THE COMMERCIAL VALUE OF (CUT) SCIENCE*

Dennis Verhoeven† Antoine Mathieu Collin‡ Reinhilde Veugelers§

July 2026

Latest version here

First version prepared for the NBER Summer Institute Science of Science Funding workshop, 2025. Very preliminary; please do not quote. Comments very welcome.

Abstract

The economic case for publicly funding science hinges on the commercial value its new knowledge creates through technological inventions. Yet we only have a bird's-eye view of which funded science creates this value, how much it is worth, to whom, and in which time frame. We measure the commercial value of science one scientific publication at a time, attributing the market value of patented inventions backward through a citation network of 521.3 million papers and patents, projecting each publication's spillover over thirty years, and attaching it to the grants that paid for the research. Per funding dollar, NSF grants return $20.18 of global technological spillover against $8.60 for the NIH; most of the value is created by indirect links, and more than half is realized on inventions made abroad. Applied to the 2,241 NIH and NSF grants terminated in 2025, every dollar of terminated grant funding forgoes $8.69 of spillover over thirty years, $5.61 in present value. The payoff arrives slowly: a research dollar does not break even for 9.8 years, beyond a typical legislative window. Among the countries that would have shared in the spillover, China loses the most: $2.46 of commercial value for every dollar of US grant funding cut, against $0.76 for the EU-27.

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

JEL codes: O31 · O33 · O38 · H50 · I23

*First version prepared for the NBER Summer Institute Science of Science Funding workshop, 2025. Very preliminary; please do not quote. Comments very welcome.

†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); Bruegel.

§KU Leuven, Department of Management, Strategy & Innovation (MSI); CEPR; Bruegel.

I. Introduction

Scientific research is a central input into technological progress, and the case for funding it publicly rests on the economic spillovers it creates for the rest of the economy. Part of these spillovers take the form of commercial value embodied in new inventions: the private value patent holders appropriate from the science they build on. In the aggregate that case is well established: because knowledge spills over, private actors appropriate only part of what their discoveries create and so underinvest (Nelson, 1959; Arrow, 1962; Griliches, 1979), and investments in science generate productivity growth and welfare gains that substantially exceed their private returns (Mansfield, 1991; Hall et al., 2010; Jones and Summers, 2022). What these aggregate rates of return do not show is how the value is created — which funded scientific discoveries become valuable inventions, over what horizon the spillovers unfold, and who captures the resulting benefits. This paper measures that value directly, one publication and one grant at a time, and applies the measure to the commercial value lost by the 2025 US science-funding cuts.

Under fiscal pressure and rising demands for accountability, governments are cutting science budgets and asking harder questions about what they buy. At the same time, policymakers face growing pressure to direct science toward the areas where the returns to society are highest — which presupposes knowing where and how those returns are realized. Underlying both is an erosion of public trust in science. This is illustrated most eminently by the United States — the largest scientific power in the world — which beginning in 2025 terminated thousands of federally funded research grants.[1] Together these pressures have strengthened calls for evidence on how much commercial value is lost when science funding is cut, and on when and where those losses materialize.

Rather than asking how much value the average research dollar creates, we ask how funded scientific knowledge travels through the innovation system, who turns science into patented invention, and how much commercial value is created along the way, and by whom. We define and measure commercial value, throughout, by the private market value of the inventions that build on the science.

[1] The terminations implemented a series of executive orders directing agencies to defund research related to diversity, equity, and inclusion, "gender ideology," and adjacent priorities (Executive Orders 14151, 14168, and 14173; 90 Fed. Reg. 8339, 8615, and 8633). Grants were flagged by screening their text for terms associated with the targeted topics — a mechanism the NIH Director acknowledged under oath before the Senate Appropriations Committee in June 2025, and that a federal district court found "arbitrary and capricious" in *American Public Health Association v. NIH* (D. Mass., June 2025); the Supreme Court stayed the part of that ruling ordering the terminated funds restored in August 2025. The parallel FY2026 proposals to cut the two agencies' budgets outright were largely rejected in the appropriations process. This paper prices the terminations, which happened, not the proposed budget cuts, which mostly did not.

To do so we use a bottom-up framework that reconstructs the knowledge network linking scientific publications to patented inventions through citations: paper-to-paper, patent-to-patent and patent-to-paper. The citation network contains 521.3 million nodes, made up of 452.1 million OpenAlex works and 69.2 million valued patent families, joined by 24.1 million patent-to-paper links from Reliance on Science (Marx and Fuegi, 2020, 2022) and 2.87 billion article-to-article links. Every patented invention is assigned a commercial value through a model that uses its patent characteristics to predict the stock-market valuation of patents granted to US listed firms (Kogan et al., 2017; Guillard et al., 2021). This approach allows extrapolating commercial values to the worldwide collection of patented inventions. A share σ of each document's value is credited to the prior art it cites and split across its references, so value propagates backward along citations and halves at each hop. Our recursive algorithm thus assigns a decayed portion of the commercial value of each downstream invention that links to it directly or indirectly. The sum of these values is an article's commercial value, i.e., its *Science Rank* (Martin et al., 2026). Any article's commercial value can be decomposed by the geography and timing of the patented invention, allowing us to trace where and with which lag science translates into commercial value. Finally, by linking each paper to the grants that funded it, we derive an approximation for the return on science-funding investment.

We start by documenting three facts about the commercial value of science. First, the value of science is extremely concentrated: the top percentile of articles holds 82.9% of the realized total. Second, a large portion of the value is assigned to science through *indirect* knowledge spillovers. Just 38.8% of the value arrives through a direct patent-to-paper citation, and the remaining 61.2% travels through chains of intervening papers or patents. Third, spillovers from science arrive slowly: half of an article's thirty-year commercial value is realized only after 18 years.

This slow arrival means that a record ending in 2022 captures only a fraction of what recently published science will earn: only 34% of an article's thirty-year spillover has arrived by age ten. We therefore project every article to age thirty with a gradient-boosted Tweedie model — a machine-learning regression suited to an outcome that is zero for most units and heavy-tailed among the rest — fitted only on what was knowable at each point in an article's life: its spillovers and citations so far, the patent-uptake trajectory of its field, its novelty (Uzzi et al., 2013; Wang et al., 2017), and its authorship. Over their first thirty years, we project the 189.0 million articles and reviews published since 1985 at $96.8 trillion. We validate the projection out of sample, refitting it as of past cutoffs and scoring it against what had materialized by 2022: it ranks science well but under-predicts levels, so we treat every projected level as a floor.

To turn value into a return we link 29.2 million grant-to-paper edges from Dimensions (a bibliometric and research-grants database) and OpenAlex, and divide a grant's projected thirty-year spillover by the money that produced it. Per funding dollar, NSF grants return $20.18 of global technological spillover against $8.60 for the NIH among grants that produced a linkable paper. These return rates do not count grants with no linked papers, assuming incomplete coverage rather than having zero output. When assuming these unlinked grants had actually zero output, we get return rates of $9.81 for NSF and $7.48 for NIH. The NSF–NIH gap reflects the cost of producing science rather than scientific merit at the paper-level: the NIH paper costs about 2.50 times more. Computer science and Engineering lead the return rate ranking for the NSF directorates; Biomedical Imaging and "General Medical" lead the NIH.

We price the 2,241 NIH and NSF grants terminated in 2025, drawn from Grant Witness. The cuts did not fall on weak science: benchmarked against contemporary survivors in the same field, the terminated investigators sit at the 47th percentile on prior productivity and the 40th on prior spillover. Because the terminated grants are too young to project directly — the median began in 2023 — we match each to similar mature donor grants within agency, field and funding band, and use their realized spillovers to stand in for what the terminated science would have produced. We find that every dollar of terminated grant funding forgoes $8.69 of global technological spillover over thirty years, $5.61 in present value at 3%. The fields cut most often and the fields worth most are not the same: the three patent-proximate divisions — biomedical, information and computing, and biological sciences — account for 70% of the forgone value. The loss is also global: well more than half of what the terminations put at risk would have accrued outside the United States. Innovations from inventors located in China, not the United States would have captured the largest share: 28.3% against the United States' 23.7%, with the EU about a tenth.

Our contribution is twofold. Methodologically, we provide a micro-level framework for measuring the commercial value of scientific knowledge by tracing the full chain: from scientific production, through diffusion, into patented invention, and on to the commercial value those inventions create. Most studies measure one link in that chain; the final step, from scientific knowledge to future commercial value, is largely missing from the empirical literature. Substantively, we quantify the economic cost of disrupting the production of science, using the 2025 US cuts as the application, and add evidence on the patterns through which scientific discoveries

create commercial value across the economy. The framework can also be used as a tool for the *ex ante* evaluation of science- policy: rather than assessing a grant's outcomes retrospectively, it can estimate the long-term cost of budget cuts, of shifts in research priorities, and of changes to funding-program design.

The next section places the paper in the literatures it speaks to, and Section III sets out our empirical methodology and the data we use. Section IV then asks what the commercial value of science is worth, how concentrated it is, how slowly it arrives, and whose invention leans on it. Section V turns to the science that public money paid for: what a funded dollar returns, when the return arrives, and how much of it lands abroad. Section VI prices the 2025 terminations against that background.

II. Contribution to the Literature

The paper speaks to several literatures: the aggregate social returns to research and development, the link from science to patented technology, the causal returns to public research funding, the geography of who captures scientific spillovers, and the estimation of the economic impact of the 2025 grant terminations themselves.

The aggregate social returns to R&D. A large literature estimates the return to research from production functions and growth accounting, and concludes that this return is high and exceeds the private return, because knowledge spills beyond the organization that pays for it (Griliches, 1958, 1979, 1992; Mansfield et al., 1977; Hall et al., 2010). Bloom et al. (2013) separate technology spillovers from product-market rivalry and put the marginal social return to R&D at 55% against a private return of 21%; Lucking et al. (2019) find the ratio has, if anything, risen since. Jones and Summers (2022) calculate that a dollar of R&D returns $13.3 in present value and argue it is difficult to construct an average return below $4. For *public* R&D specifically, Fieldhouse and Mertens (2024, 2025) exploit appropriations shocks and recover gross returns of 140–210% on US nondefense research, implying that government-funded R&D accounts for roughly a fifth of postwar business-sector productivity growth. Gazzani et al. (2026) find that government-funded but privately owned patents, though a small share of the total, explain close to a fifth of medium-term fluctuations in total factor productivity, with NIH- and NSF-funded innovations most strongly linked to subsequent productivity growth, and Dyèvre (2024) traces public-research spillovers into private-sector productivity growth.

This strand of literature uses macroeconomic methods and models to estimate the returns to science and technology, but doesn't shed light on which science produced the return (Hall et al., 2010). To fill this gap, we propose a microeconomic assessment which allows a decomposition by field, funder and region, and even at the grant level, by tracing each scientific paper's individual return.

From science to patented technology. That industrial invention rests on public science is long established, from the qualitative accounts of Rosenberg (1990) through the citation evidence of Narin et al. (1997), who find that 73% of the papers cited by US industry patents are public science, and the survey evidence of Cohen et al. (2002), who find that the channels running from public research into industrial R&D are papers and conferences rather than patents and licenses. Patent-to-paper citation data have since made the link measurable at scale (Marx and Fuegi, 2020, 2022; Bryan et al., 2020): Ahmadpoor and Jones (2017) show that 80% of cited articles link forward to a future patent and that the typical linked pair stands 2 to 4 citation degrees apart. Li et al. (2017) find that only about 8.4% of NIH grants yield a patent that directly acknowledges them, while 31% yield a patent that cites a paper acknowledging them. Fleming et al. (2019) document that up to 30% of US patents rest on federally funded research. A subset of this work attaches *value* to the link, and how it measures that value matters for what the estimates mean: Poege et al. (2019) proxy a patent's value largely by the forward citations it receives, checking the result against market-based measures, while Krieger et al. (2024) and Arora et al. (2021) price patents directly from stock-market valuations in the manner we adopt here. Poege et al. (2019) find that patents building on high-quality science are roughly twice as valuable as those building on low-quality science; Krieger et al. (2024) interact patent market values with citation distance and find patents one hop from science 26% more valuable; Arora et al. (2021) estimate an 11% value premium on science-linked patents. Masclans et al. (2025) take the further step of predicting a paper's commercial potential *ex ante* from its text, and Munari et al. (2022) trace European Research Council grants into the patents that cite them. In all of it the dollars stay attached to the *patent*, graded by its distance from science. We instead conserve those dollars and propagate them backward, so that every article receives an additive, region- and year-resolved share of the downstream invention value it enabled — a share that can then be attached to a grant. Our decomposition allow to estimate the share of the value that is attributable to direct science-to-technology links versus indirect links.

Causal returns to public research funding. A quasi-experimental literature, surveyed by Azoulay and Li (2020), identifies the effect of public funding on scientific and inventive output. Jacob and Lefgren (2011b,a) exploit NIH funding rules; Azoulay et al. (2019) find that $10 million of additional NIH funding yields 2.7 additional private-sector patents, and back out roughly $2.3 of downstream pharmaceutical sales per public dollar; Li et al. (2017) trace NIH grants into patents; Myers and Lanahan (2022) show that for every patent a Department of Energy grantee produces, about three more are produced by others; Bergeaud et al. (2025) find €0.81 of induced private R&D per public euro in French public-lab clusters; Moretti et al. (2025) document crowd-in and international spillovers from defense R&D; and Babina et al. (2023) study university funding shocks directly. Not all of it is optimistic — Kantor and Whalley (2025) find no persistent productivity spillovers from Apollo-era contract R&D — and the dispersion is itself informative. These designs recover the value of the marginal public dollar in specific agencies, fields and empirical settings. Our object is the *average* funded dollar, for science selected by topic rather than by a funder's value ranking, spanning both agencies and every field. The two turn out to be of the same order, and Section V.A sets them side by side. That the average and the marginal should be close here is not a coincidence: within the fundable range, peer-review scores barely predict what a grant will produce (Fang et al., 2016; Pier et al., 2018; Li, 2017), and reviewers penalize exactly the risky proposals whose returns are largest (Azoulay and Greenblatt, 2025; Veugelers et al., 2025). A funder that cannot rank its winners cannot select far up the value distribution. So the marginal dollar earns close to the portfolio average — which is why the terminated grants, selected on topic, fairly price the typical funded dollar.

Who captures the value. Knowledge spillovers are geographically localized (Jaffe et al., 1993; Audretsch and Feldman, 1996; Belenzon and Schankerman, 2013), though the localization is weaker than first estimated (Thompson and Fox-Kean, 2005) and is mediated by people and networks rather than by proximity as such (Agrawal et al., 2006; Kerr, 2008; Singh, 2005). They also cross borders on a large scale: Coe and Helpman (1995) find that about a quarter of the benefit of G7 research accrues to trading partners, Eaton and Kortum (1999) that research performed abroad is roughly two-thirds as potent as domestic research, and Keller (2004) that for most countries foreign sources account for the overwhelming majority of domestic productivity growth, with spillover intensity halving at roughly 1,200 kilometers (Keller, 2002). On public research specifically, Myers and Lanahan (2022) find that about 40% of the spillovers from

Department of Energy grants land outside the United States, and Moretti et al. (2025) document that one country's public R&D raises private R&D abroad. We find far more international spillovers than the aggregate literature, possibly because our approach credits indirect spillovers to a larger extent.

