAAS lines and, in cases without reported rates, the department's R&D change from CRS R48694, Table 2, with changes shown as their obligation-weighted average. Other HHS R&D: (CRS) −37.8%. Other DOE R&D: (CRS) −16.2%. Other USDA R&D: (CRS) −25.3%; APHIS 0%; ERS −12%; NASS −1%. Other Commerce R&D: (CRS) −19.2%; Census Bureau +24%. Department of Education: (CRS) −75.6%. Other Transportation R&D: (CRS) −2.9%; NHTSA 0%; FRA Railroad R&D −10%; PHMSA 0%. Other Interior R&D: (CRS) −43.9%. All other departments and agencies: USAID (CRS) 0%; DHS S&T Directorate +2%; Justice (CRS) −32.5%; Nuclear Regulatory Commission (CRS) +3.1%; Social Security Administration (CRS) 0%; other Homeland Security (CRS) +8%; Appalachian Regional Commission held flat. Sources: American Association for the Advancement of Science, "FY 2026 R&D Appropriations: The Presidential Budget Request" (June 30, 2025), `https://www.aaas.org/news/fy-2026-rd-appropriations-presidential-budget-request`; National Science Foundation, FY 2026 Budget Request to Congress; Congressional Research Service, R48694.

To compute the implied reduction in federal support for R&D at universities, I assume that the reductions would fall on performers in proportion to their existing share of federal obligations. Because a substantial portion of extramural NIH research is performed at affiliated medical centers that are typically associated with a university, I include non-profit performers in the university bucket for NIH. I then compute the weighted average reduction in support. Table 23 reports each account's proposed change, its obligations to universities, and the implied change. The resultant decline in support is approximately 41.4%. Performing an analogous exercise using agency funding from the Higher Education Research and Development Survey generates an implied cut that is slightly lower (39.9%). Because this process neglects other proposed and/or enacted policy changes (like partial upfront financing of NIH grants, indirect cost caps, and the endowment tax) that would further reduce university resources, I take the larger of the two numbers. Note that these cuts were proposed for the Fiscal Year 2026 President's budget. Congress ultimately did not enact them, instead opting to keep funding roughly constant. The Fiscal Year 2027 President's budget again proposed cuts, but those cuts were smaller—with NIH for example facing a roughly 11% reduction instead of 43.5%.[36]

Note that in the main text I refer to the R&D budget for DOD as reflecting proposed increases of 22.9%. This differs from the total in Table 24 for three reasons. First is the fact that DOD R&D spending typically excludes budget activity 6.7 (Operational systems development). Second is that the CRS numbers use a slightly different FY2025 base. Third concerns the treatment of the the 2025 reconciliation bill (the One Big Beautiful Bill Act). Table 24 breaks DOD's research, development, test and evaluation (RDT&E) request down by budget activity. The discretionary request was roughly flat (+0.5%), and the bill added $37.1 billion, for a total increase of 26.7%. Only $2.3 billion of the $37.1 billion went to science and technology (budget activities 6.1–6.3), which carries about four-fifths of DOD's support for universities. Even among the science and technology activities, nearly all of the additional funding went to advanced technology development, which has, proportionally, the

[36]See here.

least university involvement of the three. Because of this, the DOD rows of Table 23 use only the discretionary request and exclude reconciliation funding.

Table 24: DOD Research, Development, Test and Evaluation Request by Budget Activity

		FY 2026 request			Change from FY 2025	
Budget activity	FY 2025	Discretionary	Reconciliation	Total	Discretionary	Total
6.1 Basic research	2,474	2,256	10	2,266	−8.8%	−8.4%
6.2 Applied research	6,251	5,769	255	6,024	−7.7%	−3.6%
6.3 Advanced technology development	10,136	9,985	2,003	11,988	−1.5%	+18.3%
Science and technology (6.1–6.3)	18,861	18,010	2,268	20,278	−4.5%	+7.5%
6.4 Advanced component development and prototypes	35,995	30,781	9,663	40,444	−14.5%	+12.4%
6.5 System development and demonstration	24,466	27,483	12,195	39,679	+12.3%	+62.2%
6.6 Management support	9,837	9,811	293	10,104	−0.3%	+2.7%
6.7 Operational systems development	51,899	54,396	12,417	66,813	+4.8%	+28.7%
6.8 Software and digital technology pilots	295	829	233	1,062	+181.5%	+260.6%
Agile RDT&E portfolio management	0	690	0	690	–	–
Total RDT&E	141,352	142,001	37,069	179,070	+0.5%	+26.7%

Notes: Millions of dollars. All RDT&E appropriations (Army, Navy, Air Force, Space Force, Defense-Wide, and Operational Test and Evaluation). FY 2025 is the enacted level plus the supplemental. Reconciliation is the mandatory funding from the 2025 reconciliation bill (the One Big Beautiful Bill Act) that the request counts toward FY 2026, $37.1 billion in all. The largest reconciliation amounts in science and technology are Defense Innovation Unit (DIU) ($661 million, 6.3); Joint Hypersonic Technology Development & Transition ($600 million, 6.3); Test & Evaluation Science & Technology ($550 million, 6.3); Making, Maintaining, Supply Chain and Logistics ($250 million, 6.2). Agile RDT&E portfolio management is a new FY 2026 budget activity. Source: DOD Comptroller, RDT&E Programs (R-1), FY 2026 President's Budget.

