Allocating funding across research fields: Evidence from the National Institutes of Health *

François Bares
Saïd Business School,
University of Oxford

June 13, 2026

Abstract

This paper studies the allocation of a fixed budget across scientific fields. I estimate the marginal return of funding for each NIH institute by developing a micro-founded model of scientists' knowledge production. Scientists respond to changes in funding allocation by choosing where to seek funding for their ideas. I estimate the model using administrative data from the National Institutes of Health (NIH) and publication data from OpenAlex covering 2005–2018. I find substantial heterogeneity in marginal returns to funding across 19 research fields—from 5.3 additional impact-factor-weighted publications per $1 million in Nursing Research to 64.0 in Human Genome. Some fields produce more knowledge spillovers than others, helping explain this discrepancy. Scientists are responsive to changes in the funding distribution and do relocate across fields—a budget increase in one field draws scientists away from other fields—but the cost of this reallocation is negligible in equilibrium. Contrary to what is often recommended in the medical literature, distributing funding across fields proportionally to their disease burden is suboptimal. Even if the sole goal of the allocation is to minimise disease burden, an allocation that takes account of knowledge externalities and productive efficiency dominates one based purely on disease burden.

Keywords: Economics of Science, Public Research, Research Policy, Public Research Grants, Sectoral Planning

JEL Codes: O30, O32, O38, O51, O25, H4, E13, L52

*I thank Thomas Hellmann and Kyle Myers for their continuous support and feedback. I thank Seung Joo Lee and Joel Shapiro for their detailed feedback on earlier drafts. I thank the participants in seminars at the University of Oxford for their comments and suggestions.

1 Introduction

Scientific research is a central input into technological change and economic growth. Yet scientific progress appears to be slowing (Bloom et al., 2020; Jones, 2009; Jones, 2010; Park, Leahey, and Funk, 2023). At the same time, firms have retreated from upstream basic science and rely more on publicly funded science instead to support their downstream innovation (Arora, Belenzon, and Patacconi, 2018; Fleming et al., 2019; Poege et al., 2019). This leaves governments with a larger role in sustaining the scientific base just as public science budgets have tightened in many countries [1]. In this context, it is essential that governments allocate their (limited) public science budget efficiently.

A key allocation problem for every large funder of science and technology is how to distribute a fixed budget across competing lines of research. In the case of the National Institute of Health (NIH), the world's largest biomedical research funder, Congress is responsible for doing so. Each year, it determines the NIH budget and its distribution across NIH institutes (ICs). Each of these ICs is devoted to a particular body system or class of disease, such as the National Eye Institute or the National Cancer Institute. The budget split among these institutes is contentious and the subject of much lobbying (Sampat, 2012). In particular, critics argue that NIH funding is only loosely correlated with disease burden (Gross, Anderson, and Powe, 1999), and funding should instead concentrates more in institutes tackling the most burdensome disease. The NIH often responds that science developed with one disease in mind might be central to developing cures for other diseases. These long running debates lack an essential input: estimates of the efficiency with which each IC turns additional funds into scientific output. Such estimates would allow to compare institutes not simply on the social burden of the disease they focus on, but also on the cost of making further scientific progress. If these estimates include knowledge externalities, they can also capture how basic upstream research may lead to improvement in downstream applied research. This central goal of this paper is to provide such estimates by measuring the marginal return to NIH funding in each IC, taking knowledge externalities into account.

To be useful for policy evaluations and counterfactuals, these marginal-return estimates must be anchored in the preferences and behaviour of scientists. Scientists may respond strategically to funding allocations and modify the research they pursue when funding conditions change. Second, research output in one field may be an input into later research in other fields. Motivated by these challenges, I estimate a tractable micro-founded equilibrium model of scientists' field and production choices. The model studies

[1] In the United States, the proposed 2026 budget included a 40% cut to the National Institutes of Health (NIH). In Europe, changes to European Research Council grant allocation rules in 2014 reduced biomedical research funding by 25% (Adda and Ottaviani, 2023), while in the UK, the 2025–2026 UK Research and Innovation budget fell by roughly £300 million.

1) how additional funds affect output in each IC, 2) how scientists respond to changes in the funding landscape, and 3) how research in one field creates knowledge externalities for itself and for other fields. The model provides a framework to evaluate existing funding allocations and simulate counterfactual policies to identify more efficient allocations.

I estimate the model using administrative data from the NIH covering 2002–2023. The key findings are threefold. First, marginal returns vary widely across fields: an additional $1 million in funding yields anywhere from 5.3 extra impact-factor-weighted publications in Nursing Research to 64.0 in Human Genome. Second, the general-equilibrium labour reallocation effects are small: scientists do relocate towards better-funded fields, but those who replace them are equally productive once funded, so the net cost of reallocation is negligible. Third, knowledge spillovers are a key differentiator of highest-return from lowest-return institutes. Upstream institutes whose research feeds the knowledge inputs of many other fields—Human Genome, General Medicine, and Biomedical Imaging and Bioengineering—generate large cross-field spillovers; and not accounting for these spillovers would significantly underestimate their impact. The knowledge spillovers generated by one field for others is These cross-field spillovers exceed the own-field spillover for 17 of 19 institutes. Specialised institutes with narrower citation footprints—Nursing Research and Aging—sit at the bottom of the return distribution on every metric.

Existing discussions of funding allocation across NIH institutes typically argue that institutes addressing diseases with a greater social burden should receive proportionally more funds[2] (Ballreich et al., 2021; Bisias, Lo, and Watkins, 2012; Gillum et al., 2011; Gross, Anderson, and Powe, 1999; Sampat, Buterbaugh, and Perl, 2013). These studies implicitly assume that productivity is constant across fields and focus instead on differences in the social value of research—is curing cancer more pressing than preventing infectious disease? Instead I focus on field-level differences in the marginal return to funding: the rate at which additional dollars translate into research output. My estimates can be used to evaluate allocative efficiency under any arbitrary social objectives. I explore three such objectives: maximising the academic impact of the research produced, maximising the commercial potential of the research produced, and minimising disease burden. I find that the current allocation is inefficient under all three objectives. Moreover, I find that even if the sole objective is to minimise disease burden, allocating funds in proportion to disease burden is suboptimal and is dominated by an allocation that accounts for both productivity and knowledge spillovers.

Unlike in the medical literature, this allocation problem has been largely ignored in the economics of science and innovation literature. Perhaps because this question sits at the middle of two distinct programmes that feature most of the research on the public

[2]In addition to disease burden, concerns over racial (Taffe and Gilpin, 2021), gender (Mirin, 2021), and age (Lauer and Roychowdhury, 2021; Qiu, 2023) inequities also feature.

funding of science and technology. The first is concerned the design and effectiveness of public scientific funding at the micro level. For example, Boudreau et al., 2016, Lanei et al., 2022, and Li, 2017 show how study-section composition shapes which proposals are funded; Adda and Ottaviani, 2023 and Carson, Zivin, and Shrader, 2023 highlight the role of grading schemes; Howell et al., 2023 and Myers, 2020 analyse the trade-off between untargeted and targeted grants; Azoulay, Graff Zivin, and Manso, 2011 compares person- and project-based awards; and Krieger, Myers, and Stern, 2023 and Packalen and Bhattacharya, 2018 document evaluator biases. Together, these studies ask how to allocate funds among competing individual projects. By contrast, this paper takes a macro perspective and focuses on allocating funds across broad lines of research.

The other programme is concerned with funding efficiency of science and technology for the economy as a whole. In particular, this paper adds to the growing literature estimating the return on public R&D funding (Akcigit, Hanley, and Serrano-Velarde, 2021; Azoulay et al., 2019; Babina et al., 2022; Dyevre, n.d.; Myers and Lanahan, 2022; Yin et al., 2022). While these studies focus on the return on public R&D funding in general and sometimes account for heterogeneity among scientists (Bertolotti, Myers, and Tham, n.d.), I focus on the cross-field differences in the return to funding. In doing this I am perhaps closest to Liu and Ma, 2021, who study which industry and industrial policy should target, and argue that industries that generate more knowledge spillovers should be prioritised. While I also model heterogeneity in knowledge spillovers across fields, I account for an additional driver of heterogeneity: the friction scientists face when reallocating across fields. In contrast, Liu and Ma, 2021 assume a frictionless process of reallocation of scientists across fields.

The paper is structured as follows. Section 2 defines the allocation problem and links it to the NIH institute structure. Section 3 presents the model of scientist sorting, grant evaluation, and research production. Section 4 derives the comparative statics linking institute budgets to output within and across fields. Section 5 describes the sample, empirical objects, identification, and estimation. Section 6 reports model fit, production-function estimates, marginal returns, and mechanisms. Section 7 documents the budget variation that motivates a local exercise and evaluates allocation rules. Section 8 concludes.

2 The Allocation Problem

2.1 Institutional Setting

The allocation problem studied here is the distribution of a fixed NIH budget across research fields. NIH funding is primarily disbursed through institutes and centres (ICs), each focused on a body system, disease class, or broad research mission (Table 12). I define a *field* as an NIH institute and use the terms "field", "IC", and "institute" interchangeably. Congress sets both the aggregate NIH budget and IC-level allocations, informed by NIH priorities and recommendations. This institutional set-up makes the individual IC the policy-relevant way to categorise research fields for the purpose of this paper. Below I formalise the allocation problem this paper tackles and show why the key object to estimate is the marginal return to funding across fields.

2.2 The Social Planner's Problem

NIH institutes are indexed by $d \in \{1, \ldots, C\}$. In period t, the planner allocates a finite total research budget $\bar{G}_t > 0$ across institutes.

$$\mathbf{G}_t \equiv (G_{1,t}, \ldots, G_{C,t}) \in \mathbb{R}_+^C \quad \text{and} \quad \sum_{d=1}^{C} G_{d,t} = \bar{G}_t.$$

Let $\mathbf{K}_t$ denote the vector of depreciated knowledge stocks accessible at time t, summarising past scientific output through a depreciation operator:

$$\mathbf{K}_t = \mathcal{K}(\mathbf{K}_{t-1}, \mathbf{Y}_t),$$

Aggregate scientific output $\mathbf{Y}_t$ is determined in equilibrium from the mapping:

$$\mathbb{E}[\mathbf{Y}_t] = \mathbb{E}[\mathbf{F}(\mathbf{G}_t, \mathbf{K}_t)], \quad \text{where} \quad \mathbb{E}[\mathbf{F}(\cdot, \mathbf{K}_t)] : \left\{ \mathbf{G} \in \mathbb{R}_+^C : \sum_{d=1}^{C} G_d = \bar{G}_t \right\} \to \mathbb{R}_+^C.$$

For brevity, I denote this budget set by $\Delta(\bar{G}_t)$. Here $\mathbb{E}[F_h(\mathbf{G}_t, \mathbf{K}_t)]$ denotes the expected output of institute h at t given funding $\mathbf{G}_t$ and knowledge state $\mathbf{K}_t$. The planner is risk-neutral and evaluates outputs using a continuously differentiable social welfare function:

$$U : \mathbb{R}_+^C \to \mathbb{R}, \qquad U(\mathbb{E}[\mathbf{Y}_t]) \equiv U(\mathbb{E}[Y_{1,t}], \ldots, \mathbb{E}[Y_{C,t}]),$$

For a given state $\mathbf{K}_t$, the planner chooses $\mathbf{G}_t$ to maximise welfare:

$$\max_{\mathbf{G}_t \in \Delta(\bar{G}_t)} \quad U(\mathbb{E}[\mathbf{F}(\mathbf{G}_t, \mathbf{K}_t)]) \,. \tag{1}$$

The first-order conditions at an optimum $(\mathbf{G}_t^\star, \lambda^\star)$ require that the marginal welfare impact of funding be equalised across all institutes[3]:

$$\underbrace{\frac{d}{dG_{d,t}} U(\mathbb{E}[\mathbf{F}(\mathbf{G}_t, \mathbf{K}_t)])}_{\text{total marginal welfare effect of funding field } d} = \lambda, \qquad \forall d. \tag{2}$$

Combined with the budget constraint $\sum_{d=1}^{C} G_{d,t} = \bar{G}_t$, this system determines the optimal allocation.

Expanding the total derivative via the chain rule:[4]

$$\frac{d}{dG_{d,t}} U(\mathbb{E}[\mathbf{F}(\mathbf{G}_t, \mathbf{K}_t)]) = \sum_{h=1}^{C} w_h \cdot \frac{d\mathbb{E}[F_h(\mathbf{G}_t, \mathbf{K}_t)]}{dG_{d,t}}. \tag{3}$$

where $w_h \equiv \frac{\partial U(\mathbb{E}[\mathbf{Y}])}{\partial \mathbb{E}[Y_h]}$ is the social welfare weight of institute h's expected output, which can be chosen to reflect any social objective, such as prioritising certain diseases or populations.

The objective of this paper is to estimate $\frac{d\mathbb{E}[F_h(\mathbf{G}_t, \mathbf{K}_t)]}{dG_{d,t}}$ for all pairs of ICs (h, d), irrespective of the social welfare weights one might choose.

2.3 Two Mechanisms Behind Cross-Field Effects

The equilibrium mapping $\mathbb{E}[\mathbf{F}(\mathbf{G}_t, \mathbf{K}_t)]$ encodes the aggregate effect of a budget allocation without specifying the underlying behaviour of individual agents. I outline below the two mechanisms that give this mapping its shape and determine the sign and magnitude of $\frac{d\mathbb{E}[F_h]}{dG_{d,t}}$.

Mechanism 1: Scientist reallocation across fields. Research funding operates primarily by determining the number and quality of research projects undertaken in each field. When a budget shock raises $G_{d,t}$, it expands the number of grants available in institute d, thereby funding additional research projects. Because the pool of biomedical scientists is not perfectly elastic in the short run, this expansion draws talent away from other fields. Scientists at the margin—those whose comparative advantage does not strongly favour one field over another—pursue research in the better-funded field rather than in their existing fields.

[3] Assuming that the objective function is continuous and the budget set $\Delta(\bar{G}_t)$ is compact, a global maximum $\mathbf{G}_t^\star$ exists. I restrict attention to interior allocations where $G_{d,t} > 0$ for all d. In principle, corner solutions could dominate any interior solution. But a corner solution involves closing down an NIH institute, which is unlikely to be politically feasible.

[4] The FOC uses a total derivative because changes in $G_{d,t}$ affect equilibrium objects—grant application success rates, applicant sorting—and thereby all components of $\mathbb{E}[\mathbf{Y}]$, not only $\mathbb{E}[F_{d,t}]$.

This reallocation has two consequences. First, output in field d rises directly because more scientists work there. Second, output in fields $h \neq d$ may fall because those fields lose scientists who would otherwise have contributed to them. The magnitude of this cross-field effect depends on how substitutable scientific labour is across fields.

Investigator histories show that such reallocation is frequent but varies across fields. Figure 1 reports the share of PIs whose current funding institute differs from their previous funder. Broad institutes have higher mover shares than institutes tied to more specific body systems: Biomedical Imaging and Bioengineering and Human Genome record mover shares of 78% and 73%, while Eye and Deafness and Communication Disorders record mover shares of 40% and 43%.

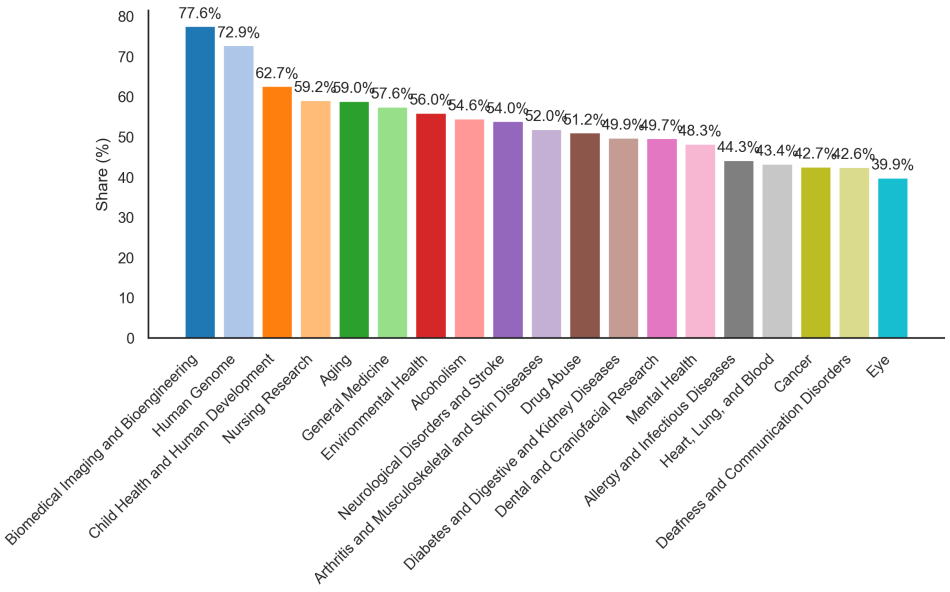


Figure 1: Share of PIs who moved between funding institutes, by destination institute.
Share of PIs whose current funding institute differs from their previous funder. Institutes whose research has broad applicability tend to have higher mover shares.

Mechanism 2: Knowledge spillovers across fields. The second mechanism operates through the knowledge stock $\mathbf{K}_t$. Scientific output is not only a final product—it is also an input into future research. A discovery in field d (e.g. a new molecular pathway identified by diabetes researchers) may spur further discoveries in other fields (e.g. cancer research, if the pathway is implicated in oncogenesis). Formally, an increase in $Y_{d,t}$ raises $\mathbf{K}_{t+1}$, which increases scientific output in other fields.

Citation patterns suggest that knowledge spillovers occur both within and across fields. Figure 2 reports the share of references made by grants in each citing field to publications associated with each cited field. The diagonal entries are large, which is consistent with large within-field knowledge accumulation. But some off-diagonal entries are also economically meaningful. For example, Alcoholism research draws heavily on Drug Abuse,

several fields cite Cancer and General Medicine, and Human Genome is cited by fields far beyond genetics. These citation patterns are asymmetric and display strong heterogeneity across fields: Some fields are more central to the biomedical research ecosystem and generate more spillovers than others.

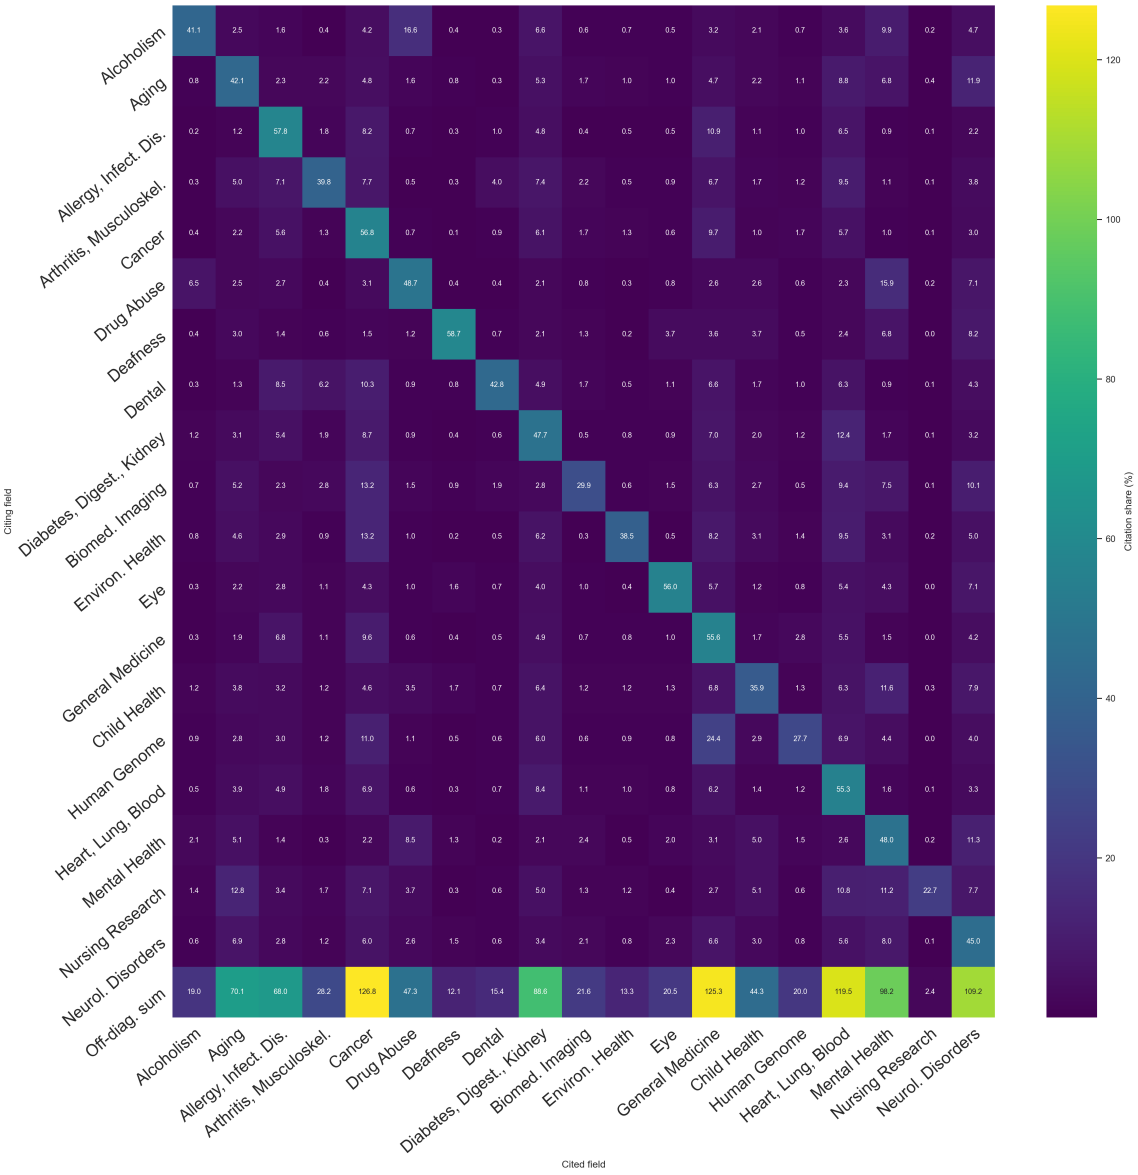


Figure 2: Citation shares across NIH fields.

Rows report the citing field and columns report the cited field. Each entry is the share of cited references in the row field that are attributed to the column field. Diagonal entries capture within-field knowledge use; off-diagonal entries capture cross-field knowledge exposure.

Together, these two mechanisms imply that the total derivative $\frac{d\mathbb{E}[F_h]}{dG_{d,t}}$ cannot be read off from production function estimates for each field in isolation. It requires modelling (i) how scientists sort across fields in response to funding, and (ii) how knowledge produced in one field feeds into the productivity of others. Both dimensions introduce cross-field dependencies that the structural model in Section 3 attempts to capture.

3 Model

Notation conventions. I use one role per index throughout, as summarised in Table 1. Unless dynamics make it necessary, the year subscript t and the outcome superscript m are suppressed in the theory sections.

Symbol	Role
i	Scientist, grant, or grant-PI observation
o	Origin field
j	Generic field index of summation
d	Destination field
h	Output or receiving field
f	Knowledge-source field
m	Outcome metric
t	Year

Table 1: Index conventions used throughout the model.

Consider an economy where a fixed number of scientists N work on research projects in C different fields, and can move from one field of origin o to a destination field d. There are N_o scientists from each origin field $o \in \{1, \ldots, C+1\}$[5].

3.1 Funding Rules and the Scientist's Problem

Each period, a scientist i from field o draws a research idea with quality a_{iod} for every field d (including her own field $o = d$). Scientist i can choose to apply for a grant in field d in order to pursue idea a_{iod}. She can apply to only one of the C funding destinations d, or she can choose not to apply. For example, we would classify a scientist whose previous grant was funded by the National Cancer Institute as a scientist with origin $o =$ Cancer. This cancer scientist may have research ideas relevant to cancer but also to any other fields. Suppose that she has a research idea that is most relevant to the National Institute of Ageing and sufficiently good that she decides to apply for funding to pursue this particular idea. Then her destination field d is classified as Ageing.

Each destination d has an exogenously given budget G_d and dispenses $G_d/\bar{g}_d$ grants, where $\bar{g}_d$ is the average grant size in field d. The grant size g_{iod} necessary to pursue idea a_{iod} is treated as exogenous to both the decision to apply and the decision to fund.[6]

Idea quality a_{iod} is private to scientist i. Grant reviewers receive grant applications with quality a_iod and give them a score x_{iod}:

[5]The total number of scientists is not fixed from one year to the next. All scientists are assigned to a field of origin o, but one origin category represents new entrants, thus allowing for entry.

[6]When applying for R01 grants, scientists must justify any funds they need in detail. Moreover, grant reviewers rank the scientific merit of proposals without necessarily taking grant size into account. Modelling grant size as a choice variable for either party would therefore be unrealistic.

$$x_{iod} = a_{iod} + \eta\,\nu_{iod}, \tag{4}$$

Grant reviewers are noisy evaluators of idea quality, and the evaluation noise $\nu_{iod} \sim \mathcal{N}(0,1)$ and $\eta > 0$ is the degree of noise in the application process.

A grant application is successful if and only if its score exceeds a threshold $\bar{x}_d$, known as the payline. The probability that an application with quality a_{iod} receives funding is therefore:

$$\Pr(x_{iod} > \bar{x}_d \mid a_{iod}, d) = \Phi\left(\frac{a_{iod} - \bar{x}_d}{\eta}\right) \tag{5}$$

Following a standard discrete choice model (McFadden, 1974), I assume that idea quality a_{iod} can be decomposed into a deterministic component—the origin-destination-specific average productivity μ_{od}—and an additive idiosyncratic draw ε_{iod}:

$$a_{iod} = \mu_{od} + \varepsilon_{iod}, \tag{6}$$

where ε_{iod} are independently drawn from a Type-I Extreme Value (Gumbel) distribution with scale parameter $\sigma > 0$.

Scientists know their idea quality draw a_{iod} and the destination paylines $\bar{x}_d$, but they do not know the exact score x_{iod} they will receive. Each scientist chooses the destination d^* that maximises her probability of being funded:

$$d^* = \arg\max_{d \in \{0,1,\ldots,C\}} \Pr(\text{funded} \mid a_{iod}, d), \tag{7}$$

where the outside option $d = 0$ yields a payoff of c. Because η is common across all origin-destination pairs,[7] the funding probability $\Phi((a_{iod} - \bar{x}_d)/\eta)$ is strictly increasing in the quality gap $a_{iod} - \bar{x}_d$. Maximising the probability of funding is therefore equivalent to maximising this gap.

Because the error terms follow a Type-I Extreme Value distribution, the probability Π_{od} that a scientist from origin o applies to field d takes a multinomial logit form:

$$\Pi_{od}(\bar{x}) \;=\; \frac{\exp\left(\frac{\mu_{od} - \bar{x}_d}{\sigma}\right)}{\exp\left(\frac{c}{\sigma}\right) + \sum_{j=1}^{C} \exp\left(\frac{\mu_{oj} - \bar{x}_j}{\sigma}\right)}. \tag{8}$$

[7] In estimation I relax this assumption and allow evaluation noise to increase with the distance between origin and destination fields. See Section 5.

The total mass of applications to field d is $A_d = \sum_{o=1}^{C} N_o \Pi_{od}(\bar{x})$, where $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_C)$ is the vector of paylines.

3.2 Equilibrium

Conditional on choosing field d, the ex ante probability of receiving funding depends on the distribution of the realised quality gap among applicants who selected destination d. Let $Z_{io} \equiv a_{iod^*} - \bar{x}_{d^*}$ denote the utility-maximising quality gap for scientist i from origin o. Let $f(z \mid \text{apply to } d)$ denote the probability density function of Z_{io} (see Appendix B for the explicit derivation based on the Gumbel structure). The expected success rate for applicants from o to d is obtained by integrating the conditional funding probability $\Phi(z/\eta)$ over this distribution:

$$S_{od}(\bar{x}) \;=\; \int_{-\infty}^{\infty} \Phi\left(\frac{z}{\eta}\right) f(z \mid \text{apply to } d)\, dz. \tag{9}$$

In equilibrium, paylines $\bar{x} = (\bar{x}_1, \ldots, \bar{x}_C)$ are such that each field fully disburses its available grants:

$$\sum_{o=1}^{C} \underbrace{N_o\, \Pi_{od}(\bar{x})}_{\substack{\text{Applications} \\ \text{from } o \text{ to } d}} \underbrace{S_{od}(\bar{x})}_{\substack{\text{Success rate of applicants} \\ \text{from } o \text{ to } d}} \;=\; \underbrace{\frac{G_d}{\bar{g}_d}}_{\text{Available grants}} \qquad \text{for all } d = 1, \ldots, C. \tag{10}$$

Equation (10) is akin to a market clearing condition where the payline $\bar{x}_d$ plays the role of a prize vector that endogenously adjusts such that the supply of funding meets demand in every fields. If d receives a budget increase (G_d rises), it must lower its threshold $\bar{x}_d$ to exhaust its funding. This raises the attractiveness of d and draws applicants away from other fields. The resulting fall in applicant volume elsewhere pushes other destinations to lower their own thresholds. Appendix C provides a proof of existence and uniqueness for the equilibrium payline vector $\bar{x}$.

3.3 Production Function and Aggregate Output

Upon receiving funding, scientists generate research output:

$$y_{iod} \;=\; a_{iod}\, k_{iod}^{\alpha_d} g_{iod}^{\beta_d} + u_{iod}, \tag{11}$$

where a_{iod} is idea quality of the research project, g_{iod} is the specific funding amount received by the project, k_{iod} is the knowledge input going into the project, α_d is the destination-specific knowledge elasticity, β_d is the destination-specific funding elasticity, and u_{iod} is a mean-zero stochastic shock representing the fundamental uncertainty of

scientific research[8].

Taking expectations over idea quality a_{iod} gives the expected production of scientists from origin o in destination d:

$$Y_{od} \;=\; \underbrace{N_o\,\Pi_{od}(\bar{x})}_{\substack{\text{Application volume}\\ \text{from } o \text{ to } d}} \;\; \mathbb{E}[k_{iod}^{\alpha_d} g_{iod}^{\beta_d}] \int_{-\infty}^{\infty} \underbrace{(z+\bar{x}_d)}_{a_{iod}} \underbrace{\Phi\left(\frac{z}{\eta}\right) f(z \mid \text{apply to } d)\, dz}_{\text{Funded component}}. \tag{12}$$

Let $W_{od}(\bar{x})$ denote the expected idea quality of a scientist from o conditional on both choosing destination d as their utility-maximising field and actually receiving funding:

$$W_{od}(\bar{x}) \;\equiv\; \mathbb{E}[a_{iod} \mid \text{chose } d \text{ and funded}] \;=\; \frac{1}{S_{od}(\bar{x})} \int_{-\infty}^{\infty} (z+\bar{x}_d)\Phi\left(\frac{z}{\eta}\right) f(z \mid \text{apply to } d)\, dz. \tag{13}$$

Total expected production in destination d is then the sum of applications, multiplied by the success rate, expected conditional quality, and expected function of inputs:

$$Y_d(\bar{x}) \;=\; \sum_{o=1}^{C} N_o\,\Pi_{od}(\bar{x})\, S_{od}(\bar{x})\, W_{od}(\bar{x})\, \mathbb{E}[k_{iod}^{\alpha_d} g_{iod}^{\beta_d}]. \tag{14}$$

I assume that the Cobb-Douglas input term $k_{iod}^{\alpha_d} g_{iod}^{\beta_d}$ is mean-independent of idea quality a_{iod}. The assumption reflects the evaluation process of R01 grants, where the scientific merit of a project is assessed independently of grant size.

[8]While an additive functional form theoretically allows negative output y_{iod} due to the infinite support of the Gumbel-distributed a_{iod}, empirically, with sensible parameter ranges, evaluating the integrals yields strictly positive expected production.

4 The marginal return to research funding

Throughout this section, all objects are evaluated in a fixed year for a fixed outcome metric, so year subscripts t and outcome superscripts m are suppressed except where intertemporal dynamics make them necessary.

Now that I have described the model and characterised its equilibrium, I characterise the main object of interest: the marginal return to research funding for each field.

How does a change in funding G_d in destination d affect scientific production in d and in other destinations $h \neq d$? The mechanism operates through endogenous paylines and applicant substitution. Consider an increase in funding G_d. To allocate the additional funding, d must lower its payline $\bar{x}_d$. This raises the success probability of grant applications to field d, making d more attractive to applicants, and changes both the number and the average quality of projects funded in d. Other fields h are affected indirectly. As applicants substitute towards d, the number of applications A_h submitted to any other field $h \neq d$ falls. Since field h still has to disburse the fixed budget G_h, it must lower its own payline $\bar{x}_h$ to clear the market. This raises the success rate in h, but it may also lower the average quality of funded projects. The net effect of a change in G_d therefore combines the own-field response in d with cross-field responses in all other destinations.

