US Science Cuts: The Cost for Europe Through Innovation Spillovers

Antoine Mathieu Collin* Dennis Verhoeven† Reinhilde Veugelers‡

March, 2026

Abstract

The Trump administration has terminated thousands of US federal research grants since early 2025. Because scientific knowledge flows across borders, these cuts threaten not only domestic innovation but also the technological progress of countries that build on US-funded research. We quantify the global and European economic value at risk using a method to value science-to-technology spillovers. Combining data on 6,034 terminated research grants across NIH, NSF, EPA, and CDC, we estimate counterfactual spillovers using PI-level track records and subfield-matched control groups, and project these forward using empirically fitted maturation curves. We estimate $577 million in global spillovers at the 2025 baseline, of which $56 million has already materialized through pre-termination publications. Projecting forward as if publications would have matured and accumulated patent citations, the cumulative foregone value reaches $6.1 billion globally and $856 million for European innovators by 2050. Since our measure captures only the patent-citation channel, these estimates represent a lower bound on the true cost of the US science cuts. Yet, our calculation of the spillover value lost hovers around the amount of budget "savings" for NSF and NIH grants.

Keywords: Knowledge spillovers · Science policy · Science funding cuts · Patent citations · Innovation · Technology transfer

JEL codes: O31 · O38

*Bruegel ; KU Leuven, Department of Management, Strategy & Innovation (MSI) | contact: antoine.mathieucollin@kuleuven.be.

†SKEMA Business School, Université Côte d'Azur (GREDEG); London School of Economics, Centre for Economic Performance (CEP); KU Leuven, Department of Management, Strategy & Innovation (MSI) | contact: dennis.verhoeven@skema.edu.

‡KU Leuven, Department of Management, Strategy & Innovation (MSI); CEPR; Bruegel | contact: reinhilde.veugelers@kuleuven.be

I. Introduction

The Trump administration has terminated or frozen thousands of US federal research grants since early 2025 (Tollefson et al., 2025; Kozlov et al., 2025). While the administration framed these actions as eliminating waste and realigning federal priorities, the affected grants span core areas of biomedical, environmental, and fundamental research across the National Institutes of Health (NIH), the National Science Foundation (NSF), the Environmental Protection Agency (EPA), and the Centers for Disease Control and Prevention (CDC). On top of these terminations, the administration proposed unprecedented cuts to science budgets for 2026, which could have affected a large share of drug approvals (Azoulay et al., 2025). Although Congress rejected the most extreme proposals, funding to several grant-awarding agencies remains disrupted (Kozlov et al., 2025).

These cuts are not just a domestic US affair. Scientific knowledge, once produced, flows freely across borders. Researchers worldwide read, cite, and build upon American science. Inventors in Europe, Asia, and elsewhere rely on US-funded discoveries when developing new technologies. This means that when the US government cuts research funding, the consequences extend well beyond American laboratories: they reduce the global stock of knowledge available to innovators everywhere. The question we address in this paper is: *if the US administration cuts science funding, how much economic value is lost globally, and how much of that loss falls on the European Union?*

Quantifying the economic value of scientific research is notoriously difficult. Research generates value through many channels: it trains skilled researchers, it produces knowledge that feeds into follow-on science, it informs public policy and regulation, and it provides the foundation for new technologies. We focus on one well-defined and measurable channel: the contribution of scientific publications to patented innovations. When a patent cites a scientific paper, it reveals that the underlying research contributed to a commercially valuable technology. By tracing these citation links and weighting them by the economic value of the citing patents, we can assign a dollar value to each scientific publication—a measure of its "spillover" to the technology sector.

We use this approach, developed by Martin et al. (2026) under the name Science Rank, to estimate the technology spillover value at risk from the Trump administration's grant terminations. Science Rank uses the full citation network between patents and scientific papers to attribute a share of each patent's private value back to the scientific publications it builds upon, producing paper-level spillover estimates in dollar terms. A key advantage of this measure is that it can be decomposed by geography: we can separately estimate how much of the spillover from a given paper flows to European innovators versus those in other regions.

Our empirical strategy proceeds in three steps. First, we link terminated grants—tracked by the Grant Witness database—to their scientific outputs using OpenAlex, and to their spillover values using Science Rank. Second, because most terminated grants are too recent to have generated measurable spillovers directly, we estimate counterfactual spillovers using two complementary control groups: the PI's own historical track record (Method B) and the average spillover of non-terminated grants in the same research subfield (Method C). Third, we account for the fact that scientific spillovers mature slowly—publications accumulate patent citations over decades—by fitting empirical maturation curves to the cross-sectional relationship between grant age and observed spillover, and projecting forward.

We find that the terminated grants represent $577 million in global technology spillovers at the 2025 baseline, of which $56 million (9.6%) has already materialized through publications produced before termination. Projecting forward to 2050, the cumulative foregone value reaches $6.1 billion globally and $856 million for European innovators. Comparing these losses to the budget savings from termination reveals stark differences across agencies: for NSF grants, the spillover value destroyed ($1.37 per dollar saved) exceeds the fiscal savings, while for NIH the ratio is $0.77—close to breakeven on the patent channel alone, before accounting for the health and welfare benefits that constitute NIH's primary mission.

These estimates should be interpreted as a lower bound on the true economic cost of the science cuts. Our measure captures only spillovers that manifest in patented technologies—one channel among many. It does not capture the value of trained researchers who leave the scientific workforce, the impact on open-source tools and datasets, the contribution to regulatory standards, or the multiplier effects of disrupted research networks. Furthermore, our projections rely on the assumption that future citation patterns will resemble historical ones, which may not hold for rapidly evolving fields. Despite these limitations, the analysis provides a concrete, empirically grounded estimate of economic value at stake—one that makes the cost of the cuts legible to policymakers on both sides of the Atlantic.

II. Data

We integrate four data sources to construct our analysis sample. The first provides the set of affected grants (our treatment group); the second enriches these grants with metadata and constructs a control group; the third links grants to their scientific outputs; and the fourth assigns dollar-valued spillover estimates to those outputs.

Our starting point is **Grant Witness**, a public database that tracks US federal grants that have been terminated, suspended, placed at risk, or reinstated under the Trump administration.[1] As of March 2026, Grant Witness records 11,597 affected grants across five agencies: NIH, NSF, CDC, EPA, and SAMHSA. These grants constitute our treatment group. We classify each grant as either a research grant or a non-research grant (training, infrastructure, service delivery) based on its NIH activity code or agency-specific identifiers. SAMHSA grants are exclusively non-research (service and community grants) and are excluded from the analysis, as are training, fellowship, career development, and infrastructure grants from other agencies. This leaves 6,034 research grants.

We enrich these grants using **Dimensions**, a commercial research intelligence database. Dimensions provides funding amounts, start and end dates, and—crucially—PI identifiers that allow us to retrieve the full grant portfolios of affected researchers. For each PI on a terminated grant, we collect all their other grants in Dimensions, forming the *control group*: the set of non-terminated grants held by the same researchers, which serves as the basis for our counterfactual spillover estimates. Of the 11,597 Grant Witness records, 10,300 (89%) match to Dimensions. Unmatched grants are predominantly CDC and EPA non-research awards that are not indexed in Dimensions.