The 2025 terminations. A fast-growing literature intends to estimate the economic impact of the the discontinuation of science funding, notably in relation to the recent science policy changes in the United States. citetliu2025characterization characterize the terminated NIH grants and the trials they disrupt; Oliveira et al. (2026) document their demographic incidence on the same Grant Witness universe we use; Azoulay et al. (2025) ask what a historically 40%-smaller NIH would have cost in drug approvals; Esposito (2026) trace citation paths from NSF-funded papers to patents; and Azoulay et al. (2026) studies the response of young scientists. A further set prices the withheld budget itself — Congressional Budget Office (2025), González et al. (2025) and United for Medical Research (2025) in drug counts or through economic-activity multipliers. Cutler and Glaeser (2025) values a proposed NIH cut through lost life expectancy, while Jalali and Hasgul (2025) and Hasgul et al. (2025) model the resulting trade-offs and the lasting "aftershock" on a shrunken research pipeline, and the scientific press chronicles the ongoing upheaval (Kozlov et al., 2026; Garisto and Lenharo, 2026; Garisto, 2026). The literature on negative funding shocks more broadly establishes that funding disruptions reduce scientific output and distort its direction (Aghion et al., 2008; Freeman and Van Reenen, 2009; Stephan, 2012). Examples include he displacement of scientific talent (Moser et al., 2014; Borjas and Doran, 2012; Doran et al., 2022; Ganguli, 2017), policy uncertainty (Baker et al., 2016), and the chilling of US–China scientific ties (Jia et al., 2024). These studies investigate the economic effects of past science funding disruptions; and for the ones focusing on the current science funding cuts in the US, they are high-level estimates. Our approach looks at the disrupted science at the grant level, which allow us to investigate where science is being cut and what is the spillover impact on commercial inventions of this cuts at a more granular level.

III. Measurement framework and data

In this paper, first we develop a microeconomic framework to measure the commercial value of science at the level of individual publications: we trace the market value of patented inventions back to science through the citation network. Second, we aggregate this measure at the grant level to estimate the commercial value derived from publicly funded science for two agencies, the National Science Foundation (NSF) and the National Institutes of Health (NIH). Third, we

apply this methodology to quantify the commercial value lost due to the 2025 cuts to NSF and NIH by estimating with a counterfactual analysis of the commercial value of the papers that have not been published and the patents that haven't been filed and granted due to lost science funding.

To value science we need a dollar figure for the spillover a scientific article creates for technologies, which represents the part of the private value of downstream patented inventions that the article plausibly enabled. To that end, we define three metrics. The commercial value of an invention (CV_p) represent the private value attached to a patented invention captured by its owner, and we estimate this value from the characteristics of the patent. (Section A). The commercial value of science (CV_s) is the part of that value a scientific article article enabled, and it should be understood as a spillover rather than as a return rate. We compute it by propagating CV_p backward along citations, following the *Science Rank* measure of Martin et al. (2026) (Section B). The commercial value of funded science (CV_{sf}) restricts CV_s to the articles linked to a grant and divides it by the grant's funding, yielding a return per dollar of funding (Section C). We now discuss each measure and its data sources in detail.

Importantly, we model the fact that science spillover take time to materialise, and estimate the commercial value of scientific article over a third-years windows. This is crucial not only to correctly model the value created by articles which are too young to have diffused entirely through science and technology, but also to estimate the impact over time of the 2025 science funding cuts. Then, the commercial value of cut science (CVsc) is the commercial value that will not materialize because of the funding cuts.

A. The commercial value of an invention (CV_p)

Only a patented invention earns a private return on the knowledge it embodies. A scientific article does not create commercial value per itself: its commercial value reaches it as a spillover from the inventions built on it. The unit of observation for CV_p is the patent family, the set of filings protecting one invention. Patented invention is also the only channel from science to commercial value that we can measure: value arriving through regulatory guidance, clinical practice, or the training of people goes uncounted. Quantifying these channels is an important avenue for other work.

Following Guillard et al. (2021), we predict an invention's value from four characteristics: filing year, technology classes, patent family size, and number of claims. The calibration comes from Kogan et al. (2017), who infer a patent's private value from the patenting firm's abnormal stock return in the three days around the news of the patent grant, scaled by the firm's market

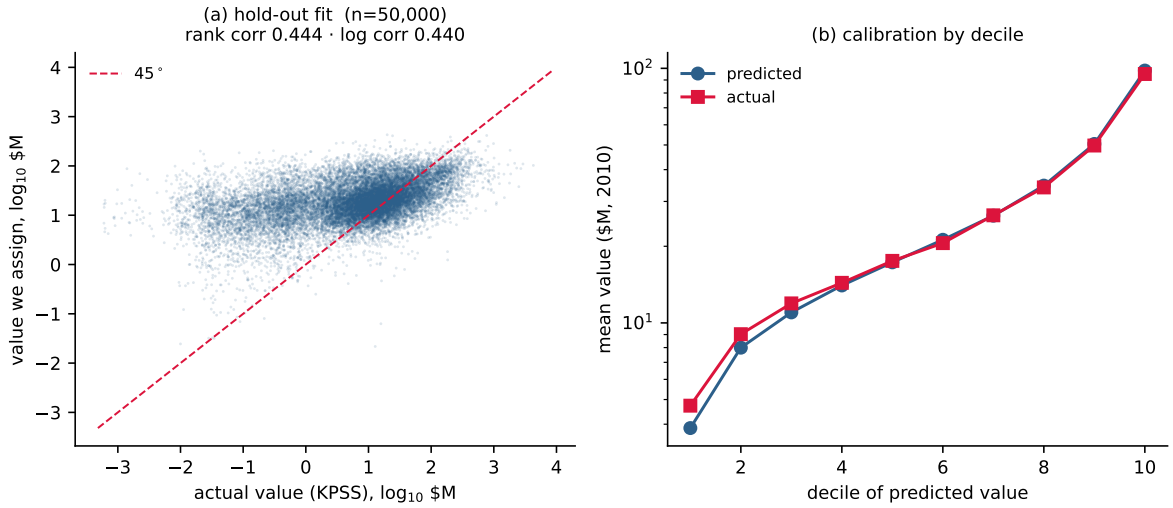
capitalization, across 3.42 million US patents in constant 2010 dollars. This market value is observed only for patents held by firms listed on a US exchange (2.9% of the world's patent families) so it serves as a training sample rather than as the measure itself. We group families into cells defined by the four characteristics and assign each family the mean market value of its cell, coarsening cells until each holds at least ten observed values.[2] Every family receives the fitted value, including those whose market reaction we observe. This estimation procedure reduces the noisiness of Kogan values: a three-day abnormal return carries the value of the patent plus everything else the market learned that week, but projecting onto patent characteristics discards the part that has nothing to do with the patent.

We hold 50,000 valued families out of the fit and compare the value we assign with the value the market paid (Figure 1). The rank correlation is 0.444, deciles of assigned value are monotone in realized value, and the assignment is unbiased in the mean: $28.34 million realized against $28.49 million assigned. A tighter family-level fit is not to be expected: a grant-window return measures average value accurately for a group of patents but noisily for any individual one, which is why Guillard et al. (2021) require ten observations per cell. This accuracy is sufficient for our purposes, because we sum and average CV_p over large sets of patents and never interpret a single family's value.[3]

[2]When a cell holds fewer than ten observed values, we coarsen the filing year first, then the technology class, and finally drop the claim count and the class rather than leave a family unvalued; a family in several classes takes the mean across them. Two kinds of family are zero by construction. Utility models (which is a lesser, largely unexamined patent right, 23.3 million of which fall in our window, the overwhelming majority Chinese) receive zero because the calibration sample contains none. Applications never granted receive their fitted value multiplied by the year-specific probability that a pending application is eventually granted, set to zero for filing years before 2004.

[3]The fitted values are more centrally distributed than the market values they are trained on: their standard deviation is 0.35 times that of the observed values. Further validation is required to establish whether this reflects our measure not picking up all of the heterogeneity in patent value, or the Kogan et al. (2017) method inflating the variance with news unrelated to the patent.

Figure 1: What the commercial value of an invention is fitted on, and how well it fits



Notes: Panel (a) plots the value we assign against the market value Kogan et al. (2017) observe, for 50,000 patent families held out of the fit. Panel (b) bins the held-out families by decile of assigned value and plots the mean realized value in each: the fit orders deciles correctly. Read the two together: the assignment is unbiased in the mean ($28.34M realized against $28.49M assigned) and correctly ordered, but it is a cell mean, so it compresses the tail and sits above the median of its cell. It is a summary of a group of patents, not an estimate of one patent's worth. *Source*: authors' calculations; market values from Kogan et al. (2017).

We locate each family's value geographically on two bases: the *inventor* basis, which records where the technology is developed, and the *assignee* basis, which records where its value is owned. On each basis, we split the family's value equally across its distinct countries and map the countries to seven recipient regions: the United States, the EU-27, the United Kingdom, China, Japan, South Korea, and the rest of the world. Families with no resolvable country form an explicit *unattributed* category rather than being added to the rest of the world. These families are disproportionately low-value: they make up 37.6% of families but only 23.4% of value.

B. The commercial value of science (CV_s)

We build a science and technology citation network comprising both academic publications and patents. The scientific side comes from OpenAlex (Priem et al., 2022): 452.1 million works linked by 2.87 billion article-to-article citations. We use 189.0 million articles and reviews published between and 1985–2022, however we allow value to travel through out-of-sample documents as long as there are citation links. The technology side comes from PATSTAT, the European Patent Office's worldwide patent statistical database (autumn 2025 edition): 69.2 million families with filing years 1985–2022, linked by 220.2 million family-to-family citations. Reliance on Science (Marx and Fuegi, 2020) joins the two sides by matching the references on patents' front pages to bibliographic records. We keep links scoring at least four on its

confidence scale, the operating point Marx and Fuegi (2020) recommend (precision above 99% at roughly 88% recall), giving 24.1 million patent-to-paper links. The full network contains 521.3 million nodes and 3.12 billion edges. Papers are 86.7% of the nodes but account for 92% of the edges: the scientific side of the network is the denser one, a paper citing nearly twice as many documents as a patent family does.

A crucial design choice in our methodology is how to attribute the commercial value of patented inventions back to the documents it builds on, and how to distribute it across different cited documents. The attribution rule follows Martin et al. (2026), who derive it from a production function for the value of a patented invention with two inputs: the inventor's own R&D, and the stock of public knowledge the invention draws on, approximated by the patents and papers it cites. The spillover a document creates, that is, the share of the invention's value that the document accounts for is the derivative of this function with respect to the stock. Because the papers in the stock are themselves produced this way, solving the recursion writes the spillover of every node in the network as a function of the structure of the network, the private value of each node lying downstream of it, and the elasticity of an invention's value with respect to its knowledge stock. We take the knowledge stock of an invention or article j to be the n_j documents it cites, weighted equally, so that each is credited a share σ/n_j of j's value. We set σ, the elasticity of value with respect to that stock, to 0.5, so that value halves at every citation hop. The calibration follows Aghion et al. (2016), whose estimates for the car industry put the elasticities of innovation value with respect to own R&D and with respect to external knowledge at similar magnitudes, implying equal weight. The evidence is from one industry and we apply it to all of them, so σ is a calibration and not an estimate. This design choice affects magnitudes but not the relative values: across a grid of σ the dollar levels we report, the rates of return included, move by more than an order of magnitude, while the ranking the measure induces across articles is close to invariant (Guillard et al., 2021).

We attribute values to the nodes of the network recursively. We collect the patent values in a vector CV_p and define the backward-citation matrix M, with $M_{ij} = 1/n_j$ when j cites i and zero otherwise. The commercial value of science — an article's *Science Rank* — is the decayed, citation-weighted sum of all private value downstream of it:

$$CV_s \;=\; \sum_{k\geq 1} (\sigma M)^k \, CV_p, \tag{1}$$

where k indexes citation hops: at $k = 1$ a patent cites the article, at $k > 1$ a chain of citations leads back to it. Because our sample is limited to article and reviews between 1985 and 2022, value leaks at every hop, because n_j counts every reference j makes, including the 4.3% that fall outside our network.

Crediting each of j's references the same σ/n_j says that a document owes as much to any one of them as to another. We admit this is a strong assumption. One way of relaxing this assumption is to use textual similarity measures to weight citations, an extention we plan to incorporate in future work.

The $k > 1$ terms separate *Science Rank* from citation-distance measures (Ahmadpoor and Jones, 2017; Poege et al., 2019; Krieger et al., 2024): basic science that patents rarely cite directly, but that applied science builds on, still accumulates commercial value. Benchmarked against scientific prizes, *Science Rank* ranks such upstream science far better than direct patent-to-paper citation counts do (Martin et al., 2026).

We build on Martin et al. (2026) closely, but extend it in four ways. First, we rebuild the network on a February 2026 OpenAlex snapshot, keeping every document type — 452.1 million works and 2.87 billion article-to-article edges, against the October 2024, article-only network of roughly 159 million works in the original and push the citing window from 2018 to 2022, which matters because the terminated grants are young. Second, where they attribute one global dollar figure to each article, we resolve each article's spillover by recipient region, priority year and basis, the partition that makes the incidence of a cut, and its time profile, observable. Third, where they report cumulative realized spillover through the end of their citing window and flag the resulting truncation as a limitation, we project every article to age thirty and validate the projection out of sample, without which grants awarded in the last five years could not be priced. Last, we supply a denominator: Martin et al. (2026) attribute value to science but never link that science to the money that paid for it, so no rate of return follows from their measure. The NSF-versus-NIH comparison, the return on a funded dollar, and the setting of foregone spillover against budget withheld all depend on that link. Appendix A describes *Science Rank* in more detail and compares it to its closest antecedents.

C. The commercial value of funded science (CV_{sf})

We complement our measure of the commercial value of science (CVf) by comparing it with the public funding received. Constructing CV_{sf} takes two inputs: what each grant paid, and which articles it paid for. The amounts come from Dimensions, which records the lifetime award of a US federal grant rather than its annual obligation. We take that lifetime award as the grant's

cost and deflate it to constant 2010 dollars, matching the patent-value calibration. The articles come from two sources: an article names the grants that supported it in its acknowledgements, which Dimensions indexes, and funders report their own outputs, which OpenAlex records. The union of the two gives 29.2 million grant-to-article edges. A grant's return is then the thirty-year commercial value of the science it produced, divided by its lifetime award. Where an article names several grants, we split its value equally across them.

We collect additional data from Dimensions and OpenAlex, including their age, their Principal Investigators (PIs) and field classifications. Of the 189.0 million articles and reviews published between 1985 and 2022, 6.5% name a grant we can trace, and those articles account for 35.3% of the commercial value we attribute to science over a thirty-year horizon (Section D). Coverage is far from complete, and its gaps are not random. Grant acknowledgements are barely indexed before 2000 and outside the funders whose awards are systematically recorded, and the share of grants with a traceable article climbs from 36% for the 1990 cohort to 75% for 2010. Appendix Table A7 and Appendix Figure A6 show what this coverage does to a naive comparison of grant-linked with unlinked science: the gap between the two tracks the spread of acknowledgement indexing rather than any change in funding, so we do not read it as a return to funding.