In equilibrium, the effect of a change in funding G_d on output Y_h (where h can be equal to d or different from d) is given by the total derivative:

$$\frac{dY_h}{dG_d} = \sum_{j=1}^{C} \underbrace{\frac{\partial Y_h}{\partial \bar{x}_j}}_{\substack{\text{Partial effect} \\ \text{of payline } j}} \underbrace{\frac{d\bar{x}_j}{dG_d}}_{\substack{\text{GE payline} \\ \text{adjustment}}} , \forall h, d. \tag{15}$$

This is the key focus of this paper[9].

4.1 General-Equilibrium Payline Adjustments

The general-equilibrium payline adjustments $\frac{d\bar{x}_j}{dG_d}$ can be obtained by differentiating the market-clearing condition (Equation 10) with respect to G_d for every field $k \in \{1, \ldots, C\}$.

[9]Production Y_h is a reduced-form object $Y_h(\bar{x}(G))$: funding does not enter production directly, but through the equilibrium payline vector induced by the full system of grant markets. The derivative with respect to G_d is therefore a *total* derivative. A partial derivative such as $\partial Y_h / \partial G_d$ would hold the equilibrium object $\bar{x}$ fixed and would therefore shut down the market-clearing response that defines the general-equilibrium effect.

This yields a system of C linear equations:

$$\sum_{j=1}^{C} \underbrace{\frac{\partial E_k}{\partial \bar{x}_j}}_{\substack{\text{Demand shift} \\ \text{in } k \text{ from } \bar{x}_j}} \frac{d\bar{x}_j}{dG_d} = \begin{cases} \frac{1}{\bar{g}_d} & \text{if } k = d, \\ 0 & \text{if } k \neq d, \end{cases} \tag{16}$$

where $E_k(\bar{x}) = \sum_{o=1}^{C} N_o \Pi_{ok}(\bar{x}) S_{ok}(M_o)$ is the expected number of funded grants in field k. When G_d changes, the right-hand side of the dth market-clearing equation shifts directly, but all C paylines generally move because application probabilities and success rates are linked across destinations through the application choices of scientists. The derivative $d\bar{x}/dG_d$ is therefore obtained by differentiating the entire equilibrium system, not by looking only at destination d in isolation.

4.2 Output Channels and Sign Ambiguity

The term $\frac{\partial Y_h}{\partial \bar{x}_j}$ is the local output response in destination h to a shift in payline j. It has two channels. First, a payline shift changes application choices across destinations. Second, it changes the expected quality of the projects funded in h:

$$\frac{\partial Y_h}{\partial \bar{x}_j} = \mathbb{E}[f(g_{ioh}, k_h)] \sum_{o=1}^{C} N_o \left[\underbrace{\frac{\partial \Pi_{oh}}{\partial \bar{x}_j}}_{\text{Volume shift}} \underbrace{W_{oh}}_{\text{Base quality}} + \underbrace{\Pi_{oh}}_{\text{Base volume}} \underbrace{\frac{\partial W_{oh}}{\partial \bar{x}_j}}_{\text{Quality shift}} \right], \tag{17}$$

where W_{oh} denotes the success-weighted expected quality contribution of a scientist from origin o applying to destination h, defined as:

$$W_{oh}(M_o, \bar{x}_h) = \int_{-\infty}^{\infty} (z + \bar{x}_h)\, \Phi\left(\frac{z}{\eta_{oh}}\right) f_Z(z; M_o)\, dz. \tag{18}$$

This term integrates over the chosen quality gap distribution $f_Z(z; M_o)$ and the probability $\Phi(z/\eta_{oh})$ of overcoming evaluation noise to secure funding, so the success-rate weighting is already included in W_{oh}. A step-by-step derivation is provided in Appendix D.

Equation (17) highlights the two channels through which funding changes affect output. An increase in funding allows to fund additional projects which increases output (volume effect). But the increase in funding also lowers the payline which means that the additional funded projects may be of lower quality compared to the baseline (quality effect). The quality effect may dampen but never dominate the volume effect, as the newly funded projects all have positive output.

4.3 Knowledge Spillovers

I extend the model to allow for knowledge spillovers within and across institutes. Knowledge production is cumulative: new discoveries build upon the existing stock of ideas, which depreciate over time (Griliches, 1979).

The knowledge stock of field f at time t accumulates past discoveries according to:

$$K_{f,t} \;=\; \sum_{s=0}^{\infty}(1-\delta_f)^s Y_{f,t-s}. \tag{19}$$

Knowledge stock for each field then enters the production function of individual grants through a grant-specific knowledge index k_{it}, defined as a Cobb–Douglas aggregator of field-specific knowledge stocks weighted by the grant's own citation profile:

$$k_{it} \;=\; \prod_{f=1}^{C} K_{f,t}^{\gamma_{fi}}, \tag{20}$$

where $\gamma_{fi} \geq 0$ is the citation share of grant i's references attributable to field f. By using citation shares as proxy for the elasticity of grant knowledge input with respect to field-level knowledge stocks, we follow Acemoglu, Akcigit, and Kerr (2016) and Liu and Ma (2021).

k_{it} captures the knowledge spillover channel. Let $Y_{h,t} = \sum_{i\in h} y_{ioh,t}$ be aggregate output in field h. I define the exposure of field h to knowledge from field d as the output-weighted average citation share:

$$\Gamma_{dh,t} \;=\; \sum_{i\in h}\omega_{i|h,t}\,\gamma_{di}, \qquad \omega_{i|h,t} \equiv \frac{y_{ioh,t}}{Y_{h,t}}. \tag{21}$$

If the production function in field h has elasticity α_h with respect to knowledge input k_{it}, then the Cobb–Douglas index implies

$$\frac{\partial \ln Y_{h,t}}{\partial \ln K_{d,t}} \;=\; \sum_{i\in h}\omega_{i|h,t}\,\alpha_h\gamma_{di} \;=\; \alpha_h\Gamma_{dh,t}. \tag{22}$$

An additional unit of expected output in field d raises the stock K_d and remains useful with geometric weight $(1-\delta_d)^s$ after s periods. Assuming that $\Gamma_{dh,t+s}$ and $\mathbb{E}[Y_{h,t+s}]/K_{d,t+s}$ are constant over time allows the series converges to $1/\delta_d$. The cumulative knowledge effect becomes:

$$\frac{d\sum_{s=0}^{\infty}\mathbb{E}[Y_{h,t+s}]}{d\mathbb{E}[Y_{d,t}]} \;=\; \frac{\alpha_h\Gamma_{dh,t}}{\delta_d}\frac{\mathbb{E}[Y_{h,t}]}{K_{d,t}}. \tag{23}$$

This gives the marginal return to funding that we have been looking for. Consistent with the allocation problem in Section 2, the object is the aggregate effect of funding field d on expected output across fields:

$$MR_{d,t} \equiv \sum_{h=1}^{C} \frac{d\mathbb{E}[F_h(\mathbf{G}_t, \mathbf{K}_t)]}{dG_{d,t}}. \tag{24}$$

Using the current-period total derivatives from Equation (15), this marginal return decomposes as

$$\begin{aligned} MR_{d,t} = & \underbrace{\frac{d\mathbb{E}[Y_{d,t}]}{dG_{d,t}}}_{\text{Own return}} \\ & + \underbrace{\sum_{h \neq d} \frac{d\mathbb{E}[Y_{h,t}]}{dG_{d,t}}}_{\text{Talent externality}} \\ & + \underbrace{\frac{d\mathbb{E}[Y_{d,t}]}{dG_{d,t}} \sum_{h=1}^{C} \frac{\alpha_h \Gamma_{dh,t}}{\delta_d} \frac{\mathbb{E}[Y_{h,t}]}{K_{d,t}}}_{\text{Knowledge externality}}. \end{aligned} \tag{25}$$

The first term is the output gained in the funded field. The second term is the effect on other fields through researcher reallocation. The third term is the cumulative output gain generated when the extra output in field d augments the knowledge stock used by all fields. Complete derivations of all comparative statics results are provided in Appendix D.

5 Empirical Strategy

This section describes how I take the model to the data. I construct a grant–investigator level panel, map each observation to an origin and destination institute, and use the resulting application numbers, cross field movement and application success-rate to estimate the structural parameters listed in Table 3.

5.1 Sample construction

Variable	Description	Source
$A_{d,t}$	Total number of applications received by institute d in year t	NIH Data Book
$\pi_{d,t}$	Aggregate success rate (funded ÷ applications) for institute d in year t	NIH Data Book
$G_{d,t}$	Total budget disbursed by institute d in year t	NIH R01 grant data
$\bar{g}_{d,t}$	Average grant size for institute d in year t	NIH Reporter
$F_{od,t}$	Number of funded grants from origin o to destination d in year t	Dimensions.ai investigator history
y_{iod}	Publication output of funded grants from o to d in year t	PubMed publications
$N_{o,t}$	Size of the potential applicant pool from origin o in year t (unique first authors of publications in year t funded by origin o grants)	Constructed from publication-grant linkages
D_{od}	Scientific distance between origin o and destination d (cosine)	Constructed from Study-section funding

Table 2: Data Variables and Their Sources

This table lists the key variables used in the empirical analysis, along with their descriptions and data sources. Origin and destination fields correspond to the $C = 19$ core NIH institutes.

The sample constructed here links public NIH grant records to publication metadata from OpenAlex and investigator histories from Dimensions.ai. The NIH Exporter provides per-grant metadata (activity codes, study-section identifiers, project identifiers) and a mapping from NIH project numbers to PubMed IDs; these mappings link grants to publications. Publication records are enriched with OpenAlex data (canonical IDs, journal names and ISSNs, citation counts, parsed authorship) and journal-level metrics (mean impact factor and h-index). Dimensions supplies investigator metadata (identifiers, affiliation, years of activity) and records of all grants received by a given investigator, which I use to identify prior funding and to expand the sample beyond the initial NIH records, including non-NIH grants.

The initial NIH sample was downloaded in December 2024 via the NIH Exporter API and is restricted to individual R01 (and R01-equivalent) awards. Grants not active as of December 2024, or with a start year after 2022, are excluded. The sample is restricted to the core NIH institutes listed in Table 12; I include only NIH institutes and exclude

smaller NIH centres (e.g., the Fogarty International Center) to improve comparability. To ensure the best institute coverage, the sample begins in 2001, as the National Institute of Biomedical Imaging and Bioengineering (NIBIB) was only established in 2000.

However, we exclude the National Institute on Minority Health and Health Disparities (NIMHD) because it was only established in 2010. For simplicity, I refer to each institute by a short research-field label rather than its official NIH name; the mapping between field labels and official institute names is provided in Table 12.

For projects funded by more than one IC, a project-level total is computed by summing per-IC contributions, and the project is assigned to the IC contributing the largest share of funding.

Dimensions is used to expand investigator-level records and to recover supporting or prior grants outside the NIH sample. When Dimensions lacks investigator metadata, the OpenAlex last-author identifier is used. Each observation in the final dataset is a unique investigator–grant pair. After cleaning, the sample contains 106,664 investigator–grant pairs, 53,175 unique investigators, and 85,290 unique grants. When a grant has multiple investigators, the award is split equally across them. All monetary values are deflated to 2017 constant dollars using the FRED GDP deflator.

Application success rates are taken from publicly available NIH Data Book tables, merged by institute and reference year. Publication stocks for knowledge-input calculations are computed from NIH grant–publication mappings: only publications linked to R01 (R01-equivalent) awards are used, and yearly stocks are computed beginning in 1980 so that the stock entering the sample window (2002) reflects prior knowledge production.

The estimation sample spans 2005–2018: this window excludes the pre-2005 NIH budget-doubling period and the post-2018 disruption attributable to the COVID-19 pandemic, although the underlying grant data cover 2001–2022.

The total number of scientist in field o at any given time $N_{o,t}$ is constructed as the number of unique first authors of papers published in year t that are linked to grants whose origin is o. This publication-based first-author definition avoids double-counting PIs who hold multiple grants and anchors the pool to researchers active at the margin of applying rather than to incumbent grant holders. This allows for many scientist never to secure a grant which is realistic[10].

Because the NIH does not publish institute-level success rates specifically for R01 grants, I proxy these with rates for all investigator-initiated mechanisms. This category includes

[10]In the small number of years where the implied aggregate pool $\sum_o N_{o,t}$ does not strictly exceed total applications A_t, I scale all origin counts $N_{o,t}$ up by a common yearly factor so that $\sum_o N_{o,t} > A_t$, preserving relative origin shares. The mass of non-applicants from origin o in year t is $N_{o,t} - \sum_d F_{od,t}/\pi_{d,t}$, where $F_{od,t}/\pi_{d,t}$ approximates the total number of applications from o to d.

R01s as well as smaller mechanisms (R03, R21) and mechanisms targeting under-served institutions (R15). The pool of non-R01 mechanisms typically exhibits a higher success rate than pure R01s, so the proxy likely overstates the true payline competitiveness. I address this by adjusting observed rates downward in a sensitivity analysis and find that the results are robust: What matters is the relative change in application success rates across institutes, not the levels.

5.2 Model change for estimation

Here I outline two modifications I made to the theoretical model presented in Section 3. The changes are designed to improve the realism of the model while preserving its core structure and tractability.

Unlike a standard location normalisation that fixes the outside-option utility at zero, I freely estimate the time-varying outside option $c_t = c_0 + c_1 t$. The time trend c_1 allows the model to accommodate secular changes in the overall propensity to apply—arising, for example, from changes in career incentives or non-NIH funding opportunities.

Secondly, I allow for the grant evaluation noise to depend on the scientific distance between origin and destination and express evaluation noise as

$$\eta_{od} = \eta_0 \exp(\gamma D_{od})$$

.

The distance-dependent evaluation noise with $\gamma \geq 0$ is economically motivated: grant reviewers in a given field may face greater difficulty evaluating proposals from distant fields. But the parametrization does not force field specific noise. If $\hat{\gamma} \approx 0$, the model then reduces to a common evaluation noise η_0 for all origin-destination pairs.

The scientific distance D_{od} is computed by emulating the technology-proximity measure of firm distance in patent space (Bloom, Schankerman, and Van Reenen, 2013), but instead of using overlap in patent portfolios between firms I use overlap in study-section funding between ICs. Let S_c be the vector of funding shares of institute c across different study sections in a given year. Specifically, we compute the distribution of funding that each institute allocates across all NIH study sections.

The overlap in these funding share distributions determines the scientific closeness between two institutes. The cosine similarity of these share vectors serves as an intuitive proximity measure: it equals 1 if two institutes distribute their funding identically across study sections, and 0 if they fund completely disjoint sets of study sections. Applying the cosine similarity to funding *shares* rather than raw volumes naturally accounts for differences in total budget sizes. The proximity between institutes o and d is the correlation

(cosine similarity) of their funding share vectors:

$$\rho_{od} = \frac{S_o \cdot S_d}{\|S_o\| \|S_d\|}. \tag{26}$$

Figure 3 displays the resulting average distance matrix over the 2005–2018 estimation window.

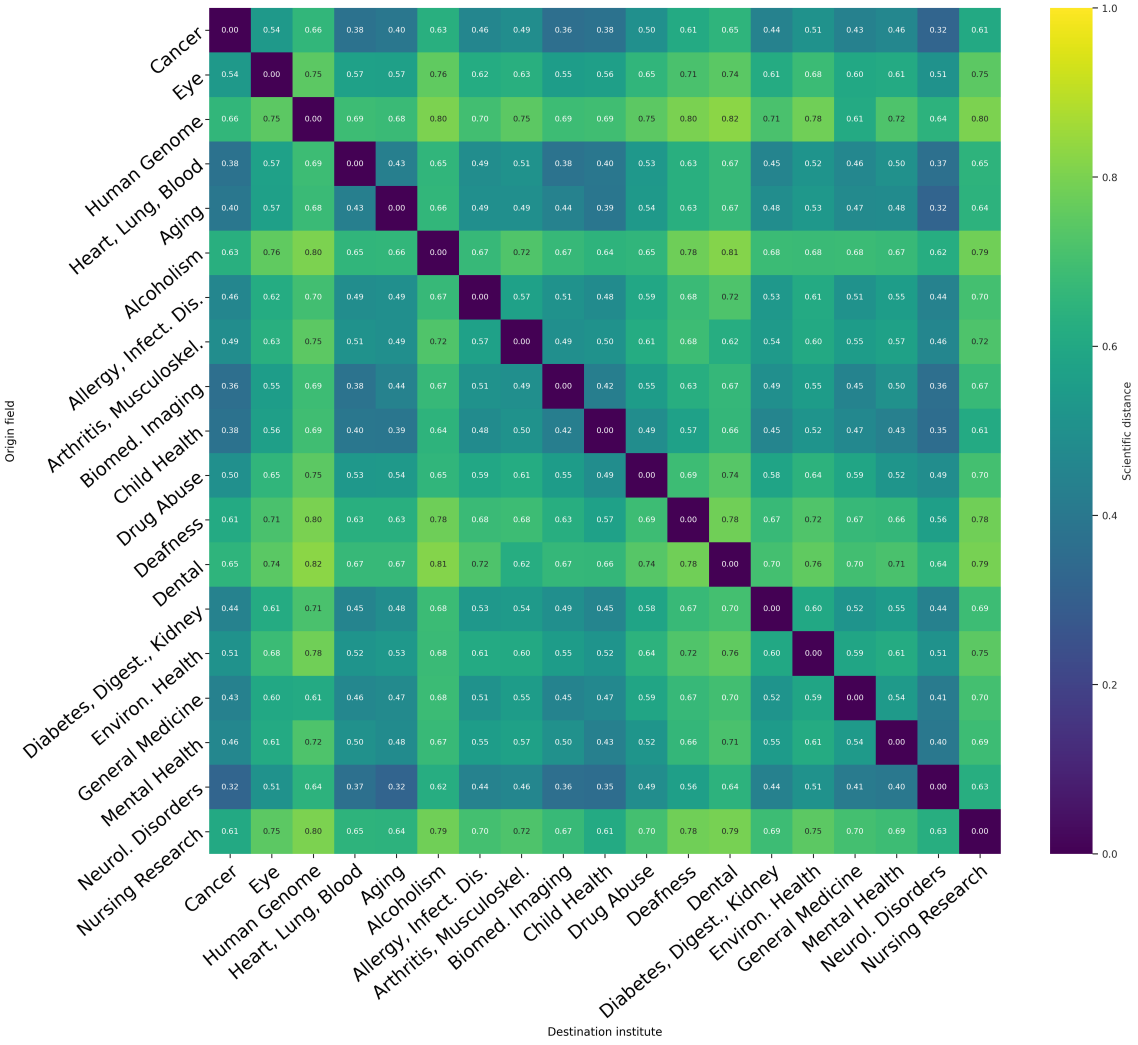


Figure 3: Scientific distance matrix, 2005–2018.

Each cell reports the average scientific distance between origin field o and destination institute d over the 2005–2018 estimation window. Larger values indicate greater scientific distance.

Parameter	Description
μ_{od}	Origin-destination mean idea quality; $C \times C$ matrix
σ	Scale parameter of the Gumbel distribution; $\sigma > 0$
c_0	Intercept of the outside-option utility
c_1	Time slope of the outside-option utility; $c_t = c_0 + c_1 \cdot t$
η_0	Baseline evaluation noise standard deviation (at zero scientific distance); $\eta_0 > 0$
γ	Distance gradient of evaluation noise; $\eta_{od} = \eta_0 \exp(\gamma D_{od})$, constrained $\gamma \geq 0$

Table 3: Structural Parameters of the Model

5.3 Moment Conditions

I match three blocks of model-implied moments to their observed analogues. All moment conditions are expressed as normalised (relative) deviations, placing observations with different scales on a common metric. Throughout, estimated model-implied quantities are written with hats and observed quantities without hats.

5.3.1 Aggregate Applications

The model predicts the total number of applications received by institute d in year t as:

$$\hat{A}_{d,t}(\theta) \;=\; \sum_{o=1}^{C} N_{o,t}\, \Pi_{od,t}(\bar{x}_t(\theta); \theta). \tag{27}$$

I match this to the observed application count $A_{d,t}$ using a relative moment condition for each valid (d, t) cell:

$$m^{A}_{d,t}(\theta) \;=\; \frac{\hat{A}_{d,t}(\theta) - A_{d,t}}{A_{d,t}} \;=\; 0. \tag{28}$$

Expressing the gap as a fraction of observed applications protects against over-weighting large institutes.

5.3.2 Bilateral Funded Flow Shares

The model predicts funded grants from origin o to destination d in year t as:

$$\hat{F}_{od,t}(\theta) \;=\; N_{o,t}\, \Pi_{od,t}(\bar{x}_t(\theta); \theta)\, S_{od}(M_{o,t}(\theta);\, \eta_{od}). \tag{29}$$

Because bilateral funded flows are sparse at the annual frequency, I pool them across all sample years and match shares of the pooled flow matrix. For each destination d, the pooled observed share of origin o is:

$$s_{od} \;=\; \frac{\sum_t F_{od,t}}{\sum_t \sum_{o'} F_{o'd,t}}, \tag{30}$$

and the corresponding moment condition is a relative (proportional) deviation:

$$m^F_{od}(\theta) \;=\; \frac{\hat{s}_{od}(\theta) - s_{od}}{s_{od}} \;=\; 0, \tag{31}$$

where $\hat{s}_{od}(\theta)$ is the analogous model-implied share using pooled model-predicted flows. I restrict to origin-destination pairs whose pooled observed share exceeds 0.1%, to avoid fitting near-zero flows dominated by measurement error. Expressing the deviation relative to the observed share places small and large bilateral flows on a comparable footing. These moments are the primary source of identification for the idea quality distribution parameter μ_{od}: conditional on total application volumes to each destination (disciplined by the application moments), and taking into account the evaluation noise, the distribution of funded grants across origins pins down the relative idea quality for any pair of origins (o, o') for the grant reviewers.

5.3.3 Success Rates

The model predicts the aggregate success rate for institute d in year t as the ratio of total funded grants to total applications:

$$\hat{\pi}_{d,t}(\theta) \;=\; \frac{\sum_o \hat{F}_{od,t}(\theta)}{\hat{A}_{d,t}(\theta)}. \tag{32}$$

I match this to the observed rate $\pi_{d,t}$ using a relative moment condition for each valid (d, t) cell:

$$m^\pi_{d,t}(\theta) \;=\; \frac{\hat{\pi}_{d,t}(\theta) - \pi_{d,t}}{\pi_{d,t}} \;=\; 0. \tag{33}$$

Together with the application moments, the success rate moments discipline the equilibrium payline level: given a payline, both total applications and total funded grants are pinned, so their ratio is determined. The success rate moments are useful relative to the application moments because they are more directly informative about how scientists respond to changes in the payline, while the application moments are more directly informative about the outside option. In addition, cross-sectional variation in success rates across institutes with different scientific-distance compositions of their applicant pool provides identification for the distance gradient γ of evaluation noise.

Table 4 summarises the identifying variation used by the estimator for each parameter block.

5.4 GMM estimation procedure

The estimator matches 881 moments to 366 free parameters: the 19×19 merit matrix μ_{od} and five scalars $(\sigma, c_0, c_1, \eta_0, \gamma)$. For each candidate θ, the equilibrium payline system

Parameter	Identifying variation	Primary moments
μ_{od} (full)	Relative bilateral funded flow shares vary with the relative merit of each origin-destination pair, conditional on paylines and pool sizes. Cross-sectional dispersion in shares across origins pins the full $C \times C$ merit matrix	Flow share moments m^F_{od}
σ	Cross-institute substitution patterns: a higher σ implies flatter, less responsive choice probabilities as paylines change over time, dampening the sensitivity of application volumes and success rates to funding shocks	Application moments $m^A_{d,t}$ and success moments $m^\pi_{d,t}$
c_0, c_1	Outside-option intercept and time trend identified from the aggregate level and time path of total applications relative to pool size $N_{o,t}$, conditional on the merit matrix and paylines	Application moments $m^A_{d,t}$
η_0, γ	Baseline noise and its distance gradient are jointly identified by the level and cross-sectional variation in success rates across institute-year cells: institutes drawing relatively more cross-field applicants face higher effective evaluation noise at a given payline, shifting their success rate away from the no-noise benchmark	Success rate moments $m^\pi_{d,t}$ conditional on paylines
$\bar{x}_{d,t}$ (solved)	Determined from the market-clearing condition given $(\mu_{od}, \sigma, c_t, \eta_{od})$; not a free parameter	Market-clearing / slots data

Table 4: Identifying Variation Used by the Estimator

is solved year by year before evaluating the moments. I estimate θ by two-stage GMM.

Stage 1. For the first stage of the GMM, I use an equal-block weighting matrix that assigns total weight one to each of the three moment blocks and distributes that weight uniformly within block:

$$W_{ii}^{(1)} \;=\; \frac{1}{n_b} \quad \text{for moment } i \text{ in block } b, \tag{34}$$

where $b \in \{\text{applications, flows, success rates}\}$ and n_b is the number of moments in block b. This weighting allows each block to contribute equally to the first-stage estimates.

Stage 2. Following the classic two-stage GMM approach, I reweight each block by the inverse of its first-stage residual variance:

$$\hat{v}_b \;=\; \mathrm{Var}_{i\in b}\Big(m_i(\hat{\theta}_1)\Big), \qquad W_{ii}^{(2)} \;=\; \frac{1}{\hat{v}_b} \quad \text{for moment } i \in \text{block } b. \tag{35}$$

The second-stage estimator minimises

$$J_2(\theta) \;=\; m(\theta)^\top W^{(2)} m(\theta), \tag{36}$$

starting from $\hat{\theta}_1$ without ridge regularisation. Appendix H provides more details on the estimation procedure.

5.5 Estimating the Production Function

5.5.1 Dependent variable construction

Y_{iodt}^m denotes the dependent variable. Here the subscript i indexes unique grants and PI pairs. The individual observation is not a unique grant, but a unique grant–investigator pair. This is because grants can have multiple investigators, although the vast majority of grants in the sample have only one investigator (See figure 19).

Y_{iodt}^m counts the number of publications that received funding from grant i. For multiple-investigator grants, the publication count for i is divided by the number of investigators to avoid double-counting publications. Certain publications received money from multiple grants, in which case the publication is split equally across the contributing grants and each grant receives $1/n$ fraction of a publication. Moreover, each publications are weighted by some weighing scheme m that captures a particular outcome of interest. As such, Y_{iodt}^m is closer to a continuous variable than a count variable.

I estimate the same production function for three outcomes $m \in \{\text{IF, CP, DB}\}$: impact-factor-weighted publications (Y_d^{IF}), commercial-potential-weighted publications (Y_d^{CP}),

and disease-burden-weighted publications (Y_d^{DB}). The impact-factor outcome weights each publication by its journal impact factor. I use the two year impact factor from 2014 as computed in the OpenAlex dataset.

The commercial-potential outcome weighs each publication by the *ex ante* probability that the paper is cited by at least one renewed patent which is estimated by Masclans, Hasan, and Cohen (2023) using the word embedings of the publication's title and abstract. Eighty-nine per cent of the publications in my sample are covered by Masclans, Hasan, and Cohen (2023). Observations without a commercial-potential score are assigned a weight of zero.

The disease-burden outcome differs from the other two. I do not observe the disease burden of each publication. Instead, I link each grant to the burden of the disease it tackles and assume that all publications funded by that grant share the same disease-burden weight.

To map grants to diseases I use the RCDC disease tags that the NIH attaches to each grant. I match these tags to the Global Burden of Disease (GBD) taxonomy using the NIH crosswalk in National Institutes of Health (n.d.), and measure burden by Disability-Adjusted Life Years (DALYs) for the US population. The same grant-level burden data serve a second purpose: aggregated to the institute level, they yield the average disease-burden profiles I use to benchmark the funding allocation in Section 7. Appendix F details the RCDC–DALY mapping.

5.5.2 Independent variable and functional form

I estimate the production function by negative binomial maximum likelihood to deal with the overdispersion in the data. Non-negativity constraints are imposed on all destination-specific knowledge and funding elasticities $\alpha_d \geq 0$ and $\beta_d \geq 0$, to avoid economically absurd results, but the constrain binds only once, for the funding elasticity of Dental and Craniofacial research. Appendix G details the constrained estimator and how standard errors are computed at the boundary. The estimating equation is:

$$\mathbb{E}[Y_{iodt}] = \exp\big(\phi_d + \tau_t + \alpha_d \ln k_{it} + \beta_d \ln g_{it} + \lambda \ln W_{od,t}\big), \tag{37}$$

where k_{it} is the knowledge stock available to grant i in year t, g_{it} is the realised grant funding, ϕ_d and τ_t are destination and year fixed effects, and α_d, β_d are destination-specific knowledge and funding elasticities.

The term $\ln W_{od,t}$ is a structural control for applicant selection. The selection object is the conditional expected idea quality among funded scientists, shifted to ensure non-

negativity:

$$W_{od,t} \equiv \mathbb{E}[a_{iod} - \underline{a} \mid \text{chose } d, \text{ funded, } t] = \frac{1}{S_{od,t}} \int (z + \bar{x}_{d,t} - \underline{a})\, \Phi\left(\frac{z}{\eta_{od}}\right) f_Z(z;\, M_{o,t})\, dz, \tag{38}$$

where $\underline{a} = \min_{o,d,z}(\mu_{od} + z)$ is a global floor computed over the quadrature support, ensuring $W_{od,t} > 0$ for all pairs. The object is computed from the estimated structural parameters. Including $\ln W_{od,t}$ absorbs cross-cell variation in average funded idea quality, so that α_d and β_d identify the causal elasticities of output with respect to knowledge and funding, conditional on project quality. This is more parsimonious than a full set of origin-destination fixed effects whilst retaining structural discipline on the selection correction. The funding elasticity β_d and the knowledge elasticity α_d are both destination-specific.

I construct the grant-level knowledge stock k_{it} using field-specific depreciation rates and grant-level citation shares. Following current research benefits from the accumulated output of related fields, depreciated over time. Rather than imposing a common depreciation rate across all fields, I estimate field-specific rates δ_f from the age distribution of cited references (details in Appendix E).

Following Acemoglu, Akcigit, and Kerr, 2016, I define the knowledge input to grant i as a citation-weighted aggregate of field-level stocks. Let $\gamma_{fi} \geq 0$ denote the share of citations to field f in the references of grant i, normalised such that $\sum_f \gamma_{fi} = 1$.[11] Each share is computed by tracing every cited publication back to the original institute that funded it. This grant-level construction follows the logic of shift-share instruments, where each unit is differentially exposed to a common shock (Borusyak, Hull, and Jaravel, 2024), and exploits within-field heterogeneity in citation profiles across grants. The knowledge stock is a Cobb-Douglas aggregate of field-level stocks:

$$k_{it} = \prod_{f=1}^{C} \left(1 + K_{f,t}\right)^{\gamma_{fi}}, \tag{39}$$

where $K_{f,t}$ is the depreciated publication stock of source field f at time t, computed using the field-specific rate δ_f. The log knowledge stock entering the production function is therefore $\ln k_{it} = \sum_{f=1}^{C} \gamma_{fi} \ln(1 + K_{f,t})$.

To address potential endogeneity, I construct citation shares using only cited publications published at least one year prior to the grant's start date. This ensures that the weights γ_{fi} are predetermined and not a functions of the grant's production, avoiding the reverse causality whereby higher-output grants mechanically cite more papers.

[11] This notation for citation-share weights γ_{fi} should not be confused with the evaluation-noise distance gradient γ introduced in Section 5.2; the two are distinguished by the presence or absence of subscripts.

6 Results

This section presents the model estimates, evaluates the model fit on both targeted and untargeted moments, and reports the comparative statics results. The marginal returns to funding are estimated as the additional scientific output generated by an extra one million dollars of funding.

6.1 Model Estimates and Fit

Table 5 reports the estimated scalar parameters and the in-sample fit; the 361 mean idea-quality terms μ_{od} are shown separately in Figure 4. The outside-option intercept is negative, $\hat{c}_0 = -2.78$, so many $\hat{\mu}_{od}$ are also negative: both are measured relative to the outside option, which serves as the numeraire. The positive time slope, $\hat{c}_1 = 0.063$, implies that the outside option became progressively more attractive over the sample, consistent with a rising cost of applying for an R01 grant.

The distance gradient is positive, $\hat{\gamma} = 1.11$, so evaluation noise $\eta_{od} = \hat{\eta}_0 \exp(\hat{\gamma} D_{od})$ rises with scientific distance: grant reviews are less accurate for cross-field applications. In the 2005–2018 distance matrix (Figure 3), the mean of D_{od} across the 361 origin-destination cells is 0.55 with a standard deviation of 0.17; a one-standard-deviation increase therefore raises η_{od} at the mean from 2.20 to 2.66, a 21 per cent increase.

The estimated mean idea-quality matrix $\hat{\mu}_{od}$ is shown in Figure 4. Diagonal entries $\hat{\mu}_{oo}$ are systematically higher than off-diagonal entries, confirming that scientists draw their best ideas in their home field. Heterogeneity is pronounced along both margins: some fields are stronger "senders" than "receivers" (e.g. Biomedical Imaging and Bioengineering) and others the reverse (e.g. Mental Health). The matrix is also markedly asymmetric. For example, scientists from Cancer applying to Aging draw lower-quality ideas on average ($\hat{\mu}_{\text{Cancer,Aging}} = 0.04$) than scientists from Aging applying to Cancer ($\hat{\mu}_{\text{Aging,Cancer}} = 0.18$).