To link grants to their scientific outputs, we use **OpenAlex**, an open bibliometric database that maps grants to published papers through funding acknowledgment metadata. OpenAlex covers approximately 250 million works and 3.4 million award records from US funders. For each grant in both the treatment and control groups, we retrieve the set of publications that acknowledge it as a funding source. This publication-level linkage is essential: it is through publications that the knowledge produced by a grant enters the broader scientific and technological ecosystem.

Finally, we assign economic values to these publications using **Science Rank**, a dataset developed by Martin et al. (2026) that quantifies the contribution of scientific papers to patented innovations. Science Rank uses the full citation network between patents and scientific papers to trace how knowledge flows from science into technology. Each patent's private value—estimated

[1] See https://grantwitness.org.

from stock market reactions to patent grants following Kogan et al. (2017)—is attributed back to the scientific papers it cites, both directly and indirectly through chains of intermediate citations. The result is a paper-level measure of technology spillovers, expressed in millions of dollars. Importantly, these spillovers can be decomposed by the geographic location of the patent applicant, allowing us to separately estimate the value flowing to European innovators (`ev_eu`) versus the global total (`ev_global`). Science Rank covers publications from approximately 1985 to 2018, reflecting the time needed for patent citations to accumulate.

A. Sample overview

Table 1 presents summary statistics for all 11,597 affected grants across five agencies. SAMHSA grants are exclusively non-research (service and community grants) and are excluded from the remainder of the analysis. Table 2 restricts the sample to the 6,057 research grants that form our analysis sample.

Table 1: Summary statistics of affected grants by agency

	Grants			Funding ($B)		Median	Data coverage	
Agency	*N*	Research	% Term.	Total	Remaining	($K)	Pubs (%)	SR (*N*)
CDC	579	33	40.8	60.3	16.6	16,052	4.3	2
EPA	671	293	90.5	29.6	28.6	1,000	3.9	0
NIH	5,418	3,735	20.6	17.2	7.2	1,629	31.5	342
NSF	1,996	1,996	68.3	1.7	0.9	500	12.3	2
SAMHSA	2,933	0	4.6	8.8	2.5	1,420	0.0	0
Total	11,597	6,057	29.8	117.7	55.8	1,296	17.3	346

Notes: Data from Grant Witness, enriched with Dimensions, OpenAlex, and Science Rank. "% Term." is the share of grants with terminated status. "Pubs (%)" is the share of grants with ≥ 1 linked OpenAlex publication. "SR (*N*)" is the number of grants with Science Rank spillover data. Funding figures include mega-grants (> $500M); 40 such grants account for $54.9B of the $117.7B total award value.

Table 2: Summary statistics of affected research grants by agency

	Grants		Funding ($B)				Data coverage	
Agency	*N*	% Term.	Total	Remaining	Median ($K)	Mega	Pubs (%)	SR (*N*)
CDC	33	60.6	0.2	0.1	2,407	0	21.2	1
EPA	293	81.9	27.4	26.9	1,217	6	7.8	0
NIH	3,735	16.5	14.8	5.9	2,113	1	35.3	230
NSF	1,996	68.3	1.7	0.9	500	0	12.3	2
Total	6,057	37.0	44.1	33.8	1,431	7	26.3	233

Notes: Research grants only ($N = 6{,}057$). Non-research grants (training, fellowships, career development, and service/community grants—including all SAMHSA grants) are excluded; see Table 1. "% Term." is the share terminated. "Mega" counts grants with award value $\geq$ $500M. "Pubs (%)" is the share with ≥ 1 linked OpenAlex publication. "SR (*N*)" is the count with Science Rank spillover data.

NIH accounts for the largest share of research grants (N = 3,735), followed by NSF (N = 1,996) and EPA (N = 293). Termination rates vary considerably: EPA has the highest rate (81.9%), followed by NSF (68.3%), CDC (60.6%), and NIH (16.5%). Total award value at stake amounts to $44.1 billion for research grants, of which $33.8 billion remains unspent. Seven "mega-grants" (award value ≥ $500M), all from EPA, account for a disproportionate share of total funding.

Figure 1 shows the distribution of award values by agency. The distribution is highly right-skewed, spanning four orders of magnitude. EPA mega-grants (red dots) are clear outliers. NSF grants are systematically smaller (median $500K) than NIH grants (median $2.1M).

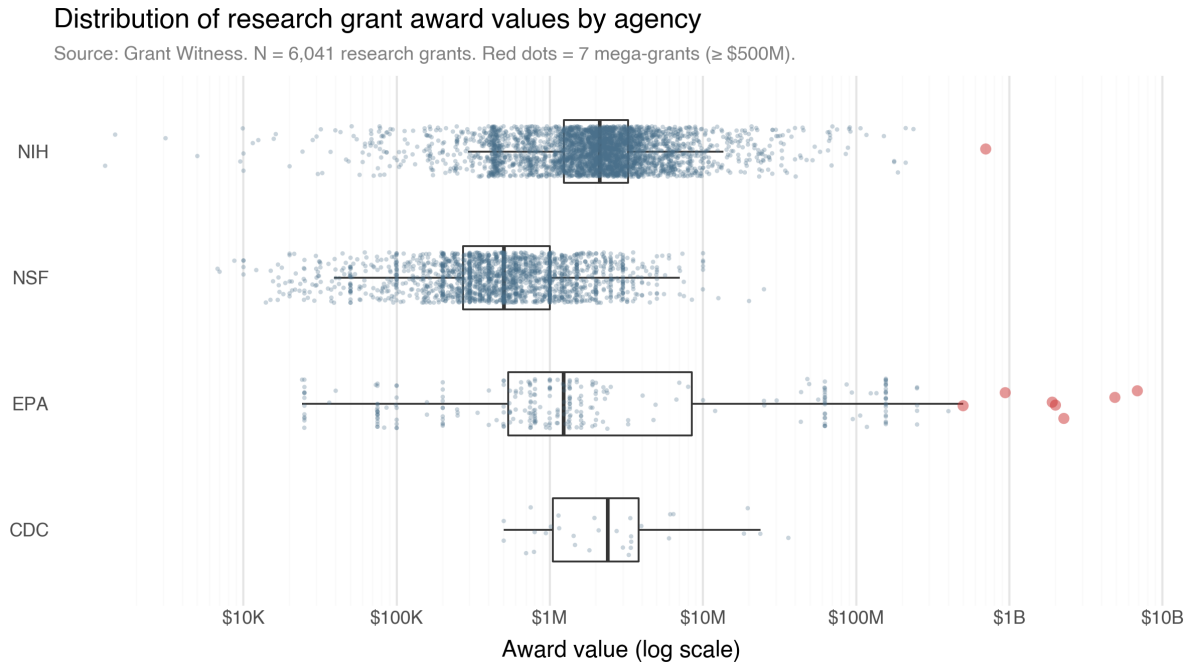


Figure 1: Distribution of research grant award values by agency (grants affected

Source: Grant Witness, enriched with Dimensions. N = 6,057 research grants. Red dots indicate mega-grants (≥ $500M).