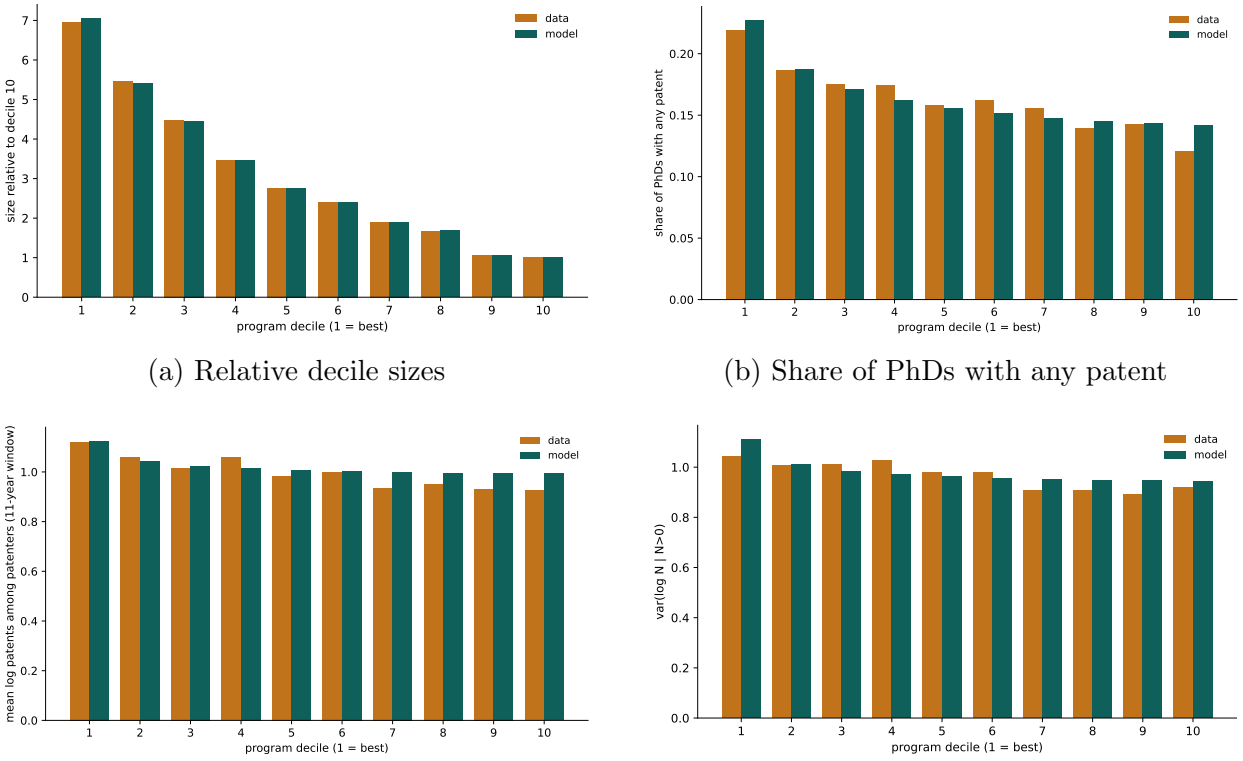
G Quantitative Model Appendix

This appendix provides additional plots from the quantitative model.

G.1 Targeted Moments

Tables 2 and 3 summarize the fit of each block of targeted moments by its mean and maximum absolute percentage error. This section plots the blocks against the data at the calibrated parameters. Figure 38 shows the four blocks constructed by program decile from the matched dissertation-inventor data: the relative sizes of the deciles, the share of PhDs with any patent, and the mean and the variance of log patent counts among those who patent. Figure 39 shows elite capture at four depths of the inventor tail. Figure 40 shows the funding shock fed to the model and the response of completions whose values at years 6 and 7 identify the admission cost elasticity ν.

Figure 38: Targeted Moments by Program Decile: Model and Data



(a) Relative decile sizes

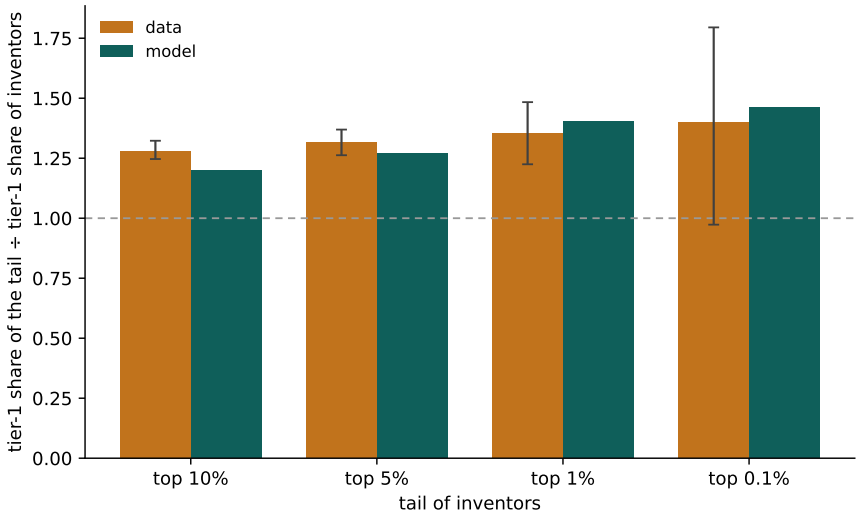
(b) Share of PhDs with any patent

(c) Mean log count among patenters

(d) Variance of log count among patenters

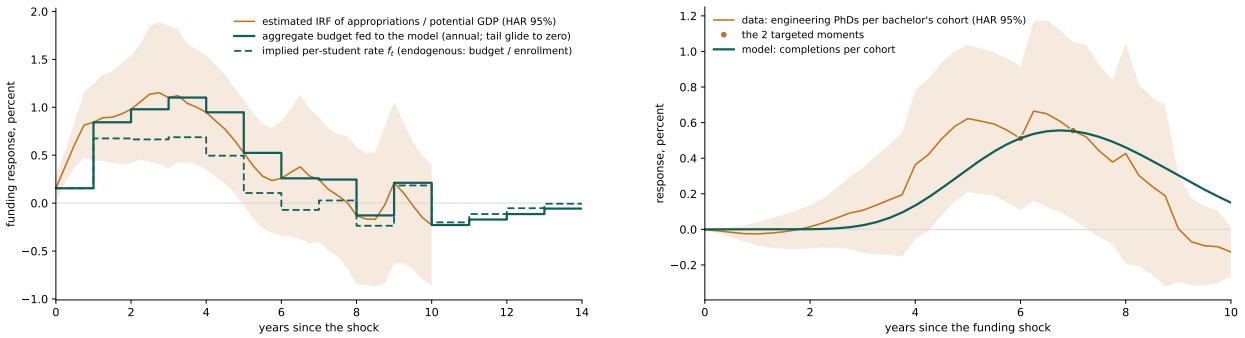
Notes: Programs are ranked within field and grouped into deciles, with decile 1 the best. Data moments pool the matched sample across fields with equal field weights; counts are patents over the 11-year window. Panel (a) normalizes the size of decile 10 to one, so nine of the ten bars are free moments. Model moments are computed on the balanced growth path through the negative binomial count layer at the calibrated parameters. The relative sizes identify the profile of program productivity $\alpha_1, \ldots, \alpha_{10}$; the within-decile variance identifies σ_R; the patenting share and its gradient across deciles identify ϱ and σ_ε; the mean log count identifies the count scale A.