Because a grant with no traceable article is not necessarily a grant that produced nothing, we report every rate two ways: treating unlinked grants as zeros and dropping them. The two can differ by at most 16.3%.[4]

A second issue is what counts as a grant. Our panel covers every funder Dimensions records, and outside the NIH and the NSF it is dominated by award vehicles our acknowledgement-based linkage cannot see: Department of Energy laboratory contracts, the European Commission, national research ministries, and private foundations, which are large, few, and rarely tied to their outputs by an acknowledgement. An aggregate over all of them measures that composition rather than the return to a research dollar: it more than doubles depending on how the largest awards are treated, while within either of our two agencies it moves by a quarter to a third, because there the large awards are linked about as well as the small ones. We therefore do not report a single global rate of return. We report rates within funders, and we restrict granters in this paper to two research agencies most affected by the 2025 US Science funding cuts, the NIH's and the NSF's.

[4]Unlinked grants make up 26.9% of the mature donor pool by count but only 14.0% by dollars, since linkage rises with award size; scoring them as zeros biases rates downward (treating unobserved spillover as no output), while dropping them biases rates upward (if funded grants are also more likely to be acknowledged), so we report both and let the reader see the bound.

D. Projecting over thirty years

So far our measures capture the commercial value an article has generated *through 2022*, which makes a young article and an old one incomparable: the science in a 1990 article has had three decades for inventors to build on, that in a 2018 article almost none, and half of an article's thirty-year spillover arrives only after 18 years. To compare them, we put every article on a common thirty-year clock and predict, for each one, how much commercial value it will have generated by the time its science is thirty years old. We fit that prediction on the articles old enough to be observed for a full thirty years, reading off how their commercial value accumulates with age and letting the profile depend on what could be seen early in an article's life: the spillovers and citations it had drawn so far, how quickly its field turns into patents, how novel it was (Uzzi et al., 2013; Wang et al., 2017), and who wrote it. An article younger than thirty is projected forward along the fitted profile; one already past thirty keeps its age-thirty value, and whatever it earned later is set aside. To check the exercise, we run it in the past: we return the model to a series of earlier years, give it only what would have been known at each date, have it forecast each article's spillover, and compare the forecast with the value the article had in fact generated by 2022.

The test tells us three things. The model is very good at telling apart the articles that will eventually have commercial value from those that never will: shown one of each, it ranks them in the right order more than nine times in ten (0.91–0.95). It is reasonably good at judging how large that value will be, ordering articles by their eventual value loosely rather than closely (a rank correlation of 0.71–0.74). Its levels are biased low: run in the past, it falls 3–35% short of the value that actually materialized, and short by more for the youngest articles — the vintage the terminated grants occupy. We therefore read every projected total in this paper as a conservative floor, and we use the model only for sums over many grants, never to value a single article.

Projected this way, the 189.0 million articles and reviews published between 1985 and 2022 generate $96.8 trillion of commercial value over their first thirty years, against the $39.4 trillion the same articles had realized by 2022. Most of the commercial value this body of research will create has yet to arrive.

Appendix B sets out the prediction panel, the model, and the out-of-sample validation in full.

IV. Results on the commercial value of science

The commercial value of science is large. For the 189.0 million research articles and reviews published between 1985 and 2022, those articles had realized $39.4 trillion of commercial value by 2022, and we project them at $96.8 trillion over their first thirty years, which means that what this body of science will be worth to invention has still to arrive.

These magnitudes pass a basic sanity check. Cumulative world output over the same 1985–2022 window was about $2,008 trillion in constant 2010 dollars, so the $39.4 trillion science had realized is about 1.96% of it, which is close to one dollar for every dollar of the roughly $41.9 trillion the world spent on research and development over the period.[5] The level itself is an order of magnitude: it scales with the citation-decay parameter σ, running from $4.7 trillion to $299 trillion across the range we consider (Section B).

Two limits apply to every figure that follows. We price only the commercial value that runs through patented invention, leaving out the non-patented and uncommercialized uses of science, and we price only its private, appropriable part read via the stock-market value of the patenting firms, a producer's return that excludes consumer surplus and other benefits such as the health, longevity, clinical and human-capital benefits of research. Except for the thirty-year projection above, every number in this section is commercial value *realized through 2022*, read off the citation network as it stands rather than projected forward. Three properties of that value organize the rest of the section: it is extremely concentrated in a few articles, it arrives only after decades, and it flows back to the invention of some countries far more than to others'.

A. The commercial value of science is very concentrated

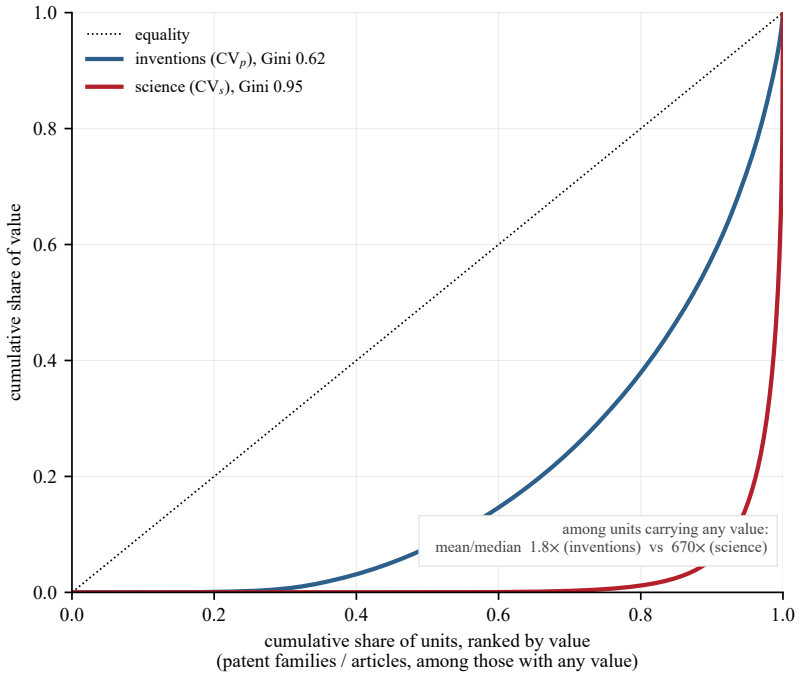
Commercial value is extremely concentrated in a few articles. Measured on the value realized through 2022, only 22.2% of articles lead to any commercial value at all, and the top percentile of articles holds 82.9% of the total.

This concentration makes identifying high-value science valuable, and difficult. With almost all of the return in the top percentile, the gain from directing funding toward it is large, but so is the precision the task demands: a selection that lands anywhere below the top percentile reaches only a small fraction of the value.

[5] World output is World Bank series NY.GDP.MKTP.CD (current US$, world aggregate) deflated year by year to constant 2010 dollars with the same deflator we apply to grant funding throughout; cumulative world R&D applies the World Bank R&D-intensity series to that GDP. Each total is cumulated over the same 1985–2022 window as the realized spillover: comparing any of these stocks to a *single* year's world output would confuse a stock with a flow.

Figure 2 plots this concentration as a Lorenz curve, and sets it against the same curve for the private value of the patented inventions themselves (CV_p). That patent values are highly skewed is well established (Hall et al., 2010; Kogan et al., 2017): the invention curve already bows well below the diagonal, its Gini reaching 0.62 among families with any value. The commercial value of *science* is more concentrated still, at a Gini of 0.95. The natural worry is that science simply inherits the skew of the patent values it is built from. It does not: assigning every patent an identical value, which strips the market valuation out entirely and leaves a pure citation-network statistic, leaves the science-side Gini essentially unchanged at 0.95. The concentration is generated by the structure of the citation network and not by the Kogan et al. (2017) calibration.

Figure 2: The concentration of commercial value across articles



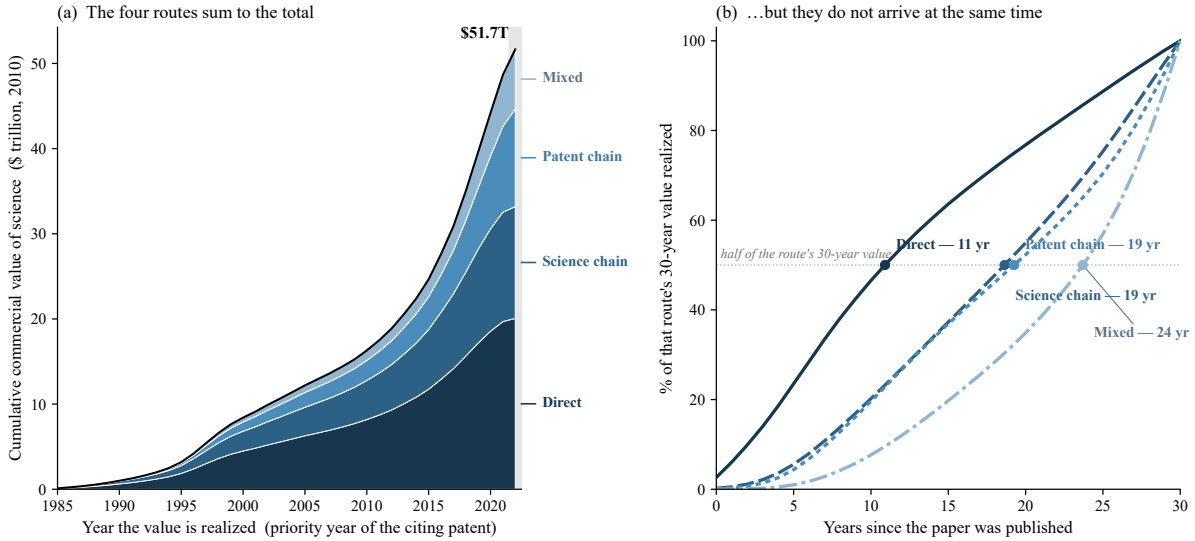
Notes: Lorenz curves for the private value of patented inventions (CV_p) and the commercial value of science (CV_s), among the units of each with any positive value. The commercial value of science is the more concentrated of the two (Gini 0.95 against 0.62). Both curves are drawn over positive-value units only, so neither speaks to the share of units with no value at all: on the invention side that share is set by our valuation rule — utility models and never-granted applications are zero by construction (Section A).

B. How value reaches science, and when

Because our measure traces each dollar of invention value back along the citation path it actually travelled, we can see not only how much commercial value science creates but how it gets there, and how long it takes. The value can reach an article by four routes: a patent may cite the article *directly*; it may reach it through a chain of intervening *papers*; it may reach it through a chain of intervening *patents*, the last of which cites the article; or it may travel through *both*. Figure 3 shows how the commercial value of science divides among them — first by the route the value took (panel a), then by how long each route took to arrive (panel b).

Most of the value reaches science indirectly. Panel (a) plots the realized value cumulated year by year, split by route: only 38.8% of it arrives through a direct patent-to-paper citation, while 25.4% travels through chains of papers, 22.2% through chains of patents, and 13.5% through both. So 61.2% of the commercial value of science is invisible to any measure that reads only a patent's direct citations to the literature. Li et al. (2017) document the same pattern at the level of grants: only about 8.4% of NIH grants are linked to a patent that cites the grant directly, against 31% linked through a patent that cites a *paper* the grant funded.

Figure 3: The four routes from an invention's value back to the science it rests on



Notes: Panel (a) plots calendar time: the *cumulative* commercial value realized through each year, split by the route the value travelled from the citing patent back to the article, and summing to the $51.7 trillion that reaches scientific articles. The four routes are exact sub-sums of the same attribution, not separate estimates, so they add to the total by construction. The final year is shaded because the patent record for priority year 2022 is still filling in. Panel (b) plots the article's own clock, for the 1985–1992 articles old enough to be observed for a full thirty years; the marked points are each route's half-life. Direct value is half-collected within 10.9 years while every indirect route takes about twice as long, so the lag reflects not only how long the science takes to be built on, but how long the resulting inventions take to be built on in turn. *Source*: authors' calculations on the 1985–2022 citation network; patent values from Kogan et al. (2017).

The indirect routes are also the slow ones, which lengthens the lag from science to commercial value. Panel (b) follows each route on the article's own clock. Direct value is half-collected within 10.9 years of publication, but every indirect route takes about twice as long — 18.6 years through the literature, 19.3 through the patent record, and 23.7 through both. The patent chain is no faster than the science chain, so the delay is not simply that science is slow to be noticed: it takes time for inventors to build on the science, and more time again for later inventors to build on those inventions in turn.

The two results compound whenever a research dollar is judged on a short horizon. Within five years of an article's publication, 75.3% of the commercial value it has generated has arrived through the direct route; over a full thirty years the direct share falls to 35.9%.[6] Said differently, a short evaluation window therefore not only understates the return to science but also distorts its composition, giving more weight to the most direct citation routes and less weight to indirect routes.

C. Whose invention leans on science

Some countries' inventions draw on science far more heavily than others'. We can see this because the measure resolves each dollar of invention value both ways: by the region whose inventors hold the patent, and by the region whose inventors hold the patents that a piece of science ultimately feeds. Comparing the two shares for a region measures how science-intensive its invention is. US inventors hold 18.9% of the private value of the world's patented inventions but account for 32.1% of the value that flows back to science: American technology leans on science about 1.70 times as heavily as the world's does.

The pattern sorts countries into two groups. On one side are the science-based economies of the West: alongside the United States, the United Kingdom's inventions lean on science just as heavily, and the EU-27's clearly more than average (10.9% of invention value against 13.5% of the science value, a ratio of 1.25). On the other are the scale-manufacturing economies of East Asia: Japan and South Korea each account for a share of the value reaching science that is only about half their share of invention value, and China (which is the largest single holder of invention value on this basis, at 29.7%) draws 26.9% of the science value, a ratio of 0.91, slightly below parity.

Again, this ordering is a property of the citation network, not of the patent valuation. Setting every patent's private value to one, so the measure becomes a pure citation statistic, reproduces 82% of the United-States-to-China regional tilt (Appendix Figure A1).

[6] Both shares are computed on the 1985–1992 articles observed for a full thirty years and capped at that horizon.

V. Results on the commercial value of funded science

Section III attached each grant's spillover to the money that bought it. We now read three things off that link: what a public research dollar returns, how much of the return is realized abroad, and when it arrives.

A. The return per funding dollar

A public research dollar returns a large multiple of itself in global technological spillover. Table 1 reports the rate at the NIH and the NSF, on the thirty-year *projected* basis of Section D: each agency's projected thirty-year spillover from the articles its grants funded, divided by the money that funded them, in constant 2010 dollars. Among grants that produced a traceable paper, a dollar of funding returns \$8.60 of global technological spillover at the NIH and \$20.18 at the NSF. A multiple remains even when the spillover is counted only where it is realized on inventions by US inventors: \$2.45 and \$4.53 respectively, under a third of the global figure. Most of what a US research dollar buys is therefore realized abroad. This is the well-known appropriability problem of Nelson (1959) and Arrow (1962) operating across borders: scientific knowledge diffuses internationally along the same citation chains that spread it across fields, so a large share of what any one country funds accrues to inventors elsewhere (Coe and Helpman, 1995; Eaton and Kortum, 1999; Keller, 2004).