Parameter	Estimate
σ (idiosyncratic shock scale)	0.2331
c_0 (outside option intercept)	-2.7767
c_1 (outside option time slope)	0.0627
η_0 (baseline evaluation noise)	1.1943
γ (distance slope in η_{od})	1.1109

Table 5: GMM scalar parameter estimates.

Two-stage GMM estimates over 2005–2018. The μ_{od} matrix ($19 \times 19 = 361$ parameters) is displayed separately in Figure 4.

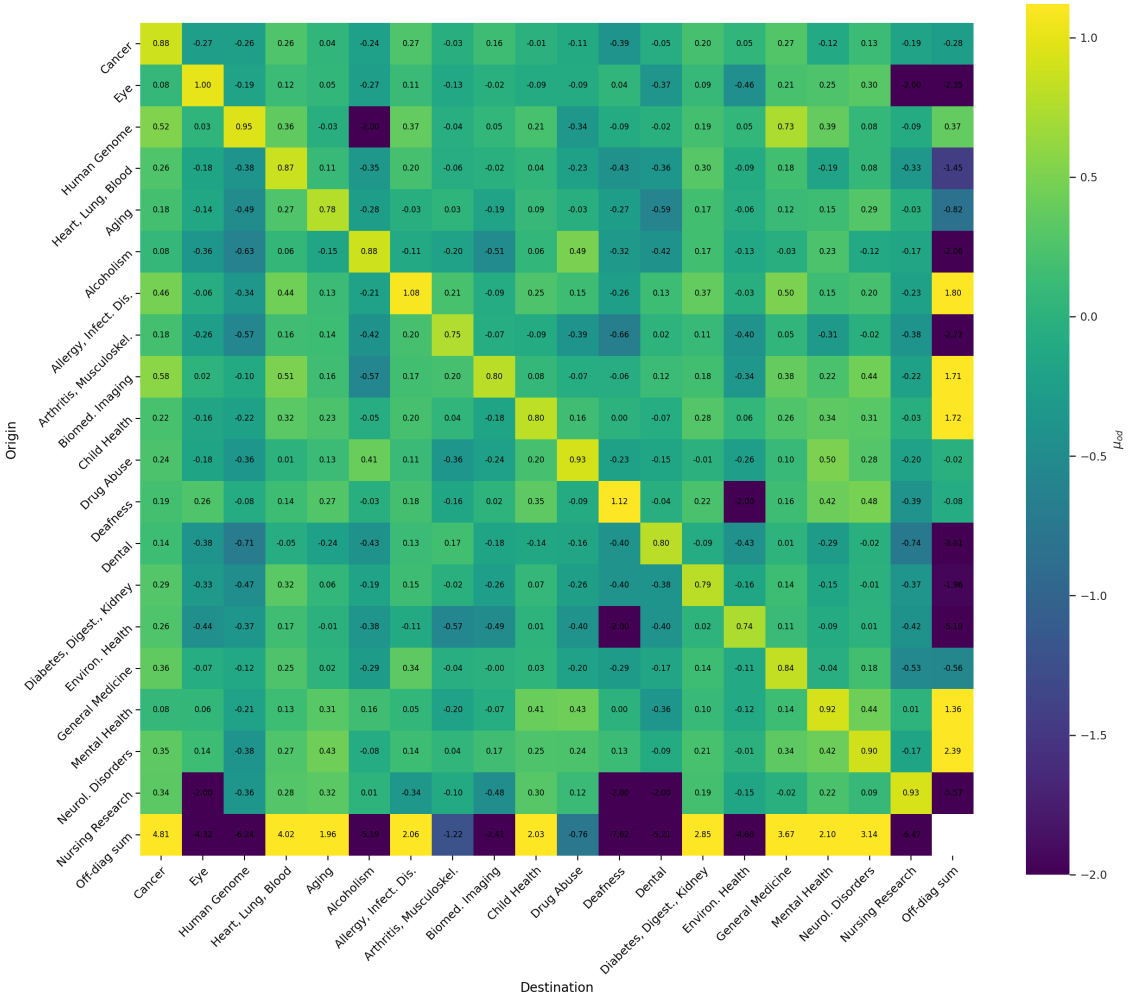
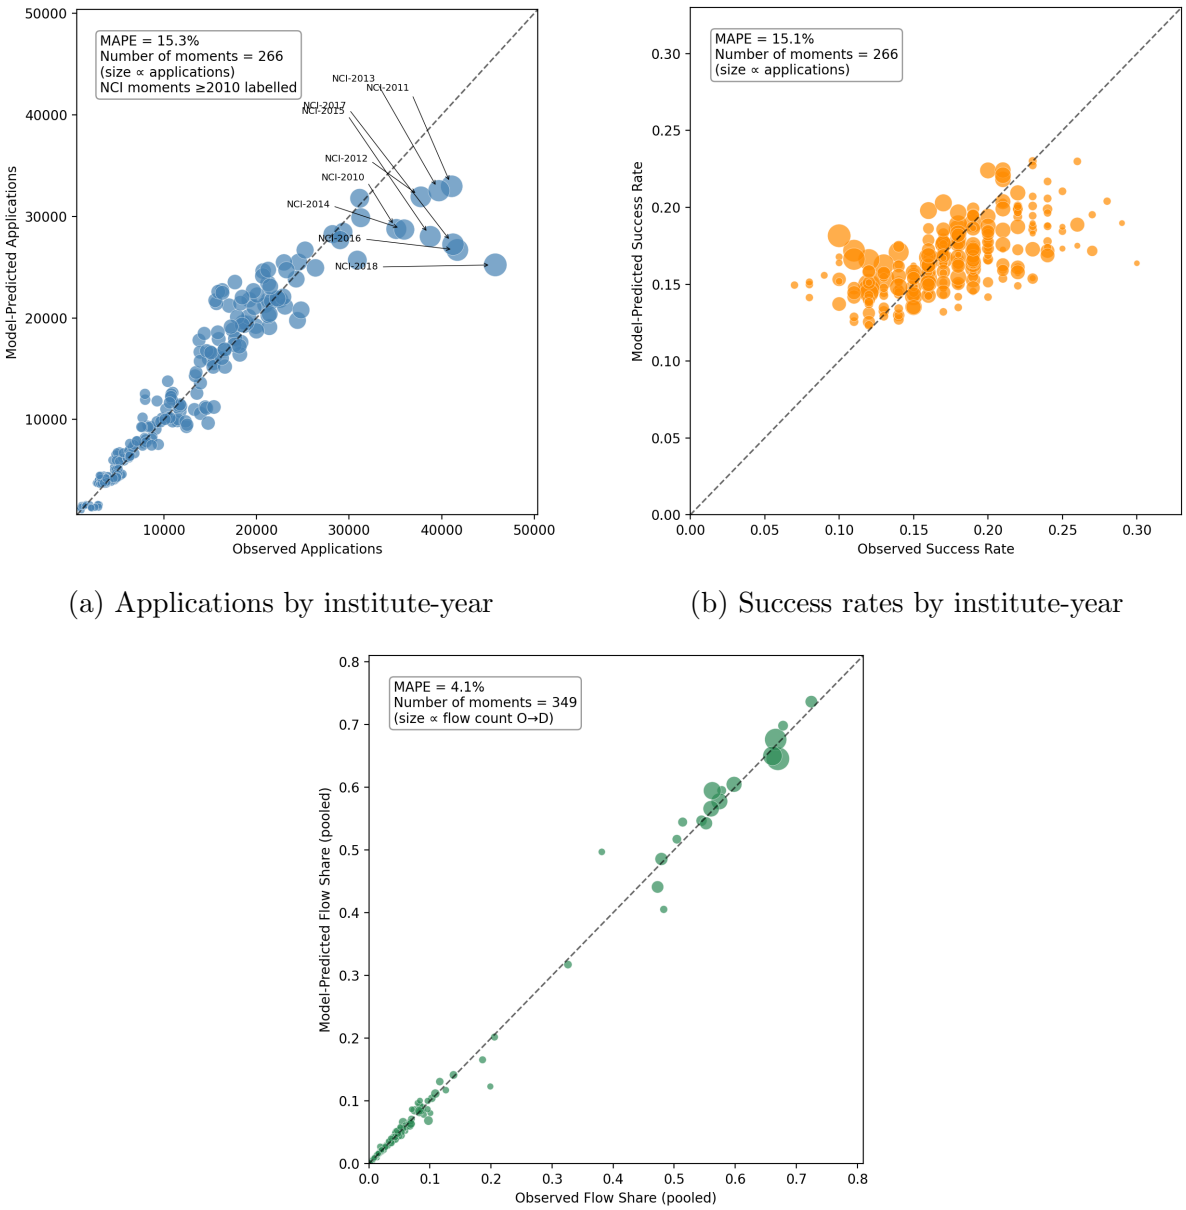


Figure 4: Estimated mean idea quality (μ_{od}).

Structurally estimated idea quality matrix $\hat{\mu}_{od}$. Rows index origin field o; columns index destination institute d. Cell entries report point estimates rounded to two decimals. Lighter green indicates higher $\hat{\mu}_{od}$: lighter rows identify origins whose scientists draw high-quality ideas when applying elsewhere; lighter columns identify destinations that attract high-quality ideas from other fields.



(a) Applications by institute-year

(b) Success rates by institute-year

(c) Bilateral flow shares (pooled)

Figure 5: Model fit: targeted moments.

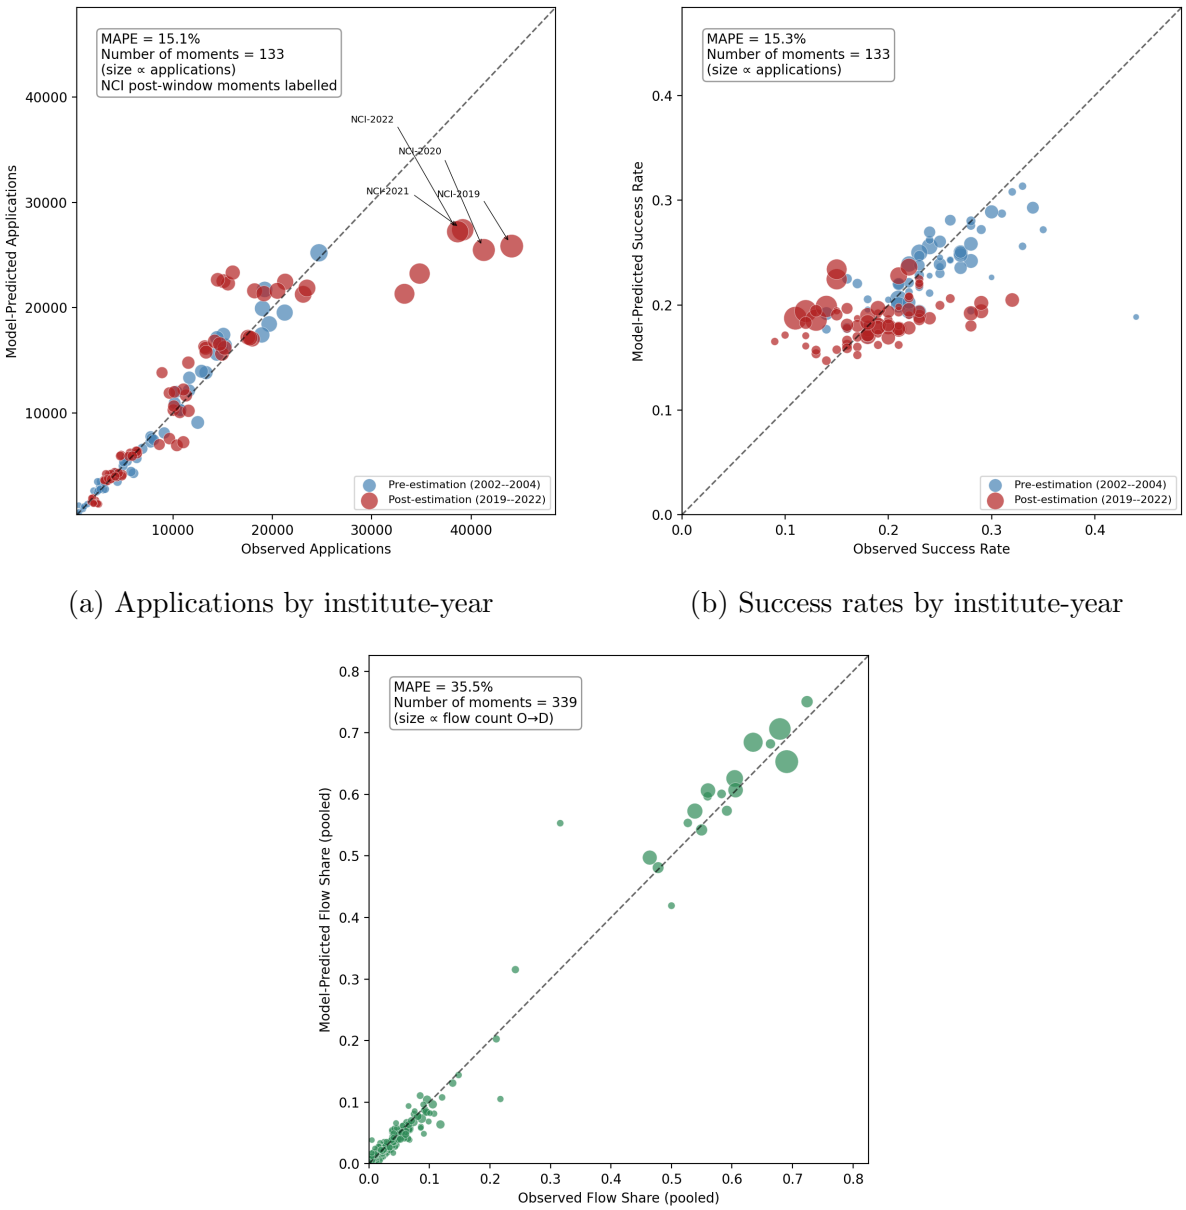
Each panel plots model-predicted versus observed values for one moment block. The diagonal line indicates perfect fit. Marker sizes are proportional to the number of observations. Panel (a): institute-year cells for total applications. Panel (b): institute-year cells for aggregate success rates. Panel (c): origin-destination pairs for pooled bilateral flow shares.

Figure 5 displays the model fit for each moment block. The model tracks application volumes well (panel a; MAPE = 15.3%, $R^2 = 0.91$), with the caveat that it underpredicts the number of applications for the National Cancer Institute after 2010. Bilateral flow shares are fitted extremely well (panel c; MAPE = 4.1%, $R^2 = 0.99$).

Success rates are not as well predicted (panel b; MAPE = 15.1%, $R^2 = 0.34$). This likely comes from the simplifications we made in the model: when a change in budget

allocation G_d occurs, paylines and scientists adjust immediately to return to equilibrium. This compresses the variation in success rates that the model can generate. In practice it is likely that a return to equilibrium takes time, and that success rates adjust with some lag, which would generate more variation in success rates: The observed success rates vary from 0.08 to 0.35 across institutes and years, but the model-predicted success rates only range from 0.12 to 0.25.

I also evaluate predictive fit outside the estimation window. I hold the estimated parameters fixed and recompute the predicted moments with data from out-of-sample years. I use the non-overlapping holdout years 2002–2004 and 2019–2022. Figure 6 shows the fit in the out-of-sample years, which is broadly similar to the fit in the in-sample years.



(a) Applications by institute-year

(b) Success rates by institute-year

(c) Bilateral flow shares (pooled)

Figure 6: Model fit: untargeted moments.

The parameters are estimated on 2005–2018 and then held fixed. The outside-option trend is evaluated with a calendar-consistent time index. Panels (a) and (b) split institute-year moments between pre-estimation years 2002–2004 and post-estimation years 2019–2022. Panel (c) pools bilateral origin-destination flow shares over both holdout windows, because rare flows are sparse within short windows. Marker sizes are proportional to applications in panels (a) and (b), and to raw flow counts in panel (c). The diagonal line indicates perfect fit.

The evolution of the model-implied equilibrium paylines ($\bar{x}_{dt}$) over the sample period is illustrated in Appendix Figure 13. The model-implied paylines generally tighten over the 2008–2013 period, reflecting an increasingly competitive environment for NIH, before easing again over 2014–2018 which is broadly consistent with observed trends in application success rates.

6.2 Production Function Estimates

The production function is estimated separately for the three outcomes defined in Section 5.5. Let $m \in \{\text{IF, CP, } b\}$ index the outcome.

Knowledge elasticities $\hat{\alpha}_d^m$ are positive for all 19 institutes and all three outcomes. Table 6 reports the mean and maximum destination-specific knowledge and funding elasticities across the 19 institutes. For impact factor (Appendix Table 19) the knowledge elasticity ranges from 0.097 (Heart, Lung, and Blood) to 0.553 (Alcoholism); for commercial potential (Appendix Table 20) from 0.112 (Eye) to 0.670 (Alcoholism); for disease burden (Appendix Table 21) from 0.055 (Aging) to 0.426 (Dental and Craniofacial Research). Funding elasticities $\hat{\beta}_d^m$ are smaller. All elasticity estimates are below one, suggesting decreasing returns to scale in scientific production. Appendix G documents the estimation procedure.

	Knowledge elasticity $\hat{\alpha}_d$		Funding elasticity $\hat{\beta}_d$	
Outcome	Mean	Max	Mean	Max
Impact factor	0.307	0.553	0.117	0.220
Commercial potential	0.384	0.670	0.086	0.244
Disease burden	0.213	0.426	0.104	0.261

Table 6: Summary of production-function elasticities across outcomes.

Each row summarises the 19 destination-specific elasticities estimated for one outcome. Reported statistics are the mean and maximum across institutes. Full destination-specific estimates are reported in Appendix Tables 19–21.

6.3 Marginal Return Estimates

Using the estimated parameters from the field-choice model and the production function, I compute the marginal return to funding for each institute d and each output m. Appendix I describes the numerical implementation of these comparative statics. The marginal return per \$1 million of funding to institute d on output $m \in \{\text{IF, CP, DB}\}$ decomposes into four components:

1. **No-externality return**: direct extensive-margin increase in output m at the funded institute.
2. **Talent-flow externality**: change in output m across all other institutes from the equilibrium reallocation of applicants as a consequence of an addition million dollar to institute d.
3. **Own knowledge spillovers**: change in output m at institute d from the augmentation of its own knowledge stock K_d^m that comes from the additional research produced by the additional funding.
4. **Knowledge spillovers to other fields**: change in output m at all other institutes

attributable to the additional research produced by the marginal million dollar in d.

I report partial effects—the change in output m for a \$1 million budget increase—evaluated at the 2018 baseline R01 allocations and application population productivity in Figure 7. Each institute has three horizontal stacked bars, one for each output metric. The length of each bar is the marginal return to funding for that output, and the bar segments show the contribution of each of the four components.

Marginal returns differ sharply across institutes and output metrics. The talent-flow externality is negligible in magnitude for every metric. Variation in the marginal return is instead driven almost exclusively by the direct effect and the knowledge externalities.

Impact-factor marginal returns. Human Genome has the highest marginal return at 64.0 impact-factor-weighted publications per \$1 million, combining a no-externality return of 35.9, own knowledge spillovers of 4.0, and cross-field knowledge spillovers of 24.2. General Medicine follows at 58.6, driven mainly by a no-externality return of 45.5. Dental and Craniofacial Research (35.0) and Diabetes and Digestive and Kidney Diseases (33.7) rank next; Nursing Research (5.3) and Aging (14.8) sit at the bottom (Appendix Table 22, Figure 7).

Commercial-potential marginal returns. General Medicine leads at 4.9 predicted patent citations per \$1 million, followed by Human Genome (4.5) and Biomedical Imaging and Bioengineering (3.4); Dental and Craniofacial Research (3.3) is close behind, while Nursing Research (0.3) is at the bottom. For most institutes the no-externality return dominates, but Human Genome is an exception: its cross-field knowledge spillovers add 2.2 predicted patent citations, slightly more than its no-externality return of 2.0 predicted patent citations. Own-plus-cross knowledge spillovers total 2.5 predicted patent citations for Human Genome, consistent with strong downstream diffusion from the most upstream fields (Appendix Table 23, Figure 7).

Disease-burden marginal returns. General Medicine again leads, at 10.3 disease-burden-weighted publications per \$1 million. Heart, Lung, and Blood (9.0) and Human Genome (7.6) follow, with Cancer close behind at 6.8. Human Genome's cross-field knowledge spillovers add 2.5 disease-burden-weighted publications. Nursing Research (1.2), Child Health and Human Development (1.4), Deafness and Communication Disorders (1.4), and Aging (1.5) have the smallest marginal returns (Appendix Table 24, Figure 7).

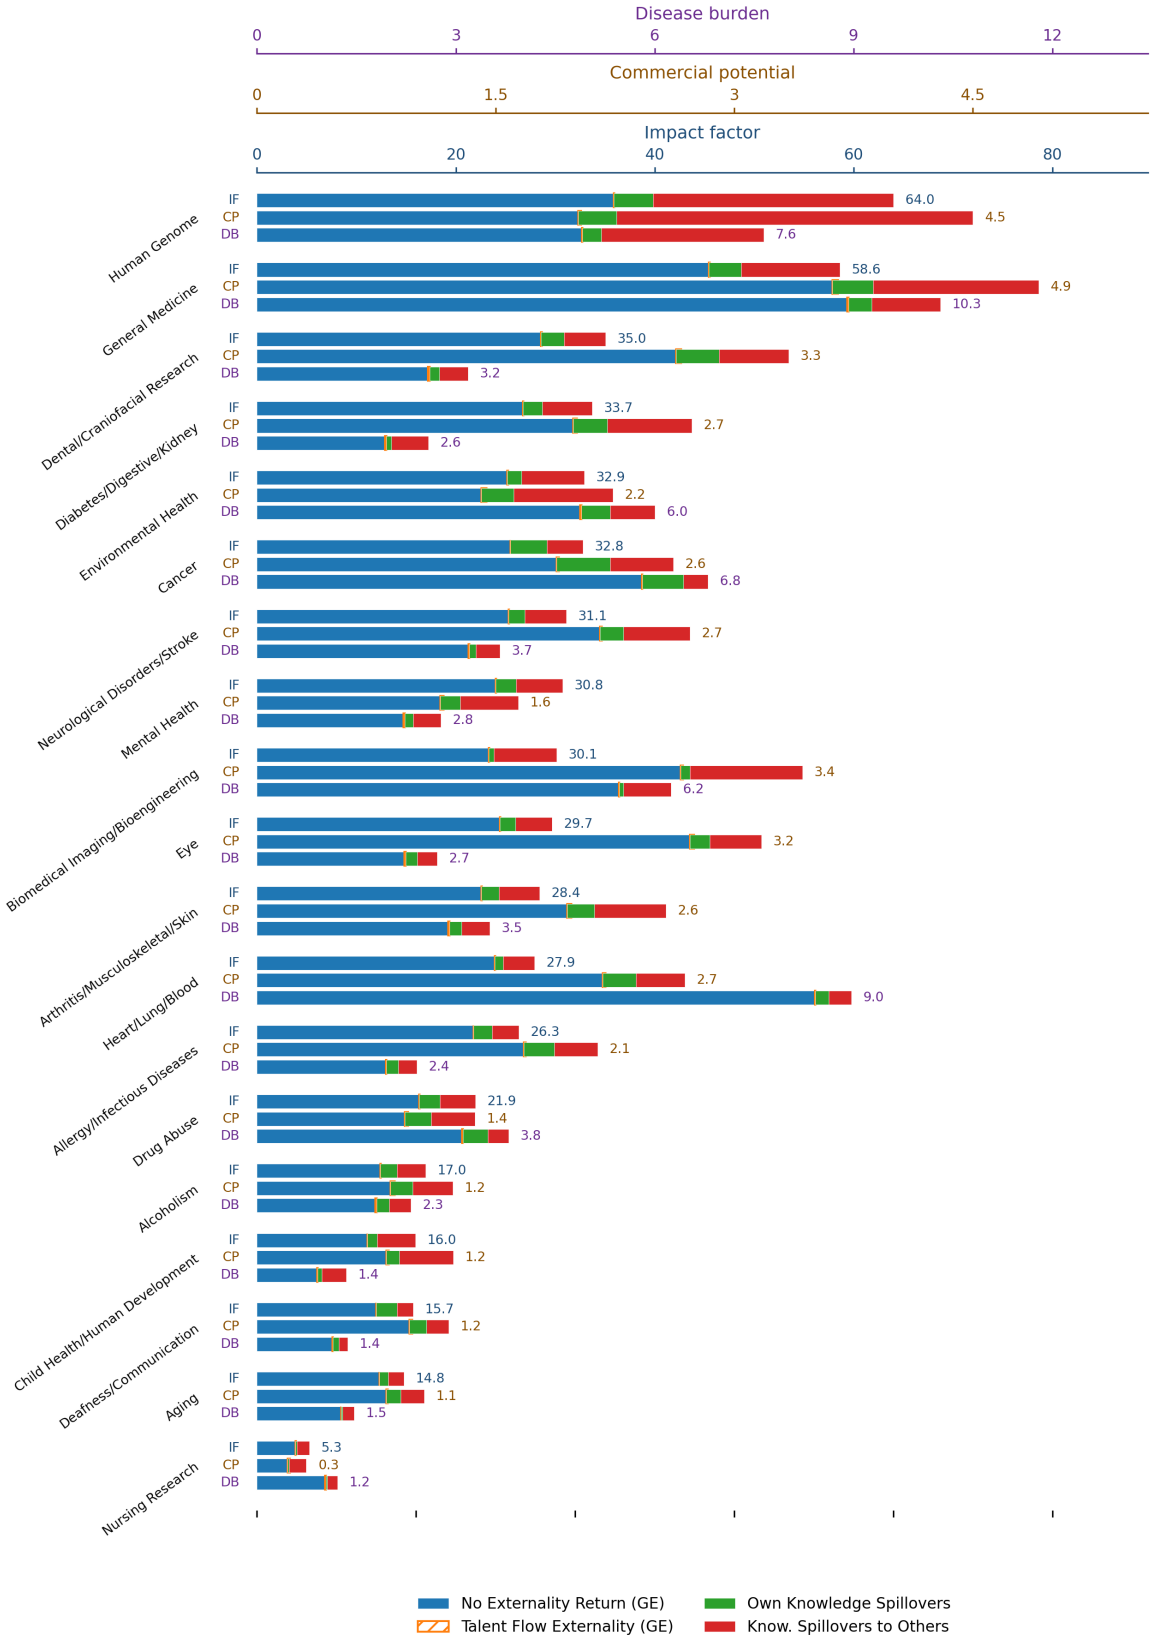


Figure 7: Decomposition of marginal returns by institute across outputs.

Institutes are ordered by the impact-factor marginal return. The three top axes give the output-specific scales, so bar lengths are comparable within outputs but not across outputs. Numbers at the end of bars report the marginal return.

Cross-outcome comparison. General Medicine ranks first on the commercial-potential and disease-burden marginal returns and second on the impact-factor marginal return. Human Genome ranks first on the impact-factor marginal return, second on the commercial-potential marginal return, and third on the disease-burden marginal return, with its position lifted on every metric by unusually large cross-field knowledge spillovers. Nursing Research and Aging sit at the bottom on every metric. The similarity between the impact-factor and commercial-potential marginal returns is not surprising: Figure 8 shows that these two publication-based outcomes are strongly correlated, much more than either is with disease burden. Publications published in high-impact journals are also more likely to be cited by patents.

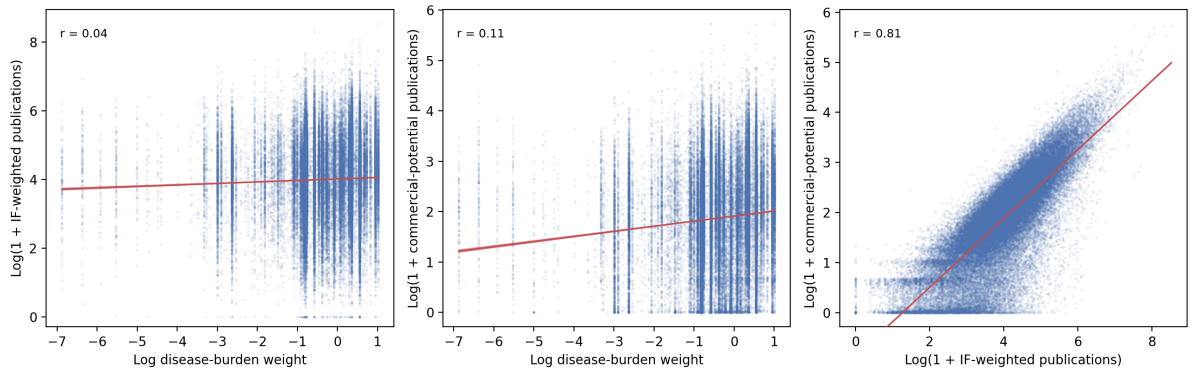


Figure 8: Grant-level correlations between disease burden, impact factor, and commercial potential.

Each point is a unique grant in the 2005–2018 analysis sample. Red lines: fitted linear trends in the plotted log variables. Shaded bands: one-standard-error bands for the conditional mean. Reported r values are log Pearson correlations.

6.4 Mechanisms: Talent Flows and Knowledge Spillovers

Talent-flow channel. Figure 7 shows that the talent-flow externality is negligible for every institute and every output metric, and yet Figure 4 suggests considerable heterogeneity in the mean idea quality μ_{od} across origin fields for any given destination field. In other words, scientists' reallocation does not affect marginal returns, even though scientists have vastly different expected productivities based on their origin field when applying to a given destination field.

One hypothesis is that scientists are not mobile, and that changes in G_d do not induce meaningful reallocation. We test this hypothesis by computing the change in the share of funded scientists at the shocked institute d whose origin field differs from d (the "off-field share"). I do so both after the change in field d's budget, which alters its payline x_d, but before the paylines in other fields adjust, referred to as the partial equilibrium response, and after the entire payline vector adjusts and every grant market clears, referred to as the general equilibrium response. Table 7 reports the PE and GE elasticities of that off-field share, together with the largest donor institute in general equilibrium. All elasticities are positive, confirming that scientists do move toward the shocked field. The elasticities range from 0.10 in Biomedical Imaging and Bioengineering to 0.68 for the National Institute of the Eye. In partial equilibrium the off-field share rises by 0.006 to 0.077 percentage points, with a median increase of 0.018 percentage points. The general-equilibrium changes are very similar, ranging from 0.005 to 0.077 percentage points, suggesting that the return to equilibrium after the shock does not attenuate or amplify the mobility response.

Scientist mobility does not explain the weakness of the talent externality. There is another hypothesis: the scientists who move toward field d are no less productive than the existing scientists in d, and the scientists that replace them in their origin fields are no less productive than them. Table 22, 23 and 24 show that the talent flow externalities are large and negative for all fields and outcomes in partial equilibrium, but close to zero in general equilibrium. In other words, when a scientist moves from field o to field d the initial opportunity cost of the project in field o that would have been completed if the scientist had stayed in o is large, but the scientist that replaces her in field o has a similar expected productivity, so the net effect on field o is close to zero. This is the case despite the replacing scientist in field o likely coming from a different field, and therefore having ideas of lower expected quality in field o. Even though scientists from other fields have lower expected idea quality *ex ante*, the replacing scientist in field o has a similar productivity *ex post* conditional on getting funded in field o, which suggests that the selection mechanism functions as intended.

Table 8 confirms this by comparing the coefficient of variation (CV) of W_{od} with the CV

of $\hat{\mu}_{od}$. The CV of W_{od} is substantially lower than the CV of $\hat{\mu}_{od}$ for every destination, indicating that the grant-evaluation mechanism compresses large differences in *ex ante* idea quality into much smaller differences in *ex post* funded idea quality.

Institute	PE Elasticity	GE Elasticity	GE largest off-field origin
Eye	0.677	0.656	Neurological Disorders and Stroke
Deafness and Communication Disorders	0.661	0.646	Neurological Disorders and Stroke
Allergy and Infectious Diseases	0.515	0.462	General Medicine
Heart, Lung, and Blood	0.361	0.311	Diabetes, Digestive, and Kidney Diseases
Cancer	0.353	0.301	General Medicine
Dental and Craniofacial Research	0.337	0.328	Cancer
General Medicine	0.305	0.258	Cancer
Alcoholism	0.278	0.266	Drug Abuse
Diabetes, Digestive, and Kidney Diseases	0.267	0.235	Heart, Lung, and Blood
Environmental Health	0.254	0.245	Cancer
Arthritis, Musculoskeletal, and Skin Diseases	0.254	0.241	Heart, Lung, and Blood
Neurological Disorders and Stroke	0.246	0.220	Mental Health
Mental Health	0.245	0.220	Drug Abuse
Drug Abuse	0.242	0.221	Mental Health
Aging	0.227	0.207	Neurological Disorders and Stroke
Nursing Research	0.179	0.176	Mental Health
Human Genome	0.171	0.170	General Medicine
Child Health and Human Development	0.154	0.140	Mental Health
Biomedical Imaging and Bioengineering	0.101	0.097	Cancer

Table 7: Off-field-share elasticities and largest donor institute after a $1 million shock.