Figure 39: Elite Capture of the Inventor Tail: Model and Data



Notes: Elite refers to the model analogue of the graduates of the top decile of a field's programs by NRC rank. This holds about 23 percent of PhDs. Each bar is the elite's share of the top 10, 5, 1 and 0.1 percent of patenters by count, pooled across fields, divided by the elite's share of all patenters, so a value of one means no over-representation in the tail beyond the participation gradient. Whiskers are 95 percent bootstrap intervals over inventors. These four moments, together with the pooled Hill tail index reported in Table 3, identify the scale of the star component τ.

Figure 40: The Funding Shock and the Completions Response: Model and Data



(a) Funding input

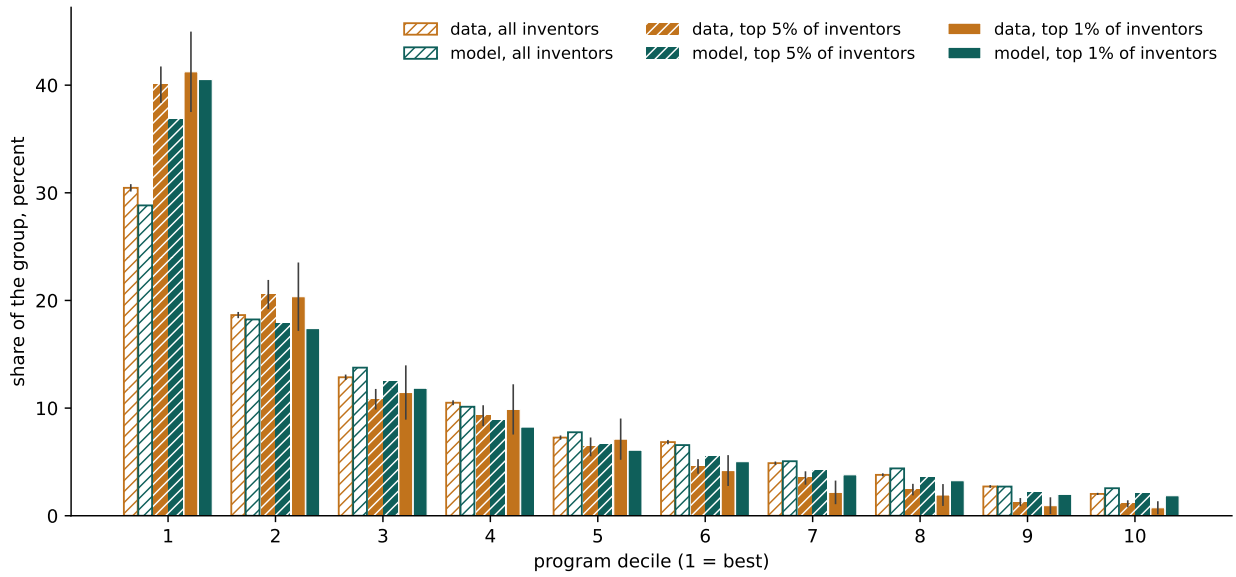
(b) Completions

Notes: Responses are in percent per one percent narrative appropriation shock, for engineering. Panel (a) shows the estimated response of research appropriations as a share of potential GDP with its 95 percent HAR band, the annual aggregate budget path fed to the model, which follows the estimate for ten years and then glides to zero, and the per-student rate f_t that the budget implies given the enrollment response. Panel (b) shows the local-projection response of PhDs per bachelor's cohort with its 95 percent HAR band and the model's completions per cohort along the transition. The two targeted moments are the responses at years 6 and 7; the rest of the path is untargeted.

G.2 Untargeted Moments

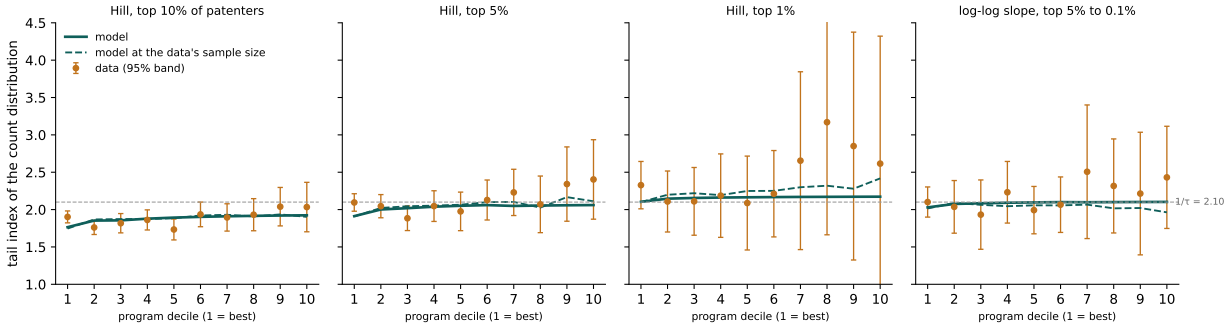
Figure 41 adds the composition of the (pooled) inventor tail, while Figure 42 plots the tail index of the count distribution by decile. In the estimation, I target only the pooled Hill index over the top 5 percent of patenters. With that being said, this is more a test of the specification of the ability distribution, as the footnote in the text about Breiman's lemma makes clear. Table 4 in the main text summarizes the untargeted moments by their values in the best and worst deciles and by their mean and maximum absolute percentage errors across deciles. Figures 44 and 43 plot the same moments by program decile. Figure 45 compares the model-implied quality impulse response to the quality impulse response estimated in the pooled design discussed in Appendix E.1.