The two agencies differ: the NSF returns about 2.35 times as much per dollar as the NIH. Part of that gap may be due to coverage. Although there are linked-grant rate conditions on a grant having produced a traceable paper, the NSF's grants link to papers appear less complete than the NIH's. We control for that by counting every grant instead with unlinked grants entered as true zeros; and that brings the two agencies closer (\$9.81 against \$7.48), yet the NSF displays higher returns on either basis.

The natural explanation for the rest of the gap is field composition: the NSF funds computer science and engineering, which are closer to the patent frontier. But reweighting the NSF's field-by-field returns to the NIH's field funding mix leaves the gap essentially unchanged, at 2.33 times.[7] This points to the idea that the gap is instead a matter of funding amounts. Splitting each agency's return into the value of a resulting article and the funding it takes to produce one,

[7] Fields are the two-digit fields-of-research divisions, a coarse control. On the 2006–2013 grant cohorts, where the two agencies' grant-to-paper linkage coverage is comparable and the comparison is cleanest, the same reweighting *widens* the gap to 3.67 times — so field mix is not the explanation on either basis.

the two agencies' articles have comparable projected value, but the NIH pays about 2.50 times as much grant money per resulting article as the NSF ($0.37 million against $0.15 million).[8] In short, the NSF returns more per dollar because it produces publishable, patent-relevant science more cheaply.

Within each agency, the return tracks proximity to the patent frontier. Computer and information science ($47.6 per dollar) and engineering lead the NSF's directorates, and the method-oriented institutes (e.g. biomedical imaging and bioengineering ($26.9)) lead the NIH's disease-oriented ones. Fields far from the patent frontier are less visible to our measure by values patented inventions (Section VII).

Table 1: Return per funding dollar, by funder

	Grants	Funding	Rate ($ / $)	
Funder	(N)	($B)	Global	US only
NSF — all directorates	619,906	282	20.2	4.5
NIH — all institutes	534,718	1,037	8.6	2.5
Highest-return NSF directorates				
Computer & Info. Sci.	69,193	30	47.6	12.5
Engineering	70,841	26	41.9	9.1
Math. & Physical Sci.	110,944	62	22.6	4.3
Technology & Innov.	20,612	6	19.7	5.0
Highest-return NIH institutes				
Biomedical Imaging (NIBIB)	6,835	8	26.9	6.5
General Medical (NIGMS)	45,726	94	14.0	3.9
Cancer (NCI)	73,567	175	11.4	3.2
Arthritis & Musculoskeletal (NIAMS)	12,698	19	10.6	2.9

Notes: Rate = projected thirty-year *Science Rank* spillover of a funder's linked articles, divided by the lifetime award value of the linked grants, in constant 2010 USD. "Global" counts spillover realized on patents anywhere; "US only" counts spillover realized on patents with US inventors. Grants and funding cover the full linked history of each agency. The rate shown is the return among grants that produced a traceable paper; counting every grant, with unlinked grants entered as true zeros, gives a conservative floor of $9.81 (NSF) and $7.48 (NIH). Because incomplete linkage only omits spillover, that floor is a genuine lower bound. NIH and NSF linked grants together account for 10.9% of all projected thirty-year spillover. Directorates and institutes are ranked within agency; no cross-agency total is reported, because the award vehicles are not comparable across the full funder population (see text).

[8] Per credit-weighted paper, splitting each paper's credit equally across the grants that funded it. On the full linked history the NSF-to-NIH value-per-paper ratio is 0.94 and the NIH pays 2.50 times as much per paper; on the comparable 2006–2013 window the value ratio is 1.00 and the price ratio 3.29.

These returns can be set against the existing estimates of the aggregate return to research, but only once the estimates are harmonised. Jones and Summers (2022) express the return as a present value per research dollar, as we do: they put it at $13.3, and argue that an average below $4 is hard to construct. Our present value at a 3% discount rate is $5.61 for the terminated grants; that figure, and the undiscounted $7.48–$9.81 we find for the average funded dollar, both clear their $4 floor and fall below their $13.3 headline. One reason we fall below the headline may be that we count less of the same thing: their estimate recovers the whole social surplus of research, including the value that reaches society and consumers and firms that never patent, while we recover only what patent-holders appropriate, and only the part of that a citation chain reveals.

Bloom et al. (2013) and Fieldhouse and Mertens (2024, 2025) report an *annual* rate of return on a stock of R&D capital rather than a present value per dollar, so a conversion is needed to compare them. Discounting a perpetual annual return at 3% and depreciating the R&D stock at the rate each study adopts,[9] Bloom, Schankerman and Van Reenen's 55% marginal social return becomes about $3.06 per dollar, which our all-grant rates exceed, and Fieldhouse and Mertens's 140–210% gross return becomes $7.37–$11.05, which brackets our $7.48–$9.81. The closest quasi-experimental benchmark, Azoulay et al. (2019), is smaller in magnitude (about $2.3 of downstream drug sales per NIH dollar) but measures a different object: the effect of funding at the *margin*, identified off NIH funding rules, for biopharmaceutical firms alone, in drug sales. The authors themselves stress that valuing those patents is difficult and report a range rather than a rate of return, so we read it as a related magnitude.

B. When the return arrives

Different public investments repay on different schedules, and governments are held to account over short and recurring windows, which biases funding toward projects that pay back inside a political horizon (Nordhaus, 1975; Rogoff, 1990). Science is at the far end of that trade-off, so how long its commercial value takes to arrive is relevant for policy-makers.

We link each grants in 5-years windows to the articles it funded (Section C with 29.2 million grant-to-article edges), and we track each cohort's cumulative realized spillover against the money that bought it. Figure 4 plots one line per grant-start cohort and displays the cumulative realized spillover as a percentage of that cohort's funding, so that break-even is the 100% line. Appendix Table A5 collects the statistics.

[9] We convert an annual rate r to a present value $r/(\rho + \delta)$, taking the discount rate ρ at 3% and the R&D depreciation rate δ from each study. The converted figure is sensitive to δ, which neither we nor the literature pin down, so we report it as ours rather than theirs.

An NIH or NSF research dollar breaks even after 9.8 years, and after 11.3 years discounted at 3%.[10] The grants we have observed for a full thirty years have realized $5.57 of commercial value per dollar of funding.

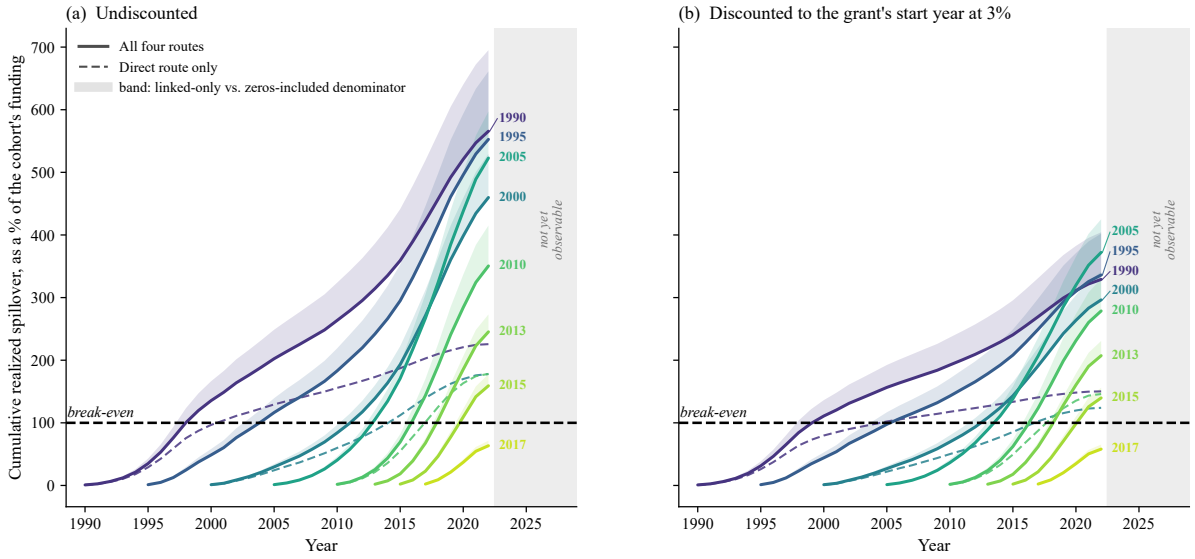
Only 2.4% of a grant's thirty-year commercial value is realized inside a three-year budget window, and 7.4% inside five years; the half-life of a grant's spillover, clocked from its start, is 20.5 years. A government books the savings from a cut immediately and in full, but inside the window on which it is judged it sees about one fortieth of what it gave up.

Appendix Table A5 decomposes the payback by the route the value travels, and the ordering restates the timing result of Section B in grant-level units. Direct patent citations alone repay a grant's funding in 13.1 years, and the science chain in 19.6 years; the patent chain (47.9% of funding recovered by year twenty) and the mixed route (24.4%) do not break even inside the window we can observe. The slow routes repay as well, on horizons beyond the observable window.

The differing slopes across cohorts in Figure 4 do not identify a trend. The recent cohorts' curves rise more steeply than the old ones, and it is tempting to read this against the literature on the declining productivity of research (Bloom et al., 2020), as though a research dollar bought more downstream value in 2010 than in 1990. Two confounds block that reading. First, linkage coverage roughly doubles across the cohorts (36% of 1990's grants have a traceable paper against 75% of 2010's). A correction for this bias is beyond our current scope, so we claim no trend in either direction and read the figure only for the level and the censoring: the young cohorts' lines are short because the data stop in 2022, not because their science stopped paying.

[10] Cohorts 1990–2002, each observed at least twenty years, undiscounted, keeping grants whose papers we cannot trace in the denominator. Every figure here is an *upper* bound on the payback time, because an untraceable grant contributes funding but no spillover and the linkage sees between 36% and 75% of grants depending on the cohort. Restricting to grants with at least one traceable paper gives 8.6 years; the highest-coverage cohorts give 10.7; the fully matured 1990–92 cohorts give 7.9.

Figure 4: How long a research dollar takes to pay for itself



Notes: NIH and NSF grants. Each line is one grant-start cohort, beginning at its own start year and stopping at 2022; the young lines are short because the data stop, not because the value does. The vertical axis is cumulative *realized* spillover as a percentage of the cohort's total funding, so break-even is the dashed 100% line. Panel (a) is undiscounted; panel (b) discounts each cohort's spillover to its start year. Cohort curves are aggregate ratios (total spillover divided by total funding), not means of per-grant ratios. The solid line divides by *all* of a cohort's funding; the band's upper edge divides only by the funding of grants whose papers we can trace, so the width of the band reflects linkage coverage, not differences in returns. Dashed lines give the direct route alone, for three cohorts. The differing slopes across cohorts do not identify a trend in science and should not be read as one (see text); this figure identifies the level and the censoring only.

VI. Results on the 2025 science cuts

Everything so far has been about science and funding in general: what the commercial value of science is worth, how concentrated it is, how long it takes to arrive, and where in the world it lands. We now use this framework to quantify the lost commercial value due to the 2025 US Science funding cuts. This section describes the cuts, builds the counterfactual that stands in for the science the terminated grants would have produced, and reports what the terminations cost, who held the terminated grants, which fields lost the most, and where the loss is borne.

A. The cuts, and the data

In early 2025 federal agencies terminated active research awards whose subject matter conflicted with the administration's stated priorities, implementing the executive orders described in Section I. Available evidence seems to indicate that agencies selected grants to stop by topic, not by scientific merit: they screened award text for diversity, "gender ideology," vaccine hesitancy, and China-related research (Liu et al., 2025; Oliveira et al., 2026). We take the terminated awards from Grant Witness, a public tracker of federal research awards ended under the 2025 executive actions, which compiles grant-level data from agency notices and litigation records. We use

the March 2026 snapshot. We restrict the headline analysis to NIH and NSF (1,116 and 1,363 terminated awards across all award types) and within them to research, career-development, and fellowship grants: 2,241 awards. Of these, 2,225 survive de-duplication and 2,221 have both positive recorded funding and a matched counterfactual. We exclude CDC, EPA, and SAMHSA terminations from the headline because our patent-citation measure cannot see put a value on these grants, not because we believe they are worthless.[11] We count grants terminated *and not reinstated*: awards restored after the June 2025 court rulings (and the partial survival of those rulings at the Supreme Court) are excluded, so the treated set is the policy's footprint as of March 2026 rather than its maximal intended scope. Counting every award ever disrupted would roughly double the size of the dataset.

The median terminated grant is a $570 thousand award (Table 2); the median start year is 2023, and 88.8% started in 2021 or later. This young age of the disrupted grants is the central measurement problem: only 197 of the 2,241 grants have even one publication in our network (4,666 articles in total), because their science was still in the pipeline when the awards were cut.

Table 2: The terminated NIH and NSF grants

	N	Median award ($k)		Median	Started
		Terminated	Project	start year	≥2021 (%)
NIH	867	931.7	938.5	2023	83.6
NSF	1,354	500.0	387.5	2023	92.1
Total	2,221	579.4	462.3	2023	88.8

Notes: NIH and NSF grants recorded as terminated (and not reinstated) in Grant Witness (March 2026 snapshot), restricted to research, career-development, and fellowship awards. The two award columns are medians in thousands of current dollars: *Terminated* is the value of the terminated award as Grant Witness records it — one segment of a multi-year NIH project, but essentially the whole of an NSF award — while *Project* is the lifetime project funding recorded in Dimensions, the denominator the per-dollar results use (Section B). "Started ≥2021 (%)" is the share of terminated grants in the row whose recorded start year is 2021 or later.

B. The counterfactual design

We estimate the expected 30-year global technology spillover embodied in the terminated research, and express it per dollar of grant budget. This is not a question that prior estimates of the marginal return to research funding can answer, because the cuts were ideological rather than made at the funding margin: a marginal funding cut removes the science a funder values least, while these cuts removed whatever matched the targeted topics. Where the cut science

[11] Most are block or community grants that fund services rather than research, so the measure would mechanically assign them near-zero spillover, and their $55.1 billion of budget would drag down the per-dollar ratio without saying anything about the science channel. For EPA the measure also misses regulatory and environmental benefit channels entirely.

falls in the value distribution is therefore an empirical question, answered directly in Section D, rather than something the design can assume. Two things bound what the resulting estimate claims. First, it is an *accounting projection*, not a causal treatment effect: we ask what spillovers the terminated research lines would plausibly have generated, not what the general-equilibrium effect of the policy on innovation will be. Second, we do not model substitution(e.g. a terminated line may be partly continued by the same investigator under other funding, or taken up by another team).

The obvious approach (projecting the terminated grants' own publications forward with the Section D model) fails for the reason Table 2 makes plain: the grants are too young to have publications, and for young articles the projection's level bias is at its worst. We instead borrow *observed* outcomes from comparable grants that ran their course. In other words, we have to estimate the papers that haven't been written, the other papers that will not build on them, and the value of the derived patents that will not be filled or granted. To do so, for each terminated grant we match *mature donor grants*, that is, non-terminated NIH and NSF awards with start years 2001–2013.