Each row corresponds to the shocked institute d. The off-field share is the share of funded grants at d awarded to scientists whose origin field differs from d. The PE columns hold all other paylines fixed. The GE columns allow all paylines to adjust. The GE largest off-field origin column reports the non-d origin institute that contributes the largest increase in funded scientists at d under the general-equilibrium response.

Institute	$CV_{\hat{\mu}_{od}}$	$CV_{W_{od}}$
Cancer	0.670	0.071
Eye	3.128	0.133
Human Genome	1.245	0.153
Heart, Lung, and Blood	0.792	0.075
Aging	1.552	0.085
Alcoholism	2.399	0.142
Allergy and Infectious Diseases	1.681	0.103
Arthritis and Musculoskeletal and Skin Diseases	10.677	0.106
Biomedical Imaging and Bioengineering	3.449	0.119
Child Health and Human Development	1.484	0.082
Drug Abuse	39.141	0.131
Deafness and Communication Disorders	2.012	0.160
Dental and Craniofacial Research	2.260	0.131
Diabetes and Digestive and Kidney Diseases	0.946	0.069
Environmental Health	2.494	0.111
General Medicine	0.998	0.086
Mental Health	1.966	0.114
Neurological Disorders and Stroke	1.128	0.087
Nursing Research	1.811	0.145

Table 8: Cross-origin dispersion in ex ante idea quality and funded idea quality.

Each row is a destination institute d. Columns report the coefficient of variation (CV = s.d. / |mean|) across origin fields o for estimated mean idea quality $\hat{\mu}_{od}$ (time-invariant GMM estimate) and idea quality *conditional* on funding W_{od}. The $\hat{\mu}_{od}$ CV denominator uses the absolute destination-specific mean because $\hat{\mu}_{od}$ is measured relative to the outside option. Destinations with $|\bar{\mu}_d| < 0.05$ (NIAMS, NIDA) have inflated $CV_{\hat{\mu}}$ values.

Knowledge-spillover channel. The knowledge channel is the main source of cross-outcome variation in rankings. It varies widely across institutes and outcomes, more so than own-field returns. Fields such as Human Genome, General Medicine, and Biomedical Imaging and Bioengineering, which are more upstream in the scientific knowledge production process, tend to have larger cross-field spillovers. Fields that are more downstream tend to generate smaller spillovers, such as Cancer or Arthritis, Musculoskeletal and Skin Diseases. For the disease-burden outcome the same upstream-to-downstream pattern survives but is muted, suggesting that the research projects benefiting most from knowledge spillovers are not necessarily those that tackle the highest disease burden conditions.

Moreover, cross-field knowledge spillovers are larger than knowledge spillovers to the same field for all fields bar Cancer and Ageing. Although the own field spillovers are larger than all individual cross-field spillovers, the sum of the cross-field spillovers is larger than the own-field spillover for 17 out of 19 institutes (see the complete knowledge spillover decomposition in Appendix Table 22). This highlights the importance of accounting for cross-field dependencies when evaluating returns to funding. The full knowledge spillover decomposition can be found in Appendix Figures 23, 24, and 25.

7 Policy Implications

This section uses the results presented in Section 6.3 to evaluate the current NIH allocation and to assess alternative allocation rules.

7.1 Implied Social Weights in the Current Allocation

The planner's optimum equalises the social value of a marginal dollar across institutes. As shown in Section 2, the first-order conditions of the planner's problem can be written

$$\omega_d \cdot MR_d^P = \lambda \quad \forall\, d,$$

where MR_d^P is institute d's marginal product of funding in raw publications and ω_d is an arbitrary social welfare weight on its marginal research. So far I have evaluated $\omega_d \cdot MR_d^P$ under three weighting schemes: impact factor, commercial potential, and US DALY burden. In every case the marginal returns are highly heterogeneous, so the current allocation is inefficient if the planner weights these three dimensions similarly across institutes.

In practice, the budget Congress sets reflects many objectives and cannot be reduced to maximising commercially relevant output or academic impact. Which preferences, then, does the current allocation reveal? The estimated model lets me recover the social welfare weights that rationalise it. I re-estimate the marginal return of funding in unweighted publication counts,[12] evaluate MR_d^P at the 2018 funding distribution by summing direct and spillover partial effects across all output destinations (see Section 4), and invert the first-order condition to recover the implied weights ω_d. Normalising the weights to sum to 19, one per institute on average, yields:

$$\omega_d = 19 \cdot \frac{1/MR_d^P}{\sum_{d'} 1/MR_{d'}^P}. \tag{40}$$

Figure 9 plots the underlying marginal returns by institute, and Table 9 reports the implied weights ω_d from equation (40). Alongside each weight I report the average academic impact, commercial potential, and disease-burden value of an institute's publications over the sample period, the three outcomes used so far. Each scheme is normalised to a mean of one to ease comparison across scales.

The implied weights do not align with the three value profiles. Their correlation with commercial potential is -0.659 and with academic impact -0.597. Alignment with US

[12] The production function and comparative-statics exercise mirror those described in Section 5 but replace the impact-factor-weighted outcome with an unweighted publication count. The estimated elasticities are qualitatively similar; see Appendix G for details.

DALY burden is closest to zero, though still negative at -0.237.

This disconnect comes from the tails of the distribution. A handful of institutes carry high implied weights but low impact factor and commercial potential, such as Nursing Research, Deafness and Communication Disorders, and Child Health and Human Development. Their research likely holds social and clinical value that the three metrics miss. At the other extreme, basic-science institutes such as Human Genome and General Medical Sciences carry low implied weights but high impact factor and commercial potential. These high values reflect the large knowledge spillovers these institutes generate, which Congress may not internalise when it allocates funding. Their focus on basic rather than readily applicable research may further explain the low weight they receive in the current allocation.

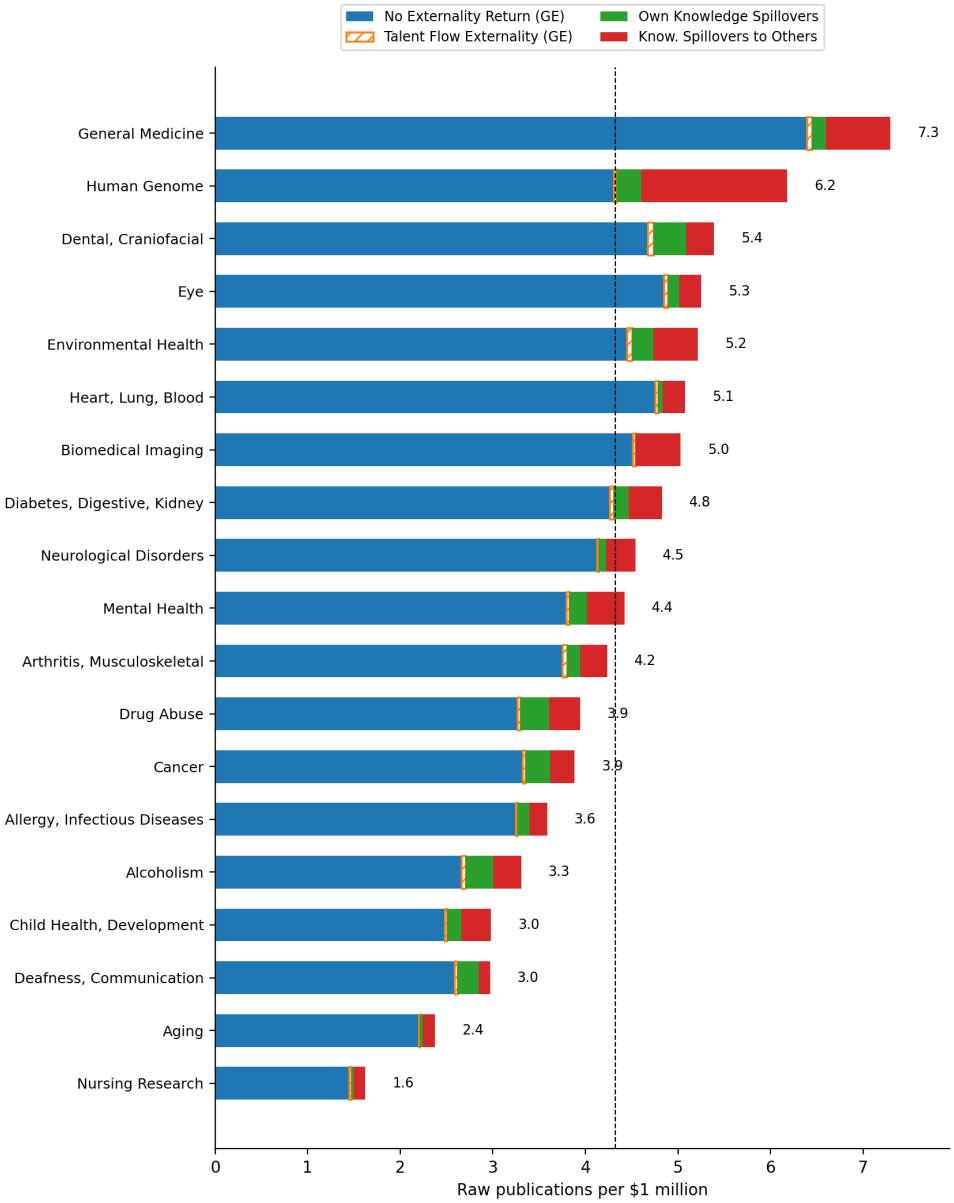


Figure 9: Raw-publication marginal returns by NIH institute, 2018.

System-wide change in unweighted publications per additional $1 million allocated to each institute, evaluated at the 2018 funding distribution. Each horizontal bar is decomposed into the no-externality return, the general-equilibrium talent-flow externality, own-field knowledge spillovers, and knowledge spillovers to other fields. The vertical dashed line marks the cross-institute mean total return.

Table 9: Implied social weights and publication-value profiles by NIH institute.

Institute	Raw pub. MR (pubs/$1M)	Implied weight ω_d	Academic impact	Commercial potential	Disease burden
Nursing Research	1.6	2.365	0.606	0.405	1.031
Aging	2.4	1.614	0.967	0.827	0.653
Deafness, Communication	3.0	1.291	0.818	0.876	0.466
Child Health, Development	3.0	1.287	0.792	0.615	0.492
Alcoholism	3.3	1.160	0.744	0.602	1.007
Allergy, Infectious Diseases	3.6	1.069	1.174	1.333	0.657
Cancer	3.9	0.987	1.150	1.293	2.008
Drug Abuse	3.9	0.973	0.829	0.583	1.073
Arthritis, Musculoskeletal	4.2	0.905	0.952	1.208	0.871
Mental Health	4.4	0.867	1.006	0.627	0.476
Neurological Disorders	4.5	0.845	1.075	1.188	0.791
Diabetes, Digestive, Kidney	4.8	0.794	1.046	1.138	0.524
Biomedical Imaging	5.0	0.763	0.967	1.384	1.547
Heart, Lung, Blood	5.1	0.756	0.895	1.092	2.044
Environmental Health	5.2	0.736	1.012	0.833	1.231
Eye	5.3	0.730	0.939	1.155	0.545
Dental, Craniofacial	5.4	0.712	0.992	1.231	0.435
Human Genome	6.2	0.621	1.719	1.245	1.464
General Medicine	7.3	0.526	1.318	1.365	1.683
Mean	4.3	1.000	1.000	1.000	1.000

ω_d is the social welfare weight that rationalises the 2018 NIH allocation under the assumption that the planner equalises weighted marginal returns across institutes (equation 40). Raw pub. MR is the system-wide change in unweighted publications per $1 million added to institute d. The three profile columns report the average value per publication for academic impact (impact factor), commercial potential, and disease burden (US DALYs); each column is normalised so the cross-institute mean equals one.

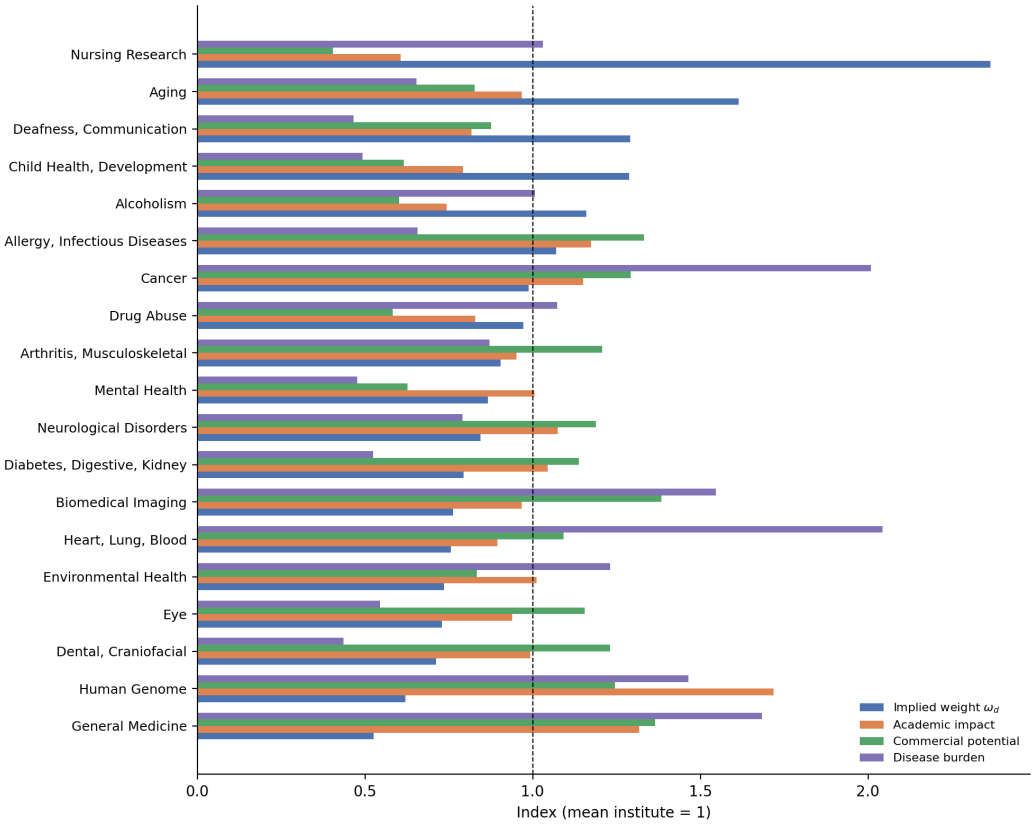


Figure 10: Implied social weights versus observed publication-value profiles, by NIH institute.

Grouped horizontal bar chart. Each series is normalised to a cross-institute mean of one. Institutes are sorted by implied social weight ω_d (highest at top). The dashed vertical line marks the common normalisation point.

7.2 Alternative Allocation Rules

The exercise asks how an extra $1 million added to the 2018 NIH R01 budget should be split across the 19 core institutes. I compare the current allocation against rules that split the increment in proportion to marginal returns or average disease burden. Two design choices need justification: the focus on a local increment rather than a full budget reallocation, and the use of a proportional split rather than a corner allocation.

The first choice is to study a local increment rather than the full budget. A wholesale reallocation is politically unrealistic. In practice, NIH reallocation usually occurs through differential growth across institutes, not through a complete redesign of the budget. IC funding is shaped by path dependence and lobbying pressures (Sampat, 2012), which produce occasional disproportionate increases for some institutes. Over the past twenty years, aggregate NIH funding has risen in most years. Nominal cuts are rare; prolonged nominal declines are absent (Figure 20 and Figure 21). Reallocation operates mainly through relative growth: some institutes stagnate in real or relative terms, while others receive discrete increases.

Reallocating from scratch would require stronger modelling assumptions as well. Currently, the model is identified from marginal variation around the observed allocation, so it speaks to the return on small reallocations, not to the output of a budget rebuilt from scratch. Evaluating a wholesale redesign would require assuming that the estimated parameter values hold globally.

The second choice is to split the increment proportionally rather than assign it all to the single highest-return institute. Optimally, the planner should place the entire $1 million in the institute with the highest marginal return, here Human Genome or General Medical Sciences. That rule needs no separate exercise, since the estimated marginal returns already report the output gain from adding $1 million to any single institute (Figure 7). Directing all the additional funding to one institute is also likely to be politically infeasible. The proportional rules below are the realistic alternative: they spread the increment across institutes in proportion to marginal returns or disease burden, preserving the ranking implied by marginal returns while avoiding a corner solution.

7.3 Benchmark and Marginal-Return Allocation Rules

The benchmark allocates the extra $1 million in proportion to current institute budget shares. Let s_d^0 denote institute d's observed 2018 R01 budget share. The benchmark gives institute d an additional s_d^0 dollars for each extra dollar of NIH R01 funding. This is the local analogue of maintaining the current allocation.

Let MR_d^m denote the institute-d marginal return of $1 million of extra funding for outcome $m \in \{\text{IF}, \text{CP}, \text{DB}\}$, evaluated at the 2018 comparative-statics baseline and summed over destination institutes as estimated in 6.3:

$$MR_d^m = \sum_h \left. \frac{\partial Y_h^m}{\partial G_d} \right|_{t^\star = 2018}. \tag{41}$$

For each outcome m, the corresponding marginal-return allocation rule splits the extra dollar in proportion to its institute marginal returns:

$$s_d^{r(m)} = \frac{MR_d^m}{\sum_{d'} MR_{d'}^m}, \qquad \sum_d s_d^{r(m)} = 1. \tag{42}$$

7.4 Average Disease-Burden Allocation

A fourth rule is to allocate the extra million dollars in proportion to *average* rather than *marginal* disease burden, echoing the medical literature that frequently benchmarks NIH funding against the burden of the diseases it targets (Gross, Anderson, and Powe, 1999). For each institute I build an average burden profile $\bar{B}_d^{\text{US DALYs}}$: the typical US-

DALY burden of the diseases its grant portfolio tackles, aggregated from the grant-level disease burden data introduced in Section 5.5, regardless of the production efficiency of the institute. The average disease-burden allocation splits the extra dollar in proportion to this profile,

$$s_d^B = \frac{\bar{B}_d^{\text{US DALYs}}}{\sum_{d'} \bar{B}_{d'}^{\text{US DALYs}}}. \tag{43}$$

Cancer, Heart, Lung, and Blood, General Medical Sciences, Biomedical Imaging, and Human Genome carry the highest burden profiles. Appendix J and Figure 22 provide more details on the construction and report the full profiles by institute.

7.5 Results

Table 10 reports $\Delta Y^{r,m}$ for the current-allocation benchmark, the average disease-burden rule, and the three marginal-return rules. Cell values are output gains from the extra \$1 million; parentheses report the ratio to the current-allocation benchmark within the same outcome. Figure 11 plots these gains relative to the benchmark for each of the three metrics. Table 11 reports the institute shares of the extra \$1 million implied by each rule.

Three results stand out. First, every combination of rule and metric delivers sizeable gains relative to the status quo, ten to thirty per cent more output for the same budget. This gap signifies either the inefficiency of the current allocation or the failure of these three metrics to internalise the social preferences Congress pursues. Second, marginal-return allocation delivers nearly identical gains for academic impact and commercial potential. But the US DALY gain is far larger under the disease-burden rule than under the commercial or academic rules. Third, allocating on average disease burden ignores production efficiency and is dominated across the board. Even on the disease-burden metric, the marginal disease-burden rule (green bar) delivers ten per cent more output than the average disease-burden rule (purple bar). This exercise highlights the importance of accounting for production efficiency in scientific research, even when one is only concerned about disease burden.

Allocation rule	IF-weighted publications	Commercial-potential publications	US DALY-weighted publications
Current allocation	28.67 (1.000)	2.35 (1.000)	4.47 (1.000)
Average disease burden	31.73 (1.107)	2.62 (1.112)	5.27 (1.177)
Impact-Factor MR rule	35.42 (1.235)	2.88 (1.223)	5.05 (1.129)
Commercial MR rule	35.21 (1.228)	2.94 (1.251)	5.10 (1.140)
US DALY MR rule	35.59 (1.241)	2.94 (1.249)	5.83 (1.304)

Table 10: Output Gain from a Local $1 Million Reallocation

Notes: Each row is an allocation rule applied to an additional $1 million addition to the 2018 R01 budget. The current allocation row splits the extra funding in proportion to observed 2018 institute budget shares. The average disease-burden row splits the extra funding in proportion to the institute's US-DALY average burden profile. The three MR rows split the extra funding in proportion to outcome-specific institute marginal returns. Each column is an outcome metric. Cell values are output gains $\Delta Y^{r,m}$ in units of the column outcome. Parentheses report the ratio to the current-allocation benchmark within the same column.

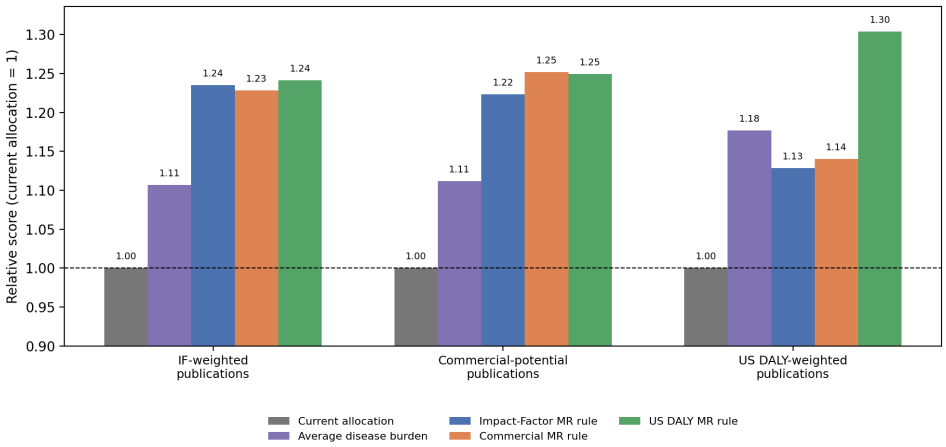


Figure 11: Output Gain from a Local $1 Million Reallocation, Relative to Current Allocation

Notes: Bars report $\Delta Y^{r,m}$ from Table 10 divided, within each metric column, by the gain from allocating the extra $1 million according to observed 2018 institute budget shares. Numbers above bars report the same ratios. The dashed line marks parity with the current-allocation benchmark.

Institute	Current allocation (%)	Average burden (%)	Impact-Factor MR rule (%)	Commercial MR rule (%)	US DALY MR rule (%)
Heart, Lung, Blood	12.57	10.32	5.05	5.96	11.45
Allergy, Infectious	12.43	3.00	4.77	4.75	3.08
Cancer	12.05	10.52	5.94	5.80	8.68
Aging	10.70	3.51	2.67	2.33	1.87
Neurological	8.18	4.19	5.64	6.03	4.67
Diabetes, Digestive, Kidney	7.08	2.73	6.10	6.06	3.30
General Medical Sci.	6.58	8.54	10.62	10.89	13.16
Mental Health	5.63	2.29	5.57	3.64	3.54
Child Health, Development	5.38	2.64	2.89	2.74	1.72
Drug Abuse	3.71	5.63	3.98	3.04	4.84
Arthritis, Musculoskeletal	2.74	4.92	5.15	5.70	4.49
Eye	2.54	3.10	5.38	7.03	3.47
Deafness, Communication	2.26	2.50	2.85	2.67	1.75
Alcohol Abuse	2.03	5.36	3.08	2.73	2.97
Biomedical Imaging	1.62	8.28	5.46	7.60	7.97
Environmental Health	1.54	6.72	5.97	4.96	7.66
Dental, Craniofacial	1.37	2.25	6.35	7.41	4.07
Nursing	0.94	5.53	0.96	0.69	1.55
Human Genome	0.65	7.96	11.59	9.98	9.76

Table 11: Institute Shares of the Local $1 Million Reallocation (%)

Notes: Institute shares of the extra $1 million under the current-allocation benchmark, the average disease-burden rule, and each marginal-return rule. The current allocation column is the observed 2018 R01 budget share. The average burden column is proportional to the US-DALY profile in Table **??**. Each column sums to 100%. Institutes are sorted by the current-allocation share.

8 Conclusion

This paper provides a structural framework for evaluating the allocative efficiency of scientific funding in biomedical research. Using administrative data on NIH R01 grants and their associated publications, I estimate a micro-founded model of how scientists choose which institute to seek funding from. Scientists respond to the funding distribution by changing where they apply, and research output in one field generates knowledge spillovers both to the same field and to other fields. By taking into account the behavioural response of scientists and the knowledge externalities of the research they produce, I can evaluate the efficiency of the current NIH funding allocation and provide counterfactual allocations.

I find that the marginal return to a research dollar varies by an order of magnitude across fields: a $1 million expansion yields 5.3 additional impact-factor-weighted publications in Nursing Research and 64.0 in Human Genome. Knowledge spillovers, not talent reallocation, drive this heterogeneity. Scientists do relocate towards better-funded fields, but those who replace them are equally productive once funded, so the net cost of reallocation is negligible. The gap between high-return and low-return institutes instead reflects their position in the research network: upstream institutes such as Human Genome and General Medical Sciences provides larger knowledge spillovers to other institutes compared to other downstream fields. Cross-field knowledge spillovers are quantitatively large: they exceed own-field knowledge spillover for all but two fields.

This heterogeneity in marginal returns does not seem to be guiding the existing funding allocation. The social welfare weights that rationalise the 2018 budget correlate negatively with the weights a social planner would assign if she were to maximise academic impact, commercial impact, or minimise disease burden. This suggests that the current allocation is either inefficient or reflects other social objectives that are not captured by these metrics.

The medical literature typically recommends funding institutes in proportion to the *average* burden of the diseases they tackle. This approach is suboptimal. Even a planner who values only disease burden, setting aside academic and commercial impact, could improve on this approach by allocating instead on *marginal* burden—which accounts for how efficiently each institute converts dollars into output and for the knowledge spillovers it generates—raising output by at least ten per cent.

This paper is a first step in providing a structural economic framework for discussions surrounding the allocation of scientific funding. As such, many extensions naturally follow. First, the model assumes scientists within a field share the same average productivity. But scientific output is in fact highly skewed, with a small share of researchers produc-

ing the majority of the output. Allowing productivity to vary within fields would refine the reallocation estimates, since the scientists a funding shock attracts need not match those they displace, even when coming from the same field of origin. Second, the model collapses a scientist's outside option into a single reduced-form alternative. Attracting researchers to the NIH carries an opportunity cost at other funders—the National Science Foundation, private foundations, or industry—whose research output the model does not track. Third, the model is static. It abstracts from the inter-temporal trade-off between funding established investigators with higher productivity and training younger scientists, including PhD students, and improving their future productivity Qiu, 2023. While the metrics this paper considers to evaluate the efficiency of funding allocation cover a wide range of social objectives, they are not exhaustive. For example, the model does not capture clinical outcomes beyond disease burden.

References

Acemoglu, Daron, Ufuk Akcigit, and William R. Kerr (Oct. 2016). "Innovation network". en. In: *Proceedings of the National Academy of Sciences* 113.41, pp. 11483–11488. ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.1613559113. URL: https://pnas.org/doi/full/10.1073/pnas.1613559113 (visited on 09/26/2025).

Adda, Jérôme and Marco Ottaviani (Sept. 2023). "Grantmaking, Grading on a Curve, and the Paradox of Relative Evaluation in Nonmarkets*". In: *The Quarterly Journal of Economics*, qjad046. ISSN: 0033-5533. DOI: 10.1093/qje/qjad046. URL: https://doi.org/10.1093/qje/qjad046 (visited on 10/12/2023).

Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde (Jan. 2021). "Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth". In: *The Review of Economic Studies* 88.1, pp. 1–43. ISSN: 0034-6527. DOI: 10.1093/restud/rdaa061. URL: https://doi.org/10.1093/restud/rdaa061 (visited on 10/03/2023).

Arora, Ashish, Sharon Belenzon, and Andrea Patacconi (Jan. 2018). "The decline of science in corporate R&D". en. In: *Strategic Management Journal* 39.1, pp. 3–32. ISSN: 01432095. DOI: 10.1002/smj.2693. URL: https://onlinelibrary.wiley.com/doi/10.1002/smj.2693 (visited on 07/23/2023).

Azoulay, Pierre, Joshua S. Graff Zivin, and Gustavo Manso (2011). "Incentives and creativity: evidence from the academic life sciences". en. In: *The RAND Journal of Economics* 42.3. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1756-2171.2011.00140.x, pp. 527–554. ISSN: 1756-2171. DOI: 10.1111/j.1756-2171.2011.00140.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1756-2171.2011.00140.x (visited on 05/23/2024).

Azoulay, Pierre et al. (Jan. 2019). "Public R&D Investments and Private-sector Patenting: Evidence from NIH Funding Rules". In: *The Review of Economic Studies* 86.1, pp. 117–152. ISSN: 0034-6527. DOI: 10.1093/restud/rdy034. URL: https://doi.org/10.1093/restud/rdy034 (visited on 09/11/2023).

Babina, Tania et al. (Aug. 2022). *Cutting the Innovation Engine: How Federal Funding Shocks Affect University Patenting, Entrepreneurship, and Publications*. en. SSRN Scholarly Paper. Rochester, NY. DOI: 10.2139/ssrn.3560195. URL: https://papers.ssrn.com/abstract=3560195 (visited on 08/31/2023).

Ballreich, Jeromie M. et al. (Jan. 2021). "Allocation of National Institutes of Health Funding by Disease Category in 2008 and 2019". In: *JAMA Network Open* 4.1, e2034890. ISSN: 2574-3805. DOI: 10.1001/jamanetworkopen.2020.34890. URL: https://doi.org/10.1001/jamanetworkopen.2020.34890 (visited on 12/08/2024).

Bertolotti, Fabio, Kyle Myers, and Wei Yang Tham (n.d.). "Productivity Beliefs and Efficiency in Science". en. In: ().

Bisias, Dimitrios, Andrew W. Lo, and James F. Watkins (May 2012). "Estimating the NIH Efficient Frontier". en. In: *PLoS ONE* 7.5. Ed. by Eshel Ben-Jacob, e34569. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0034569. URL: https://dx.plos.org/10.1371/journal.pone.0034569 (visited on 12/08/2024).

Bloom, Nicholas, Mark Schankerman, and John Van Reenen (2013). "Identifying Technology Spillovers and Product Market Rivalry". en. In: *Econometrica* 81.4. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA9466, pp. 1347–1393. ISSN: 1468-0262. DOI: 10.3982/ECTA9466. URL: https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA9466 (visited on 12/21/2024).

Bloom, Nicholas et al. (Apr. 2020). "Are Ideas Getting Harder to Find?" en. In: *American Economic Review* 110.4, pp. 1104–1144. ISSN: 0002-8282. DOI: 10.1257/aer.20180338. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20180338 (visited on 09/01/2023).

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (Dec. 2024). *A Practical Guide to Shift-Share Instruments*. Working Paper. DOI: 10.3386/w33236. URL: https://www.nber.org/papers/w33236 (visited on 12/09/2024).

Boudreau, Kevin J. et al. (Sept. 2016). "Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance, Novelty, and Resource Allocation in Science". In: *Management Science* 62.10. Publisher: INFORMS, pp. 2765–2783. ISSN: 0025-1909. DOI: 10.1287/mnsc.2015.2285. URL: https://pubsonline.informs.org/doi/abs/10.1287/mnsc.2015.2285 (visited on 03/04/2024).

Carson, Richard T., Joshua S. Graff Zivin, and Jeffrey G. Shrader (July 2023). *Choose Your Moments: Peer Review and Scientific Risk Taking*. en. Tech. rep. w31409. National Bureau of Economic Research. DOI: 10.3386/w31409. URL: https://www.nber.org/papers/w31409 (visited on 11/18/2025).

Dyevre, Arnaud (n.d.). "Public R&D Spillovers and Productivity Growth". en. In: ().

Fleming, L. et al. (June 2019). "Government-funded research increasingly fuels innovation". en. In: *Science* 364.6446, pp. 1139–1141. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aaw2373. URL: https://www.science.org/doi/10.1126/science.aaw2373 (visited on 11/06/2022).

Gillum, Leslie A. et al. (2011). "NIH Disease Funding Levels and Burden of Disease". en. In: *PLOS ONE* 6.2. Publisher: Public Library of Science, e16837. ISSN: 1932-6203. DOI: 10.1371/journal.pone.0016837. URL: https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0016837 (visited on 01/24/2025).

Griliches, Zvi (1979). "Issues in Assessing the Contribution of Research and Development to Productivity Growth". In: *The Bell Journal of Economics* 10.1, pp. 92–116. ISSN: 0361-915X. DOI: 10.2307/3003321. JSTOR: 3003321. URL: https://www.jstor.org/stable/3003321 (visited on 05/31/2026).