Figure 41: Composition of the Inventor Tail by Program Decile: Model and Data



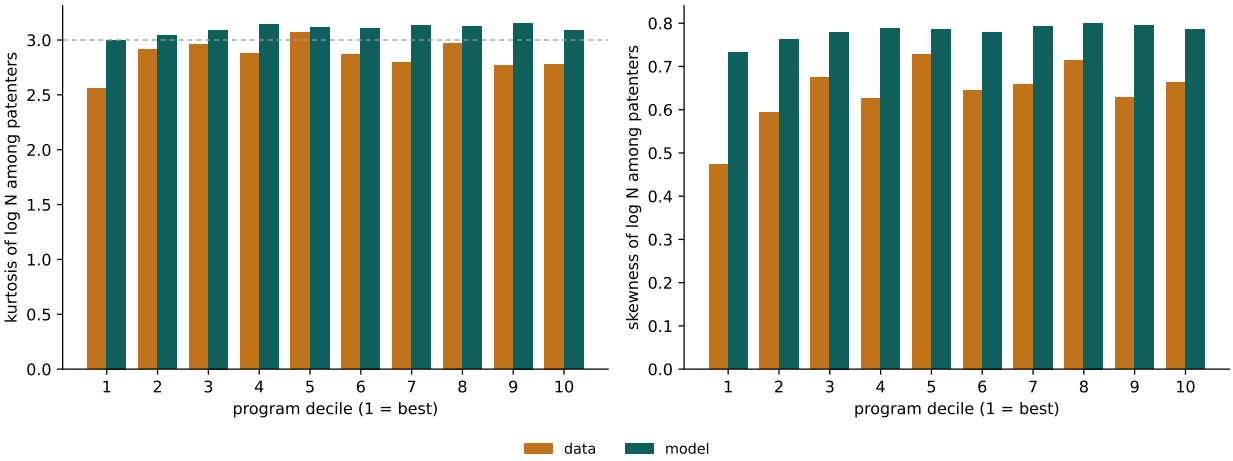
Notes: Each bar is the share of the group, in percent, that graduated from programs in each within-field decile. I pool across fields and include all inventors (outlined bars), the top 5 percent of inventors (hatched) and the top 1 percent (solid). Whiskers are 95 percent bootstrap intervals on the data shares. Inventors are PhDs with at least one patent filed within ten years of the degree. Tails are by patent count, pooled, with tied counts weighted so that each tail is exactly the top fraction. Model shares come from the balanced growth path's researcher densities by decile following the same rule for constructing the tail. The share of all inventors is a product of two targeted blocks, (decile sizes and participation by decile). Shares of the top 5 and top 1 percent are not targeted.

Figure 42: Tail Index of the Count Distribution by Program Decile: Model and Data



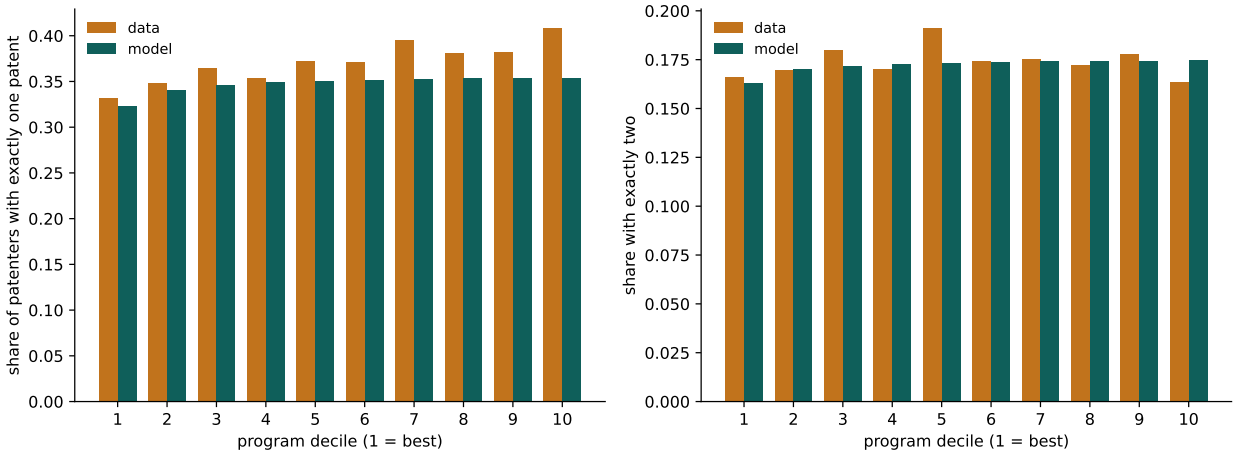
Notes: Each panel estimates the tail index of the distribution at various depths, in order of the Hill estimator over the top 10, 5 and 1 percent of patenters, and the slope of the log survival function against the log count between the top 5 and the top 0.1 percent. Data points pool the decile's patenters across fields and reportsh 95 percent bootstrap bands. Outcome is patents filed within ten years of the degree. The solid line is the model's population value from the count distribution on the balanced growth path. The dashed line applies the data's estimator to samples drawn from the model with each decile's number of patenters. The grey line is $1/\tau$, the asymptotic tail index implied by estimated exponential component of ability. Estimation targets only the pooled Hill index over the top 5 percent of patenters (Table 3).

Figure 43: Shape of Log Output among Patenters by Program Decile: Model and Data

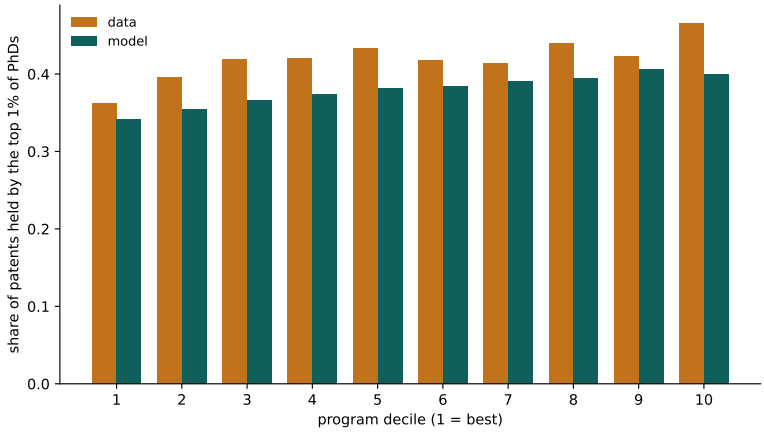


Notes: Kurtosis (left) and skewness (right) of log patent counts among patenters, by program decile. Log output is centred on the field × cohort cell's geometric mean and pooled by decile. The dashed line marks the kurtosis of a normal distribution. I generate a sample comparable to the data's field × cohort cells. The model is more right-skewed than the data in every decile and slightly more leptokurtic, and the cross-decile gradient in both moments is flatter in the model than in the data.