The donors are the population of comparable non-terminated grants: 403,134 NIH and NSF awards drawn from Dimensions with start years 2001–2016, of which 323,134 are mature and enter the matching.[12] A terminated research line is thus measured against the science that similar grants (same agency, field and funding band) actually produced. Matching is coarsened exact matching (Iacus et al., 2012) on agency × field × funding band. More precisely, field at the level of Dimensions' two-digit research divisions, funding bands on the pooled treated-and-donor distribution. The counterfactual assigned to a terminated grant is the mean donor spillover vector in its cell: 100% of treated grants match in the finest cell, and the median treated grant draws on 2,943 donors. Section D discusses the principal investigators in charge of the disrupted research grants.

The design rests on one key assumption: that mature donors' realized spillovers transport across a vintage gap to younger grants. A pseudo-cutoff backtest inside the donor universe supports it, and we apply no transport correction.[13]

[12] Whether the population pool is the safer choice hinges on how to read its zeros. On the canonical grant-to-article linkage, 72.85% of the mature population pool links to at least one article (235,394 grants, 1,851,823 articles), so most of the pool's zeros are *real* zeros — grants whose science generated no traceable downstream invention value. A design that discarded them would condition on success; keeping them is the conservative choice.

[13] Using only *earlier* donors to predict the realized spillover of *later* grants, through the same matching pipeline: compared over windows of the *same length* — so both sides have had the same number of years to accumulate spillover — the prediction recovers 0.94–0.98 of the later grants' realized aggregate spillover, and the transport does not decay as the vintage gap widens. Comparing a short prediction window against a long truth window would manufacture an apparent decay, because the window mismatch grows with the vintage gap by construction; equalizing the windows removes it.

The headline statistic sets each grant's counterfactual foregone spillover against the money that bought the science. The denominator is the lifetime project funding recorded in Dimensions, for the terminated grant and for the donor grants whose realized spillover stands in for it. The total sums to $3.00 billion across the 2,221 grants.[14]

Because a dollar of spillover decades out is not a 2025 dollar, we discount each grant's projected flow along the empirical development curve from the backtest (Figure A3 in the Appendix), giving a present-value multiplier of 0.65 at a 3% discount rate.[15]

C. What it cost

We estimate as a result that every dollar of terminated grant funding forgoes $8.69 of global technology spillover, $5.61 in present value at 3% (Table 3). That means that the 2,221 terminated NIH and NSF grants would have generated $26.1 billion of expected spillover over thirty years ($16.9 billion in present value) against $3.00 billion of funding.[16]

The resulting ratio is a floor, for two reasons. The projection under-predicts aggregate spillovers, especially for younger articles that the terminated grants would have produced (Section B). And the grant-to-article linkage coverage is incomplete, so a terminated grant's counterfactual omits the spillover that a fuller linkage would have credited to it. More fundamentally, we price only the patent channel; Section VII returns to what that leaves out.

[14] We divide by lifetime project funding rather than the terminated-award value recorded by the tracker: an award record is one segment of a multi-year NIH project but essentially the whole of an NSF one, so dividing by it would deflate the NIH denominator and reverse the agency ordering. Appendix C.1 sets out that segment-versus-whole wedge and shows it is mechanical. Table A2 also reports the ratio per dollar of terminated *award* and per dollar actually *withheld*, which answer different questions rather than variants of this one.

[15] 0.74 at 2% and 0.41 at 7%; 18 years to the halfway point. Starting the clock at the 2025 cut slightly front-loads the flow, since the terminated grants' articles would have appeared with a lag of a few years, so a publication-lag variant would lower the multiplier; the undiscounted ratio is unaffected. Confidence intervals resample donor *investigators* — the level at which donor information clusters — holding the treated census fixed.

[16] Drawing the donors instead from the grants held by the terminated investigators themselves — the more established among them, since a donor grant exists only where an investigator holds more than one award — raises the ratio to $11.31 per dollar, about 30% higher; those grants earn more per dollar than population grants in the same agency, field and funding band, so the population pool is the conservative choice (Appendix C).

A second force can push the headline the other way: acknowledgement coverage itself has risen sharply over time, so part of the measured increase in spillover per dollar may just reflect better linkage, not more valuable science. Because this drift could compound to a factor of 1.4 over the decade between cohorts, we report a band from \$6.21 (treating the rise as pure artifact) to \$8.69 (treating it as fully real), since the bias can be bounded but not pinned down exactly (Appendix C).[17][18]

Table 3: Foregone spillover per dollar of grant budget

	N	Funding (\$B)	Foregone (\$B)	Spillover per \$	Spillover per \$ (PV@3%)
NIH	867	2.03	15.2	7.48	4.83
NSF	1,354	0.97	10.9	11.20	7.23
Total	2,221	3.00	26.1	8.69	5.61
				[8.42, 8.98]	

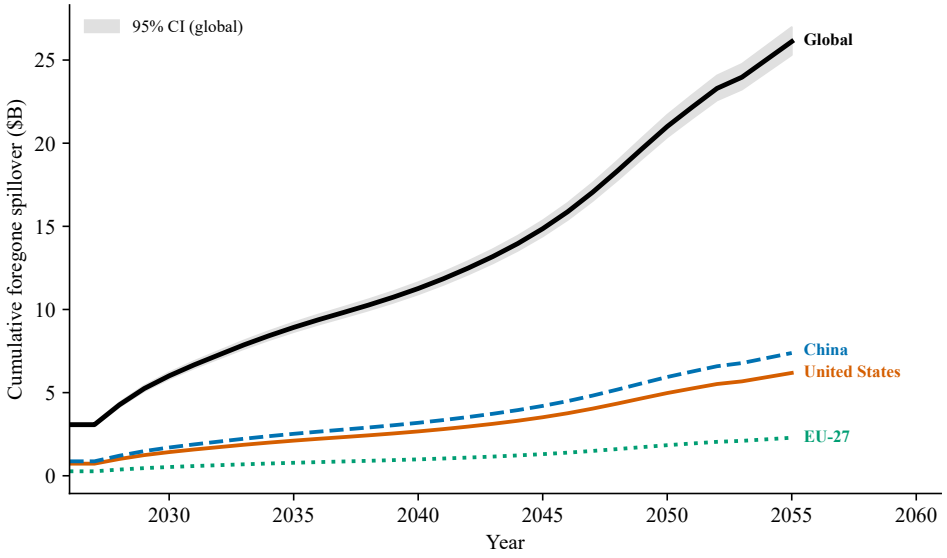
Notes: Terminated NIH and NSF grants with positive recorded funding and a matched counterfactual (2,221 grants). Foregone spillover is the sum of matched-donor counterfactual 30-year global spillovers (Section VI), in millions of constant 2010 dollars under the Kogan et al. (2017) calibration. PV@3% discounts the spillover flow along the empirical development curve from the 2025 termination date (present-value multiplier 0.65). The per-dollar columns divide by lifetime project funding as recorded in Dimensions, the same object on the treated and the donor side — not by the value of the terminated award, which is a segment of an NIH project but essentially the whole of an NSF one, and which for that reason reverses the agency ordering (Section B). Each article's spillover is split equally across the grants that supported it. The donor-sampling 95% confidence interval on the total foregone spillover is [25,307, 26,977] million dollars, clustered at the donor-investigator level.

Because science matures into patented technology slowly, the loss quantified above is not a 2025 event but a decades-long flow. Only 34% of the 30-year loss has accrued ten years after the cut, and half arrives only after 18 years: the empirical development curve applied to the totals of Table 3 (Figure 5). The bulk of the cost therefore lands in the 2040s, long after the 2025 savings were booked.

[17] Our per-dollar measure relies on grant-to-article linkage, and the share of papers acknowledging a grant rose several-fold from the oldest to the youngest donor cohorts, inflating measured spillover under our equal-split convention even absent any real change.

[18] The accounting choices that define a matching cell are second-order by comparison: re-banding the same donors on alternative edges moves the peer-pool construct from \$11.31 to \$12.17 per dollar (a 7.6% change), and restricting the donor window to its oldest or youngest cohorts moves it by a further 5.9% and 5.8%. Appendix Table A2 collects these accounting variants.

Figure 5: The 30-year trajectory of the foregone spillover



Notes: Cumulative foregone spillover (in billions of 2010 dollars) over the 30 years following the 2025 terminations, globally and for the three largest recipient regions (inventor basis). We build the time profile by applying the empirical development curve to the totals of Table 3; the shaded band applies the donor-sampling 95% confidence interval to the global trajectory. *Source*: authors' calculations from the held-out backtest development curve applied to the Table 3 totals.

D. Who held the terminated grants

A political justification and natural defence of the grant terminations is that they removed weak or unproductive science. Our matched counterfactual already sheds light on this: the cut grants are priced by ordinary donor grants in the same agency, field and funding band. A limitation is that it matches on only three characteristics and on relatively few terminated grants, so we add an independent piece of evidence that uses no spillover model: the publication records of the people whose grants were cut. To that end, we match each principal investigator to their publication record, strictly as of the grant year, for 87% of the terminated grants.[19] On that record, evidence suggests that the cut grants are held by established, actively publishing scientists (Table A3): the median terminated investigator had 47 prior publications over a 21-year career, and only 6% had fewer than five prior papers. Only the fellowship awards diverge, but these are by design attributed to trainee grants whose holders are early-career (median 5 prior papers, 129 grants). For the rest, the research grants that account for almost all of the foregone spillover are held by senior people (median 54 papers, 1,742 grants). On aggregate, the terminations did not select on inexperience.

[19]Principal investigators are resolved from Dimensions to their OpenAlex author record by ORCID (Open Researcher and Contributor ID, a persistent per-researcher identifier), DOI match, and author-name matching against the authors of their grant-linked papers (Section VI); the unresolved remainder are disproportionately early-career fellowship holders, whose thin publication records reflect genuinely early careers rather than a resolution failure.

A fair comparison with the investigators who survived must hold across vintages. We therefore compare the terminated investigators to the holders of grants awarded in the same years (2021 to 2024) that were not terminated, resolving both cohorts to OpenAlex by ORCID (the same resolution method on both sides, reaching 76% of the terminated research investigators). Because the two cohorts are contemporaries, we can use a measure closer to scientific impact than raw output: the *Science Rank* spillover each investigator's prior research had already generated, per career-year. We restrict the terminated side to research grants (to exclude fellowships) and rank each terminated investigator among the 62,471 surviving contemporaries, deduplicated to unique people.

Table 4 narrows the comparison from all survivors down to the investigator's own agency and field. Against the national pool the terminated investigators rank below the median grantholder — the 42th percentile on productivity and the 29th on prior spillover. Within the same agency and field, the finest cell that still covers the whole cohort, they reach the 47th percentile on productivity (which is essentially the peer median) and the 40th on prior spillover (modestly below the peer median). The layered comparison could in principle produce a below-median result on its own, so we check it against a placebo: relabelling random surviving grants as terminated and running the same layers returns the median percentile at every one, as a neutral procedure should.[20]

All in all, the terminated investigators are established and productive; they are at the median of their same-field contemporaries on output and somewhat below it on prior spillover; and per research dollar their grants rank at the 38th percentile of comparable grants, modestly below the middle rather than at the bottom.[21] They are not the markedly weaker scientists, nor, on our measure, the best of their peers. Termination appears uncorrelated with the future commercial value of the science it removed.

[20] A second placebo confirms the same for the counterfactual positioning machinery: relabelling random donor grants as pseudo-terminated and running the positioning pipeline recovers a median percentile of 50.15, against the 50 a null pipeline yields by construction.

[21] Unconditional on grant size their foregone spillover ranks lower, at the 20th percentile, that is, the gap a portfolio of below-median-*sized* grants produces. What the terminations removed were the *smaller* of the grants on offer (Appendix Figure A5).

Table 4: Terminated investigators' standing among their surviving contemporaries

	Percentile among surviving contemporaries	
Comparison set	Productivity	Prior spillover
All 2021–2024 survivors	41.6	29.2
Same agency and field	46.9	40.2
Same field and institution	49.7	45.1
Placebo (random survivors)	49.9	50.6

Notes: Median percentile of the 1,234 terminated research-grant investigators within the distribution of 62,471 surviving contemporary investigators, on two as-of measures of prior track record: publication productivity (papers per career year) and Science-Rank spillover per career year, both computed strictly from pre-grant publications. Both cohorts are 2021–2024 NIH and NSF grant-holders, resolved to their OpenAlex record identically by ORCID and deduplicated to unique investigators, so the comparison has no vintage gap and no resolution asymmetry. Each row restricts the comparison set to survivors in the same cell; a percentile of 50 is the peer median. The all and field rows cover the whole resolved cohort; the institution row covers only the 34.6% of terminated grants that fall in a field-and-institution cell with at least twenty surviving investigators — those at large research universities. Its near-median standing therefore describes that subset, not additional conditioning on the whole cohort. The field row is the headline conditional comparison. ORCID resolves 76% of the terminated research investigators; the unresolved remainder, disproportionately those with thinner records, are omitted, so the resolved comparison is if anything favourable to the terminated side. The placebo relabels random survivors as terminated and recovers the median at every layer.

E. The fields that lost the most, and whether the cut was selective

Importantly, the fields that lost the most grants are not the fields that lost the most value. By count, Education is the most-terminated division (742 grants, reflecting the topic-driven targeting of diversity- and education-related work) yet it accounts for only 8.5% of the foregone spillover. This is expected by our measure that is based on patentable inventions, excluding the wider societal benefits of education.[22] By value, the loss concentrates in patent-proximate science: biomedical and clinical sciences (29.8%), information and computing (23.6%) and biological sciences (17.0%) jointly account for 70% of the total (Appendix Table A4, which reports the dollar loss and grant count of each field).

A raw count of terminations cannot show whether the cuts singled out particular fields or recipient countries, because it ignores how large each field or region is in the funded portfolio. Table 5 supplies the missing benchmark: it compares each field's and each region's share of the foregone spillover with its share of the funded NIH and NSF portfolio, and calls a field or region over-represented, or spared, only when the ratio stays on one side of parity across every vintage window we can construct. On that standard the selectivity is narrow. Only Education and Commerce are robustly over-represented (Education by a factor of 2.9 to 14.1) and only

[22]Grant Witness flags diversity-related language for a subset of terminations. Read strictly as a proxy as the flags are coded from award text and termination lists, not from official rationales, these diversity-, equity-, and inclusion- (DEI-) flagged grants account for 10% of the foregone spillover: frequent among the terminations by count, but a modest share of the value put at risk.

Biological Sciences robustly spared. Biomedicine and every recipient region, China included, are *untargeted* in this sense: their ratios cross parity as the benchmark's vintage changes, so an apparent over-representation of China, or under-representation of biomedicine, on any single window is a composition artifact rather than evidence that the cut selected on geography or on those fields.