Gross, C. P., G. F. Anderson, and N. R. Powe (June 1999). "The relation between funding by the National Institutes of Health and the burden of disease". eng. In: *The New England Journal of Medicine* 340.24, pp. 1881–1887. ISSN: 0028-4793. DOI: 10.1056/NEJM199906173402406.

Howell, Sabrina T. et al. (Mar. 2023). *Opening up Military Innovation: Causal Effects of 'Bottom-Up' Reforms to U.S. Defense Research*. en. SSRN Scholarly Paper. Rochester, NY. DOI: 10.2139/ssrn.3825034. URL: https://papers.ssrn.com/abstract=3825034 (visited on 10/13/2023).

Jones, Benjamin F. (Jan. 2009). "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?" In: *The Review of Economic Studies* 76.1, pp. 283–317. ISSN: 0034-6527. DOI: 10.1111/j.1467-937X.2008.00531.x. URL: https://doi.org/10.1111/j.1467-937X.2008.00531.x (visited on 09/24/2023).

Jones, Benjamin F (Feb. 2010). "Age and Great Invention". In: *The Review of Economics and Statistics* 92.1, pp. 1–14. ISSN: 0034-6535. DOI: 10.1162/rest.2009.11724. URL: https://doi.org/10.1162/rest.2009.11724 (visited on 09/24/2023).

Krieger, Joshua L., Kyle R. Myers, and Ariel D. Stern (May 2023). "How Important Is Editorial Gatekeeping? Evidence from Top Biomedical Journals". In: *The Review of Economics and Statistics*, pp. 1–33. ISSN: 0034-6535. DOI: 10.1162/rest_a_01340. URL: https://doi.org/10.1162/rest_a_01340 (visited on 10/05/2023).

Lanei, Jacqueline N. et al. (June 2022). "Conservatism Gets Funded? A Field Experiment on the Role of Negative Information in Novel Project Evaluation". In: *Management science* 68.6, pp. 4478–4495. ISSN: 0025-1909. DOI: 10.1287/mnsc.2021.4107. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC9531843/ (visited on 11/18/2025).

Lauer, Michael S and Deepshikha Roychowdhury (Sept. 2021). "Inequalities in the distribution of National Institutes of Health research project grant funding". In: *eLife* 10. Ed. by Carlos Isales, Mone Zaidi, and Mark Peifer. Publisher: eLife Sciences Publications, Ltd, e71712. ISSN: 2050-084X. DOI: 10.7554/eLife.71712. URL: https://doi.org/10.7554/eLife.71712 (visited on 11/17/2025).

Li, Danielle (Apr. 2017). "Expertise versus Bias in Evaluation: Evidence from the NIH". en. In: *American Economic Journal: Applied Economics* 9.2, pp. 60–92. ISSN: 1945-7782. DOI: 10.1257/app.20150421. URL: https://www.aeaweb.org/articles?id=10.1257/app.20150421 (visited on 09/28/2024).

Liu, Ernest and Song Ma (Dec. 2021). *Innovation Networks and R&D Allocation*. Working Paper. DOI: 10.3386/w29607. URL: https://www.nber.org/papers/w29607 (visited on 09/18/2023).

Masclans, Roger, Sharique Hasan, and Wesley M. Cohen (June 16, 2023). *The Commercial Potential of Science and Its Realization: Evidence From a Measure Using a Large Language Model*. DOI: 10.2139/ssrn.4481928. URL: https://papers.ssrn.com/abstract=4481928 (visited on 08/10/2023). Pre-published.

McFadden, Daniel (Nov. 1, 1974). "The Measurement of Urban Travel Demand". In: *Journal of Public Economics* 3.4, pp. 303–328. ISSN: 0047-2727. DOI: 10.1016/0047-2727(74)90003-6. URL: https://www.sciencedirect.com/science/article/pii/0047272774900036 (visited on 03/18/2026).

Mirin, Arthur A. (July 2021). "Gender Disparity in the Funding of Diseases by the U.S. National Institutes of Health". In: *Journal of Women's Health* 30.7, pp. 956–963. ISSN: 1540-9996. DOI: 10.1089/jwh.2020.8682. URL: https://pmc.ncbi.nlm.nih.gov/articles/PMC8290307/ (visited on 11/17/2025).

Myers, Kyle (Oct. 2020). "The Elasticity of Science". en. In: *American Economic Journal: Applied Economics* 12.4, pp. 103–134. ISSN: 1945-7782. DOI: 10.1257/app.20180518. URL: https://www.aeaweb.org/articles?id=10.1257/app.20180518 (visited on 09/26/2023).

Myers, Kyle R. and Lauren Lanahan (July 2022). "Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy". en. In: *American Economic Review* 112.7, pp. 2393–2423. ISSN: 0002-8282. DOI: 10.1257/aer.20210678. URL: https://www.aeaweb.org/articles?id=10.1257/aer.20210678 (visited on 04/09/2024).

National Institutes of Health (n.d.). *Report on NIH Funding vs. Global Burden of Disease*. URL: https://report.nih.gov/report-nih-funding-vs-global-burden-disease (visited on 04/17/2026).

Packalen, Mikko and Jay Bhattacharya (July 2018). *Does the NIH Fund Edge Science?* Working Paper. DOI: 10.3386/w24860. URL: https://www.nber.org/papers/w24860 (visited on 03/03/2024).

Park, Michael, Erin Leahey, and Russell J. Funk (Jan. 2023). "Papers and patents are becoming less disruptive over time". en. In: *Nature* 613.7942. Number: 7942 Publisher: Nature Publishing Group, pp. 138–144. ISSN: 1476-4687. DOI: 10.1038/s41586-022-05543-x. URL: https://www.nature.com/articles/s41586-022-05543-x (visited on 09/15/2023).

Poege, Felix et al. (Dec. 2019). "Science quality and the value of inventions". en. In: *Science Advances* 5.12, eaay7323. ISSN: 2375-2548. DOI: 10.1126/sciadv.aay7323. URL: https://www.science.org/doi/10.1126/sciadv.aay7323 (visited on 11/17/2022).

Qiu, Y. J. Jeff (2023). "The Matthew effect, research productivity, and the dynamic allocation of NIH grants". en. In: *The RAND Journal of Economics* 54.1. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1756-2171.12433, pp. 135–164. ISSN: 1756-2171. DOI: 10.1111/1756-2171.12433. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12433 (visited on 05/10/2025).

Sampat, Bhaven N. (Dec. 2012). "Mission-oriented biomedical research at the NIH". In: *Research Policy*. The need for a new generation of policy instruments to respond to the

Grand Challenges 41.10, pp. 1729–1741. ISSN: 0048-7333. DOI: 10.1016/j.respol.2012.05.013. URL: https://www.sciencedirect.com/science/article/pii/S0048733312002168 (visited on 01/30/2024).

Sampat, Bhaven N., Kristin Buterbaugh, and Marcel Perl (Mar. 2013). "New Evidence on the Allocation of NIH Funds across Diseases". en. In: *The Milbank Quarterly* 91.1, pp. 163–185. ISSN: 0887-378X, 1468-0009. DOI: 10.1111/milq.12005. URL: https://onlinelibrary.wiley.com/doi/10.1111/milq.12005 (visited on 12/08/2024).

Taffe, Michael A. and Nicholas W. Gilpin (Jan. 2021). "Racial inequity in grant funding from the US National Institutes of Health". eng. In: *eLife* 10, e65697. ISSN: 2050-084X. DOI: 10.7554/eLife.65697.

Yin, Yian et al. (Oct. 2022). "Public use and public funding of science". en. In: *Nature Human Behaviour* 6.10. Number: 10 Publisher: Nature Publishing Group, pp. 1344–1350. ISSN: 2397-3374. DOI: 10.1038/s41562-022-01397-5. URL: https://www.nature.com/articles/s41562-022-01397-5 (visited on 10/05/2023).

A Supplementary Figures and Tables

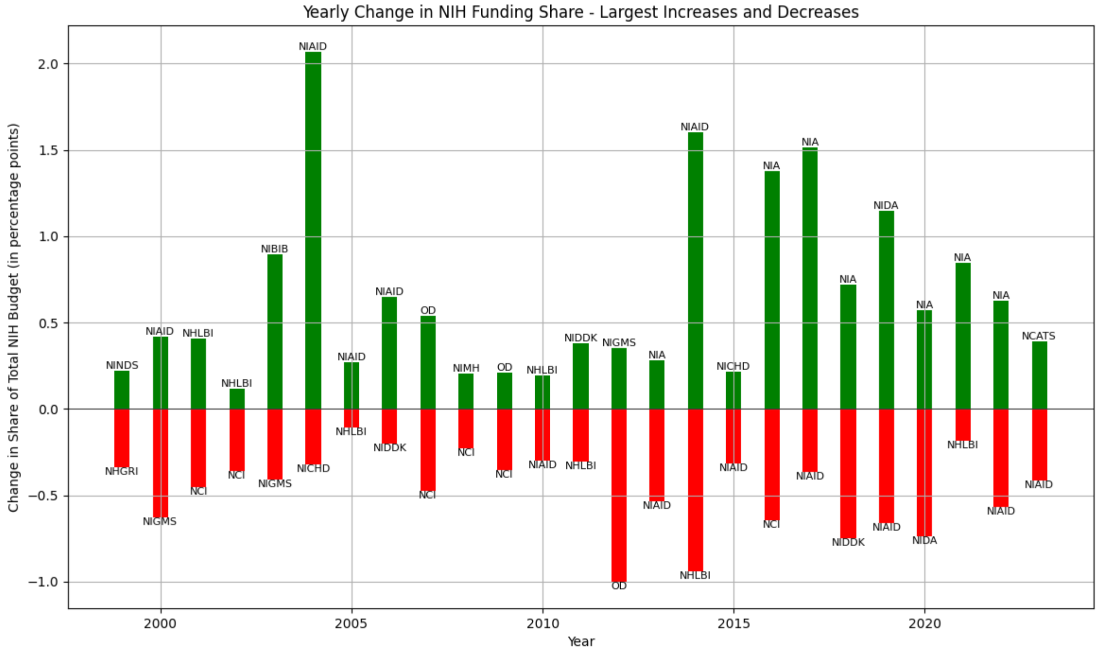


Figure 12: Largest annual shocks in budget share (positive and negative).
Largest positive and negative changes in any institute's budget share in each year. Large drops are both rarer and smaller than large increases.

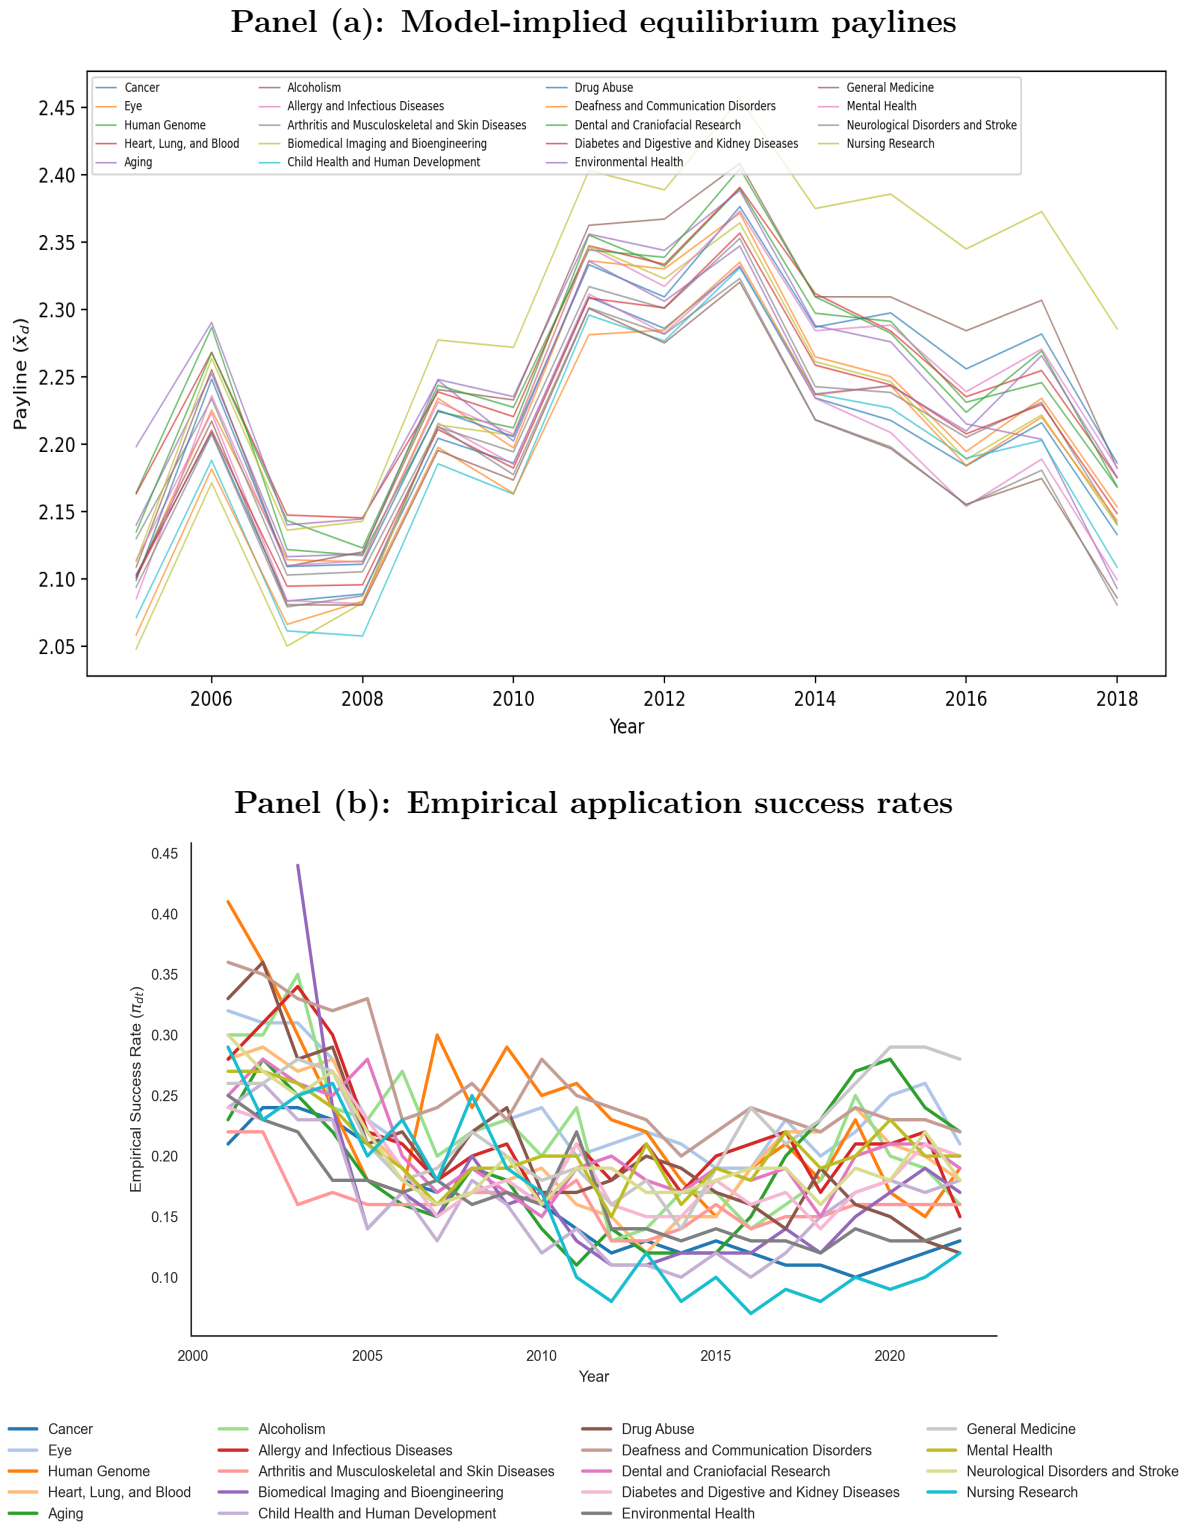


Figure 13: Model-implied equilibrium paylines and empirical success rates over time.

Panel (a) plots the estimated equilibrium paylines $\bar{x}_{dt}$ for each institute across the sample period 2005–2018. Higher values indicate more stringent funding thresholds required to clear the market given the institute's budget. Panel (b) plots the corresponding empirical grant application success rates π_{dt} over the wider 2001–2022 window for comparison.

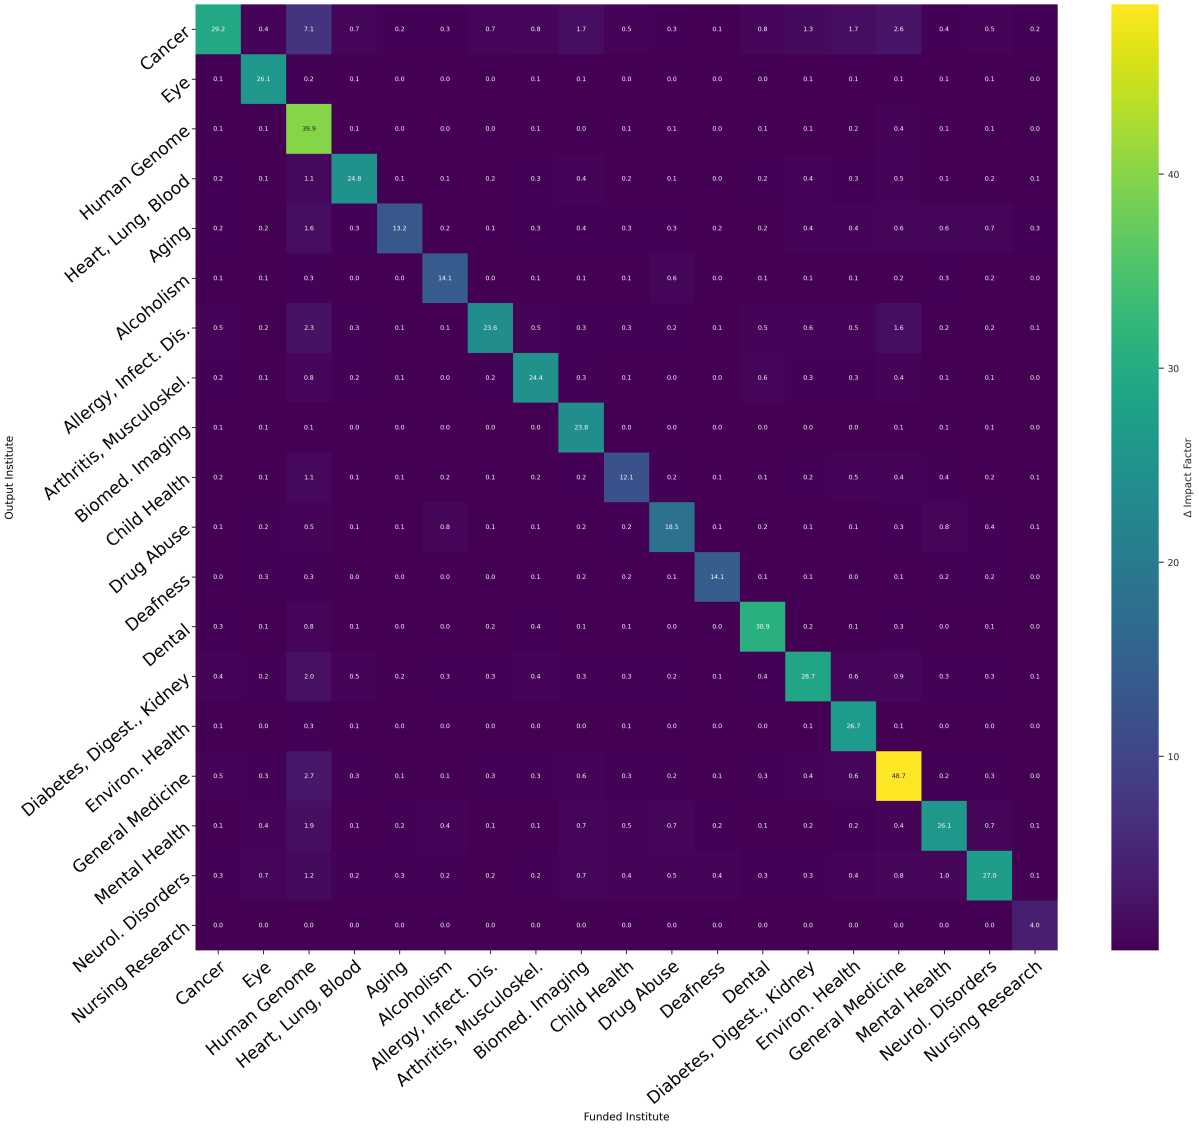


Figure 14: Heatmap of total net partial effects per $1 million.

Each cell evaluates the total absolute change in impact-factor-weighted output (including no externality return, talent flow externality, and knowledge spillovers) of the row institute resulting from a $1 million funding increase to the column institute.

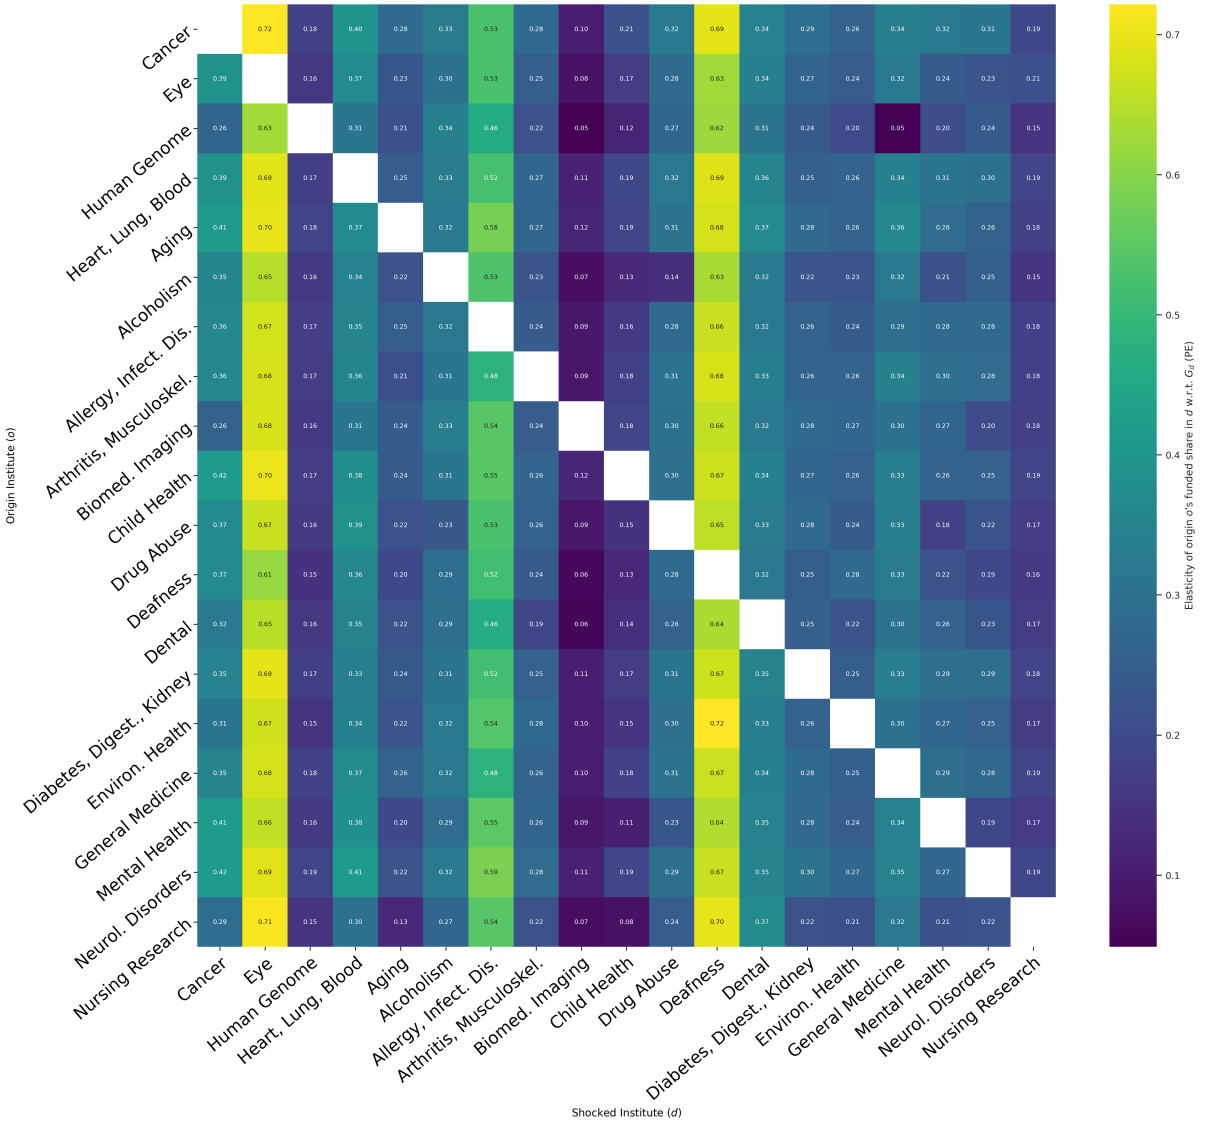


Figure 15: Partial-equilibrium elasticity of origin funded-scientist share with respect to destination funding.

Each cell reports the partial-equilibrium (PE) elasticity of the funded-scientist share of origin o (row) inside destination d (column) with respect to a 1% increase in d's budget. The diagonal is masked to focus on cross-field inflows. A positive value indicates that more funding in d raises the fraction of d's funded scientists that come from field o. The bar chart in Figure 16 summarises the total off-field share elasticity for each destination by summing over all cross-field origins.

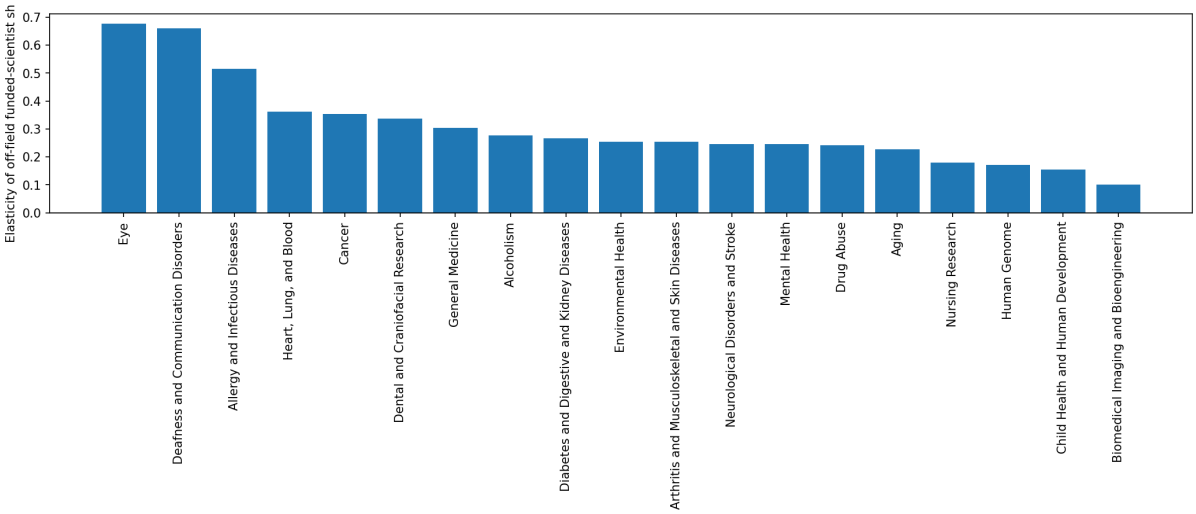


Figure 16: Partial-equilibrium elasticity of the off-field funded-scientist share.

For each funded institute d, the bar reports the elasticity of the share of d's funded scientists coming from other fields with respect to a 1% increase in d's budget G_d, computed in partial equilibrium (only d's payline adjusts). Positive values confirm that higher funding in d attracts scientists from outside fields. The origin breakdown is provided in Figure 15.

Table 12: List of NIH institutes and centres

Acronym	Official Name	Name in this paper
FIC	Fogarty International Center	NA
NCCIH	National Center for Complementary and Integrative Health	Complementary and Integrative Health Medicine
NCI	National Cancer Institute	Cancer
NCRR	National Center for Research Resources	NA
NIDCD	National Institute on Deafness and Other Communication Disorders	Deafness and Communication disorders
NEI	National Eye Institute	Eye
NHGRI	National Human Genome Research Institute	Human Genome
NHLBI	National Heart, Lung, and Blood Institute	Heart, Lung, and Blood
NIA	National Institute on Aging	Aging
NIAAA	National Institute on Alcohol Abuse and Alcoholism	Alcoholism
NIAID	National Institute of Allergy and Infectious Diseases	Allergy and Infectious Diseases
NIAMS	National Institute of Arthritis and Musculoskeletal and Skin Diseases	Arthritis and Musculoskeletal and Skin Diseases
NIBIB	National Institute of Biomedical Imaging and Bioengineering	Biomedical Imaging and Bioengineering
NICHD	National Institute of Child Health and Human Development	Child Health and Human Development
NIDA	National Institute on Drug Abuse	Drug Abuse
NIDCR	National Institute of Dental and Craniofacial Research	Dental and Craniofacial Research
NIDDK	National Institute of Diabetes and Digestive and Kidney Diseases	Diabetes and Digestive and Kidney Diseases
NIEHS	National Institute of Environmental Health Sciences	Environmental Health
NIGMS	National Institute of General Medical Sciences	General Medicine
NIMH	National Institute of Mental Health	Mental Health
NINDS	National Institute of Neurological Disorders and Stroke	Neurological Disorders and Stroke
NINR	National Institute of Nursing Research	Nursing Research

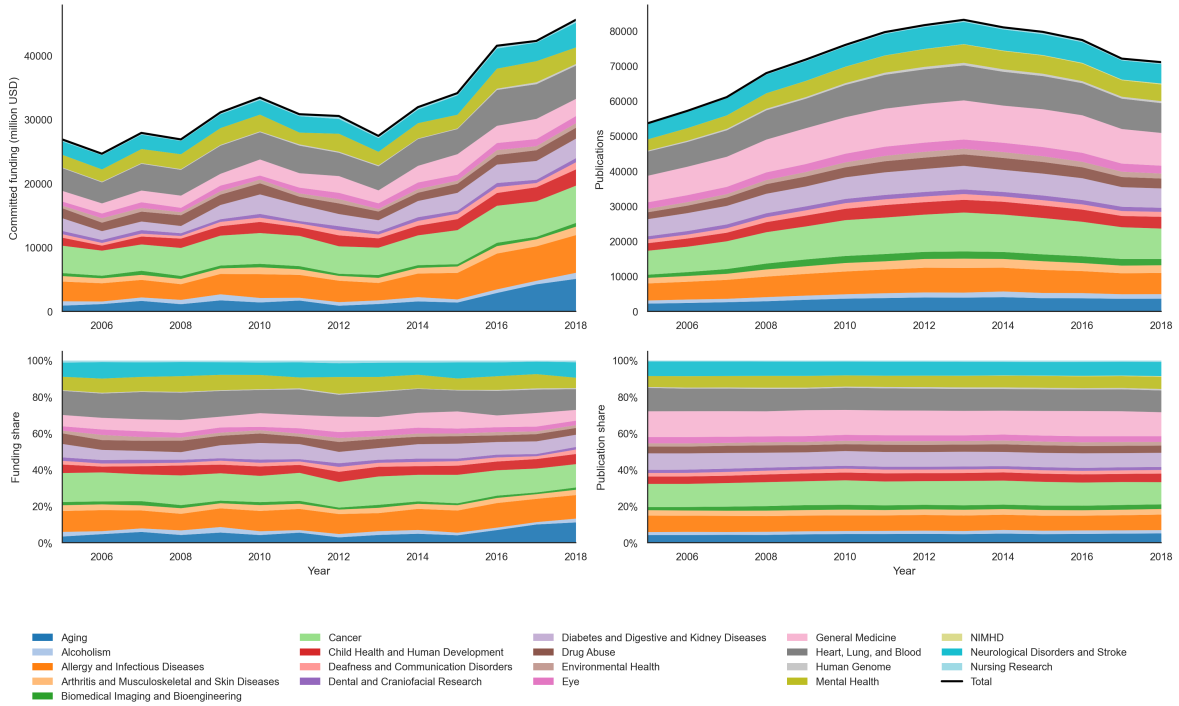


Figure 17: Funding and publications over time, stacked by institute.

The figure reports funding levels and publication levels (top row), and corresponding shares (bottom row), stacked by NIH institute. Funding is measured in 2017 real USD and publication counts are based on grant-linked outputs.