Figure 44: Untargeted Moments of the Count Distribution by Program Decile: Model and Data



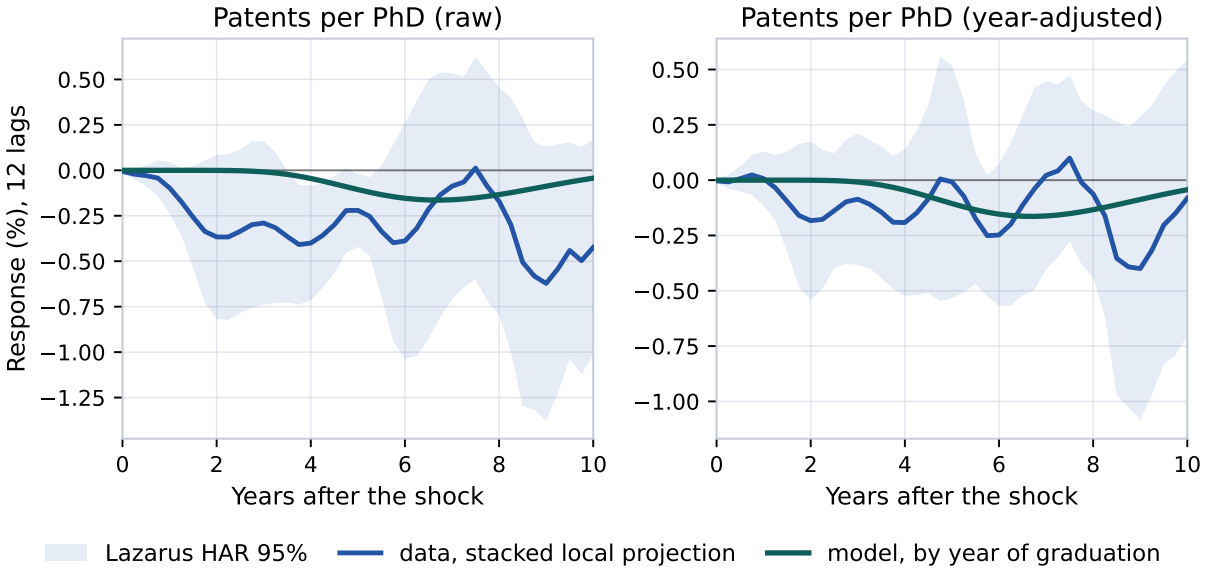
(a) Share of patenters with exactly one patent (left) and with exactly two (right)



(b) Share of a decile's patents held by its top 1 percent of PhDs

Notes: Programs are ranked within field and grouped into deciles, with decile 1 the best. Data moments pool the matched sample across fields with equal field weights. I compute counts as patents filed within ten years of the degree. In panel (a) the denominator is the decile's PhDs with at least one patent. In panel (b) the top 1 percent is taken over all of a decile's PhDs, regardless of patenting. The data take the top $\lceil 0.01\,n \rceil$ PhDs of each field $\times$ decile cell. I interpolate tails so that an integer count at the 99th percentile does not move the effective threshold. Table 4 reports the decile 1 and decile 10 values and the mean and maximum absolute percentage errors across deciles.

Figure 45: Response of Patents per PhD to a Funding Shock: Model and Data



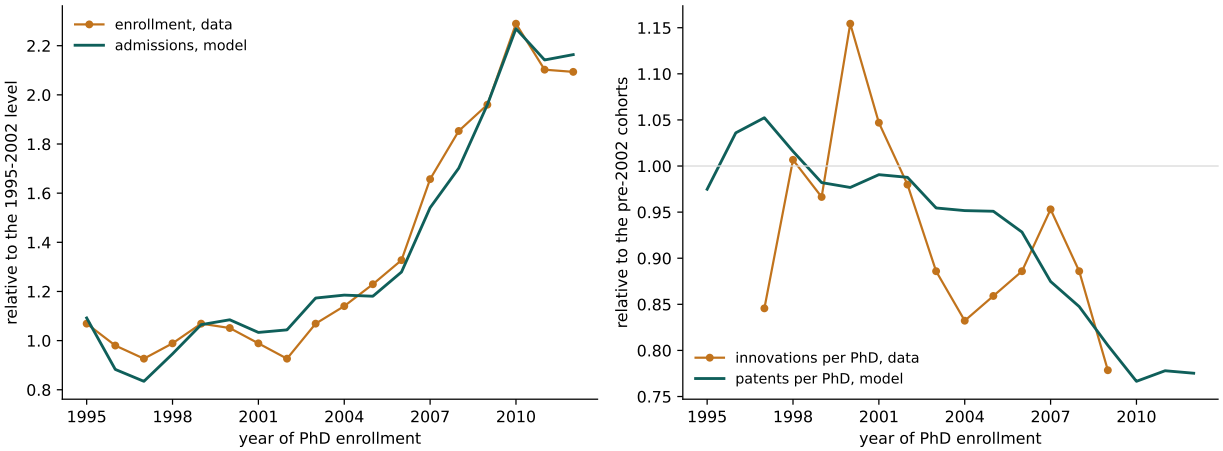
Notes: Response of patents per PhD to a 1% narrative appropriation shock, in percent, by years after the shock. The data are stacked local projections that pool nine fields with field fixed effects and common coefficients on 12 lags of log appropriations, the outcome and cyclical controls, for cohorts graduating from 1980 to 2010. The left panel uses the log of patents filed within 10 years of the degree per PhD in the cohort. The right panel uses the cohort effect from a Poisson regression of each cohort's patents on cohort and field-by-filing-year fixed effects, which removes calendar-year swings in patenting. Shaded bands are 95% intervals with Lazarus et al. (2018) standard errors. The model line is patents per PhD by year of graduation, when the model is fed the appropriation response used to estimate the quantity response.