Table 5: Was the cut field- or region-selective? Over-representation relative to the funded portfolio

Panel A. By field (ratio of cut share to portfolio share)					
	cut %	matched	all vintages	funding	verdict
Biomedical	29.8	1.03	0.89	0.73	not identified
Information & computing	23.6	1.47	1.97	2.81	not identified
Biological sciences	17.0	0.65	0.55	0.70	spared
Education	8.5	11.6	14.1	2.94	over-cut
Health sciences	5.6	2.68	2.03	0.83	not identified
Engineering	3.9	0.30	0.42	1.11	not identified
Chemical sciences	3.5	0.48	0.56	2.01	not identified
Commerce & management	2.2	11.2	13.7	3.01	over-cut
Psychology	1.3	2.73	2.68	0.88	not identified
Panel B. By recipient region, inventor basis					
	cut %	matched	all vintages	cut-era	verdict
China	28.3	0.98	1.11	1.16	not identified
United States	23.7	1.04	0.88	0.73	not identified
Unattributed	16.3	0.97	1.11	1.20	not identified
European Union	8.8	1.01	0.90	0.95	not identified
Rest of World	8.2	1.02	0.99	1.05	not identified
Japan	7.2	1.00	0.97	1.19	not identified
South Korea	5.9	0.98	1.07	1.19	not identified
United Kingdom	1.7	1.02	0.88	0.92	not identified

Notes: Each ratio compares a field's or region's share of the cut's foregone spillover to its share of the NIH and NSF portfolio's spillover, which is the benchmark the raw field counts of Table A4 do not supply. Because the cut grants are priced by mature donors started 2001–2013, the portfolio denominator is reported on several vintage windows; a field or region is called over-represented or spared only when the ratio stays on the same side of parity across all of them. On that standard, only Education and Commerce are robustly over-represented (Education by a factor of 2.9 to 14.1) and Biological Sciences robustly spared. Biomedicine and every recipient region are *untargeted* in this sense: their ratios cross parity as the benchmark's vintage changes, so the apparent under-representation of biomedicine and over-representation of China on any single window is a composition artifact rather than evidence of selective cutting.

F. Where it is borne

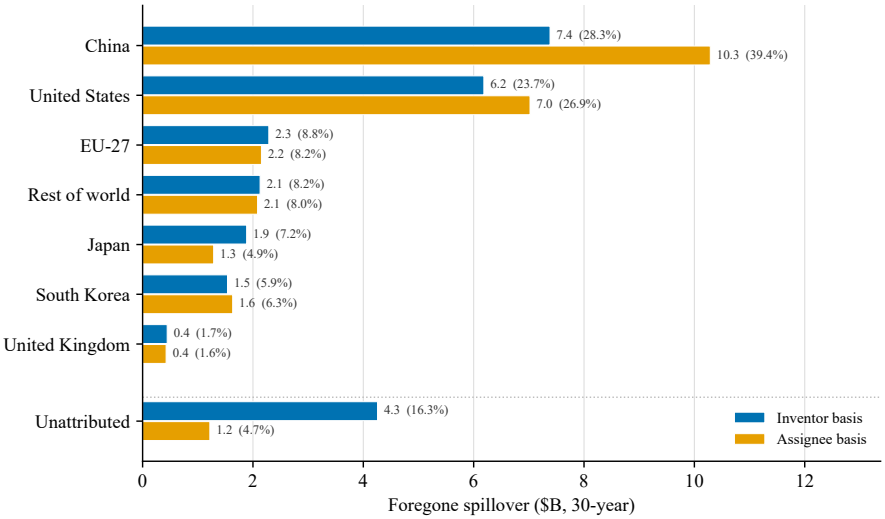
The commercial value of this science does not stay where it was funded. The free-riding pattern of Section A, which means a national funder pays the whole bill and collects a minority of the value, also holds for the specific grants that were cut. Of the foregone spillover, 60.0% would have been realized on inventions made outside the United States, against 23.7% at home.[23] Among the foreign locations China is the largest, at 28.3% on the inventor basis (which locates where the downstream technology is developed) and 39.4% on the assignee basis, which locates who owns it; the EU-27 is a distant third at about 8.8% (Figure 6).

We read these shares as accounting incidence of where the value would have surfaced, and not as value destroyed, transferred, or taken from anyone as science is not a zero-sum game. Even read as pure incidence, the asymmetry is large: the country that bore the entire fiscal cost of this science would have realized under a quarter of the technological value built on it. Two measurement features pull on the Chinese number in opposite directions. The cells that assign CV_p have no country dimension, so a Chinese family is assigned the value of US-listed families sharing its technology, vintage, family size and claim count (Section A), which likely overstates it. Pulling the other way, the geographic tilt is not only due to patent valuation: setting every patent's private value to one, so the measure becomes a pure citation-network statistic, reproduces most of the pattern.

The geographical location of commercial value has a policy corollary. A region has a private incentive to fund research only up to the value it expects to appropriate from it, and on that criterion the country with the most to gain from this particular science is not the one that paid for it: China would capture a larger share of its commercial value (28.3%) than the United States would (23.7%). No government funds another's terminated grants, and we do not suggest one would. But the number makes the structural point precise: research whose benefits diffuse across borders is underfunded by any single national budget, and a unilateral cut transfers, in the first instance, future commercial value to the countries that build on the science it stops producing.

[23] A further 16.3% falls on inventions we cannot place in any country. The shares reported here are shares of the full total, so this unplaceable residual is stated rather than reallocated across the regions.

Figure 6: Foregone spillover by recipient region, inventor and assignee bases



Notes: The foregone 30-year global spillover of Table 3 across recipient regions, on the inventor and assignee bases, ordered largest to smallest. The inventor basis locates the downstream technology's development; the assignee basis locates the ownership of its value. The two bases are alternative partitions of the same total and are never summed. "Unattributed" collects patent families with no resolvable country — a residual and a coverage diagnostic, not a region. The held-out backtest cannot sign the direction of any error in the regional split, because the held-out predictions do not include region components (Appendix B), so the shares are reported without a bias correction.

G. Cross-cutting findings

In this final results section, we illustrate the advantage of our bottom-up approach to measure the lost commercial value from the science cuts, by comparing the distribution of commercial value at each stage, from the invention that carries it to the terminated grant whose science would have produced it. Table 6 sets the four stages side by side:

- the commercial value of patented inventions (CV_p),
- the commercial value of science from the patented inventions building on it (CV_s),
- the commercial value carried by grant-linked science (CV_{sf}),
- the commercial value of the cut science (CV_{sc}).

Each of these stages is a successively narrower restriction of the same downstream dollars, from $868.39 trillion to $26.1 billion, as illustrated in Table 6.

Table 6: The four commercial-value objects, defined and described

	$\mathbf{CV}_p$ inventions	$\mathbf{CV}_s$ science	$\mathbf{CV}_{sf}$ funded science	$\mathbf{CV}_{sc}$ the cut
Unit	patent family	article	grant	terminated grant
Sample	global, priority 1985–2022	global	global, all funders	US, NIH + NSF, 2025
Basis	realized	realized through 2022	realized through 2022	30-year projection[a]
N units	69,215,274	452,057,476 (78,438,549 with value)	12,336,794 linked articles; 29,214,405 edges	2,221
Total	$868.39T	$51.66T	$8.92T (allocated to grants)	$26,112M
Mean per unit	$12.55M all; $22.58M among positive	$0.11M all; $0.66M among positive	–	$11.76M
Share carrying any value	55.6%	22.2%	–	–
Concentration: top 1% of units hold	13.2%	82.9%	–	–
Gini (among units with any value)	0.62	0.95	–	–
Rate ($ per $ of funding)	–	–	2.91 all funders; 4.62 NIH+NSF[f]	8.69 undisc.; 5.61 at 3%
Share of the object above	–[b]	5.9% of CV_p	19.1% of CV_s[c]	–[d]
Region: US / CN / EU[e]	US 14.5 / CN 22.8 / EU 8.3	US 27.5 / CN 23.1 / EU 11.6	–	US 23.7 / CN 28.3 / EU 8.8

Notes: Each object is the same downstream technological value, restricted to a narrower set: the private value of a patented invention (CV_p), the value that reaches scientific articles (CV_s), the value on grant-linked articles (CV_{sf}), and the value of the terminated grants' science (CV_{sc}). CV_p, CV_s and CV_{sf} are realized through 2022; CV_{sc} is a thirty-year projection (keyed cell [a]). Moments and concentration are computed on the network window, *not* on the separate patent-value table, which covers a different population. Table 7 breaks each column down by macro-field. Science's commercial value is far more concentrated than invention's (Gini 0.95 against 0.62), a gap generated by the citation network rather than by the patent-value calibration. Keyed cells: [a] CV_{sc} is a thirty-year projection — an estimator, not a sum over an observed set: the cut grants are too young to have produced the science whose value we price. [b] CV_p is the source of the downstream value, not a superset of the science objects, so it has no share of an object above it. [c] a coverage floor, not the funded share of science: grant linkage is incomplete and collapses before 2000. [d] no like-for-like share exists — CV_{sc} is US, NIH and NSF only, while CV_{sf} is world, all-funder — so dividing one by the other is not printed. [e] region shares are value-weighted; the unattributed residual is a coverage gap, not a region, and is omitted from these three. The CV_s split is the paper-side global partition at the baseline $\sigma = 0.5$. [f] both rates are realized through 2022 on grants started by 2013, so neither is comparable to the projected thirty-year rates of Table 1. They differ only in the funders they cover, and the gap is the reason we report rates within funder (Section C): the all-funder figure pools award vehicles whose outputs our acknowledgement-based linkage cannot trace — national-laboratory contracts, the European Commission, research ministries, foundations — and is dragged below *both* of the agencies it contains. We pool the NIH and the NSF here rather than splitting them, because on this vintage window the split is a linkage artifact and not a fact about the agencies.

Several observations are worth highlighting. First, commercial value is scarce and extreme-tailed. Table 6 allows us to compare this across the stages, decomposing the contribution of each stage to the skew. Table 6 confirms the finding reported in Section A, that the value of science is far more concentrated than the value of the inventions it feeds: a Gini of 0.95 against 0.62. It is science more than inventions that skews the distribution of commercial value created. Only 22.2% of articles carry any commercial value, and the top percentile holds 82.9% of the total commercial value from science, resulting in the extremely high Gini value of 0.95, much higher than the Gini for the total commercial value among patents. This gap is generated by the knowledge linkage networks rather than by the patent value calibration.

Second, science functions as a global public good. The commercial value of science falls where science's downstream users are, and they are geographically dispersed. Table 6 shows this geographic distribution for the various stages. The dominance of the US is not particularly because of its dominance in creating patenting value, but in creating patenting value out of science. Among the foreign countries, China is the strongest country on both, while the EU is weak on both. This translates into the highest losses outside the US from the 2025 funding cut falling on China.

Finally, we look at the distribution of the commercial value of science, split up by scientific discipline (Table 7). When looking from the patent perspective (CV_p), the most important scientific field is by far physics, engineering and energy. But looking from the publication perspective, life sciences and biomedicine dominate by far, reflecting the larger size of the scientific discipline. This dominance of the life sciences is further exacerbated when looking at the value of funded science, with the high share of the NIH in total US funding. All this results in the life sciences being also the field most represented in the loss of commercial value from the 2025 US science funding cuts, although somewhat less than their commercial value in funded science would have predicted. The physical sciences and engineering seem to be less affected in the commercial value loss than would have been predicted from their value in funded science, with the exception of information and computing, which seems to be overly affected. These numbers, however, need to be taken very carefully, as the number of grants in this field is very small.

Table 7: Commercial value by macro-field (share of each object's total, %)

	$\mathbf{CV}_p$ inventions	$\mathbf{CV}_s$ science	$\mathbf{CV}_{sf}$ funded science	$\mathbf{CV}_{sc}$ the cut
Life sciences & biomedicine	12.8	44.5	62.1	52.5
Chemistry & materials	16.7	16.4	10.8	3.5
Information & computing	20.6	13.8	8.6	23.6
Physics, engineering & energy	34.9	21.5	16.8	6.0
Social sciences, humanities & education	–	2.0	1.8	14.4
Other & unclassified	15.1	1.9	0.0	0.0

Notes: Share of each object's total commercial value held by each of six macro-fields, aggregated from the OpenAlex concept and patent-class hierarchies. The four objects (CV_p, the private value of a patented invention; CV_s, the value reaching articles; CV_{sf}, grant-linked science; CV_{sc}, the value of the cut) are defined, with their unit counts and dollar totals, in Table 6; each cell is that field's share of the corresponding column total there. The patent channel does not price the social sciences, so that cell is empty for CV_p by construction. Column shares need not sum to exactly 100 because of rounding.

VII. Conclusions

Recent policy developments have made the measuring and understanding of the returns to public investments in science more important than ever. This is in particular relevant in the major scientific powerhouse, the United States, whose public funding of science is significantly changing course.

This paper focuses on the commercial value created from patented inventions using science. To do so, we develop a bottom-up framework that reconstructs the knowledge network linking scientific publications to patented inventions. Rather than asking how much commercial value is created on average from investing in science, we look bottom-up at how funded scientific knowledge travels through the innovation system. We decompose the sources and users of science by geography, by field and by quality segment, at both the paper and the patent level. We identify which scientific disciplines generate the largest commercial spillovers. We trace these patterns internationally to assess how much of the commercial value is created outside the national funder's territory.

We demonstrate the usefulness of this measurement framework by studying the consequences of the 2025 US executive actions that terminated large numbers of federally funded research grants. We quantify the downstream commercial spillovers that are put at risk by these funding cuts, compared to the fiscal savings achieved through grant termination. Terminating 2,241 NIH and NSF research grants gave up an expected $26.1 billion of global technology spillover over thirty years against $3.00 billion of funding, which is $8.69 per dollar, $5.61 in present value. An interesting finding is that the cuts did not fall on weak science. And the loss arrives on a schedule no budget window can see: a research dollar breaks even only after 9.8 years, and 2.4% of its eventual value lands within three.

Five limitations bound what these estimates can claim. First, they are accounting attributions, not causal effects: we do not model substitution when terminated lines partly continued by the same investigators under other funding, or taken up elsewhere. Second, every monetary estimate of commercial value (the headline per-dollar ratio included) scales with the Kogan et al. (2017) market-value calibration and with the decay σ. The paper's comparative structure is however invariant to that calibration: the rankings of per-dollar returns across agencies and fields, the regional shares, and the positioning percentiles. Third, the patent-citation channel misses economic spillovers of science that never pass through a patent, for example because it is not captured by firms or takes the form of wider societal benefits such as clinical practice, regulation, environmental and public-health benefits, and the training of people. Fourth, the seed values themselves are noisy in a way that biases our partitions. Arora et al. (2024) show that the

measurement error in Kogan et al. (2017) values is negatively correlated with true patent value, which leaves the values usable as a regressor but biases comparisons across groups of patents. We have no correction for this bias, and we flag it as an unresolved bias of the composition results. Last, the data carry known coverage gaps: investigator disambiguation is imperfect, a share of patent geography is unattributable, grant-to-article links may be incomplete, and the diversity-related flags are coded proxies rather than official rationales. The grant-to-article linkage reaches only 35.3% of projected spillover, and roughly a third of patents built on federally funded research fail to disclose that support at all (Kwon, 2024), so acknowledgement-based linkage undercounts by construction. Except for the measurement-error threat, each limitation we can sign pushes the true loss above our estimate.

Taking these caveats into account, we see our contribution as twofold. More broadly, the paper contributes to a growing literature seeking to understand not only the overall returns to science funding, but also the patterns through which scientific discoveries create commercial value throughout the economy. Methodologically, we provide a new micro-level framework for measuring the commercial value from scientific knowledge by micro-tracing the pathways through which science generates economically valuable inventions. Most studies only measure one of the steps in the chain. Our framework offers not only a new metric for the returns on science but can also be used as a new tool for the *ex ante* evaluation of science policy. Instead of measuring the outcomes of a grant retrospectively, the method can be used to estimate the long-term costs of budget cuts, shifts in research priorities, or changes to funding programs.