B. Research fields affected

We classify terminated grants by research field using OpenAlex topic metadata. For each grant, we retrieve the publications linked to it via the OpenAlex Works API and assign the grant's field as the *modal* primary topic across its publications, following the four-level OpenAlex taxonomy (topic, subfield,field, domain). This covers 95.8% of terminated research grants (5,780 out of 6,057). Figure 2 shows the distribution of terminated grants across OpenAlex fields. Medicine and Biochemistry/Genetics/Molecular Biology dominate, reflecting the large NIH share.

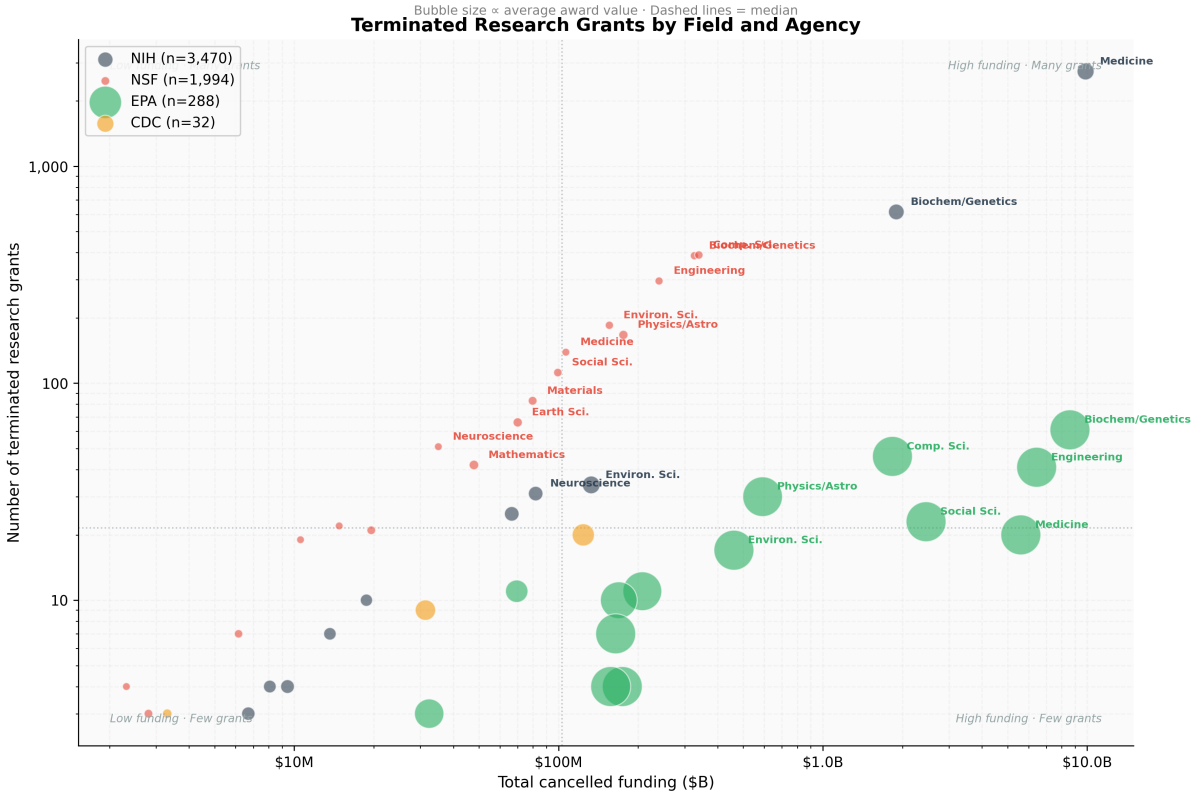


Figure 2: Terminated research grants by OpenAlex field and funding agency. Each bubble represents a field–agency combination; x-axis = total cancelled funding, y-axis = number of grants, bubble size $\propto$ average award value. Dashed lines mark medians.

III. Methodology

This section describes our approach to estimating the technology spillover value at risk from terminated grants. We first present the three complementary estimation methods, then describe the S-curve projection used to account for citation lag.

A. Spillover estimation

We estimate the technology spillover value at risk from terminated grants using three complementary methods (A, B, C) that provide base estimates, plus a projection layer (P) that accounts for citation lag:

- **Method A** — Direct observation from ScienceRank for grants whose publications are already indexed. These capture spillovers that have *already materialized* despite termination.

- **Method B** — PI-level historical track record as a *within-PI control group* (1,198 grants, 19.9%). For each terminated grant whose PI can be identified and has ScienceRank-covered prior grants, we assign the PI's historical average spillover, age-adjusted to the grant's current maturity using the S-curve[2].

- **Method C** — Subfield-matched conditional spillover from *non-terminated grants* in the PI portfolio as an *alternative control group* (4,836 grants, 80.1%). For each terminated grant, we assign the mean spillover of all PI-portfolio grants in the same OpenAlex subfield, age-adjusted to the grant's current maturity[3]. Of these: 4,650 matched by subfield, 8 by field-level fallback, 178 by overall portfolio mean.

- **P (age adjustment)** — A, B and C estimates are age-adjusted using a logistic S-curve: $EV(T) = \text{raw_estimate} \times f(\text{future_age}) \,/\, f(\text{reference_age})$, where $f(t) = S(t)/L$ is the maturity fraction.

We present in Annex A more details about our methodology.

The total *counterfactual* spillover (i.e. what the terminated grants would have generated as spillovers from the generated research, had they continued) combines methods B and C, projected forward via S-curves:

$$EV^{\text{counterfactual}}(T) = \sum_{g \in \mathcal{G}_B} P\left(\hat{EV}_g^B, T\right) + \sum_{g \in \mathcal{G}_C} P\left(\hat{EV}_g^C, T\right)$$

where $\mathcal{G}_B$ and $\mathcal{G}_C$ are mutually exclusive and $P(\cdot, T)$ denotes S-curve projection to horizon T. Method A identifies spillovers that have already materialized through publications produced before termination. These already-published papers will continue accumulating patent citations regardless of the grant's termination, so the realized portion is also projected forward:

$$EV^{\text{realized}}(T) = \sum_{g \in \mathcal{G}_A} P\left(EV_g^{\text{obs}}, T\right)$$

The *genuinely lost* spillover—the research that will now not take place—is:

$$EV^{\text{lost}}(T) = EV^{\text{counterfactual}}(T) - EV^{\text{realized}}(T)$$

We repeat the same procedure for `ev_eu`, the spillover value accruing to European innovators.

[2] Age-adjusted to the grant's current maturity using the global logistic S-curve (L=\$1.49M, k=0.20, t_0=18 years).

[3] Same age-adjustment method

The estimates capture spillovers at their *current* level, which underestimates the true long-run impact of younger grants. Scientific publications accumulate patent citations over decades, so a grant started in 2020 has had far less time to generate measurable spillover than one started in 2005. The S-curve projection (P) is applied as a *layer on top of* to project each grant's base estimate forward to its long-run asymptote.