G.3 The Danish PhD Expansion

There is a remarkable lack of systematic evidence on the effects of expansions of post-secondary education on quality of individual applicants. What evidence there is typically comes from the Danish PhD expansion, which was a rough doubling of PhD slots in Denmark that occurred in the 21st century. Akcigit et al. (2025) study this expansion and find that it led to a substantial increase in the number of PhDs enrolled, but that the quality of the marginal PhD, as measured by patents per PhD, declined somewhat.[37] Figure 46 simulates an expansion of comparable size and plots the resultant change in patents per PhD. The left panel shows the expansion in PhD slots, and the right panel shows the decline in patents per PhD. The model is able to match the expansion in PhD slots and generates a decline in patents per PhD that is comparable to the Danish time-series evidence.

[37] See Figures C1 and C2 in their online appendix. Their outcome is "share of patents filed per student within three years of PhD enrollment from 1995–2015."

Figure 46: The Danish PhD Expansion: Data and Model



Notes: Left panel: PhD enrollment in the data and the model's admissions along the appropriation path, relative to the 1995–2002 level. Right panel: innovations per PhD in the data and the model's patents per PhD, by year of enrollment, relative to the pre-2002 cohorts.

H A Taste for Research

This appendix documents an extension wherein I assume that individuals' choices are driven not just by an individual's ability at research but also by their taste for it. As in the baseline model, learning is public. Lacking evidence on initial application behavior, I assume that taste is revealed only upon graduation: no one, including the individual, observes it before then, and the public signal drawn at birth is informative about taste only through taste's correlation with research ability. In effect, only through doing research can an individual learn whether or not they like it.

H.1 Researchers

As in the baseline model, a fraction of individuals are eligible to become researchers. These eligible individuals are now indexed by their ability at research, z_R, and their taste for research, ψ_R. As in the main text, log ability is the sum of a normal component $\eta \sim \mathcal{N}(0, \sigma_R^2)$ and an independent exponential component γ with scale parameter τ, so that $\mathbb{E}[\ln z_R] = \tau$ and $\text{sd}(\ln z_R) = \sqrt{\sigma_R^2 + \tau^2}$. Log taste is normally distributed conditional on log ability, with a mean that is linear in log ability:

$$\ln \psi_R \,\big|\, z_R \;\sim\; \mathcal{N}\Big(m_\psi(z_R),\; \sigma_\psi^2(1-\omega^2)\Big), \qquad m_\psi(z_R) = \mu_\psi + \omega\,\sigma_\psi\,\frac{\ln z_R - \tau}{\sqrt{\sigma_R^2+\tau^2}}, \tag{76}$$

where $\omega \in (-1, 1)$. Log taste therefore has mean μ_ψ and variance σ_ψ^2, and its correlation with log ability is $\text{Corr}(\ln z_R, \ln \psi_R) = \omega$. When $\omega = 0$, taste is independent of ability and log taste is distributed $\mathcal{N}(\mu_\psi, \sigma_\psi^2)$. Note that μ_ψ is not a normalization.

Individuals in the model have flow preferences given by:

$$u(c) = \begin{cases} \ln(c) & \text{if not a researcher,} \\ \ln(c) + \ln(\psi_R) & \text{if a researcher.} \end{cases} \tag{77}$$

The value of a match between worker and innovative firm depends now on both the ability of the researcher to do research z_R and their taste for research ψ_R. The Hamilton-Jacobi-Bellman equation of a researcher with state (z_R, ψ_R) in the continuation region is given by:

$$(\rho+\delta)\hat{\mathcal{R}}_t(z_R,\psi_R;\{\hat{\mathcal{P}}_s(z_R,\psi_R)\}_{s\geq t}) - \dot{\hat{\mathcal{R}}}_t = \ln\big((1-\tau_W)(\hat{\mathcal{P}}_t(z_R,\psi_R)+\hat{\Pi}_t)\big) + \ln(\psi_R) \tag{78}$$

The value of an innovative firm is then given by:

$$(r^*+\delta-g_t)\hat{\mathcal{V}}_t(z_R,\psi_R;\{\hat{\mathcal{P}}_s(z_R,\psi_R)\}_{s\geq t}) - \dot{\hat{\mathcal{V}}}_t = \frac{\hat{V}_t\varphi_R z_R}{1-s_R} - \hat{\mathcal{P}}_t(z_R,\psi_R) \tag{79}$$

As in the baseline, the R&D payroll credit enters as the factor $1/(1-s_R)$ on the value of an innovation. Notice that here, the value of a firm does not depend on the value of taste,

except through the firm's ability to extract additional pecuniary value as compensation for the ability to do research. As in the baseline, the Nash Bargaining problem at every instant solves:

$$\max_{\hat{\mathcal{P}}_t(z_R,\psi_R)} \left(\hat{\mathcal{R}}_t\big(z_R,\psi_R;\hat{\mathcal{P}}_t(z_R,\psi_R)\big) - \hat{\mathcal{W}}_t\right)^{\theta} \left(\hat{\mathcal{V}}_t\big(z_R,\psi_R;\hat{\mathcal{P}}_t(z_R,\psi_R)\big)\right)^{1-\theta} \tag{80}$$

H.1.1 Characterization

To characterize the decisions of individuals in this economy, I proceed in two parts. First, I characterize the beliefs of an individual at entry, and then I characterize the exit decisions as a function of an individual's taste draw.

An individual at entry must form expectations over both their true research productivity and their taste for research, which are potentially correlated. The public signal is (potentially) informative about taste only through its correlation with ability. As a consequence, taste is independent of the signal conditional on ability, and one can factor the joint density:

$$f(z_R,\psi_R \mid n_R) = f(z_R \mid n_R)\, f(\psi_R \mid z_R), \tag{81}$$

where $f(\psi_R \mid z_R)$ is given by Equation (76) and $f(z_R \mid n_R)$ is the baseline posterior over ability. With reliability $\beta = \sigma_R^2/(\sigma_R^2+\sigma_\varepsilon^2)$, the normal component has posterior $\eta \mid n_R \sim \mathcal{N}\big(\beta n_R, \sigma_R^2(1-\beta)\big)$, while the tail component γ, which the signal does not load on, retains its prior.