Year	Unique grants	Grant size (2017 USD)	Publications (IFW)	Unique PIs	Mover share	Success rate
2001	3946	1742601.47	82.53	4002	0.53	0.27
2002	3833	1784467.65	81.97	3923	0.55	0.27
2003	4230	1817123.93	88.07	4332	0.54	0.27
2004	4018	1757951.74	86.51	4098	0.52	0.26
2005	3603	1702198.57	89.25	3685	0.52	0.21
2006	3456	1630520.90	92.70	3586	0.52	0.19
2007	3878	1656982.07	103.95	4017	0.51	0.17
2008	3628	1683337.67	98.14	3970	0.52	0.19
2009	4344	1727089.78	85.14	4896	0.51	0.19
2010	3673	2022403.47	95.58	4300	0.53	0.17
2011	3333	2049200.67	100.86	4061	0.52	0.18
2012	3470	1962765.32	95.66	4300	0.50	0.15
2013	3181	1928402.13	95.32	3963	0.51	0.16
2014	3549	1998690.28	86.25	4481	0.49	0.15
2015	3845	1970543.45	82.13	4858	0.49	0.17
2016	4303	2129289.38	73.42	5493	0.49	0.17
2017	4344	2149465.87	67.44	5574	0.49	0.18
2018	4661	2145106.50	55.37	6027	0.47	0.17
2019	4544	2253655.13	46.97	5971	0.48	0.20
2020	4474	2310500.06	37.09	5893	0.47	0.20
2021	3865	2323408.33	24.57	5234	0.47	0.20
2022	3112	2221194.66	15.35	4258	0.48	0.18

Table 13: Descriptive statistics: annual summary.

Annual summary of key aggregate statistics. Award sizes are in thousands of 2017 real USD. Mover share and success rate are in percentage points.

Institute	Number of grants	Grant size (2017 USD)	Publications (IFW)	Share of movers	Success rate
Cancer	12078	1947150.48	80.55	0.43	0.15
Eye	2598	1761559.87	69.23	0.40	0.23
Human Genome	610	1830909.29	131.05	0.73	0.22
Heart, Lung, and Blood	11075	2114738.26	74.01	0.43	0.20
Aging	4599	2566409.59	76.64	0.59	0.21
Alcoholism	1521	1971762.21	49.43	0.55	0.21
Allergy and Infectious Diseases	8638	2148098.08	77.18	0.44	0.21
Arthritis and Musculoskeletal and Skin Diseases	2367	1869750.50	75.02	0.52	0.16
Biomedical Imaging and Bioengineering	1559	1850124.02	87.83	0.78	0.17
Child Health and Human Development	3826	2035851.47	47.64	0.63	0.16
Drug Abuse	3334	2225904.73	55.74	0.51	0.21
Deafness and Communication Disorders	1638	1891300.68	55.93	0.43	0.26
Dental and Craniofacial Research	1352	1912768.42	70.89	0.50	0.20
Diabetes and Digestive and Kidney Diseases	6894	1732563.51	75.35	0.50	0.19
Environmental Health	1419	1937666.67	78.56	0.56	0.16
General Medicine	8832	1313299.09	101.30	0.58	0.22
Mental Health	5351	2284202.49	72.45	0.48	0.20
Neurological Disorders and Stroke	6883	1869267.66	74.75	0.54	0.20
Nursing Research	716	2144937.56	27.56	0.59	0.15

Table 14: Descriptive statistics by funding institute.

Descriptive statistics by funding institute. Award sizes are in thousands of 2017 real USD. Mover share and success rate are in percentage points.

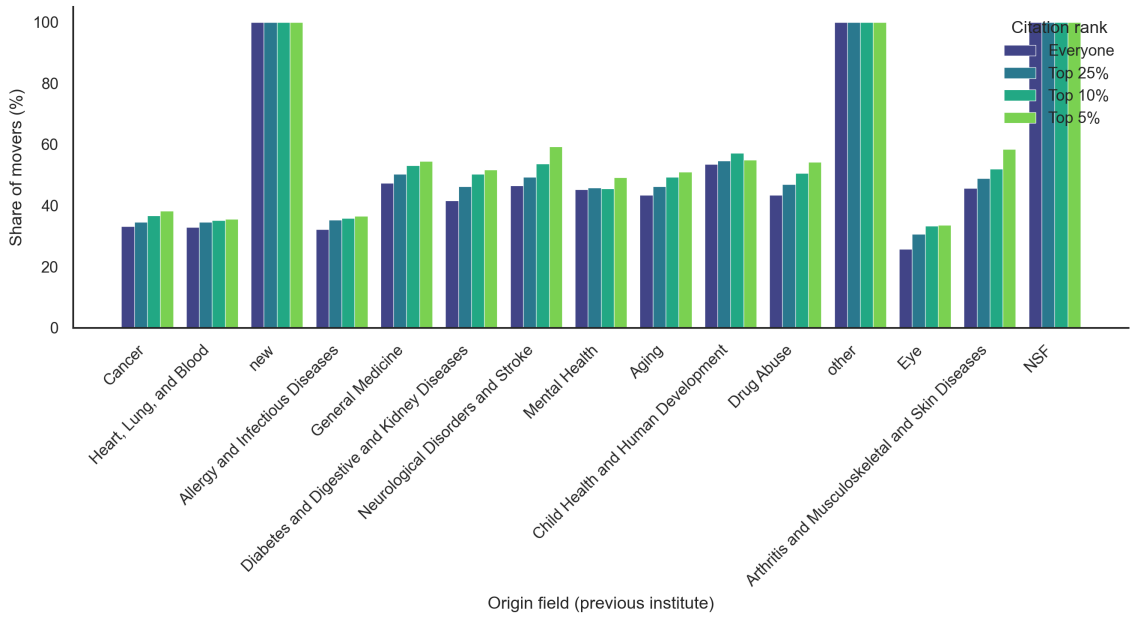
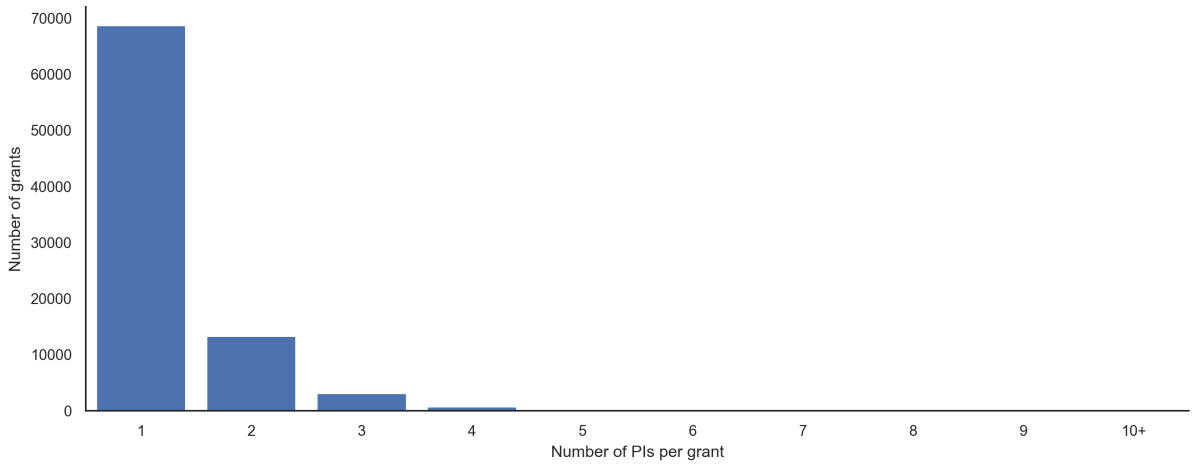
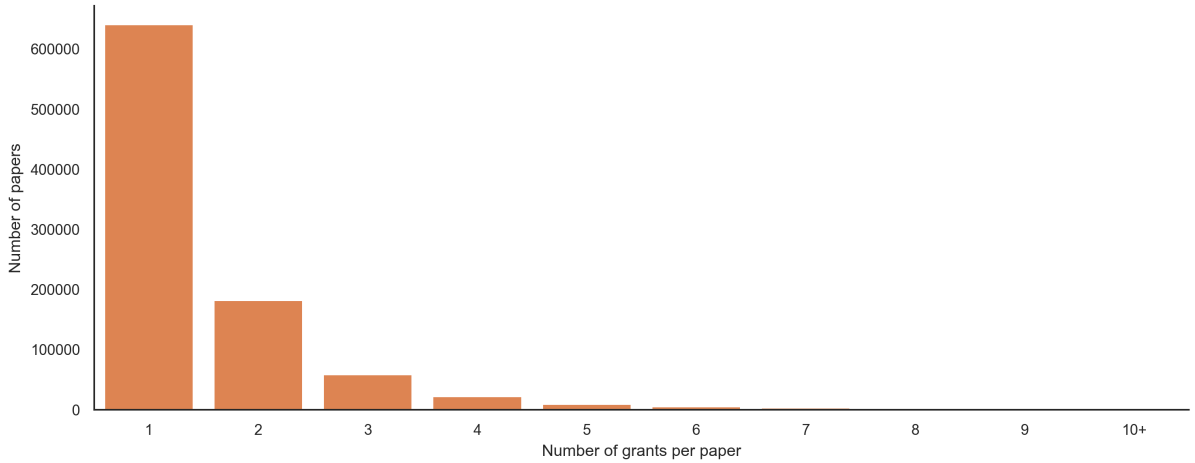


Figure 18: Share of PIs who moved between funding institutes, by citation rank.

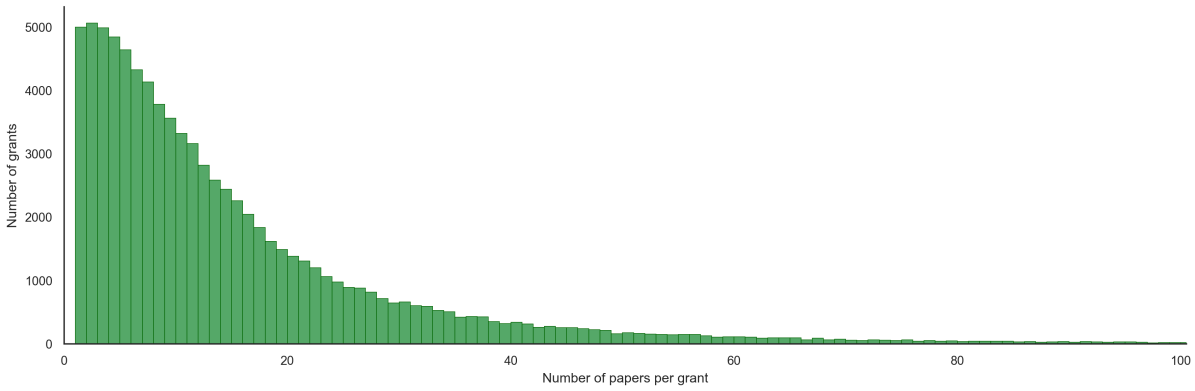
Share of investigators whose current funding institute differs from their previous funder, computed separately for the full sample and for investigators in the top 25%, top 10% and top 5% of lifetime citations within their field-year. The citation percentile is the within-(primary institute, year) rank of the investigator's pre-grant citation stock. Only cells with at least 20 observations are displayed.



(a) PIs per grant



(b) Grants per paper



(c) Papers per grant

Figure 19: Many-to-many links between PIs, grants, and publications.
Panel (a): distribution of the number of PIs per grant. Panel (b): distribution of the number of grants acknowledged per paper. Panel (c): distribution of the number of papers produced per grant, truncated at 100 for display.

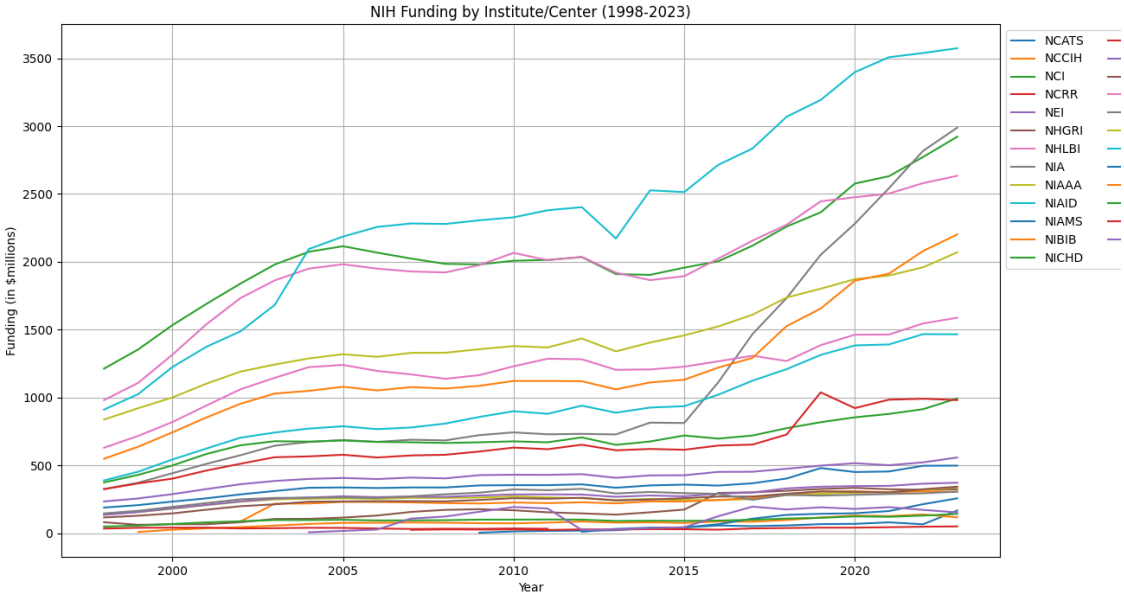


Figure 20: Nominal funding trends by NIH institute.

Nominal funding by NIH institute. Source: NIH budget data.

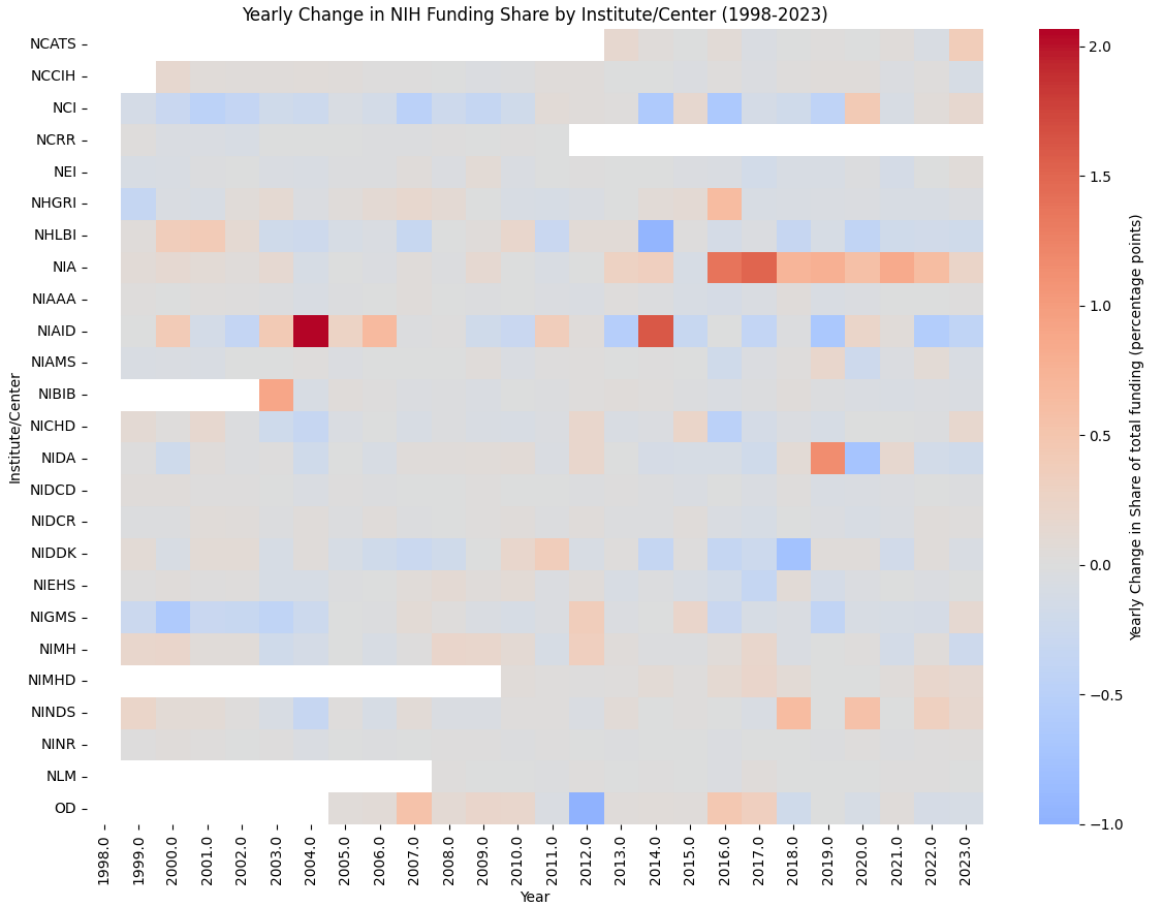


Figure 21: Yearly change in institutes' share of total NIH funding.

Yearly change in each institute's share of the total NIH budget. Colour intensity indicates the magnitude and direction of the change.

B Derivation of the Realised Quality Gap Distribution

This section derives the distribution of the realised quality gap z_{io}, conditional on a scientist from origin o choosing destination d, and defines the inclusive value M_o.

Assuming the idiosyncratic shocks ε_{iod} are drawn independently from a Type-I Extreme Value (Gumbel) distribution with location 0 and scale $\sigma > 0$, standard discrete choice properties imply that the distribution of the realised shock, conditional on choice d, is identical to the unconditional distribution of the maximum shock across all alternatives.

Consequently, the realised quality gap $z_{io} \equiv \mu_{od} - \bar{x}_d + \varepsilon_{iod}$, conditional on choosing d, follows a Gumbel distribution with scale σ and location parameter given by the inclusive value M_o:

$$M_o \;=\; \sigma \ln\left(\exp\left(\frac{c}{\sigma}\right) + \sum_{j=1}^{C} \exp\left(\frac{\mu_{oj} - \bar{x}_j}{\sigma}\right)\right), \tag{44}$$

where c represents the outside option. The conditional density of the realised quality gap z is therefore:

$$f_Z(z; M_o) \;=\; \frac{1}{\sigma} \exp\left(-\frac{z - M_o}{\sigma} - \exp\left(-\frac{z - M_o}{\sigma}\right)\right). \tag{45}$$

This density is shifted right when M_o increases (i.e. when scientists from origin o face a more favourable landscape overall, their chosen quality gaps tend to be larger).

C Equilibrium Existence and Uniqueness

C.1 Existence

Proposition 1 (Existence of an Equilibrium)**.** *Let $s_d = G_d/\bar{g}_d$ denote the number of available grant slots in destination d. Suppose $s_d > 0$ for all d and the aggregate slot vector is feasible, so that $\sum_{d=1}^{C} s_d < \sum_{o=1}^{C} N_o$. Suppose further that there exists a compact rectangle $\mathcal{X} = \prod_{d=1}^{C} [L_d, U_d]$ such that, for every destination d,*

$$E_d(L_d, \bar{x}_{-d}) \geq s_d \quad \textit{and} \quad E_d(U_d, \bar{x}_{-d}) \leq s_d \quad \textit{for all } \bar{x}_{-d} \in \prod_{c \neq d} [L_c, U_c]. \tag{46}$$

Then there exists an equilibrium payline vector $\bar{x}^* \in \mathbb{R}^C$ solving (10).

Proof. Let the excess demand function for destination d be $Z_d(\bar{x}) = E_d(\bar{x}) - s_d$, where $E_d(\bar{x}) = \sum_{o=1}^{C} N_o \Pi_{od}(\bar{x}) S_{od}(M_o(\bar{x}))$. The equilibrium condition is $Z(\bar{x}) = \mathbf{0}$. The aggregate feasibility condition is necessary because each scientist can apply to at most one

destination and $S_{od} \leq 1$, so $\sum_d E_d(\bar{x}) \leq \sum_o N_o$ for any payline vector.

The stated face conditions are the uniform version of the one-dimensional limit argument. On an upper face, the application probability to destination d vanishes as $\bar{x}_d$ becomes large, uniformly over bounded values of $\bar{x}_{-d}$. Since $0 \leq S_{od} \leq 1$, this makes E_d arbitrarily small. The lower-face condition imposes the corresponding capacity requirement: when destination d's payline is sufficiently loose relative to the other coordinates in $\mathcal{X}$, expected funded mass in d weakly exceeds its slot capacity.

The functions $Z_d(\bar{x})$ are continuous because the logit probabilities, the inclusive value, and the success integrals are continuous in $\bar{x}$. By assumption, $Z_d(L_d, \bar{x}_{-d}) \geq 0$ and $Z_d(U_d, \bar{x}_{-d}) \leq 0$ for every d and every $\bar{x}_{-d}$ on the corresponding face of $\mathcal{X}$. The Poincaré-Miranda theorem therefore implies that there is some $\bar{x}^* \in \mathcal{X}$ such that $Z(\bar{x}^*) = \mathbf{0}$. This vector clears all grant markets. □

C.2 Uniqueness

To establish uniqueness of the equilibrium, we impose a parameter restriction on the magnitude of the indirect quality effect across all alternatives.

Proposition 2 (Global Uniqueness of Equilibrium)**.** *Under Assumption 1 (Dominant Substitution Effect), which requires $\frac{1}{\sigma} S_{od}(M_o) > S'_{od}(M_o)$ for all o, d throughout the relevant equilibrium domain, the equilibrium $\bar{x}^*$ is globally unique.*

Proof. By the Gale–Nikaido univalence theorem, a continuously differentiable mapping defined on a rectangular region is globally injective if its Jacobian matrix is a P-matrix everywhere on that region. We apply this theorem to the mapping $-Z$, not to Z itself. Since the Jacobian of Z is $J(\bar{x})$, the Jacobian of $-Z$ is $-J(\bar{x})$. We show that $-J(\bar{x})$ is an M-matrix, which is a subclass of P-matrices.

Recall the elements of the Jacobian J derived in Section D:

$$J_{dd} = -\sum_{o=1}^{C} N_o \Pi_{od} \left[\frac{1}{\sigma}(1 - \Pi_{od}) S_{od} + \Pi_{od} S'_{od} \right]$$

$$J_{dj} = \sum_{o=1}^{C} N_o \Pi_{od} \Pi_{oj} \left[\frac{1}{\sigma} S_{od} - S'_{od} \right] \quad \text{for } j \neq d.$$

Under Assumption 1, $J_{dj} > 0$ for all strictly off-diagonal elements, making $-J_{dj} < 0$. Thus, $-J$ has positive diagonal elements and non-positive off-diagonal elements.

To determine if $-J$ is strictly diagonally dominant, we sum the entries of a row of J:

$$\begin{aligned}\sum_{j=1}^{C} J_{dj} &= J_{dd} + \sum_{j \neq d} J_{dj} \\ &= \sum_{o=1}^{C} N_o \Pi_{od} \left[-\frac{1}{\sigma}(1 - \Pi_{od}) S_{od} - \Pi_{od} S'_{od} + \sum_{j \neq d} \Pi_{oj} \left(\frac{1}{\sigma} S_{od} - S'_{od} \right) \right] \\ &= \sum_{o=1}^{C} N_o \Pi_{od} \left[\frac{1}{\sigma} S_{od} \left(\sum_{j \neq d} \Pi_{oj} - (1 - \Pi_{od}) \right) - S'_{od} \left(\Pi_{od} + \sum_{j \neq d} \Pi_{oj} \right) \right].\end{aligned}$$

Let Π_{o0} denote the probability of choosing the outside option. By the properties of the multinomial logit, $\sum_{j=1}^{C} \Pi_{oj} + \Pi_{o0} = 1$. Therefore, $\sum_{j \neq d} \Pi_{oj} = 1 - \Pi_{od} - \Pi_{o0}$. Substituting this yields:

$$\sum_{j=1}^{C} J_{dj} = \sum_{o=1}^{C} N_o \Pi_{od} \left[-\frac{1}{\sigma} S_{od} \Pi_{o0} - S'_{od}(1 - \Pi_{o0}) \right]. \tag{47}$$

Since $\Pi_{o0} > 0$, $S_{od} > 0$, and $S'_{od} > 0$, we have $\sum_{j=1}^{C} J_{dj} < 0$ for every row d. This implies that $-J$ is strictly diagonally dominant by rows.

Any strictly diagonally dominant matrix with non-positive off-diagonal elements and positive diagonal elements is a non-singular M-matrix, and consequently a P-matrix. Because $-J(\bar{x})$ is a P-matrix everywhere, Gale–Nikaido implies that $-Z$ is injective on any rectangle in $\mathbb{R}^C$. Since injectivity of $-Z$ is equivalent to injectivity of Z, the system $Z(\bar{x}) = \mathbf{0}$ admits at most one solution. Combined with existence, this establishes that $\bar{x}^*$ is globally unique. □

D Comparative Statics Derivation

This section derives every component of $\text{CS1}_{hd} = \frac{d \ln Y_h}{d \ln G_d}$ step-by-step, starting from the primitives of the model.

Step 1: Total derivative via the chain rule. Total production Y_h depends on the budget vector G only through the equilibrium paylines $\bar{x}(G)$. By the chain rule:

$$\frac{dY_h}{dG_d} = \sum_{j=1}^{C} \frac{\partial Y_h}{\partial \bar{x}_j} \cdot \frac{d\bar{x}_j}{dG_d}. \tag{48}$$

Converting to an elasticity by multiplying both sides by G_d/Y_h:

$$\frac{d \ln Y_h}{d \ln G_d} = \frac{G_d}{Y_h} \sum_{j=1}^{C} \frac{\partial Y_h}{\partial \bar{x}_j} \cdot \frac{d\bar{x}_j}{dG_d}. \tag{49}$$

We now derive each of the two components in the sum: the partial derivatives $\frac{\partial Y_h}{\partial \bar{x}_j}$ (Steps 2–5) and the GE payline adjustments $\frac{d\bar{x}_j}{dG_d}$ (Steps 6–7).

Step 2: Differentiate total production with respect to paylines. Using the W_{oh} convention in Section 4, where W_{oh} already includes the success-probability weighting from peer review, total production in h can be written as:

$$Y_h = \mathbb{E}[f(g_{ioh}, k_h)] \sum_{o=1}^{C} N_o \, \Pi_{oh}(\bar{x}) \, W_{oh}(M_o(\bar{x}), \bar{x}_h).$$

The prefactor $\mathbb{E}[f(g_{ioh}, k_h)]$ does not depend on $\bar{x}$. Inside the sum, for each origin o, the term $\Pi_{oh} \cdot W_{oh}$ is a product of two functions that both depend on $\bar{x}_j$. Applying the product rule:

$$\frac{\partial Y_h}{\partial \bar{x}_j} = \mathbb{E}[f(g_{ioh}, k_h)] \sum_{o=1}^{C} N_o \left[\frac{\partial \Pi_{oh}}{\partial \bar{x}_j} \cdot W_{oh} + \Pi_{oh} \cdot \frac{\partial W_{oh}}{\partial \bar{x}_j} \right]. \tag{50}$$

We need to compute $\frac{\partial \Pi_{oh}}{\partial \bar{x}_j}$ and $\frac{\partial W_{oh}}{\partial \bar{x}_j}$ separately.

Step 3: Differentiate the application probabilities (logit derivatives). The application probability from origin o to destination h is given by the multinomial logit formula (8):

$$\Pi_{oh}(\bar{x}) = \frac{\exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right)}{\exp\left(\frac{c}{\sigma}\right) + \sum_{k=1}^{C} \exp\left(\frac{\mu_{ok} - \bar{x}_k}{\sigma}\right)}.$$

Write the denominator as $D_o = \exp(c/\sigma) + \sum_k \exp((\mu_{ok} - \bar{x}_k)/\sigma)$, so that $\Pi_{oh} = \exp((\mu_{oh} - \bar{x}_h)/\sigma)/D_o$.

Case 1: Own derivative ($j = h$). Both the numerator and denominator depend on $\bar{x}_h$. Using the quotient rule:

$$\begin{aligned}
\frac{\partial \Pi_{oh}}{\partial \bar{x}_h} &= \frac{1}{D_o} \frac{\partial}{\partial \bar{x}_h} \exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right) - \frac{\exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right)}{D_o^2} \frac{\partial D_o}{\partial \bar{x}_h} \\
&= \frac{1}{D_o} \cdot \left(-\frac{1}{\sigma}\right) \exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right) - \frac{\exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right)}{D_o^2} \cdot \left(-\frac{1}{\sigma}\right) \exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right) \\
&= -\frac{1}{\sigma} \Pi_{oh} + \frac{1}{\sigma} \Pi_{oh}^2 \\
&= -\frac{1}{\sigma} \Pi_{oh} (1 - \Pi_{oh}).
\end{aligned}$$

Case 2: Cross derivative ($j \neq h$). The numerator $\exp((\mu_{oh} - \bar{x}_h)/\sigma)$ does not depend on $\bar{x}_j$, so only the denominator changes:

$$\begin{aligned}
\frac{\partial \Pi_{oh}}{\partial \bar{x}_j} &= -\frac{\exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right)}{D_o^2} \cdot \frac{\partial D_o}{\partial \bar{x}_j} \\
&= -\frac{\exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right)}{D_o^2} \cdot \left(-\frac{1}{\sigma}\right) \exp\left(\frac{\mu_{oj} - \bar{x}_j}{\sigma}\right) \\
&= \frac{1}{\sigma} \cdot \frac{\exp\left(\frac{\mu_{oh} - \bar{x}_h}{\sigma}\right)}{D_o} \cdot \frac{\exp\left(\frac{\mu_{oj} - \bar{x}_j}{\sigma}\right)}{D_o} \\
&= \frac{1}{\sigma} \Pi_{oh} \Pi_{oj}.
\end{aligned}$$

Step 4: Differentiate the inclusive value M_o. Before differentiating W_{oh}, we need $\frac{\partial M_o}{\partial \bar{x}_j}$. Recall from (44):

$$M_o(\bar{x}) = \sigma \ln D_o = \sigma \ln \left(\exp\left(\frac{c}{\sigma}\right) + \sum_{k=1}^{C} \exp\left(\frac{\mu_{ok} - \bar{x}_k}{\sigma}\right) \right).$$

Differentiating with respect to $\bar{x}_j$:

$$\begin{aligned}\frac{\partial M_o}{\partial \bar{x}_j} &= \sigma \cdot \frac{1}{D_o} \cdot \frac{\partial D_o}{\partial \bar{x}_j} \\ &= \sigma \cdot \frac{1}{D_o} \cdot \left(-\frac{1}{\sigma}\right) \exp\left(\frac{\mu_{oj} - \bar{x}_j}{\sigma}\right) \\ &= -\frac{\exp\left(\frac{\mu_{oj} - \bar{x}_j}{\sigma}\right)}{D_o} \\ &= -\Pi_{oj}.\end{aligned}$$

Step 5: Differentiate the funded-quality term W_{oh}. The funded-quality contribution from o in h is:

$$W_{oh}(M_o, \bar{x}_h) = \int_{-\infty}^{\infty} (z + \bar{x}_h)\, \Phi\left(\frac{z}{\eta_{oh}}\right) f_Z(z; M_o)\, dz.$$

This depends on $\bar{x}_j$ through two channels: (i) directly via $\bar{x}_h$ in the integrand (only when $j = h$), and (ii) indirectly via $M_o(\bar{x})$ in the density f_Z. We treat each in turn.

Step 5a: Direct effect of $\bar{x}_h$. Since $\bar{x}_h$ appears linearly inside the integrand:

$$\begin{aligned}\left.\frac{\partial}{\partial \bar{x}_h}\right|_{\text{direct}} W_{oh} &= \frac{\partial}{\partial \bar{x}_h} \int_{-\infty}^{\infty} (z + \bar{x}_h)\, \Phi\left(\frac{z}{\eta_{oh}}\right) f_Z(z; M_o)\, dz \\ &= \int_{-\infty}^{\infty} 1 \cdot \Phi\left(\frac{z}{\eta_{oh}}\right) f_Z(z; M_o)\, dz \\ &= S_{oh}(M_o).\end{aligned}$$

Step 5b: Indirect effect through M_o. Recall that $f_Z(z; M_o) = f_{Z_0}(z - M_o)$ where $Z_0 \sim \text{Gumbel}(0, \sigma)$. Using the change of variables $z_0 = z - M_o$, rewrite W_{oh} as:

$$W_{oh} = \int_{-\infty}^{\infty} (z_0 + M_o + \bar{x}_h)\, \Phi\left(\frac{z_0 + M_o}{\eta_{oh}}\right) f_{Z_0}(z_0)\, dz_0.$$

Since f_{Z_0} does not depend on M_o, differentiating under the integral sign yields:

$$\begin{aligned}\frac{\partial W_{oh}}{\partial M_o} &= \int_{-\infty}^{\infty} \Phi\left(\frac{z_0 + M_o}{\eta_{oh}}\right) f_{Z_0}(z_0)\, dz_0 + \int_{-\infty}^{\infty} (z_0 + M_o + \bar{x}_h)\, \frac{1}{\eta_{oh}}\, \phi\left(\frac{z_0 + M_o}{\eta_{oh}}\right) f_{Z_0}(z_0)\, dz_0 \\ &= S_{oh}(M_o) + \int_{-\infty}^{\infty} (z + \bar{x}_h)\, \frac{1}{\eta_{oh}}\, \phi\left(\frac{z}{\eta_{oh}}\right) f_Z(z; M_o)\, dz,\end{aligned}$$

where the last line reverts to the original variable $z = z_0 + M_o$. The first term arises because a rightward shift in the quality gap distribution increases the mass of funded

projects. The second term captures the change in the success-weighted quality contribution.