The base spillover estimates reflect the mature PI portfolio (median grant age 19 years), but the terminated grants are much younger (median 3 years). We therefore age-adjust by rescaling each grant's base estimate to its current maturity:

$$EV_{\text{adj}}(t) = EV_{\text{base}} \times \frac{f(t)}{f(t_{\text{ref}})}$$

where $f(t) = S(t)/L$ is the maturity fraction at grant age t, and $t_{\text{ref}} = 19$ years is the median age of the reference portfolio. At the 2025 baseline, a 3-year-old grant receives approximately 9% of its base estimate; by 2050 (age 28), it reaches 170%. The maximum multiplier is empirically bounded at $f(50)/f(t_{\text{ref}}) = 1.91\times$.

Not all publications will yield a lot of patent citation, and this is mainly driven by field specificity. For that reason, we fit a logistic (S-curve) model to the cross-sectional relationship between grant age and mean observed spillover, using the full population of PI portfolio grants with ScienceRank coverage ($N \approx 18{,}000$):

$$S(t) = \frac{L}{1 + e^{-k(t-t_0)}}$$

where t is grant age, L is the asymptotic mature spillover, k is the growth rate, and t_0 is the half-max age. The global fitted parameters are $L = \$1.49\text{M}$, $k = 0.20/\text{yr}$, and $t_0 = 18$ years (Figure 3).

We fit the logistic model separately for the 15 largest subfields. Four subfields yield reliable fits (Molecular Biology, Infectious Diseases, Genetics, Epidemiology), and five additional fits are obtained one level up, at the field level (Medicine, Environmental Science, Engineering, Neuroscience, Physics & Astronomy). The remaining subfields exhibit three failure patterns (Figure 9):

1. **Humped** (Cognitive Neuroscience, Physiology): spillover peaked for 1990s–2000s grants then declined following a past "golden era."

2. **Front-loaded** (AI, Health/Toxicology): younger grants generate higher spillovers than older ones, reflecting field transformation and a recent technology (e.g. the deep learning revolution).

3. **Blockbuster-driven** (Oncology): a few outliers dominate the distrubution, producing implausible asymptotes that cannot reasonably be used for prediction.

For these reasons, in our method, each grant is assigned an S-curve: (1) subfield curve if available, (2) field curve, (3) static estimate for low-patent subfields, (4) global curve as fallback. Figures in Annex shows different curves for some fields.

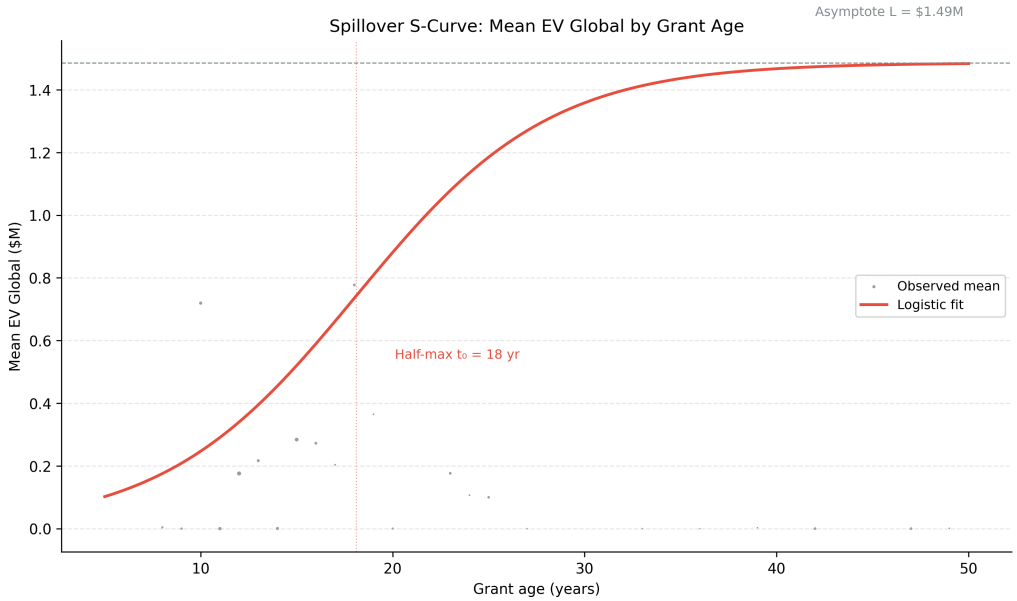


Figure 3: Global S-curve: mean EV Global by grant age with fitted logistic ($L = \$1.49\text{M}$, $t_0 = 18$ years). Dot size $\propto$ sample count.

IV. Preliminary results

We present our estimates in two parts: first the spillover values by estimation method, then a comparison with the budget savings from grant termination.

A. Spillover estimates

Of the 6,034 terminated research grants, 1,198 (19.8%) receive a Method B estimate based on PI track records, while the remaining 4,836 (80.2%) are covered by Method C using subfield-matched conditional means. The combined baseline estimate is $577 million in age-adjusted global spillovers at 2025, growing to $6.1 billion by 2050 as publications accumulate patent citations along grant-level S-curves.

Table 3: Projected spillover at five-year horizons: counterfactual, realized, and lost

Year	Counterfactual (B+C)		Realized (A)		Genuinely Lost	
	Global ($M)	EU ($M)	Global ($M)	EU ($M)	Global ($M)	EU ($M)
2025	577	80	56	5	520	75
2030	1,170	164	85	9	1,085	155
2035	2,251	316	116	14	2,135	302
2040	3,780	530	141	18	3,640	513
2050	6,254	876	160	21	6,094	856

Counterfactual = age-adjusted B+C estimates projected forward via hierarchical S-curves. Realized = already-materialized spillovers (Method A) also projected forward, since existing papers continue accumulating patent citations over time. Lost = Counterfactual − Realized. The realized portion shrinks from 9.8% (2025) to 2.6% (2050) as the counterfactual grows rapidly while realized spillovers grow more slowly.

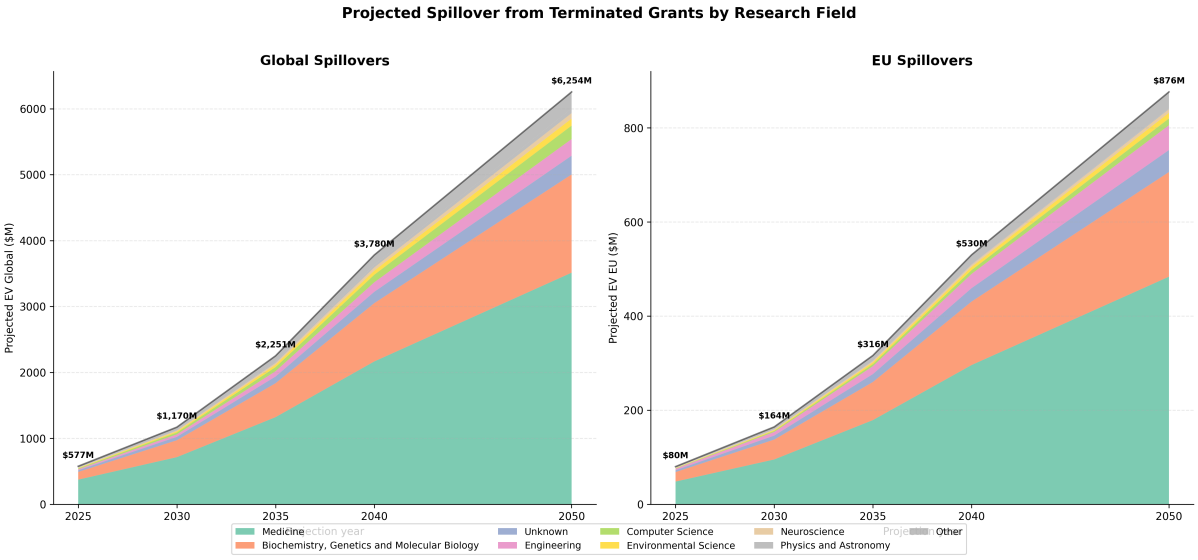


Figure 4: Projected spillover over time by OpenAlex field (stacked area). Left: global. Right: EU. Medicine dominates (∼60%).