As in the main text, a graduate student passes through $\mathcal{O}$ stages, each completed at rate $\mathcal{O}\mu$, and receives the stipend $\hat{s}_{t_0}(n_R)$ contracted at the admission date t_0. The value of a graduate student with signal n_R, admitted at t_0, in stage k is thus given by:

$$\begin{aligned}
(\rho+\delta+\mathcal{O}\mu)\hat{E}_t^{(k)}(n_R;t_0) - \dot{\hat{E}}_t^{(k)} &= \ln\big((1-\tau_W)(\hat{s}_{t_0}(n_R)+\hat{\Pi}_t)\big) + \mathcal{O}\mu\, \hat{E}_t^{(k+1)}(n_R;t_0), \qquad k=1,\dots,\mathcal{O}-1,\\
(\rho+\delta+\mathcal{O}\mu)\hat{E}_t^{(\mathcal{O})}(n_R;t_0) - \dot{\hat{E}}_t^{(\mathcal{O})} &= \ln\big((1-\tau_W)(\hat{s}_{t_0}(n_R)+\hat{\Pi}_t)\big) + \mathcal{O}\mu\, \mathbb{E}_{z_R,\psi_R}\Big[\max\big\{\hat{\mathcal{W}}_t, \hat{\mathcal{R}}_t(z_R,\psi_R)\big\}\Big| n_R\Big].
\end{aligned} \tag{82}$$

Taste enters the problem of a graduate student only through the expected value of research at graduation.

The next step is to characterize the problem of an individual who has realized their research ability and taste for research. The production of a match is given by:

$$I_t = \varphi_R z_R$$

A match ends when there are no longer gains from trade. In particular, a researcher exits when the continuation value of research lies strictly below the value of becoming a production worker:

$$\hat{\mathcal{W}}_t > \hat{\mathcal{R}}_t(z_R,\psi_R;\{\hat{\mathcal{P}}_t^*(z_R,\psi_R)\}) \tag{83}$$

Define the researcher's surplus over production, $\mathcal{D}_t(z_R, \psi_R) \equiv \hat{\mathcal{R}}_t(z_R, \psi_R) - \hat{\mathcal{W}}_t$. Subtracting the HJB equation of a production worker with a doctorate from Equation (78), it satisfies:

$$(\rho + \delta)\mathcal{D}_t - \dot{\mathcal{D}}_t = \ln\left(\frac{\psi_R(\hat{\mathcal{P}}_t(z_R, \psi_R) + \hat{\Pi}_t)}{\hat{w}_t \xi_2 z_P + \hat{\Pi}_t}\right) \tag{84}$$

This ODE governs the relative benefit for a researcher of type (z_R, ψ_R) of being in research as opposed to production. As in the baseline, the split of flow surplus is determined by

$$\theta \hat{\mathcal{V}}_t = (1 - \theta)(\hat{\mathcal{P}}_t(z_R, \psi_R) + \hat{\Pi}_t)\mathcal{D}_t \tag{85}$$

Solve this expression for $\hat{\mathcal{P}}_t(z_R, \psi_R)$ to recover that:

$$\hat{\mathcal{P}}_t(z_R, \psi_R) = \left(\frac{\theta}{1 - \theta}\right)\left[\frac{\hat{\mathcal{V}}_t}{\mathcal{D}_t}\right] - \hat{\Pi}_t \tag{86}$$

Given this expression, one can substitute into Equations (78) and (79) and recover a coupled system of HJB equations. The match persists if and only if there are nonnegative gains from trade, so that $\hat{\mathcal{R}}_t(z_R, \psi_R) \geq \hat{\mathcal{W}}_t$ and $\hat{\mathcal{V}}_t(z_R, \psi_R) \geq 0$.

Lemma 4. *For each taste ψ_R and date t, there exists a threshold $z_t^*(\psi_R)$ so that researchers of taste ψ_R and ability $z_R < z_t^*(\psi_R)$ choose to exit. The threshold $z_t^*(\psi_R)$ is strictly decreasing in ψ_R.*

Proof. I first show that a match of type (z_R, ψ_R) is viable at t only if there exists $s \geq t$ such that

$$\psi_R\big[\frac{\hat{V}_s \varphi_R z_R}{1 - s_R} + \hat{\Pi}_s\big] \geq \hat{w}_s \xi_2 z_P + \hat{\Pi}_s.$$

Suppose the match is viable at t and let $T \in (t, \infty]$ be the date at which the match ends. If $T < \infty$, continuation values at T are zero, so the firm's participation constraint $\hat{\mathcal{V}}_T \geq 0$ requires $\hat{\mathcal{P}}_T \leq \hat{V}_T \varphi_R z_R/(1 - s_R)$ and the researcher's participation constraint $\hat{\mathcal{R}}_T \geq \hat{\mathcal{W}}_T$ requires $\psi_R(\hat{\mathcal{P}}_T + \hat{\Pi}_T) \geq \hat{w}_T \xi_2 z_P + \hat{\Pi}_T$. A payment satisfying both exists only if

$$\psi_R\big[\frac{\hat{V}_T \varphi_R z_R}{1 - s_R} + \hat{\Pi}_T\big] \geq \hat{w}_T \xi_2 z_P + \hat{\Pi}_T.$$

If $T = \infty$, the match is viable on the terminal balanced growth path, where both participation constraints hold with stationary values and the same inequality follows. Either way, the inequality holds at some $s \geq t$.

Let

$$z_t^*(\psi_R) = \inf_{s \geq t} \frac{1 - s_R}{\hat{V}_s \varphi_R}\left[\frac{\hat{w}_s \xi_2 z_P + \hat{\Pi}_s}{\psi_R} - \hat{\Pi}_s\right].$$

For $z_R < z_t^*(\psi_R)$ the inequality fails at every $s \geq t$, so a researcher of type (z_R, ψ_R) exits at t.

Suppose $\psi_1 < \psi_2$. For every $s \geq t$,

$$\frac{1-s_R}{\hat{V}_s\varphi_R}\left[\frac{\hat{w}_s\xi_2 z_P + \hat{\Pi}_s}{\psi_1} - \hat{\Pi}_s\right] - \frac{1-s_R}{\hat{V}_s\varphi_R}\left[\frac{\hat{w}_s\xi_2 z_P + \hat{\Pi}_s}{\psi_2} - \hat{\Pi}_s\right] = \frac{(1-s_R)(\hat{w}_s\xi_2 z_P + \hat{\Pi}_s)}{\hat{V}_s\varphi_R}\left(\frac{1}{\psi_1} - \frac{1}{\psi_2}\right) > 0,$$

and the first factor is bounded away from zero along the path, so taking the infimum over s gives $z_t^*(\psi_2) < z_t^*(\psi_1)$. ■

Lemma 4 establishes a necessary condition for a match to persist. The final step to characterize the new properties under the taste-based economy is to derive aggregate quantities.