Step 5c: Combine using the chain rule.

For $j = h$ (own-payline):

$$\frac{\partial W_{oh}}{\partial \bar{x}_h} = \underbrace{S_{oh}(M_o)}_{\text{direct effect (Step 5a)}} + \underbrace{\frac{\partial W_{oh}}{\partial M_o}}_{\text{from Step 5b}} \cdot \underbrace{\frac{\partial M_o}{\partial \bar{x}_h}}_{=-\Pi_{oh}\text{ (Step 4)}}$$
$$= S_{oh}(M_o) - \Pi_{oh}\frac{\partial W_{oh}}{\partial M_o}.$$

For $j \neq h$ (cross-payline), the direct effect is zero because $\bar{x}_j$ does not appear in the integrand of W_{oh}:

$$\frac{\partial W_{oh}}{\partial \bar{x}_j} = \frac{\partial W_{oh}}{\partial M_o} \cdot \frac{\partial M_o}{\partial \bar{x}_j}$$
$$= -\Pi_{oj}\frac{\partial W_{oh}}{\partial M_o}.$$

Step 6: Differentiate the market-clearing condition. We now derive the GE payline adjustments $\frac{d\bar{x}_j}{dG_d}$. The market-clearing condition for field k is:

$$E_k(\bar{x}) \;\equiv\; \sum_{o=1}^{C} N_o\, \Pi_{ok}(\bar{x})\, S_{ok}(M_o(\bar{x})) \;=\; \frac{G_k}{\bar{g}_k}.$$

To find how $\bar{x}$ responds to a marginal change in G_d, totally differentiate each equation with respect to G_d:

$$\sum_{j=1}^{C} \frac{\partial E_k}{\partial \bar{x}_j}\frac{d\bar{x}_j}{dG_d} = \begin{cases} \frac{1}{\bar{g}_d} & \text{if } k = d, \\ 0 & \text{if } k \neq d. \end{cases}$$

We now compute $\frac{\partial E_k}{\partial \bar{x}_j}$. Using the product rule on $\Pi_{ok} \cdot S_{ok}$:

$$\frac{\partial E_k}{\partial \bar{x}_j} = \sum_{o=1}^{C} N_o \left[\frac{\partial \Pi_{ok}}{\partial \bar{x}_j} \cdot S_{ok} + \Pi_{ok} \cdot \frac{\partial S_{ok}}{\partial \bar{x}_j}\right].$$

The first term uses the logit derivatives from Step 3. The second term requires $\frac{\partial S_{ok}}{\partial \bar{x}_j}$, which operates through M_o:

$$\frac{\partial S_{ok}}{\partial \bar{x}_j} = \frac{\partial S_{ok}}{\partial M_o} \cdot \frac{\partial M_o}{\partial \bar{x}_j} = S'_{ok}(M_o) \cdot (-\Pi_{oj}) = -\Pi_{oj}\, S'_{ok}(M_o),$$

where $S'_{ok}(M_o) = \int_{-\infty}^{\infty} \frac{1}{\eta_{ok}} \phi\Big(\frac{z}{\eta_{ok}}\Big)\, f_Z(z; M_o)\, dz > 0$.

Case 1: Own derivative ($j = k$). Substituting the logit own-derivative from Step 3 and the S_{ok} derivative:

$$\begin{aligned}
\frac{\partial E_k}{\partial \bar{x}_k} &= \sum_{o=1}^{C} N_o \left[-\frac{1}{\sigma} \Pi_{ok}(1 - \Pi_{ok}) \cdot S_{ok} + \Pi_{ok} \cdot \left(-\Pi_{ok}\, S'_{ok} \right) \right] \\
&= \sum_{o=1}^{C} N_o \left[-\frac{1}{\sigma} \Pi_{ok} S_{ok} + \frac{1}{\sigma} \Pi_{ok}^2 S_{ok} - \Pi_{ok}^2 S'_{ok} \right] \\
&= -\sum_{o=1}^{C} N_o \Pi_{ok} \left[\frac{1}{\sigma}(1 - \Pi_{ok}) S_{ok} + \Pi_{ok} S'_{ok} \right].
\end{aligned}$$

This is strictly negative: raising the payline in k both (i) drives applicants away (Π_{ok} falls) and (ii) lowers M_o, reducing the success rate.

Case 2: Cross derivative ($j \neq k$). Substituting the logit cross-derivative from Step 3:

$$\begin{aligned}
\frac{\partial E_k}{\partial \bar{x}_j} &= \sum_{o=1}^{C} N_o \left[\frac{1}{\sigma} \Pi_{ok} \Pi_{oj} \cdot S_{ok} + \Pi_{ok} \cdot \left(-\Pi_{oj}\, S'_{ok} \right) \right] \\
&= \sum_{o=1}^{C} N_o \Pi_{ok} \Pi_{oj} \left[\frac{1}{\sigma} S_{ok} - S'_{ok} \right].
\end{aligned}$$

The results from Step 6 are collected in the following lemma.

Lemma 1 (Jacobian of the Equilibrium System)**.** *Let J denote the $C \times C$ Jacobian matrix of the equilibrium system (10) with respect to $\bar{x}$, with representative element $J_{dj} = \frac{\partial E_d}{\partial \bar{x}_j}$. Then:*

$$J_{dd} = -\sum_{o=1}^{C} N_o \Pi_{od} \left[\frac{1}{\sigma}(1 - \Pi_{od}) S_{od}(M_o) + \Pi_{od} S'_{od}(M_o) \right] < 0, \tag{51}$$

$$J_{dj} = \sum_{o=1}^{C} N_o \Pi_{od} \Pi_{oj} \left[\frac{1}{\sigma} S_{od}(M_o) - S'_{od}(M_o) \right] \quad \text{for } j \neq d. \tag{52}$$

Assumption 1 (Dominant Substitution Effect)**.** *For all o, d and every payline vector in the relevant equilibrium domain, $\frac{1}{\sigma} S_{od}(M_o) > S'_{od}(M_o)$.*

This assumption states that, when a payline $\bar{x}_j$ falls, the direct substitution effect pulling applicants away from destination d dominates the indirect effect operating through the inclusive value M_o. Under Assumption 1, $J_{dj} > 0$ for $j \neq d$ and $\sum_j J_{dj} < 0$, so $-J$ is a strictly diagonally dominant M-matrix.

Remark. For vanishing evaluation noise ($\eta_{od} \to 0$), we have $S_{od} \to \Pr(Z_0 > -M_o)$ and $S'_{od} \to f_{Z_0}(-M_o)$. For the Gumbel distribution in (45), this inequality holds strictly for all relevant values of M_o. Hence the regularity condition is naturally satisfied in low-noise environments.

Step 7: Solve the linear system. The system from Step 6 is:

$$\sum_{j=1}^{C} \frac{\partial E_k}{\partial \bar{x}_j} \frac{d\bar{x}_j}{dG_d} = \mathbf{1}_{k=d} \cdot \frac{1}{\bar{g}_d}, \quad k = 1, \ldots, C.$$

To see the logic cleanly, define the equilibrium mapping

$$F_k(\bar{x}, G) \equiv E_k(\bar{x}) - \frac{G_k}{\bar{g}_k}, \qquad k = 1, \ldots, C.$$

An equilibrium is any vector $\bar{x}^*(G)$ such that $F(\bar{x}^*(G), G) = 0$. The paylines are therefore not free arguments once G is given; they are endogenous functions pinned down implicitly by the full system. Differentiating the identity $F(\bar{x}^*(G), G) = 0$ with respect to G_d gives

$$\sum_{j=1}^{C} \frac{\partial F_k}{\partial \bar{x}_j} \frac{d\bar{x}_j}{dG_d} + \frac{\partial F_k}{\partial G_d} = 0.$$

Since $\partial F_k / \partial \bar{x}_j = \partial E_k / \partial \bar{x}_j$ and $\partial F_k / \partial G_d = -\mathbf{1}_{k=d} / \bar{g}_d$, this reduces to the system above. This is why the object of interest is a total derivative $d\bar{x}_j / dG_d$: when G_d changes, every equilibrium payline may move in order to keep all markets clearing simultaneously. A partial derivative $\partial \bar{x}_j / \partial G_d$ holding the other equilibrium conditions fixed would not correspond to any feasible counterfactual equilibrium.

This is a $C \times C$ linear system in the unknowns $\left\{ \frac{d\bar{x}_j}{dG_d} \right\}_{j=1}^{C}$. Solving it yields the local GE payline adjustments around a given equilibrium. Analytically, the procedure has two steps: first solve the original nonlinear market-clearing system numerically for the baseline equilibrium $\bar{x}^*$; then evaluate the Jacobian $J(\bar{x}^*)$ and solve the linear system for the derivative vector. The empirical comparative-statics implementation does not invert J at $\bar{x}^*$; it instead re-solves the nonlinear market-clearing system at the perturbed budget and reports finite-difference responses, as documented in Section I.

Proposition 3 (General Equilibrium Payline Adjustments)**.** *Suppose there exists an equilibrium $\bar{x}^*$ solving (10), and suppose the Jacobian evaluated at $\bar{x}^*$ is invertible. Under Assumption 1, the inverse Jacobian satisfies $J^{-1} \leq 0$ elementwise. The Implicit Function Theorem then implies the local comparative statics*

$$d\bar{x} = J^{-1} \mathrm{diag}(1/\bar{g}_d)\, dG. \tag{53}$$

Hence an increase in the budget of destination h weakly lowers the payline of every destination:

$$\frac{\partial \bar{x}_j}{\partial G_h} = (J^{-1})_{jh} \frac{1}{\bar{g}_h} \leq 0 \quad \textit{for all } j. \tag{54}$$

Proof. The uniqueness proof shows that, under Assumption 1, $-J(\bar{x}^*)$ is a non-singular M-matrix. A non-singular M-matrix has a non-negative inverse, so $(-J(\bar{x}^*))^{-1} \geq 0$ elementwise. Since $J(\bar{x}^*)^{-1} = -(-J(\bar{x}^*))^{-1}$, it follows that $J(\bar{x}^*)^{-1} \leq 0$ elementwise. The Implicit Function Theorem applied to $F(\bar{x}, G) = E(\bar{x}) - \text{diag}(1/\bar{g}_d)G$ gives

$$d\bar{x} = J(\bar{x}^*)^{-1} \text{diag}(1/\bar{g}_d)\, dG.$$

Because $\bar{g}_d > 0$ for every destination, the sign of each payline response is inherited from $J(\bar{x}^*)^{-1}$. □

The expression above is local. It describes the slope of the equilibrium correspondence at the baseline allocation, not the exact response to a large discrete funding reform. For small counterfactual reallocations this linearisation is the correct object. For larger changes, one should re-solve the full nonlinear equilibrium system at the new budget vector rather than extrapolating from J^{-1}.

Clarifying PE vs. GE. In partial equilibrium, holding $\bar{x}_j$ fixed for $j \neq h$, an increase in G_h lowers $\bar{x}_h$. This creates *cross-field crowding out*: applications to other fields fall because Π_{oj} declines. It also raises M_o, shifting the conditional quality distribution upwards.

In general equilibrium, the paylines $\bar{x}_j$ must also adjust to clear the remaining markets. Since applicant volume in field j falls while G_j is fixed, field j must lower $\bar{x}_j$ to fund a larger share of the remaining applicants. The response of applicant volumes depends on the full matrix J^{-1}, but the payline effect is weakly negative under the stated conditions.

Step 8: Assemble the elasticity. Substituting the results from Steps 2–5 into Step 2 gives the partial derivatives $\frac{\partial Y_h}{\partial \bar{x}_j}$. Substituting the solution from Step 7 gives $\frac{d\bar{x}_j}{dG_d}$. The elasticity is then:

$$\text{CS1}_{hd} = \frac{G_d}{Y_h} \sum_{j=1}^{C} \frac{\partial Y_h}{\partial \bar{x}_j} \cdot \frac{d\bar{x}_j}{dG_d}. \tag{55}$$

E Field-Specific Knowledge Depreciation Rates

This section documents the construction of the field-specific depreciation rates δ_d entering the knowledge stock in the production function.

E.1 Reference Age Distribution

For every publication i with observed publication year t_i, I iterate over all its references j in the OpenAlex citation network. A reference j is retained if (i) its publication year t_j can be recovered from OpenAlex metadata, (ii) it can be matched to a NIH grant via a unique PubMed identifier (PMID), and (iii) the matched grant is attributed to a single NIH institute d.[13] For each retained reference pair (i, j), the *reference age* is

$$a_{ij} = t_i - t_j, \quad a_{ij} \geq 1.$$

Restricting citing publications to those from 2002 onwards and aggregating over all such pairs, I obtain for each institute d and age a a citation count C_{da}—the total number of times a paper attributed to institute d was cited at age a.

E.2 Half-Life Estimation

The empirical cumulative citation distribution for institute d is

$$\hat{F}_d(a) = \frac{\sum_{a'=1}^{a} C_{da'}}{\sum_{a'=1}^{A} C_{da'}},$$

where A is the maximum observed age. The *citation half-life* h_d is the median of this distribution:

$$h_d = \min\{a : \hat{F}_d(a) \geq 0.5\}.$$

A short half-life indicates that citations to papers in field d are concentrated in the years immediately following publication—i.e. the field's stock of relevant knowledge turns over rapidly.

E.3 Implied Depreciation Rate

Under a geometric decay model, the fraction of a knowledge unit still in productive use after a periods is $(1 - \delta_d)^a$. Imposing that half of each knowledge unit has depreciated by the empirical half-life gives $(1 - \delta_d)^{h_d} = 0.5$, which yields

$$\delta_d = 1 - 0.5^{1/h_d}. \tag{56}$$

This rate then governs knowledge accumulation. Starting from $K_{d,0} = 0$, the depreciated stock evolves as

$$K_{d,t} = (1 - \delta_d)\, K_{d,t-1} + P_{d,t},$$

where $P_{d,t}$ is the impact-factor-weighted publication count of institute d in year t.

[13] References matched to multiple institutes are discarded to ensure unambiguous field attribution.

F Outcome Construction

This section defines the three output variables used in the production-function estimates. Let $\mathcal{P}_g$ denote the set of OpenAlex publications linked to grant g, and let n_g^{PI} denote the number of investigators on that grant. A linked publication is split equally across the grant's investigators. If a publication is linked to several grants, each grant-publication link is retained. For investigator i on grant g, the metric-m output is

$$Y_{ig}^{m} = \sum_{p \in \mathcal{P}_g} \frac{w_{pg}^{m}}{n_g^{\mathrm{PI}}}, \tag{57}$$

where w_{pg}^{m} is the value weight attached to publication p under metric m.

F.1 Impact Factor

The impact-factor outcome measures publication quantity weighted by journal impact factor. For each publication p, I match the journal metric file to OpenAlex publication records using the journal ISSN-L. Let IF_p denote the matched journal impact factor. The publication weight is

$$w_{pg}^{\mathrm{IF}} = IF_p, \tag{58}$$

and the grant-PI outcome is

$$Y_{ig}^{\mathrm{IF}} = \sum_{p \in \mathcal{P}_g} \frac{IF_p}{n_g^{\mathrm{PI}}}. \tag{59}$$

This measure uses journal-level impact rather than realised citations, avoiding citation-truncation differences across publication years. Publications without a matched journal impact factor contribute zero to the weighted sum.

F.2 Commercial Potential

The commercial-potential outcome measures publication quantity weighted by predicted commercial relevance. For each publication p, I merge the DOI-level score from Masclans, Hasan, and Cohen (2023). Let $c_p \in [0, 1]$ denote the commercial-potential score. The publication weight is

$$w_{pg}^{\mathrm{CP}} = c_p, \tag{60}$$

and the grant-PI outcome is

$$Y_{ig}^{\mathrm{CP}} = \sum_{p \in \mathcal{P}_g} \frac{c_p}{n_g^{\mathrm{PI}}}. \tag{61}$$

The score is interpreted as the publication's patentability or downstream commercial potential. Publications without a matched DOI-level score receive weight zero. In the analysis sample, the score is observed for 89 per cent of grant-publication observations.

F.3 Disease Burden

The disease-burden outcome measures publication quantity weighted by the disease burden of the funding grant. The weight is attached at the grant level because the NIH RCDC categories describe the grant's disease area, whereas publication titles and abstracts may reflect intermediate or basic-science content. Let T_g denote the RCDC tags of grant g, let C denote the set of RCDC tags matched to the IHME burden workbook, and let B_ℓ denote US disability-adjusted life years for matched tag ℓ. For grants with $M_g = T_g \cap C \neq \emptyset$, the raw matched burden is

$$\tilde{b}_g = \frac{1}{|M_g|} \sum_{\ell \in M_g} B_\ell. \tag{62}$$

If M_g is empty, I impute the raw burden with the mean matched burden among grants funded by the same institute. If that institute has no matched grant, I use the overall mean matched burden. Let b_g denote this observed-or-imputed raw grant burden. The normalised grant weight is

$$w_g^B = \frac{b_g}{|G|^{-1} \sum_{g' \in G} b_{g'}}, \tag{63}$$

so the average grant has weight one. Each publication funded by grant g receives this grant weight. The disease-burden-weighted grant-PI outcome is therefore

$$Y_{ig}^{\mathrm{DB}} = \sum_{p \in \mathscr{P}_g} \frac{w_g^B}{n_g^{\mathrm{PI}}} = \frac{|\mathscr{P}_g| w_g^B}{n_g^{\mathrm{PI}}}. \tag{64}$$

This outcome measures publications weighted by the disease burden of the funding grant, not by disease content inferred from the publication. The RCDC–DALY match diagnostics are as follows.

Coverage is reported in Tables 15 and 16. In the analysis sample, 45.4 per cent of grants have a direct DALY match, 33.8 per cent have no RCDC tag, and 20.5 per cent have RCDC tags absent from the workbook. The grant-level disease-burden weight is weakly correlated with the impact-factor and commercial-potential outcomes: Pearson correlations are 0.048 and 0.099 respectively (Table 17); the corresponding scatter plots are reported in Figure 8.

Unit	Total	No RCDC	RCDC no DALY	Direct DALY	IC avg	Overall avg
Grants	65,349	21,859 (33.4)	13,837 (21.2)	29,653 (45.4)	35,947 (55.0)	0 (0.0)
Publication links	1,179,568	398,211 (33.8)	241,882 (20.5)	539,475 (45.7)	643,645 (54.6)	0 (0.0)

Table 15: RCDC–DALY coverage summary.

2002–2018 analysis sample. "No RCDC": observations with no RCDC tag. "RCDC no DALY": observations with RCDC tags absent from the US-DALY workbook. "Direct DALY": observations with at least one matched RCDC–DALY tag. "IC avg" and "Overall avg" record the imputation source when no direct match exists. Publication links are unique grant-publication pairs.

Institute	Grants	G direct	G no RCDC	G no DALY	Pub. links	P direct	P no RCDC	P no DALY	Mean weight
Aging	3,015	1,339	937	739	65,480	28,919	21,418	15,143	0.567
Alcoholism	1,217	823	394	0	19,105	13,199	5,906	0	0.858
Allergy and Infectious Diseases	6,477	2,445	2,116	1,916	105,208	39,329	34,344	31,535	0.552
Arthritis and Musculoskeletal and Skin Diseases	1,847	520	633	694	35,800	10,173	12,659	12,968	0.778
Biomedical Imaging and Bioengineering	1,223	371	481	371	30,151	8,862	12,477	8,812	1.355
Cancer	9,366	6,116	3,249	1	162,875	107,350	55,513	12	1.735
Child Health and Human Development	3,034	1,262	927	845	45,936	19,057	13,849	13,030	0.433
Deafness and Communication Disorders	1,305	224	391	690	21,408	3,749	6,728	10,931	0.405
Dental and Craniofacial Research	1,050	684	366	0	19,095	12,671	6,424	0	0.380
Diabetes and Digestive and Kidney Diseases	5,160	3,403	1,488	269	91,922	61,695	25,854	4,373	0.452
Drug Abuse	2,688	1,683	1,004	1	44,024	26,174	17,807	43	0.927
Environmental Health	1,089	513	336	240	21,279	10,474	6,047	4,758	1.100
Eye	1,961	364	553	1,044	35,899	6,923	11,078	17,898	0.475
General Medicine	7,198	1,259	2,763	3,176	125,779	21,011	50,530	54,238	1.431
Heart, Lung, and Blood	8,350	4,491	2,714	1,145	174,637	96,962	55,385	22,290	1.720
Human Genome	457	41	136	280	9,087	850	2,055	6,182	1.227
Mental Health	4,115	1,833	1,397	885	74,977	32,511	26,902	15,564	0.416
Neurological Disorders and Stroke	5,215	2,022	1,771	1,422	89,921	36,396	31,183	22,342	0.688
Nursing Research	582	260	203	119	6,985	3,170	2,052	1,763	0.909

Table 16: RCDC–DALY coverage by institute.

"G": grants. "P": grant-publication links. "Direct": observations with at least one matched RCDC–DALY tag. "No RCDC": observations without an RCDC tag. "No DALY": observations with RCDC tags not in the workbook. "Mean weight": mean normalised grant-level US-DALY burden weight within the institute.

Pair	N	Pearson	Spearman	Log Pearson
Burden–IF	65,349	0.048	0.053	0.040
Burden–commercial	65,349	0.099	0.144	0.110
IF–commercial	65,349	0.846	0.819	0.809

Table 17: Grant-level correlations between disease burden and publication outcomes.

Unique grants, 2002–2018. Pearson and Spearman correlations in levels; log Pearson correlations use log disease-burden weight and $\log(1 + y)$ for impact-factor and commercial-potential publications.

G Production Function Estimation Details

This section documents aspects of the production function estimation not covered by the main text.

G.1 Estimator

The production function in equation (37) is estimated by constrained negative binomial maximum likelihood. The negative binomial model allows for over-dispersion in publication counts relative to the Poisson benchmark, which is empirically important given the heavy right tail of the impact-factor-weighted output distribution. Non-negativity of all destination-specific elasticities, $\alpha_d \geq 0$ and $\beta_d \geq 0$, is imposed as a bound constraint throughout the optimisation, ruling out economically inadmissible negative elasticities. The constrained log-likelihood is maximised by L-BFGS-B with bound constraints; gradients are numerically approximated by SciPy for the NB objective.

G.2 Standard Errors

Standard errors reported in Table 19 use a quasi-Newton GLM negative-binomial working-weight covariance proxy evaluated at the constrained maximum: $\widehat{\text{Var}}(\hat{\theta}) = (X'WX)^{+}$, with weights $w_i = 1/(1/\hat{\mu}_i + \hat{\alpha})$, where $\hat{\mu}_i$ is the fitted mean and $\hat{\alpha}$ the estimated dispersion. The pseudo-inverse $(\cdot)^{+}$ accommodates the rank deficiency induced by binding non-negativity constraints. For parameters at an interior solution, this proxy approximates the conventional observed-information standard error. For a parameter at a boundary (e.g. $\hat{\beta}_{\text{NIDCR}} = 0$), the reported value reflects local curvature of the working likelihood at the boundary and should be interpreted as a lower bound on parameter uncertainty at the constraint.

G.3 Relationship to Comparative Statics

The estimated knowledge elasticities $\hat{\alpha}_d$, the common selection coefficient $\hat{\lambda}$, and the destination-level baseline productivity $\hat{f}_d$—the mean fitted production per PI at the 2018 baseline year, with the $W_{od,t}^{\hat{\lambda}}$ quality term factored out—are the production-side inputs to the comparative statics in Section 4. The funding elasticities $\hat{\beta}_d$ are reported in Table 19 for completeness, but they do not enter the comparative-statics calculations directly. Table 19 is drawn from the same estimation script/output dataframe as the CSV consumed by comparative statics, ensuring full reproducibility.

G.4 Full Destination-Specific Estimates

The main text reports summary statistics across destinations. The full destination-specific production-function estimates for four outcomes are reported below. The raw-

publication specification used to construct MR_d^P in Section 7.1 is reported first.

Institute	$\hat{\alpha}_d$	$\hat{\beta}_d$
Aging	0.044	0.246
	[0.023]	[0.025]
Alcoholism	0.409	0.041
	[0.053]	[0.035]
Allergy and Infectious Diseases	0.098	0.068
	[0.018]	[0.015]
Arthritis and Musculoskeletal and Skin Diseases	0.168	0.119
	[0.033]	[0.030]
Biomedical Imaging and Bioengineering	0.000†	0.208
	[0.043]	[0.048]
Cancer	0.200	0.110
	[0.015]	[0.015]
Child Health and Human Development	0.249	0.082
	[0.022]	[0.018]
Deafness and Communication Disorders	0.230	0.050
	[0.057]	[0.036]
Dental and Craniofacial Research	0.332	0.000†
	[0.044]	[0.035]
Diabetes and Digestive and Kidney Diseases	0.130	0.051
	[0.016]	[0.013]
Drug Abuse	0.301	0.112
	[0.032]	[0.029]
Environmental Health	0.190	0.216
	[0.039]	[0.042]
Eye	0.061	0.230
	[0.045]	[0.039]
General Medicine	0.067	0.031
	[0.019]	[0.016]
Heart, Lung, and Blood	0.030	0.124
	[0.013]	[0.015]
Human Genome	0.307	0.143
	[0.059]	[0.078]
Mental Health	0.152	0.158
	[0.024]	[0.024]
Neurological Disorders and Stroke	0.067	0.126
	[0.023]	[0.022]
Nursing Research	0.100	0.192
	[0.049]	[0.051]
Selection control		
$\hat{\lambda}$ ($\ln W_{od,t}$)		0.200 [0.037]
Sample		
Observations		55,710

Standard errors in brackets. † Non-negativity constraint binding. Destination and year fixed effects included but not reported.

Table 18: Production function estimates: raw-publication outcome.

Outcome Y_d^P: unweighted publications per PI. Specification identical to Table 19. Standard errors in brackets; † non-negativity constraint binding.

Institute	$\hat{\alpha}_d$	$\hat{\beta}_d$
Aging	0.244	0.219
	[0.022]	[0.024]
Alcoholism	0.553	0.042
	[0.051]	[0.034]
Allergy and Infectious Diseases	0.219	0.080
	[0.017]	[0.015]
Arthritis and Musculoskeletal and Skin Diseases	0.392	0.120
	[0.032]	[0.029]
Biomedical Imaging and Bioengineering	0.134	0.180
	[0.042]	[0.046]
Cancer	0.356	0.107
	[0.014]	[0.015]
Child Health and Human Development	0.390	0.082
	[0.021]	[0.017]
Deafness and Communication Disorders	0.523	0.061
	[0.055]	[0.035]
Dental and Craniofacial Research	0.403	0.000†
	[0.042]	[0.034]
Diabetes and Digestive and Kidney Diseases	0.257	0.052
	[0.016]	[0.013]
Drug Abuse	0.449	0.082
	[0.031]	[0.028]
Environmental Health	0.198	0.220
	[0.037]	[0.041]
Eye	0.171	0.215
	[0.043]	[0.037]
General Medicine	0.191	0.040
	[0.018]	[0.016]
Heart, Lung, and Blood	0.097	0.125
	[0.013]	[0.014]
Human Genome	0.543	0.156
	[0.056]	[0.076]
Mental Health	0.278	0.146
	[0.024]	[0.023]
Neurological Disorders and Stroke	0.238	0.101
	[0.022]	[0.020]
Nursing Research	0.199	0.190
	[0.048]	[0.048]
Selection control		
$\hat{\lambda}$ ($\ln W_{od,t}$)		-0.035 [0.036]
Sample		
Observations		55,710

Standard errors in brackets. † Non-negativity constraint binding. Destination and year fixed effects included but not reported.

Table 19: Production function estimates: impact-factor outcome.

Outcome Y_d^{IF}: impact-factor-weighted publications per PI. Estimator: constrained negative binomial with non-negativity constraints on $\hat{\alpha}_d$ and $\hat{\beta}_d$. Knowledge input k_{it}: PI-level Cobb–Douglas aggregate of field-specific knowledge stocks, weighted by grant-level citation shares with field-specific depreciation (Appendix E). Funding input g_{it}: real grant funding per PI. Destination (ϕ_d) and year (τ_t) fixed effects. $\ln W_{od,t}$: structural selection control from the first-stage GMM (equation (38)); coefficient $\hat{\lambda}$ common across destinations. $n = 55{,}710$ grant-PI-year observations, 2005–2018. Standard errors in brackets; † non-negativity constraint binding.

Institute	$\hat{\alpha}_d$	$\hat{\beta}_d$
Aging	0.388	0.206
	[0.025]	[0.027]
Alcoholism	0.670	0.000†
	[0.060]	[0.038]
Allergy and Infectious Diseases	0.292	0.047
	[0.019]	[0.016]
Arthritis and Musculoskeletal and Skin Diseases	0.402	0.082
	[0.035]	[0.030]
Biomedical Imaging and Bioengineering	0.121	0.244
	[0.044]	[0.050]
Cancer	0.435	0.075
	[0.015]	[0.016]
Child Health and Human Development	0.430	0.017
	[0.026]	[0.020]
Deafness and Communication Disorders	0.275	0.056
	[0.060]	[0.037]
Dental and Craniofacial Research	0.490	0.000†
	[0.046]	[0.036]
Diabetes and Digestive and Kidney Diseases	0.370	0.028
	[0.018]	[0.013]
Drug Abuse	0.648	0.001
	[0.037]	[0.032]
Environmental Health	0.623	0.100
	[0.043]	[0.044]
Eye	0.112	0.199
	[0.046]	[0.040]
General Medicine	0.197	0.022
	[0.019]	[0.017]
Heart, Lung, and Blood	0.263	0.104
	[0.014]	[0.015]
Human Genome	0.566	0.106
	[0.064]	[0.082]
Mental Health	0.391	0.084
	[0.027]	[0.026]
Neurological Disorders and Stroke	0.241	0.133
	[0.024]	[0.023]
Nursing Research	0.387	0.125
	[0.061]	[0.062]
Selection control		
$\hat{\lambda}$ ($\ln W_{od,t}$)		-0.235 [0.039]
Sample		
Observations		55,710

Standard errors in brackets. † Non-negativity constraint binding. Destination and year fixed effects included but not reported.

Table 20: Production function estimates: commercial-potential outcome.

Outcome Y_d^{CP}: commercial-potential-weighted publications per PI. Specification identical to Table 19. Standard errors in brackets; † non-negativity constraint binding.

Institute	$\hat{\alpha}_d$	$\hat{\beta}_d$
Aging	0.055	0.261
	[0.023]	[0.026]
Alcoholism	0.377	0.034
	[0.053]	[0.035]
Allergy and Infectious Diseases	0.233	0.049
	[0.019]	[0.016]
Arthritis and Musculoskeletal and Skin Diseases	0.270	0.076
	[0.034]	[0.030]
Biomedical Imaging and Bioengineering	0.067	0.235
	[0.042]	[0.048]
Cancer	0.261	0.103
	[0.014]	[0.015]
Child Health and Human Development	0.284	0.082
	[0.024]	[0.019]
Deafness and Communication Disorders	0.195	0.097
	[0.060]	[0.038]
Dental and Craniofacial Research	0.426	0.000†
	[0.048]	[0.037]
Diabetes and Digestive and Kidney Diseases	0.143	0.038
	[0.017]	[0.014]
Drug Abuse	0.365	0.094
	[0.032]	[0.029]
Environmental Health	0.355	0.153
	[0.039]	[0.042]
Eye	0.186	0.225
	[0.047]	[0.041]
General Medicine	0.110	0.016
	[0.019]	[0.016]
Heart, Lung, and Blood	0.064	0.124
	[0.013]	[0.014]
Human Genome	0.274	0.111
	[0.058]	[0.078]
Mental Health	0.164	0.084
	[0.026]	[0.025]
Neurological Disorders and Stroke	0.117	0.096
	[0.024]	[0.022]
Nursing Research	0.093	0.103
	[0.050]	[0.050]
Selection control		
$\hat{\lambda}$ ($\ln W_{od,t}$)		0.086 [0.038]
Sample		
Observations		55,710

Standard errors in brackets. † Non-negativity constraint binding. Destination and year fixed effects included but not reported.

Table 21: Production function estimates: disease-burden outcome.

Outcome Y_d^b: US-DALY-weighted publications per PI. Specification identical to Table 19. Standard errors in brackets; † non-negativity constraint binding.

H Implementation Details

This section documents the numerical implementation of the estimation procedure described in Section 5.