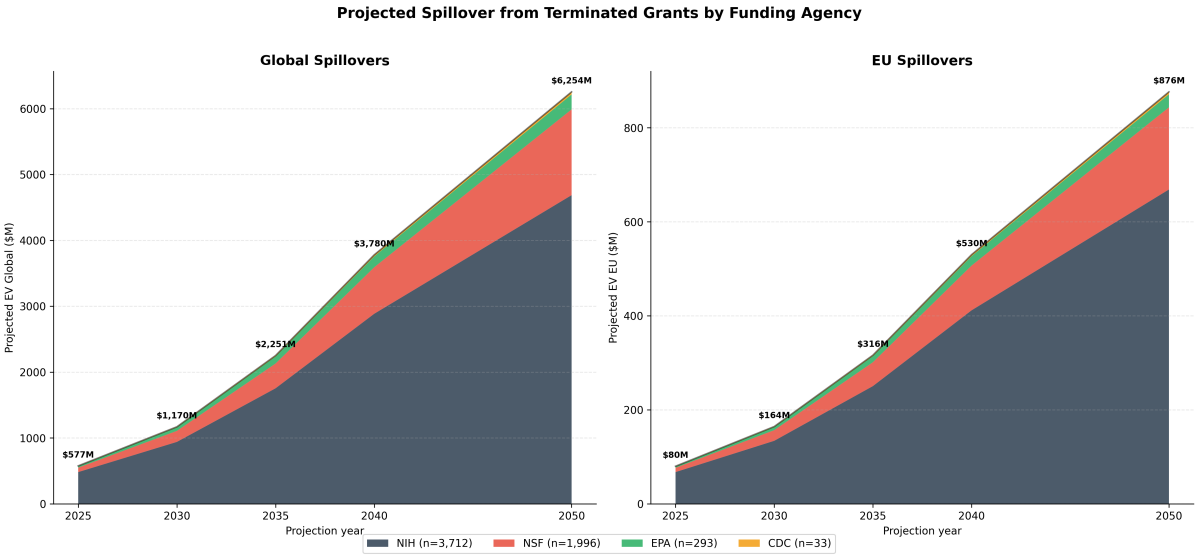


Figure 5: Projected spillover over time by funding agency (stacked area). Left: global. Right: EU.

B. Budget savings vs. foregone spillover

A central policy question is whether the fiscal savings from terminating research grants justify the innovation value destroyed. Table 4 compares the unspent funds returned to the US Treasury (budget savings) with the projected foregone spillover, by agency. The terminated grants represent $33.7 billion in unspent funds. Against this, we estimate $6.1 billion in foregone global spillovers by 2050—a ratio of **$0.18 in spillover destroyed per $1 saved**.

Table 4: US budget savings vs. foregone global spillover by agency

Agency	N	Savings ($B)	EV 2025 ($M)	EV 2050 ($M)	EU 2050 ($M)	EV/$ saved
NIH	3,712	5.9	463	4,500	642	0.769
NSF	1,996	0.9	64	1,250	168	1.367
EPA	293	26.9	23	218	27	0.008
CDC	33	0.1	3	37	5	0.671
Total	**6,034**	**33.7**	**554**	**6,005**	**842**	**0.178**

For every $1 in US budget savings, $0.18 in global technology spillovers by 2050 is lost. NSF grants yield $1.37 in spillover per dollar saved—terminating them costs more in spillovers than it saves in budget. EPA's low ratio ($0.008) reflects ScienceRank's bias toward patent-based spillovers; EPA grants may generate value through regulatory and environmental channels not captured here.

This aggregate ratio masks substantial heterogeneity across agencies (Figure 7):

- **NSF**: \$1.37 in spillover per \$1 saved. NSF grants are small (median \$500K) but in high-spillover fields (computer science, physics, mathematics). Terminating NSF grants *costs more in foregone spillovers than it saves in budget.*

- **NIH**: \$0.77 per \$1 saved—close to breakeven on patent spillovers alone, before accounting for non-patent channels (health outcomes, clinical guidelines).

- **CDC**: \$0.67 per \$1 saved, reflecting a mix of epidemiological research with moderate patent impact.

- **EPA**: \$0.008 per \$1 saved. This very low ratio reflects two factors: EPA's large "mega-grants" (7 grants $\geq$ \$500M) inflate budget savings, while EPA research generates value primarily through regulatory standards and environmental policy rather than patents.

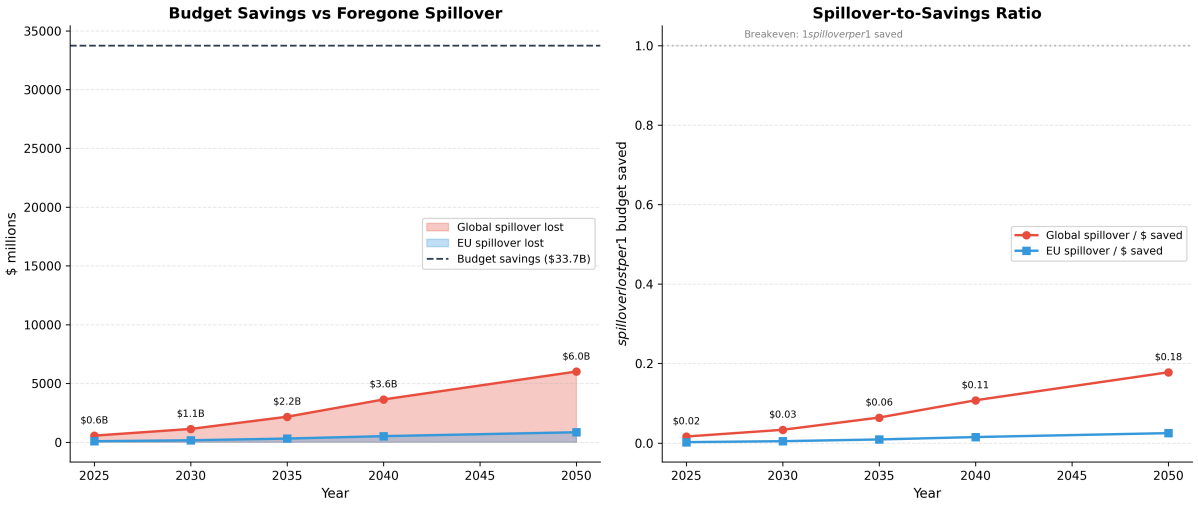


Figure 6: Left: cumulative spillover lost (global and EU) compared to the one-time budget savings from grant termination. The dashed line shows the \$33.7B in unspent funds. Right: spillover-to-savings ratio over time.