Let $\Gamma_t^R(z_R, \psi_R)$ denote the measure of employed researchers of type (z_R, ψ_R), and let $dF_t(n_R)$ denote the flow of graduates with signal n_R, as in the baseline. Let $f(z_R, \psi_R \mid n_R)$ denote the density of (z_R, ψ_R) among graduates with signal n_R, given by Equation (81).

Along the balanced growth path, the threshold in Lemma 4 is exact. That is to say that a graduate of type (z_R, ψ_R) enters research if and only if $z_R \geq z^*(\psi_R)$, and a researcher of that type remains in research if and only if the same inequality holds. Along a transition, the threshold in the Lemma is a lower bound. To see why this is, imagine a case where the threshold is low at calendar time 20 years in the future. Because the threshold computed in the Lemma is the lowest static threshold, being above the lowest static threshold far in the future is generally not sufficient to guarantee one would want to remain in research. The correct threshold for determining exit today directly compares continuation values. This threshold is weakly decreasing in taste. With some abuse of notation, I let $z_t^*(\psi_R)$ denote this exact threshold in what follows. The measure of researchers therefore satisfies

$$\dot{\Gamma}_t^R(z_R, \psi_R) = \int \mathbf{1}\{z_R \geq z_t^*(\psi_R)\}\, f(z_R, \psi_R \mid n_R)\, dF_t(n_R) - \delta\, \Gamma_t^R(z_R, \psi_R) \tag{87}$$

for all (z_R, ψ_R) with $z_R \geq z_t^*(\psi_R)$, and $\Gamma_t^R(z_R, \psi_R) = 0$ otherwise. Aggregate research skill and the research headcount are

$$H_{Rt} = \iint z_R\, d\Gamma_t^R(z_R, \psi_R), \qquad R_t = \iint d\Gamma_t^R(z_R, \psi_R). \tag{88}$$

Aggregate researcher payroll is given by:

$$\iint \hat{\mathcal{P}}_t(z_R, \psi_R)\, d\Gamma_t^R(z_R, \psi_R), \tag{89}$$

which replaces $\int \hat{\mathcal{P}}_t(z_R)\, d\Gamma_t^R(z_R)$ in the government budget constraint. Along the balanced growth path the stationary measure of researchers is

$$\Gamma^R(z_R, \psi_R) = \frac{1}{\delta}\int \mathbf{1}\{z_R \geq z^*(\psi_R)\}\, f(z_R, \psi_R \mid n_R)\, dF(n_R). \tag{90}$$

There are three kinds of production workers: ineligible individuals, with productivity $\xi_1 z_P$; eligible individuals who did not enroll, with productivity z_P; and doctorate holders who are not in research, with productivity $\xi_2 z_P$. Let L_t^I, L_t^E and X_t denote the number of people of each type. Of the δ births that occur at any given instant, a fraction π_R are

eligible. Letting $x_t = \sum_j x_{jt}$ denote the admission flow, we have:

$$\dot{L}_t^I = \delta(1 - \pi_R) - \delta L_t^I, \qquad \dot{L}_t^E = \delta\pi_R - x_t - \delta L_t^E. \tag{91}$$

Let $\Gamma_t^G = \int dF_t(n_R)$ denote the flow of graduates and e_t the flow of graduates who enter research,

$$e_t = \int \mathbb{E}\Big[\mathbf{1}\big\{z_R \geq z_t^*(\psi_R)\big\} \,\Big|\, n_R\Big]\, dF_t(n_R), \tag{92}$$

and let $\mathcal{E}_t$ denote the flow of researchers removed from Γ_t^R at t because there are no longer gains from trade. There are two kinds of PhDs outside research: those who did not enter research at graduation and those who exited after entering, so

$$\dot{X}_t = \big(\Gamma_t^G - e_t\big) + \mathcal{E}_t - \delta X_t, \qquad \dot{R}_t = e_t - \mathcal{E}_t - \delta R_t, \tag{93}$$

and the population constraint is $L_t^I + L_t^E + N_t + R_t + X_t = 1$. Production labor in efficiency units is

$$L_t = z_P\big(\xi_1 L_t^I + L_t^E + \xi_2 X_t\big), \tag{94}$$

so that $Y_t = Q_t L_t$ as in Equation (20) and the labor income of production workers in the government budget constraint is $w_t L_t$. Relative to the baseline, taste changes e_t, through the exit probability at graduation, and introduces $\mathcal{E}_t$, which is zero along the balanced growth path and nonzero along a transition only when the viable set contracts. Along the balanced growth path,

$$e = \Gamma^G \int \big(1 - q(z_R)\big)\, f^G(z_R)\, dz_R, \qquad X = \frac{\Gamma^G - e}{\delta}, \qquad R = \frac{e}{\delta}. \tag{95}$$

where f^G is the density of ability among graduates and $q(z_R)$ is the probability that a graduate of ability z_R does not enter research, as derived in Equation (96) below.

H.1.2 Measurement

Along the balanced growth path, since $z_t^*(\psi_R)$ is strictly decreasing, it has an inverse $\bar{\psi}_t(z_R)$, strictly decreasing in z_R, and a graduate of ability z_R enters research if and only if $\psi_R \geq \bar{\psi}_t(z_R)$. Define the exit probability conditional on ability among graduates, $q_t(z_R) = \Pr\big(\psi_R < \bar{\psi}_t(z_R) \mid z_R\big)$. Because taste is conditionally normal given ability (Equation (76)), this probability has a closed form for any ω:

$$q_t(z_R) = \Phi\left(\frac{\ln \bar{\psi}_t(z_R) - m_\psi(z_R)}{\sigma_\psi\sqrt{1 - \omega^2}}\right). \tag{96}$$

When $\omega = 0$, $m_\psi(z_R) = \mu_\psi$, and the exit probability depends on ability only through the boundary $\bar{\psi}_t(z_R)$.