Parameter vector. The structural parameter vector is $\theta = (\text{vec}(\mu_{od}),\ \log\sigma,\ c_0,\ c_1,\ \log\eta_0,\ \gamma)$. The merit matrix μ_{od} is $C \times C$ with all entries structurally free (C^2 parameters). The scale $\sigma > 0$ and baseline noise $\eta_0 > 0$ are reparametrised as $\log\sigma$ and $\log\eta_0$ so the optimiser operates on the real line. The distance slope γ is constrained to $[0, \infty)$ to ensure evaluation noise is non-decreasing in scientific distance. Both outside-option parameters c_0 and c_1 are freely estimated scalars. The total free parameter count is therefore $C^2 + 5 = 366$ for $C = 19$.

Equilibrium solver. The nonlinear system of market-clearing conditions is solved using a hybrid Powell root-finding algorithm (`scipy.optimize.root` with `method='hybr'`), trying multiple starting points to avoid spurious local solutions. Each year's solve is warm-started from the previous year's payline vector, exploiting temporal continuity.

Quadrature. The integrals S_{od} and $Q_{od}(M_{o,t}; \eta_{od}) = \int z\,\Phi(z/\eta_{od})\,f_Z(z; M_{o,t}, \sigma)\,dz$ are computed by deterministic quadrature over the Gumbel$(0, \sigma)$ density, evaluating 100 evenly-spaced quantile points with normalised PDF weights. All (o, d) pairs are evaluated simultaneously via array broadcasting.

Initialisation. The merit matrix is initialised from the logarithm of empirical bilateral funded flow shares, centred to have mean zero and rescaled so that the maximum absolute entry is 2. The remaining parameters are initialised at: $\sigma = 1$, $c_0 = -3$, $c_1 = 0$ (no trend), $\eta_0 = 1$, $\gamma = 0$ (no distance effect).

Optimiser. Both GMM stages use the limited-memory BFGS algorithm with box constraints (`L-BFGS-B`). The first-stage ridge penalty is $\lambda_1 = 10^{-4}$. Non-convergence of the equilibrium solver for a given year incurs a smooth penalty proportional to the mean squared market-clearing residual, steering the optimiser away from parameter regions where the system has no solution.

I Numerical Comparative-Statics Implementation

Section 4 presents the analytical decomposition. In the quantitative implementation, I compute local comparative statics by finite differences around the final estimation year $t^* = T - 1$ (2018 in the current sample window). Let baseline slots be $B_{d,t^*} \equiv G_{d,t^*}/\bar{g}_{d,t^*}$.

For each shocked destination d, I apply a proportional slots perturbation $H = 0.01$:

$$B'_{d,t^*} = (1 + H)B_{d,t^*}, \tag{65}$$

re-solve equilibrium paylines, and evaluate model-implied output at baseline and shocked equilibria.

Baseline destination-level output is evaluated from structural primitives as

$$\hat{Y}^{\text{base}}_{h,t^*} = \hat{f}_h \sum_{o=1}^{C} N_{o,t^*} \Pi_{oh,t^*} S_{oh,t^*} W^{\hat{\lambda}}_{oh,t^*}, \tag{66}$$

where W_{oh,t^*} is defined in (38). The extensive-margin elasticity matrix is then computed as

$$\varepsilon^{\text{ext}}_{hd} \approx \frac{\hat{Y}_{h,t^*}(B'_{d,t^*}) - \hat{Y}_{h,t^*}(B_{d,t^*})}{H\, \hat{Y}_{h,t^*}(B_{d,t^*})}. \tag{67}$$

Knowledge spillovers are implemented as a static local feedback term. Using production-function elasticities $\hat{\alpha}_h$ and citation-share exposures Γ_{hd}, the pairwise knowledge-elasticity object is

$$E^{\text{know}}_{hd} = \hat{\alpha}_h \Gamma_{hd}. \tag{68}$$

For a shock to destination d, the code maps this into

$$\varepsilon^{\text{know}}_{hd} = E^{\text{know}}_{hd}\, \varepsilon^{\text{ext}}_{dd}, \tag{69}$$

and sets intensive-margin elasticities to zero in the current implementation. Hence total elasticities are

$$\varepsilon^{\text{tot}}_{hd} = \varepsilon^{\text{ext}}_{hd} + \varepsilon^{\text{know}}_{hd}. \tag{70}$$

Elasticities $\varepsilon^{\text{tot}}_{hd}$ are computed from model-implied baseline output. The reported partial effect of a \$1M increase in G_d then scales these elasticities by observed baseline output and funding:

$$\Delta Y^{(\$1\text{M})}_{h \leftarrow d} = \varepsilon^{\text{tot}}_{hd} \left(\frac{Y_{h,t^*}}{G_{d,t^*}} \right) 10^6, \tag{71}$$

where Y_{h,t^*} denotes the observed (empirical) output of institute h at the baseline year, while the elasticity itself is evaluated at the model-predicted output. The code also computes institute-specific impact-factor conversion factors IF_{d,t^*} at the same baseline year for diagnostics and output reporting.

J Policy Counterfactual Construction

J.1 US DALY Outcome and Institute Weight

The US-DALY-weighted production outcome uses the normalised grant weights defined in Appendix F. The budget-allocation rule uses a different object: an institute-level disease-burden profile used only for policy counterfactuals. Let M_g denote the subset of grant g's 2017 RCDC tags that match the burden workbook. For each matched tag ℓ, grant funding is split equally across tags. If $s_{h\ell}$ is institute h's funding share in RCDC category ℓ, the destination burden proxy is

$$\widetilde{W}^{B,\text{dest}}_{h,\text{US DALYs}} = \sum_{\ell} s_{h\ell} B_{\ell}. \tag{72}$$

Using the pairwise impact-factor marginal-return matrix MR_{dh}, this destination burden is mapped back to source institutes through

$$\Psi_{d,\text{US DALYs}} = \sum_{h} \widetilde{W}^{B,\text{dest}}_{h,\text{US DALYs}} MR_{dh}, \qquad W^{B,\text{spill}}_{d,\text{US DALYs}} \propto \Psi_{d,\text{US DALYs}}/MR_d, \qquad MR_d = \sum_{h} MR_{dh}. \tag{73}$$

The resulting equivalent source weights are normalised to sum to 19 and then converted into allocation shares. In the 2017 unique-grant sample, 2,979 of 4,344 grants, or 68.6%, have at least one matched RCDC tag.

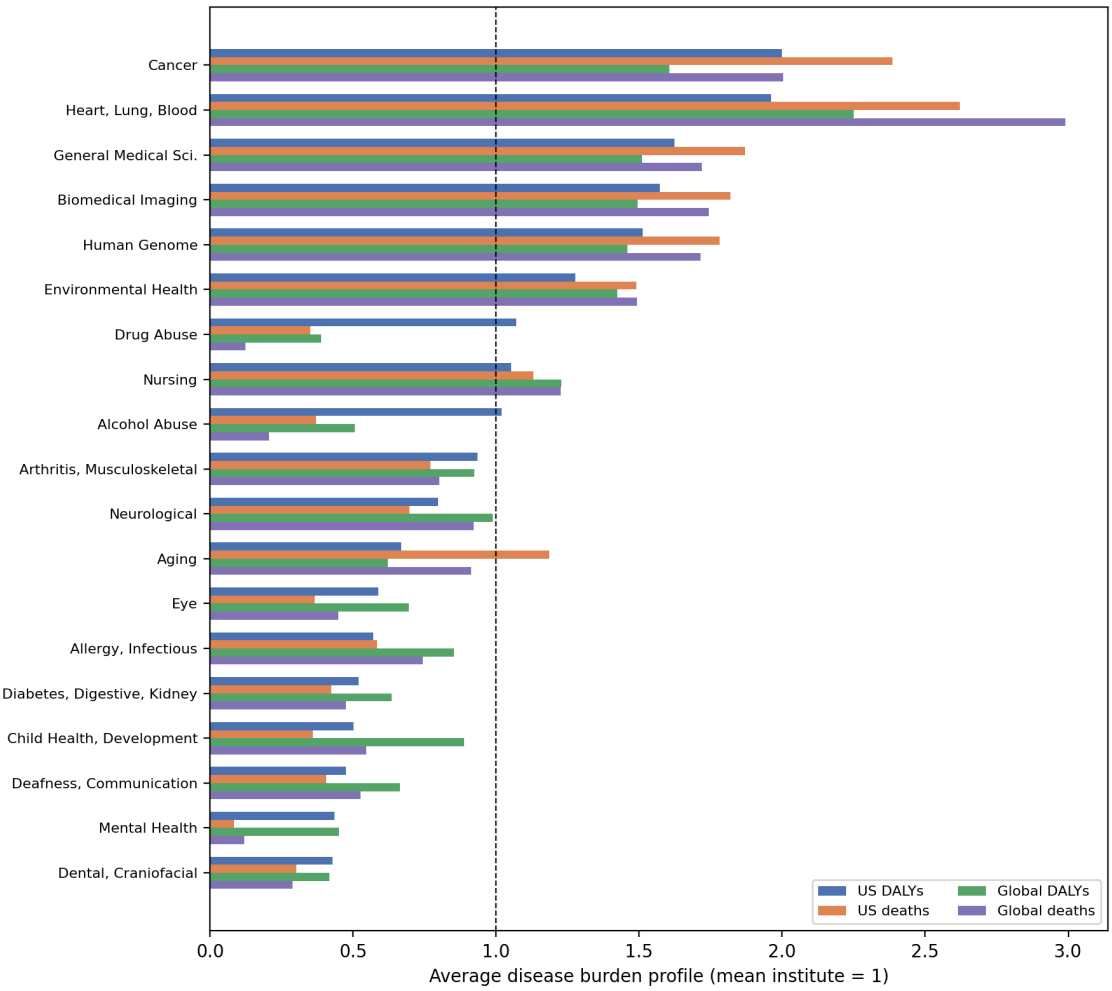


Figure 22: Average disease-burden profiles by NIH institute

Grouped horizontal bars report the institute-level average disease-burden profiles under four burden measures: US DALYs, US deaths, global DALYs, and global deaths. Each series is normalised to a cross-institute mean of one. The dashed vertical line marks the common normalisation point. Institutes are ordered by the US-DALY profile.

K Additional Results

K.1 Detailed Return Decompositions

Tables 22–24 report the full institute-level decomposition of the absolute marginal returns discussed in Section 6.3. To keep the layout readable, each table is split into an own-institute panel and a cross-institute panel.

Institute	No Ext. Return (PE)	No Ext. Return (GE)	Own Know. Spillovers	Net Own Return
Human Genome	35.915	35.914	3.955	39.869
General Medicine	45.525	45.521	3.227	48.748
Dental and Craniofacial Research	28.633	28.632	2.301	30.934
Diabetes and Digestive and Kidney Diseases	26.787	26.785	1.943	28.728
Environmental Health	25.248	25.247	1.411	26.658
Cancer	25.509	25.507	3.663	29.170
Neurological Disorders and Stroke	25.343	25.342	1.629	26.971
Mental Health	24.050	24.049	2.034	26.083
Biomedical Imaging and Bioengineering	23.362	23.362	0.480	23.841
Eye	24.469	24.467	1.585	26.053
Arthritis and Musculoskeletal and Skin Diseases	22.630	22.629	1.752	24.381
Heart, Lung, and Blood	23.973	23.971	0.794	24.766
Allergy and Infectious Diseases	21.804	21.801	1.840	23.641
Drug Abuse	16.364	16.364	2.091	18.454
Alcoholism	12.465	12.466	1.646	14.112
Child Health and Human Development	11.132	11.132	0.997	12.129
Deafness and Communication Disorders	12.013	12.012	2.095	14.107
Aging	12.332	12.331	0.869	13.200
Nursing Research	3.921	3.921	0.103	4.024

Institute	Talent Flow Ext. (PE)	Talent Flow Ext. (GE)	Know. Spill. to Others	Net Cross Return
Human Genome	-24.932	-0.034	24.160	24.126
General Medicine	-33.077	-0.060	9.907	9.847
Dental and Craniofacial Research	-19.513	-0.062	4.165	4.102
Diabetes and Digestive and Kidney Diseases	-17.187	-0.043	5.001	4.958
Environmental Health	-20.202	-0.060	6.327	6.267
Cancer	-14.368	-0.029	3.647	3.618
Neurological Disorders and Stroke	-16.920	-0.025	4.159	4.135
Mental Health	-15.770	-0.035	4.704	4.669
Biomedical Imaging and Bioengineering	-23.724	-0.030	6.338	6.307
Eye	-19.138	-0.051	3.671	3.620
Arthritis and Musculoskeletal and Skin Diseases	-19.462	-0.050	4.071	4.021
Heart, Lung, and Blood	-12.266	-0.031	3.167	3.136
Allergy and Infectious Diseases	-14.413	-0.020	2.695	2.675
Drug Abuse	-15.561	-0.038	3.530	3.493
Alcoholism	-14.919	-0.045	2.910	2.865
Child Health and Human Development	-15.846	-0.027	3.856	3.828
Deafness and Communication Disorders	-15.498	-0.031	1.652	1.621
Aging	-6.739	-0.013	1.573	1.560
Nursing Research	-11.463	-0.028	1.282	1.254

Table 22: Total returns per $1 million: impact-factor outcome.

Change in aggregate NIH-wide impact-factor-weighted output per $1 million budget increase. The top panel reports the direct and own-institute knowledge components: "No Ext. Return (PE)", "No Ext. Return (GE)", "Own Know. Spillovers", and "Net Own Return". The bottom panel reports the off-institute components: "Talent Flow Ext. (PE)", "Talent Flow Ext. (GE)", "Know. Spill. to Others", and "Net Cross Return". Institutes are sorted by total return.

Institute	No Ext. Return (PE)	No Ext. Return (GE)	Own Know. Spillovers	Net Own Return
General Medicine	3.653	3.651	0.254	3.906
Human Genome	2.038	2.038	0.240	2.278
Biomedical Imaging and Bioengineering	2.678	2.678	0.060	2.738
Dental and Craniofacial Research	2.667	2.667	0.272	2.938
Eye	2.748	2.747	0.126	2.873
Diabetes and Digestive and Kidney Diseases	2.012	2.011	0.216	2.226
Neurological Disorders and Stroke	2.170	2.169	0.147	2.316
Heart, Lung, and Blood	2.192	2.191	0.208	2.398
Cancer	1.900	1.899	0.338	2.237
Arthritis and Musculoskeletal and Skin Diseases	1.978	1.977	0.169	2.147
Environmental Health	1.443	1.442	0.203	1.645
Allergy and Infectious Diseases	1.690	1.688	0.192	1.880
Mental Health	1.173	1.173	0.123	1.295
Drug Abuse	0.949	0.949	0.168	1.117
Child Health and Human Development	0.829	0.829	0.080	0.909
Alcoholism	0.865	0.865	0.140	1.005
Deafness and Communication Disorders	0.975	0.975	0.106	1.081
Aging	0.822	0.821	0.088	0.909
Nursing Research	0.208	0.208	0.010	0.218

Institute	Talent Flow Ext. (PE)	Talent Flow Ext. (GE)	Know. Spill. to Others	Net Cross Return
General Medicine	-2.597	-0.032	1.040	1.008
Human Genome	-1.948	-0.018	2.240	2.222
Biomedical Imaging and Bioengineering	-1.872	-0.015	0.706	0.691
Dental and Craniofacial Research	-1.563	-0.033	0.437	0.404
Eye	-1.481	-0.027	0.325	0.298
Diabetes and Digestive and Kidney Diseases	-1.375	-0.023	0.530	0.507
Neurological Disorders and Stroke	-1.265	-0.013	0.419	0.406
Heart, Lung, and Blood	-0.947	-0.016	0.309	0.292
Cancer	-1.157	-0.016	0.396	0.381
Arthritis and Musculoskeletal and Skin Diseases	-1.555	-0.027	0.451	0.425
Environmental Health	-1.584	-0.033	0.624	0.591
Allergy and Infectious Diseases	-1.144	-0.010	0.270	0.260
Mental Health	-1.224	-0.019	0.367	0.348
Drug Abuse	-1.121	-0.021	0.274	0.253
Child Health and Human Development	-1.189	-0.015	0.341	0.326
Alcoholism	-1.054	-0.025	0.250	0.225
Deafness and Communication Disorders	-1.204	-0.016	0.139	0.123
Aging	-0.530	-0.007	0.149	0.142
Nursing Research	-0.851	-0.015	0.107	0.092

Table 23: Total returns per $1 million: commercial-potential outcome.

Change in aggregate NIH-wide commercial-potential-weighted output per $1 million budget increase. The top panel reports the own-institute components and the bottom panel reports the cross-institute components, using the same decomposition as Table 22. Institutes are sorted by total return.

Institute	No Ext. Return (PE)	No Ext. Return (GE)	Own Know. Spillovers	Net Own Return
General Medicine	8.900	8.902	0.347	9.249
Heart, Lung, and Blood	8.412	8.414	0.203	8.618
Human Genome	4.901	4.901	0.283	5.185
Cancer	5.805	5.807	0.615	6.422
Biomedical Imaging and Bioengineering	5.458	5.458	0.065	5.523
Environmental Health	4.874	4.875	0.429	5.304
Drug Abuse	3.093	3.093	0.371	3.465
Neurological Disorders and Stroke	3.190	3.191	0.104	3.295
Arthritis and Musculoskeletal and Skin Diseases	2.883	2.883	0.182	3.065
Dental and Craniofacial Research	2.581	2.582	0.147	2.729
Mental Health	2.212	2.212	0.128	2.339
Eye	2.223	2.223	0.178	2.401
Diabetes and Digestive and Kidney Diseases	1.934	1.934	0.080	2.015
Allergy and Infectious Diseases	1.945	1.946	0.177	2.123
Alcoholism	1.786	1.786	0.186	1.972
Aging	1.267	1.267	0.021	1.288
Deafness and Communication Disorders	1.135	1.135	0.086	1.221
Child Health and Human Development	0.907	0.907	0.060	0.967
Nursing Research	1.028	1.028	0.018	1.046

Institute	Talent Flow Ext. (PE)	Talent Flow Ext. (GE)	Know. Spill. to Others	Net Cross Return
General Medicine	-5.489	0.023	1.036	1.060
Heart, Lung, and Blood	-1.697	0.011	0.336	0.347
Human Genome	-4.509	0.011	2.450	2.460
Cancer	-2.380	0.010	0.368	0.379
Biomedical Imaging and Bioengineering	-4.445	0.009	0.713	0.722
Environmental Health	-3.645	0.024	0.674	0.698
Drug Abuse	-2.143	0.016	0.312	0.328
Neurological Disorders and Stroke	-2.717	0.011	0.355	0.366
Arthritis and Musculoskeletal and Skin Diseases	-3.200	0.021	0.429	0.449
Dental and Craniofacial Research	-3.208	0.025	0.432	0.457
Mental Health	-2.350	0.016	0.415	0.431
Eye	-3.020	0.020	0.295	0.316
Diabetes and Digestive and Kidney Diseases	-3.385	0.017	0.553	0.570
Allergy and Infectious Diseases	-2.734	0.007	0.285	0.292
Alcoholism	-2.329	0.018	0.333	0.351
Aging	-1.151	0.005	0.171	0.177
Deafness and Communication Disorders	-2.183	0.013	0.133	0.146
Child Health and Human Development	-2.516	0.012	0.372	0.383
Nursing Research	-1.845	0.012	0.153	0.166

Table 24: Total returns per $1 million: disease-burden outcome.

Change in aggregate NIH-wide US-DALY-weighted output per $1 million budget increase. The top panel reports the own-institute components and the bottom panel reports the cross-institute components, using the same decomposition as Table 22. The knowledge channels use the spillover estimates from Table 21. Institutes are sorted by total return.

K.2 Pairwise Knowledge-Spillover Heatmaps

Appendix Figures 23–25 report the full pairwise knowledge-spillover matrices for the three output metrics discussed in Section 6.4.

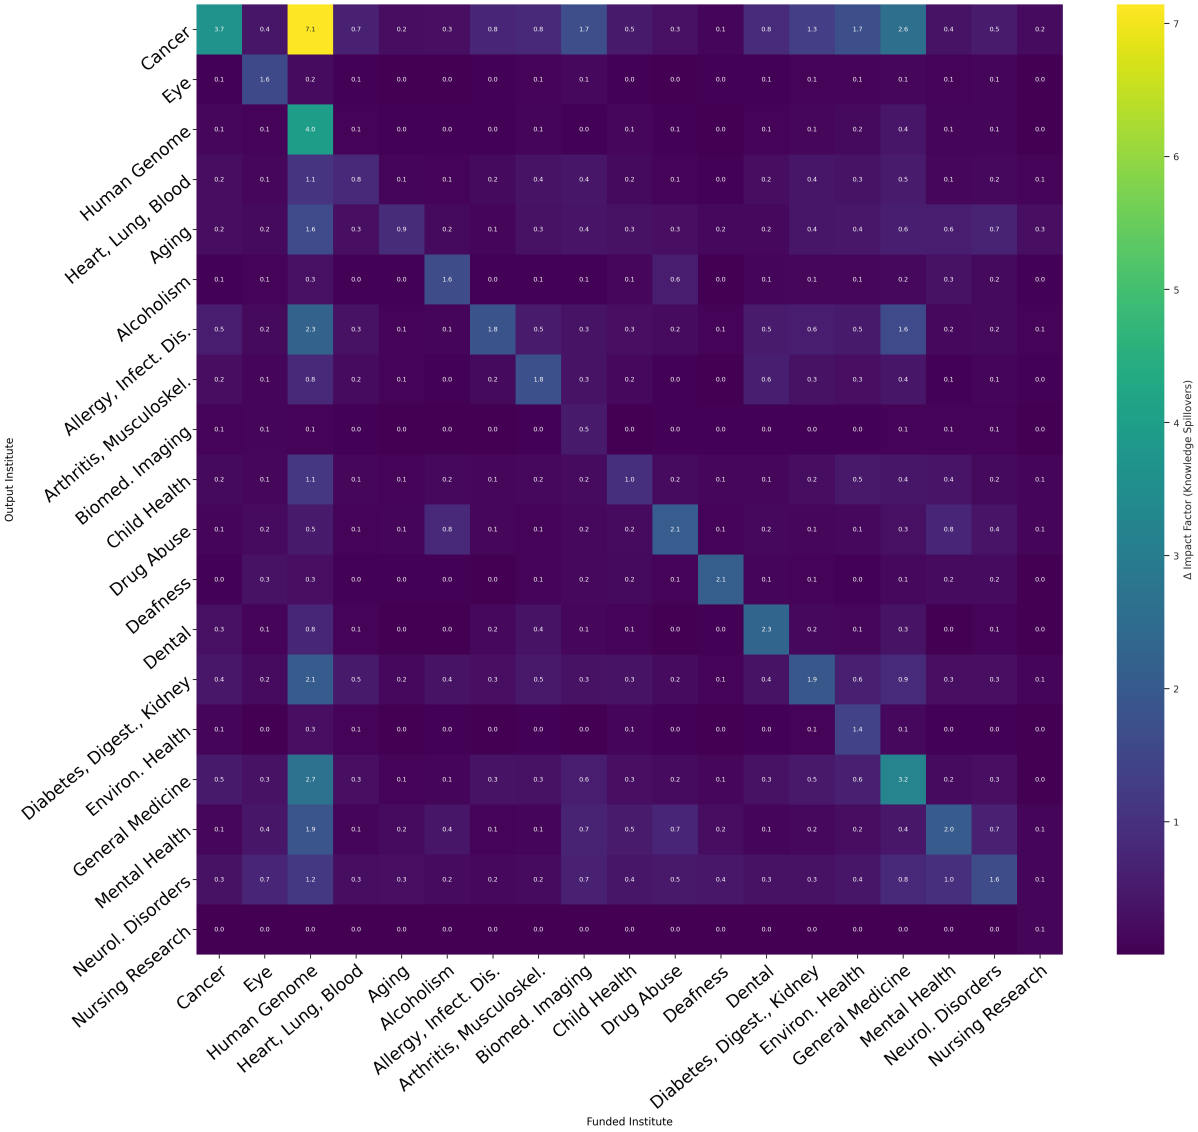


Figure 23: Pairwise knowledge spillover partial effects: impact factor.

Each cell (o, d) reports the change in impact-factor-weighted publications at output institute o per \$1 million budget increase at funded institute d, operating through the knowledge spillover channel, evaluated under the local log-output approximation at the 2018 baseline. Colour scale centred at zero.

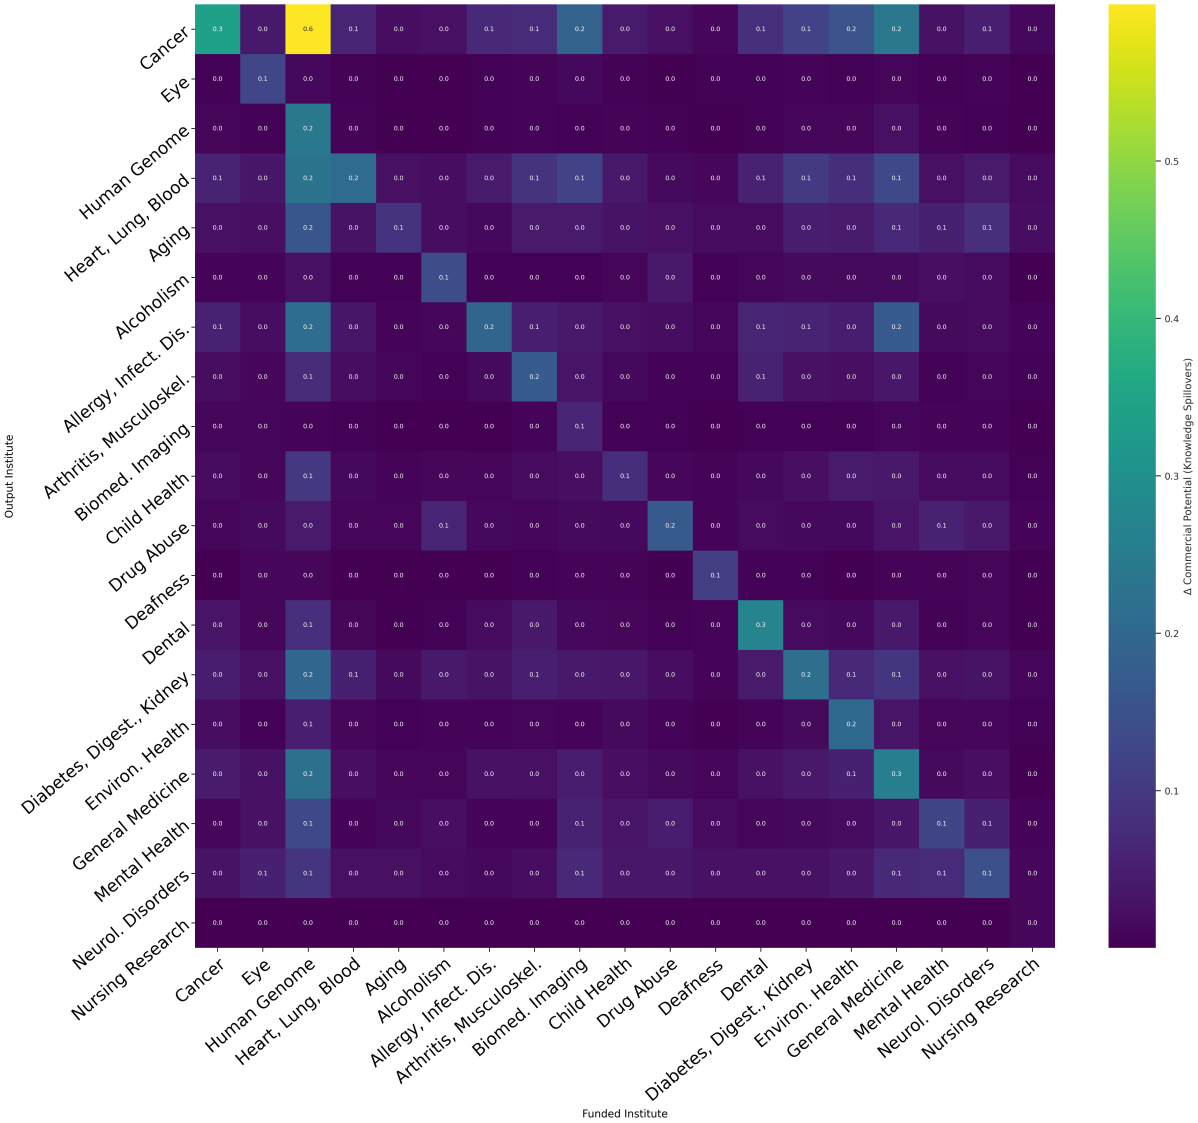


Figure 24: Pairwise knowledge spillover partial effects: commercial potential.

Each cell (o, d) reports the change in predicted patent citations at output institute o per \$1 million budget increase at funded institute d, operating through the knowledge spillover channel.

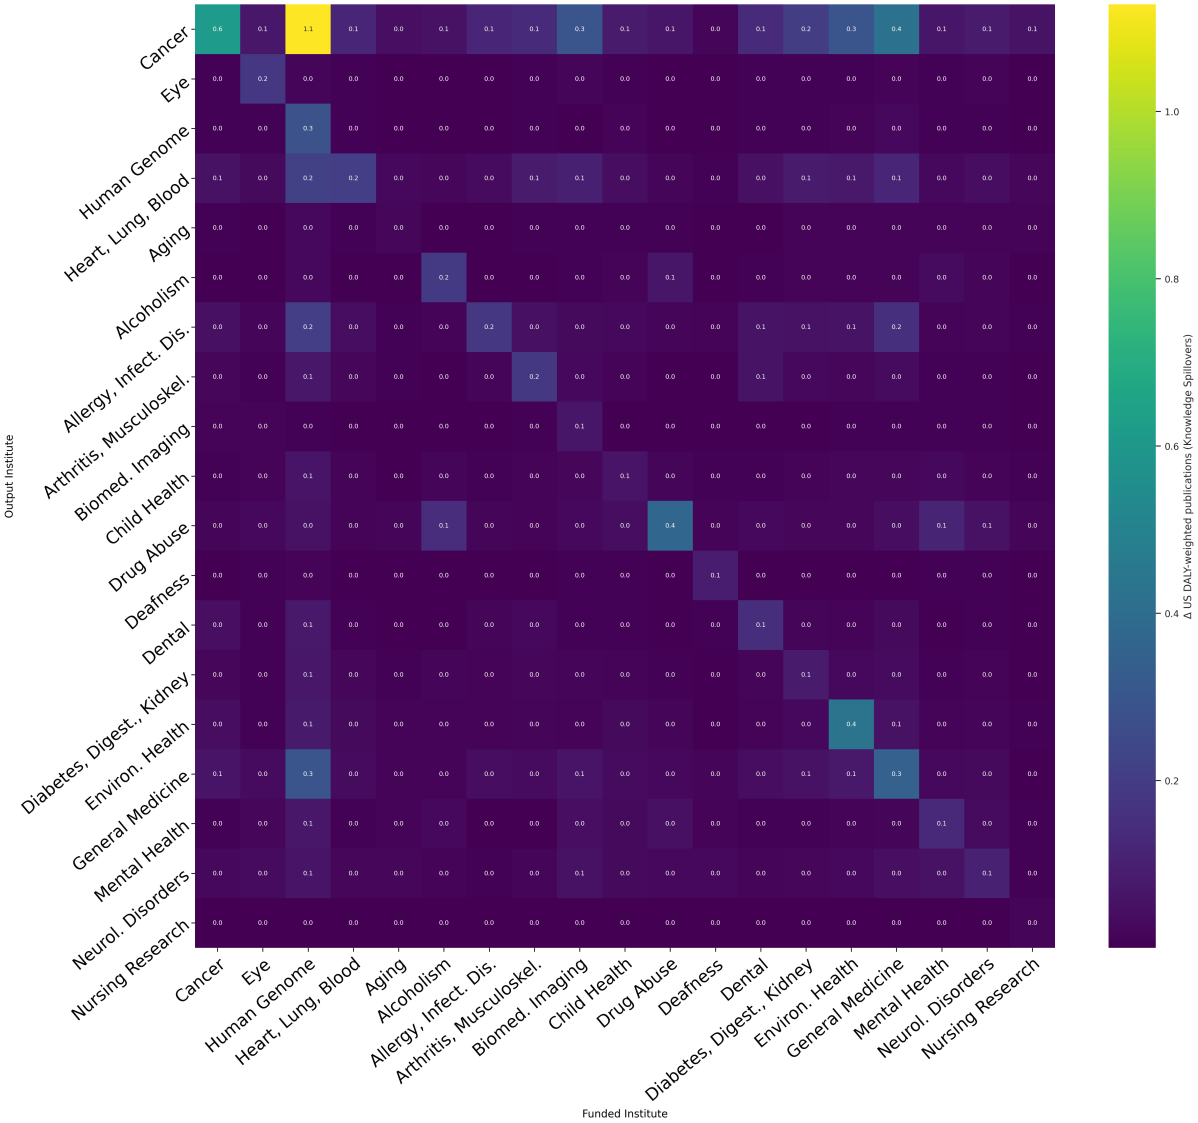


Figure 25: Pairwise knowledge spillover partial effects: disease burden.

Each cell (o, d) reports the change in US-DALY-weighted publications at output institute o per \$1 million budget increase at funded institute d, operating through the knowledge spillover channel.