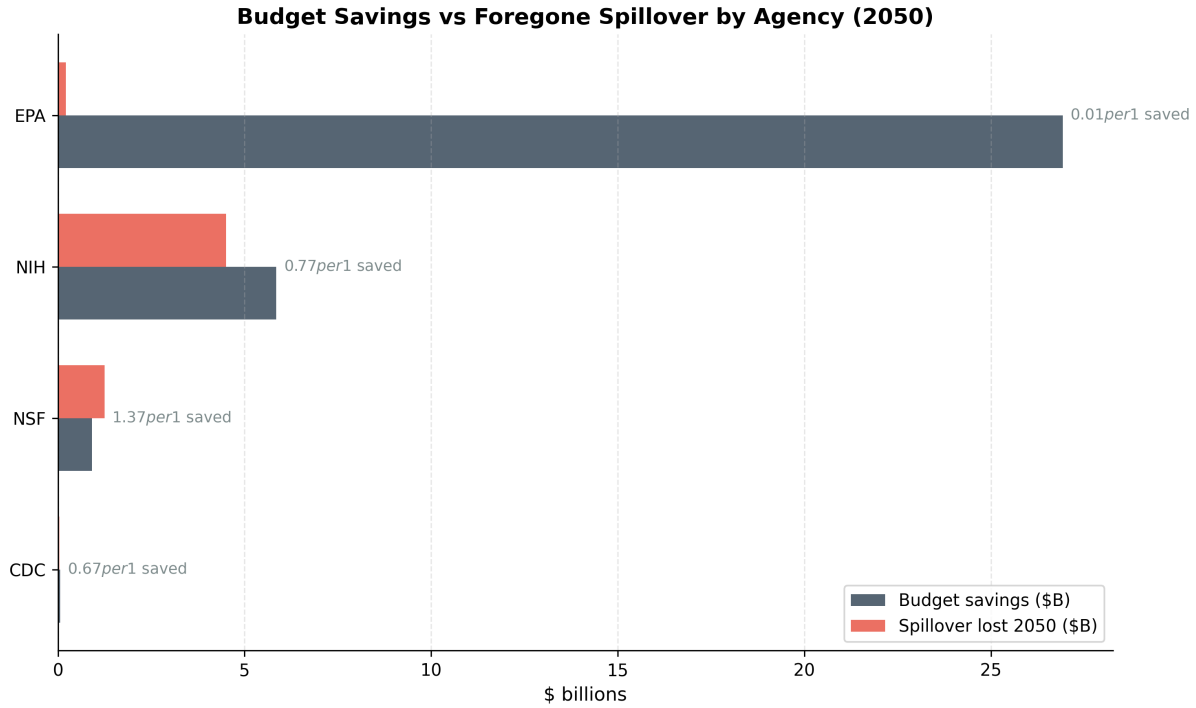


Figure 7: Budget savings (blue) vs. foregone spillover by 2050 (red) by funding agency. Labels show spillover per dollar saved. NSF grants exceed breakeven ($1.37/$1).

Table 5 and Figure 8 present the same comparison by research field. Medicine generates the largest absolute spillover loss ($3.4B by 2050), but Materials Science ($0.47/$1), Neuroscience ($0.38/$1), and Medicine ($0.36/$1) have the highest per-dollar ratios. Engineering ($0.04/$1) and Social Sciences ($0.02/$1) show low ratios, reflecting either low patent-citation intensity or the dominance of non-patent spillover channels in these fields.

Table 5: US budget savings vs. foregone global spillover by research field

Field	N	Savings ($B)	EV 2025 ($M)	EV 2050 ($M)	EU 2050 ($M)	EV/$ saved
Medicine	2,900	9.4	359	3,373	464	0.357
Biochemistry, Genetics	1,072	9.3	112	1,431	214	0.154
Engineering	340	6.5	15	242	51	0.037
Computer Science	436	2.0	11	196	14	0.100
Social Sciences	139	2.5	4	50	7	0.020
Environmental Science	235	0.6	7	102	12	0.173
Physics and Astronomy	197	0.7	4	68	2	0.101
Neuroscience	92	0.2	5	79	7	0.381
Materials Science	101	0.1	3	61	2	0.468
Unknown / Other	255	1.1	24	273	45	0.241
Total	**6,034**	**33.7**	**554**	**6,005**	**842**	**0.178**

Medicine generates the highest absolute spillover ($3.4B by 2050), while Materials Science and Neuroscience have the highest spillover-to-savings ratios ($0.47 and $0.38 per $1 saved respectively). Engineering and Social Sciences show the lowest ratios, reflecting either low patent-citation intensity or large grant sizes with modest spillover. "EV/$ saved" is the ratio of projected 2050 spillover to unspent cancelled funding.

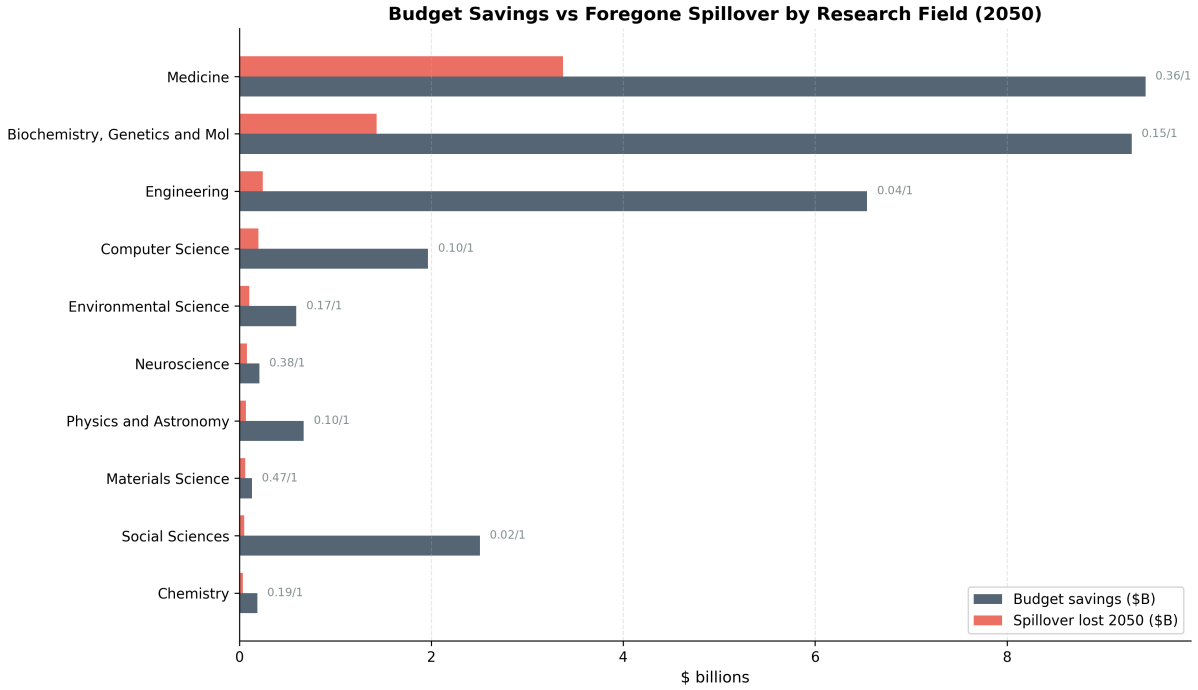


Figure 8: Budget savings (blue) vs. foregone spillover by 2050 (red) by OpenAlex research field. Labels show spillover per dollar saved.