Finally, an important conclusion is that the United States' public science functions as a global public good. The loss falls where science's downstream users are, and they are not mostly based in the US. More than half of the value would have surfaced outside the country that paid for the science. Among the non-US destinations, China would have captured the largest part.

References

Acemoglu, D., Akcigit, U., and Kerr, W. R. (2016). Innovation network. *Proceedings of the National Academy of Sciences*, 113(41):11483–11488. [Cited on page 47.]

Aghion, P., Dewatripont, M., and Stein, J. C. (2008). Academic freedom, private-sector focus, and the process of innovation. *The RAND Journal of Economics*, 39(3):617–635. [Cited on page 8.]

Aghion, P., Dechezleprêtre, A., Hemous, D., Martin, R., and Van Reenen, J. (2016). Carbon taxes, path dependency, and directed technical change: Evidence from the auto industry. *Journal of Political Economy*, 124(1):1–51. [Cited on page 12.]

Agrawal, A., Cockburn, I., and McHale, J. (2006). Gone but not forgotten: Knowledge flows, labor mobility, and enduring social relationships. *Journal of Economic Geography*, 6(5):571–591. [Cited on page 7.]

Ahmadpoor, M. and Jones, B. F. (2017). The dual frontier: Patented inventions and prior scientific advance. *Science*, 357(6351):583–587. [Cited on pages 6, 13, and 47.]

Arora, A., Belenzon, S., Ferracuti, E., and Nagar, J. P. (2024). Revisiting stock market signals as a lens for patent valuation. NBER Working Paper 33056, National Bureau of Economic Research. Revised April 2026. [Cited on page 37.]

Arora, A., Belenzon, S., and Sheer, L. (2021). Knowledge spillovers and corporate investment in scientific research. *American Economic Review*, 111(3):871–898. [Cited on page 6.]

Arrow, K. J. (1962). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press, Princeton, NJ. [Cited on pages 2 and 20.]

Audretsch, D. B. and Feldman, M. P. (1996). R&D spillovers and the geography of innovation and production. *American Economic Review*, 86(3):630–640. JSTOR stable ID 2118216, https://www.jstor.org/stable/2118216. [Cited on page 7.]

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 8.]

Azoulay, P., Graff Zivin, J. S., Li, D., and Sampat, B. N. (2019). Public R&D Investments and Private-sector Patenting: Evidence from NIH Funding Rules. *The Review of Economic Studies*, 86(1):117–152. [Cited on pages 7 and 22.]

Azoulay, P. and Greenblatt, W. H. (2025). Does peer review penalize scientific risk taking? Evidence from NIH grant renewals. NBER Working Paper 33495, National Bureau of Economic Research, Cambridge, MA. [Cited on page 7.]

Azoulay, P., Sadun, R., and Scur, D. (2026). Before the Exodus? Young Scientists and the Future of US Science. Working Paper 35330, National Bureau of Economic Research, Cambridge, MA. [Cited on page 8.]

Azoulay, P. and Li, D. (2020). Scientific grant funding. NBER Working Paper 26889, National Bureau of Economic Research, Cambridge, MA. [Cited on page 7.]

Babina, T., He, A. X., Howell, S. T., Perlman, E. R., and Staudt, J. (2023). Cutting the Innovation Engine: How Federal Funding Shocks Affect University Patenting, Entrepreneurship, and Publications. *The Quarterly Journal of Economics*, 138(2):895–954. [Cited on page 7.]

Baker, S. R., Bloom, N., and Davis, S. J. (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics*, 131(4):1593–1636. [Cited on page 8.]

Belenzon, S. and Schankerman, M. (2013). Spreading the word: Geography, policy, and knowledge spillovers. *The Review of Economics and Statistics*, 95(3):884–903. [Cited on page 7.]

Bergeaud, A., Guillouzouic, A., Henry, E., and Malgouyres, C. (2025). From Public Labs to Private Firms: Magnitude and Channels of Local R&D Spillovers. *The Quarterly Journal of Economics*, 140(4):3233–3282. [Cited on page 7.]

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393. [Cited on pages 5, 22, and 47.]

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4):1104–1144. [Cited on page 23.]

Bonacich, P. (1987). Power and centrality: A family of measures. *American Journal of Sociology*, 92(5):1170–1182. [Cited on page 47.]

Borjas, G. J. and Doran, K. B. (2012). The collapse of the Soviet Union and the productivity of American mathematicians. *The Quarterly Journal of Economics*, 127(3):1143–1203. [Cited on page 8.]

Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117. [Cited on page 47.]

Bryan, K. A., Ozcan, Y., and Sampat, B. (2020). In-text patent citations: A user's guide. *Research Policy*, 49(4):103946. [Cited on page 6.]

Chen, P., Xie, H., Maslov, S., and Redner, S. (2007). Finding scientific gems with Google's PageRank algorithm. *Journal of Informetrics*, 1(1):8–15. [Cited on page 47.]

Coe, D. T. and Helpman, E. (1995). International R&D spillovers. *European Economic Review*, 39(5):859–887. [Cited on pages 7 and 20.]

Cohen, W. M., Nelson, R. R., and Walsh, J. P. (2002). Links and impacts: The influence of public research on industrial r&d. *Management Science*, 48(1):1–23. [Cited on page 6.]

Cutler, D. M. and Glaeser, E. (2025). Cutting the NIH—the $8 Trillion Health Care Catastrophe. *JAMA Health Forum*, 6(5):e252791. [Cited on page 8.]

Congressional Budget Office (2025). How Changes to Funding for the NIH and Changes in the FDA's Review Times Would Affect the Development of New Drugs. Report, Congressional Budget Office. [Cited on page 8.]

Doran, K., Gelber, A., and Isen, A. (2022). The effects of high-skilled immigration policy on firms: Evidence from visa lotteries. *Journal of Political Economy*, 130(10):2501–2533. [Cited on page 8.]

Dosi, G., Llerena, P., and Sylos Labini, M. (2006). The relationships between science, technologies and their industrial exploitation: An illustration through the myths and realities of the so-called 'European Paradox'. *Research Policy*, 35(10):1450–1464. [Not cited.]

Dyèvre, A. (2024). Public R&D spillovers and productivity growth. Working paper, ECB Forum on Central Banking, Sintra. [Cited on page 5.]

Eaton, J. and Kortum, S. (1999). International technology diffusion: Theory and measurement. *International Economic Review*, 40(3):537–570. [Cited on pages 7 and 20.]

Esposito, C. R. (2026). Global Science Sustains U.S. Innovation. arXiv preprint arXiv:2605.30435. [Cited on page 8.]

Fang, F. C., Bowen, A., and Casadevall, A. (2016). NIH peer review percentile scores are poorly predictive of grant productivity. *eLife*, 5:e13323. [Cited on page 7.]

Fieldhouse, A. and Mertens, K. (2025). The Social Returns to Public R&D. Working Paper 33780, National Bureau of Economic Research, Cambridge, MA. [Cited on pages 5 and 22.]

Fieldhouse, A. J. and Mertens, K. (2024). The Returns to Government R&D: Evidence from U.S. Appropriations Shocks. Working Paper 2305 (revised), Federal Reserve Bank of Dallas. [Cited on pages 5 and 22.]

Fleming, L., Greene, H., Li, G.-C., Marx, M., and Yao, D. (2019). Government-funded research increasingly fuels innovation. *Science*, 364(6446):1139–1141. [Cited on page 6.]

Freeman, R. B. and Van Reenen, J. (2009). What if Congress doubled R&D spending on the physical sciences? In Lerner, J. and Stern, S., editors, *Innovation Policy and the Economy, Volume 9*, pages 1–38. University of Chicago Press, Chicago. [Cited on page 8.]

Ganguli, I. (2017). Saving soviet science: The impact of grants when government R&D funding disappears. *American Economic Journal: Applied Economics*, 9(2):165–201. [Cited on page 8.]

Garisto, D. (2026). Exclusive: NSF puts new research grants to top universities on hold. *Nature*. News. [Cited on page 8.]

Garisto, D. and Lenharo, M. (2026). White house proposes vast overhaul of US science funding: what you need to know. *Nature*. News. [Cited on page 8.]

Gazzani, A. G., Martinez, J., Natoli, F., and Surico, P. (2026). The Public Origins of American Innovation. Temi di Discussione (Working Paper) 1521, Bank of Italy, Rome. [Cited on page 5.]

González, I., Montecino, J., and Ramaswamy, V. (2025). Preliminary Estimates of the Macroeconomic Costs of Cutting Federal Funding for Scientific Research. Report, Institute for Macroeconomic & Policy Analysis, American University. [Cited on page 8.]

Griliches, Z. (1958). Research costs and social returns: Hybrid corn and related innovations. *Journal of Political Economy*, 66(5):419–431. [Cited on page 5.]

Griliches, Z. (1979). Issues in assessing the contribution of research and development to productivity growth. *The Bell Journal of Economics*, 10(1):92–116. [Cited on pages 2 and 5.]

Griliches, Z. (1992). The search for R&D spillovers. *The Scandinavian Journal of Economics*, 94(Supplement):S29–S47. [Cited on page 5.]

Guillard, C., Martin, R., Mohnen, P., Thomas, C., and Verhoeven, D. (2021). Efficient industrial policy for innovation: Standing on the shoulders of hidden giants. CEP Discussion Paper 1813, Centre for Economic Performance, London School of Economics and Political Science, London. Also circulated as CEPR Discussion Paper No. 16850 and POID Working Paper No. 054. [Cited on pages 3, 9, 10, 12, and 47.]

Hall, B. H., Mairesse, J., and Mohnen, P. (2010). Measuring the returns to R&D. In Hall, B. H. and Rosenberg, N., editors, *Handbook of the Economics of Innovation*, volume 2, chapter 24, pages 1033–1082. North-Holland/Elsevier, Amsterdam. [Cited on pages 2, 5, 6, and 17.]

Hasgul, Z., Lee, H., and Jalali, M. S. (2025). The Aftershock: The Impact of Budget Reductions at the US National Institutes of Health on Research and Innovation. SSRN Working Paper 5399303. [Cited on page 8.]

Iacus, S. M., King, G., and Porro, G. (2012). Causal Inference without Balance Checking: Coarsened Exact Matching. *Political Analysis*, 20(1):1–24. [Cited on page 26.]

Jacob, B. A. and Lefgren, L. (2011a). The impact of NIH postdoctoral training grants on scientific productivity. *Research Policy*, 40(6):864–874. [Cited on page 7.]

Jacob, B. A. and Lefgren, L. (2011b). The impact of research grant funding on scientific productivity. *Journal of Public Economics*, 95(9–10):1168–1177. [Cited on page 7.]

Jaffe, A. B. (1986). Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value. *American Economic Review*, 76(5):984–1001. [Cited on page 47.]

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics*, 108(3):577–598. [Cited on page 7.]

Jalali, M. S. and Hasgul, Z. (2025). Potential Trade-Offs of Proposed Cuts to the US National Institutes of Health. *JAMA Health Forum*, 6(7):e252228. [Cited on page 8.]

Jia, R., Roberts, M. E., Wang, Y., and Yang, E. (2024). The impact of US–China tensions on US science: Evidence from the NIH investigations. *Proceedings of the National Academy of Sciences*, 121(19):e2301436121. [Cited on page 8.]

Jones, B. F. and Summers, L. H. (2022). A calculation of the social returns to innovation. In Goolsbee, A. and Jones, B. F., editors, *Innovation and Public Policy*, chapter 1, pages 13–59. University of Chicago Press, Chicago. Earlier version: NBER Working Paper 27863, September 2020. [Cited on pages 2, 5, and 22.]

Kantor, S. and Whalley, A. (2025). Moonshot: Public R&D and Growth. *American Economic Review*, 115(9):2891–2925. [Cited on page 7.]

Katz, L. (1953). A new status index derived from sociometric analysis. *Psychometrika*, 18(1):39–43. [Cited on page 47.]

Keller, W. (2002). Geographic localization of international technology diffusion. *American Economic Review*, 92(1):120–142. [Cited on page 7.]

Keller, W. (2004). International technology diffusion. *Journal of Economic Literature*, 42(3):752–782. [Cited on pages 7 and 20.]

Kerr, W. R. (2008). Ethnic scientific communities and international technology diffusion. *The Review of Economics and Statistics*, 90(3):518–537. [Cited on page 7.]

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 pages 3, 9, 10, 11, 17, 18, 28, 37, 38, 53, and 54.]

Kozlov, M., Tollefson, J., and Garisto, D. (2026). US science after a year of Trump: what has been lost and what remains. *Nature*, 649(8098):812–815. [Cited on page 8.]

Krieger, J. L., Schnitzer, M., and Watzinger, M. (2024). Standing on the shoulders of science. *Strategic Management Journal*, 45(9):1670–1695. [Cited on pages 6, 13, and 47.]

Kwon, S. (2024). Underappreciated government research support in patents. *Science*, 385(6712):936–938. [Cited on page 38.]

Li, D. (2017). Expertise versus bias in evaluation: Evidence from the NIH. *American Economic Journal: Applied Economics*, 9(2):60–92. [Cited on page 7.]

Li, D., Azoulay, P., and Sampat, B. N. (2017). The applied value of public investments in biomedical research. *Science*, 356(6333):78–81. [Cited on pages 6, 7, and 18.]

Liu, M., Kadakia, K. T., Patel, V. R., and Krumholz, H. M. (2025). Characterization of Research Grant Terminations at the National Institutes of Health. *JAMA*, 334(6):534–537. [Cited on page 24.]

Lucking, B., Bloom, N., and Van Reenen, J. (2019). Have R&D spillovers declined in the 21st century? *Fiscal Studies*, 40(4):561–590. Earlier circulated as NBER Working Paper 24622 and CEP Discussion Paper 1548 (May 2018) under the title "Have R&D Spillovers Changed?". [Cited on page 5.]

Mansfield, E. (1991). Academic research and industrial innovation. *Research Policy*, 20(1):1–12. [Cited on page 2.]

Mansfield, E., Rapoport, J., Romeo, A., Wagner, S., and Beardsley, G. (1977). Social and private rates of return from industrial innovations. *The Quarterly Journal of Economics*, 91(2):221–240. [Cited on page 5.]

Martin, R., Shah, A., Valero, A., and Verhoeven, D. (2026). Spillovers from Science. CEP Discussion Paper 2165, Centre for Economic Performance, LSE, London. ISSN 2042-2695. [Cited on pages 3, 9, 12, 13, 45, and 47.]

Martin, R. and Verhoeven, D. (2023). Knowledge Spillovers from Clean Innovation: A Tradeoff Between Growth and Climate? CEP Discussion Paper, Centre for Economic Performance, LSE, London. [Cited on page 47.]

Marx, M. and Fuegi, A. (2020). Reliance on science: Worldwide front-page patent citations to scientific articles. *Strategic Management Journal*, 41(9):1572–1594. [Cited on pages 3, 6, 11, and 12.]

Marx, M. and Fuegi, A. (2022). Reliance on science by inventors: Hybrid extraction of in-text patent-to-article citations. *Journal of Economics & Management Strategy*, 31(2):369–392. [Cited on pages 3 and 6.]