These ratios represent a *lower bound* on the true cost of termination, for three reasons: (i) ScienceRank captures only patent-citation spillovers, missing channels such as open-source software, trained researchers, regulatory standards, and public health improvements; (ii) the S-curve projection uses the global maturation profile, which may underestimate fields with longer citation lags; (iii) we do not account for the multiplier effects of lost research on subsequent generations of innovation.

References

Azoulay, P., Clancy, M., Li, D., and Sampat, B. N. (2025). What if nih had been 40% smaller? *Science*, 389(6767):1303–1305. [Cited on page 2.]

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The quarterly journal of economics*, 132(2):665–712. [Cited on page 5.]

Kozlov, M., Witze, A., and Garisto, D. (2025). White House stalls release of approved US science budgets. *Nature*. [Cited on page 2.]

Martin, R., Shah, A., Valero, A., and Verhoeven, D. (2026). Spillovers from science. Technical report, Centre for Economic Performance, LSE. [Cited on pages 3 and 4.]

Tollefson, J., Witze, A., and Kozlov, M. (2025). Science under Trump 2.0: the likely winners and losers. *Nature*, 637:533–536. [Cited on page 2.]

Appendix

A. Methods

Method A — Direct estimation (pre-2018 grants)

For grants with publications covered by Science Rank, we directly observe spillover values. Let $g \in \mathcal{G}_A$ denote the set of terminated grants with at least one publication indexed in Science Rank. For each such grant, we sum spillover values across its linked publications:

$$EV_g^{\text{obs}} = \sum_{p \in \mathcal{P}_g} ev_global_p$$

where $\mathcal{P}_g$ is the set of Science Rank-covered publications linked to grant g. Note that Method A captures spillovers that have *already materialized*—publications produced before the grant was terminated, whose citations to patents have already been recorded. For the 228 Method A grants, $2.4 billion in global spillover was realized despite termination, dominated by a small number of NIH Cancer Center grants (P30) with thousands of publications.

Method B — PI-level track record (control group 1)

Method B uses the PI's own research history as a *within-PI control group*: for each terminated grant whose PI can be identified and has ScienceRank-covered grants, we compute the PI's historical average spillover and assign it to the terminated grant. This approach controls for unobservable researcher quality and is our preferred estimator when available.

We link each terminated grant to its PI using OpenAlex award records. The OpenAlex awards dataset ($N \approx 3.4$M) contains PI names alongside normalized grant identifiers. For each PI identified on a terminated grant, we retrieve all their other grants in the PI portfolio that have Science Rank coverage, and compute their historical average spillover per grant:

$$\bar{EV}_\pi = \frac{1}{|\mathcal{G}_\pi^{\text{hist}}|} \sum_{g \in \mathcal{G}_\pi^{\text{hist}}} EV_g^{\text{obs}}$$

where $\mathcal{G}_\pi^{\text{hist}}$ is the set of PI π's prior grants with observed spillover data.

Funding adjustment. The simple PI average does not account for differences in grant size between the terminated grant and the PI's historical portfolio. Since funding and spillover are positively correlated ($\rho = 0.56$ in the portfolio data), we scale the PI's average by the ratio of terminated grant funding to PI average funding, with an elasticity of 0.5:

$$\hat{EV}_g^B = \bar{EV}_\pi \cdot \left(\frac{F_g}{\bar{F}_\pi}\right)^{0.5}$$

where F_g is the terminated grant's award value and $\bar{F}_\pi$ is the mean award value across the PI's historical portfolio. The ratio is clipped to $[0.1, 10]$ to prevent extreme outliers.

PI disambiguation. OpenAlex grants data does not provide ORCIDs (the field exists in the schema but is entirely empty as of March 2025). We therefore disambiguate PIs using a composite key: (PI name, affiliation). This is a pragmatic trade-off: homonymous researchers at the *same* institution will be treated as one person (averaging their spillover histories), while a single researcher who changed institution will appear as two distinct PIs (each with a partial portfolio). In practice this resolves the vast majority of name collisions (1,061 unique names → 1,151 unique name-institution pairs).

Age adjustment. The raw PI average $\bar{EV}_\pi$ is computed across grants at various ages (5, 10, 25 years old), so it reflects a mix of maturation stages. The average portfolio age is $\bar{t} \approx 19$ years, corresponding to a maturity fraction $f(\bar{t}) = 0.52$. A terminated grant at age 3 has a maturity fraction of only $f(3) = 0.05$. We therefore age-adjust:

$$\hat{EV}_g^B(T) = \bar{EV}_\pi \times \frac{f(T_{\text{age}})}{f(\bar{t})}$$

where T_{age} is the grant's age at year T. At 2025, a 3-year-old grant receives $\bar{EV}_\pi \times 0.05/0.52 \approx 9\%$ of the raw PI average—appropriate because it has barely started generating spillovers. By 2050 (age 28), it receives $\bar{EV}_\pi \times 0.88/0.52 \approx 1.7\times$ the PI average, reflecting full maturation. The maximum multiplier at the asymptote is $1/f(\bar{t}) \approx 1.91\times$.

Method B coverage. Of the 6,034 terminated research grants, 1,198 (19.9%) receive a Method B estimate. Table 6 shows the breakdown.

Table 6: Method B estimate breakdown

Category	N	%	Mean EV ($M)	Treatment
B > $0.001M (spillover history)	508	42.4%	0.999	Trust B
B ≈ 0, ≥5 SR grants (genuine zero)	185	15.4%	0.000	Trust B
B ≈ 0, 1–4 SR grants (thin history)	505	42.2%	0.000	max(B, C)
Total Method B	1,198	100%	0.424	

Thin-history PIs explain 98.2% of the B–C gap ($366M of $373M). For PIs with positive B, C *under*estimates B ($381M vs. $507M).

Thin-history PI handling. Table 6 reveals that 505 Method B grants (42.2%) have B ≈ 0 with only 1–4 ScienceRank-covered grants in the PI's portfolio. Cross-referencing with subfield data shows that the B=0 rate drops sharply with portfolio size: 71–84% of thin-history PIs in high-spillover subfields (Molecular Biology, Infectious Diseases) show B=0, compared to 36–44% for PIs with ≥5 SR grants. This suggests that thin-history B=0 reflects *data insufficiency* rather than genuine zero spillover.

We address this with a hybrid rule: for PIs with ≤4 SR grants, we assign $\max(B, C)$—preserving B when it is positive but falling back to the subfield mean when B=0. For PIs with ≥5 SR grants, we trust the B estimate. This hybrid approach adds $79M to the 2025 baseline (+16%) and $452M to the 2050 projection (+7.8%).

Importantly, thin-history PIs explain 98.2% of the aggregate B–C gap ($366M of $373M total). For PIs with positive B (>$0.001M), Method C actually *underestimates* B ($381M vs. $507M). The entire B–C discrepancy is driven by the thin-history zeroes receiving the positive subfield average under C.

The 505 PIs with thin history (1–4 SR grants) may simply lack sufficient data; for these, Method C provides a complementary estimate.

Method C — Subfield-matched conditional spillover (control group 2)

For grants where Method B is not available—either because the PI could not be matched in OpenAlex or because the PI has no prior grants with ScienceRank coverage—we construct an alternative control group by matching terminated grants to *non-terminated grants in the same OpenAlex subfield*. For each subfield s, we compute the unconditional mean spillover from the PI control group:

$$\hat{EV}_g^C = E[EV \mid \text{subfield}(g) = s] = \frac{1}{|\mathcal{G}_s|} \sum_{g' \in \mathcal{G}_s} EV_{g'}^{\text{obs}}$$

where $\mathcal{G}_s$ is the set of PI portfolio grants in subfield s with ScienceRank coverage. For grants whose subfield is not available, we fall back to the field-level mean, then the overall portfolio mean. As with Method B, the subfield mean is age-adjusted: $\hat{EV}_g^C(T) = \bar{EV}_s \times f(T_{\text{age}})/f(\bar{t}_s)$, where $\bar{t}_s$ is the mean portfolio age in subfield s.

S-curves

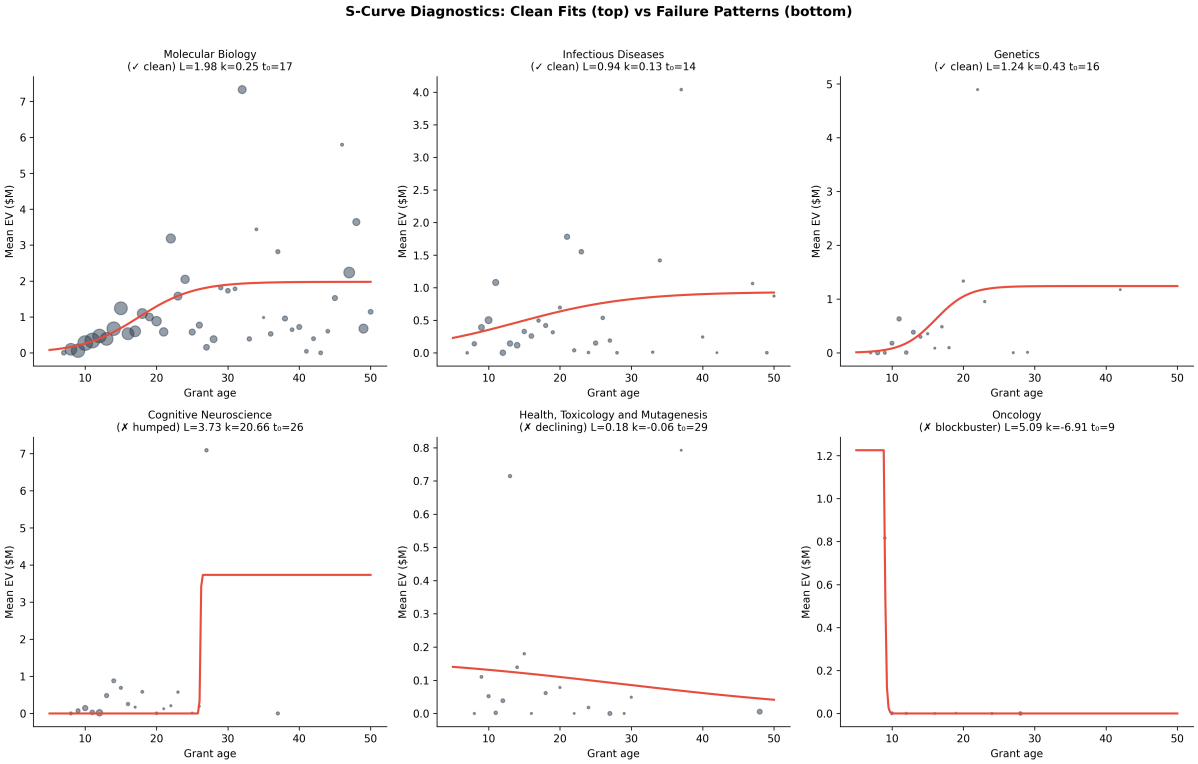


Figure 9: S-curve diagnostics. Top: clean fits (Molecular Biology, Infectious Diseases, Genetics). Bottom: failure patterns—humped, declining, blockbuster-driven.

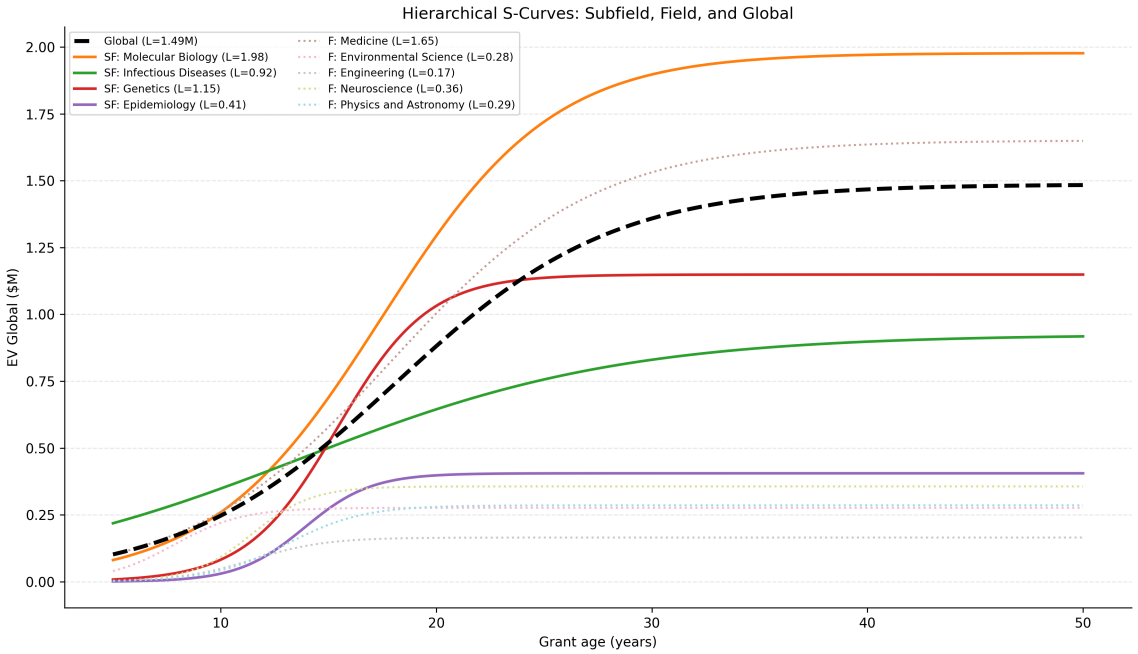


Figure 10: Hierarchical S-curves: 4 subfield (solid), 5 field (dotted), and global (black dashed).