Masclans, R., Hasan, S., and Cohen, W. M. (2025). Measuring the Commercial Potential of Science. *Strategic Management Journal*, 46(9):2199–2236. [Cited on page 6.]

Moretti, E., Steinwender, C., and Van Reenen, J. (2025). The Intellectual Spoils of War? Defense R&D, Productivity, and International Spillovers. *Review of Economics and Statistics*, 107(1):14–27. [Cited on pages 7 and 8.]

Munari, F., Righi, H. M., Sobrero, M., Toschi, L., Leonardelli, E., Mainini, S., and Tonelli, S. (2022). Assessing the Influence of ERC-Funded Research on Patented Inventions. Report, European Research Council, Brussels. [Cited on page 6.]

Moser, P., Voena, A., and Waldinger, F. (2014). German Jewish émigrés and US invention. *American Economic Review*, 104(10):3222–3255. [Cited on page 8.]

Myers, K. R. and Lanahan, L. (2022). Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy. *American Economic Review*, 112(7):2393–2423. [Cited on page 7.]

Narin, F., Hamilton, K. S., and Olivastro, D. (1997). The increasing linkage between u.s. technology and public science. *Research Policy*, 26(3):317–330. [Cited on page 6.]

Nelson, R. R. (1959). The Simple Economics of Basic Scientific Research. *Journal of Political Economy*, 67(3):297–306. [Cited on pages 2 and 20.]

Nordhaus, W. D. (1975). The Political Business Cycle. *The Review of Economic Studies*, 42(2):169–190. [Cited on page 22.]

Oliveira, D. F. M., Huang, Q., Woodruff, T. K., and Uzzi, B. (2026). How the 2025 NIH grant terminations varied by researchers' demographic groups. *Proceedings of the National Academy of Sciences*, 123(13):e2527755123. [Cited on pages 8 and 24.]

Palacios-Huerta, I. and Volij, O. (2004). The Measurement of Intellectual Influence. *Econometrica*, 72(3):963–977. [Cited on page 47.]

Pier, E. L., Brauer, M., Filut, A., Kaatz, A., Raclaw, J., Nathan, M. J., Ford, C. E., and Carnes, M. (2018). Low agreement among reviewers evaluating the same NIH grant applications. *Proceedings of the National Academy of Sciences*, 115(12):2952–2957. [Cited on page 7.]

Pinski, G. and Narin, F. (1976). Citation influence for journal aggregates of scientific publications: Theory, with application to the literature of physics. *Information Processing & Management*, 12(5):297–312. [Cited on page 47.]

Poege, F., Harhoff, D., Gaessler, F., and Baruffaldi, S. (2019). Science quality and the value of inventions. *Science Advances*, 5(12):eaay7323. [Cited on pages 6, 13, and 47.]

Priem, J., Piwowar, H., and Orr, R. (2022). OpenAlex: A fully-open index of scholarly works, authors, venues, institutions, and concepts. arXiv preprint arXiv:2205.01833. [Cited on page 11.]

Rogoff, K. (1990). Equilibrium Political Budget Cycles. *American Economic Review*, 80(1):21–36. [Cited on page 22.]

Rosenberg, N. (1990). Why do firms do basic research (with their own money)? *Research Policy*, 19(2):165–174. [Cited on page 6.]

Singh, J. (2005). Collaborative networks as determinants of knowledge diffusion patterns. *Management Science*, 51(5):756–770. [Cited on page 7.]

Stephan, P. E. (2012). *How Economics Shapes Science*. Harvard University Press, Cambridge, MA. [Cited on page 8.]

Thompson, P. and Fox-Kean, M. (2005). Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, 95(1):450–460. [Cited on page 7.]

United for Medical Research (2025). NIH's Role in Sustaining the U.S. Economy: 2025 Update. Report, United for Medical Research. [Cited on page 8.]

Uzzi, B., Mukherjee, S., Stringer, M., and Jones, B. (2013). Atypical Combinations and Scientific Impact. *Science*, 342(6157):468–472. [Cited on pages 3, 15, and 48.]

Veugelers, R., Wang, J., and Stephan, P. (2025). Do funding agencies select and enable risky research: evidence from ERC using novelty as a proxy of risk taking. *Economics of Innovation and New Technology*, pages 1–18. [Cited on page 7.]

Wang, J., Veugelers, R., and Stephan, P. (2017). Bias against novelty in science: A cautionary tale for users of bibliometric indicators. *Research Policy*, 46(8):1416–1436. [Cited on pages 3, 15, and 48.]

Appendix

A. *Science Rank*: the measure in full

[back to main text]

Every node in our network is an *innovation*: a patented invention or a scientific article. Innovation j cites a set of prior innovations — its knowledge stock, of size B_j — and is cited in turn by the set F_j. Following Martin et al. (2026), we model the value of j as Cobb–Douglas in the knowledge it inherits and the effort it adds:

$$V_j \;=\; A_j \, B_j^{\sigma} \, RD_j^{1-\sigma}, \tag{2}$$

where RD_j is j's own research effort, A_j an innovation-specific shifter, and $\sigma \in [0,1]$ the elasticity of value with respect to prior knowledge. A higher σ means inherited knowledge matters more — equivalently, that replacing it would take more own effort.

The spillover an innovation i creates for a downstream innovation j that cites it is the value j loses when i is taken out of its knowledge stock. To a first-order approximation,

$$\Delta V_j \;\approx\; \frac{\partial V_j}{\partial B_j}\,\Delta B_j \;=\; \sigma\,\frac{V_j}{B_j}\,\Delta B_j. \tag{3}$$

Weighting each cited innovation equally makes the knowledge stock proportional to the number of references, $B_j = n_j$, and removing one of them $\Delta B_j = 1$. The spillover from i to j is then $\sigma V_j / n_j$, which is the rule Section B states.

An innovation contributes to everything that cites it, so summing over F_i gives the spillover i creates in total, and adding what i captures privately gives its total value:

$$V_i \;=\; CV_{p,i} \;+\; \sigma \sum_{j \in F_i} \frac{V_j}{n_j}. \tag{4}$$

The object that propagates matters: i is credited a share of the *total* value of each innovation citing it, not merely of the value that innovation captures privately. This is how value travels along chains — an article cited by an article cited by a patent is credited through the middle term. Equation 4 is one linear equation per node. Collecting them gives $V = CV_p + \sigma M V$,

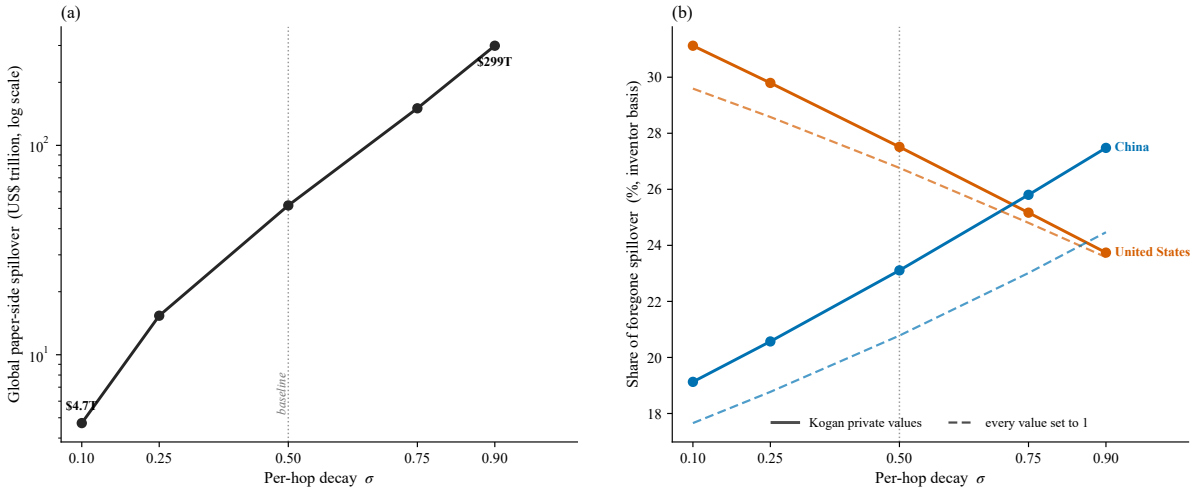
whose solution is $V = (I - \sigma M)^{-1} CV_p = \sum_{k \geq 0} (\sigma M)^k CV_p$. The $k = 0$ term is the private value a node already held; everything above it is the spillover it receives from downstream, which is equation 1. We solve by iteration rather than inversion, starting from CV_p, applying σM repeatedly, and truncating at $K = 24$ hops.[24]

Three inputs have to be supplied: the seed CV_p, the commercial value of every patent (Section A); the citation network from which M and the reference counts n_j are built (Sections B and C); and σ. Only the last is free, and we calibrate rather than estimate it. Figure A1 recomputes the whole measure over a grid of σ: two things move, and one does not. The *level* of attributed value runs from $4.7 trillion at the bottom of the grid to $299 trillion at the top, so every dollar figure in this paper is contingent on the calibration. The *regional split* tilts monotonically: the United States falls from 31.1% of the value reaching science to 23.7%, while China rises from 19.1% to 27.5%, and the two cross just above our baseline. A higher σ weights long citation chains more heavily and Chinese invention reaches science along longer chains, so we do not claim the regional shares are invariant to σ; Section F discusses the movement directly. The ranking the measure induces across articles does not move. Setting every patent's private value to one, which strips the valuation out entirely and leaves a pure citation-network statistic, reproduces 82% of the regional tilt: it is a property of the citation structure, not of what we think patents are worth.

[24] Because n_j counts references that leave the network, M is column-substochastic and the Neumann series converges geometrically at rate σ. At $\sigma = 0.5$ the truncated tail, of order σ^K, is below 10^{-7} of the total.

Figure A1: Sensitivity of the spillover level and its regional shares to the per-hop decay σ

[back to main text]



Notes: The *Science Rank* attribution recomputed over a grid of the per-hop citation decay σ on the same network (Section III); the baseline is $\sigma = 0.5$, discussed in the text above. Panel (a): the global paper-side spillover level, on a log axis, from \$4.7T at the sparsest decay through \$51.7T at the baseline to \$299T at the densest. Panel (b): the regional shares (inventor basis), which tilt monotonically and cross just above the baseline — the United States 31.1% to 23.7% (27.5% at the baseline), China 19.1% to 27.5% (23.1% at the baseline). Dashed: the same shares with every patent's private value set to one, which reproduces 82% of the United-States-to-China tilt, so the tilt is structural to the citation network rather than a product of the valuation.

Our attribution rule is a backward propagation of dollars on a citation network, and its ancestry runs through two families. The first is influence weights and eigenvector centrality (Katz, 1953; Pinski and Narin, 1976; Bonacich, 1987; Brin and Page, 1998) — a document's importance derives from being cited by important documents, the principle behind PageRank — which economists have used to measure intellectual influence (Palacios-Huerta and Volij, 2004) and physicists to find undervalued science (Chen et al., 2007). The second is the innovation-network and technology-proximity approach to spillovers in economics (Jaffe, 1986; Bloom et al., 2013; Acemoglu et al., 2016). From the first family we take the recursion; the difference is that *dollars* propagate rather than influence scores, decaying by a fixed factor at each hop. Guillard et al. (2021) apply that logic to the patent network alone and Martin and Verhoeven (2023) to clean-technology spillovers; Martin et al. (2026) extend it across the patent–paper boundary and introduce *Science Rank*, the measure we adopt. Closest in purpose, though not in method, are the citation-distance measures (Ahmadpoor and Jones, 2017; Poege et al., 2019; Krieger et al., 2024), which locate science by how many citation steps separate it from a patent rather than by how much value flows back along them. Section B sets out why the difference matters and what we add to *Science Rank* itself — above all the funding denominator that turns an attribution into a return.

B. The prediction panel and the backtest

[back to main text]

The modeling table has one row per article per year: 898.9 million rows with 159 predictors, each frozen at what was knowable in that year. Time-varying predictors — citations, spillover-so-far, the realized regional split — use running values truncated at the row's year, never totals that peek ahead. Field- and journal-level patent propensities are estimated leave-one-out from older, fully observed cohorts, so an article never helps predict itself. Static predictors use only pre-publication information.

Most article-years have no spillover at all, and a panel holding every one of them would be an order of magnitude larger than the one we can compute on. We therefore keep *every* article-year with positive spillover and a 5% random sample of those with none, and attach to each row the inverse of its sampling probability — a weight of one for a positive row, twenty for a sampled zero. We sample for tractability, not because the data force it, and the sample is unbiased provided the weights are used: every mean, total, error metric and fitted model in this paper is computed on the weighted panel, so each quantity describes the population of articles rather than the sample drawn from it. Dropping the weights silently discards 95% of the zero stratum, which biases every aggregate downward and, worse, mimics a real phenomenon: recent cohorts, where the zeros are most numerous, look as though their science had stopped generating value.

Four families of predictors provide most of the signal. First, spillover and attention so far: cumulative and recent patent-linked value and scholarly citations, the autoregressive core and the strongest predictor where any history exists. Second, patent proximity: how patent-prone the article's field, topic, and journal are, measured as of the row's year and leave-one-out; patent proximity drives the base rate. Third, novelty: atypical combinations of referenced journals (Uzzi et al., 2013) and new-to-science combinations weighted by their distance (Wang et al., 2017); novel recombination is the kind of science that occasionally seeds an entire patent lineage, so it lifts the predicted tail. Last, authors and geography: the team's prior output, reach, and inventor overlap, and the article's own country mix, the best available predictor of where its spillover will land.

Arranged as publication-year cohorts (rows) against age since publication (columns), the data form a lower-left triangle — old cohorts observed at all ages, young cohorts only young — and the task is to fill the upper right out to age 30. We learn each horizon directly rather than chaining one-year-ahead forecasts. A training pair is an article observed at anchor age a and again at age $a + h$; the features are the article's observed state at a (plus a and h); the

target is cumulative spillover at $a + h$. One pooled model across cohorts, ages, and horizons lets data-starved long horizons borrow shape from data-rich short ones; recency weighting sets how hard we lean on recent cohorts. No prediction is ever fed back into another: the model is indexed by age and horizon rather than calendar year, and every feature is an *observed* anchor state.

Cumulative spillover is zero for most articles and heavily fat-tailed among the rest — the top percentile of articles holds 82.9% of the realized total — and the object we need is the *mean*, since dollars must sum across articles. We therefore fit a gradient-boosted Tweedie regression, whose compound Poisson–Gamma structure matches this shape exactly, and constrain predictions to be monotone in the horizon so no article's projected trajectory decreases.

We model regional composition separately: we predict the global total with the full machinery, then the sixteen regional *shares* (eight regions on two bases, each set summing to one), so that small regions borrow strength from the well-estimated total and the parts add up by construction.

The projection reaches thirty years into the future, so the only way to test it is against the past: stand at a past year, use only what was known then, project forward, and compare with what actually happened by 2022. For pseudo-cutoffs 1997, 2002, 2007, and 2012, we refit the entire pipeline using only information available at the cutoff — between 10.8 and 36.7 million articles per cutoff, at horizons of ten to twenty-five years. Table A1 reports the results; Figure A2 shows the three headline diagnostics.

Table A1: Rolling-origin backtest of the spillover projection

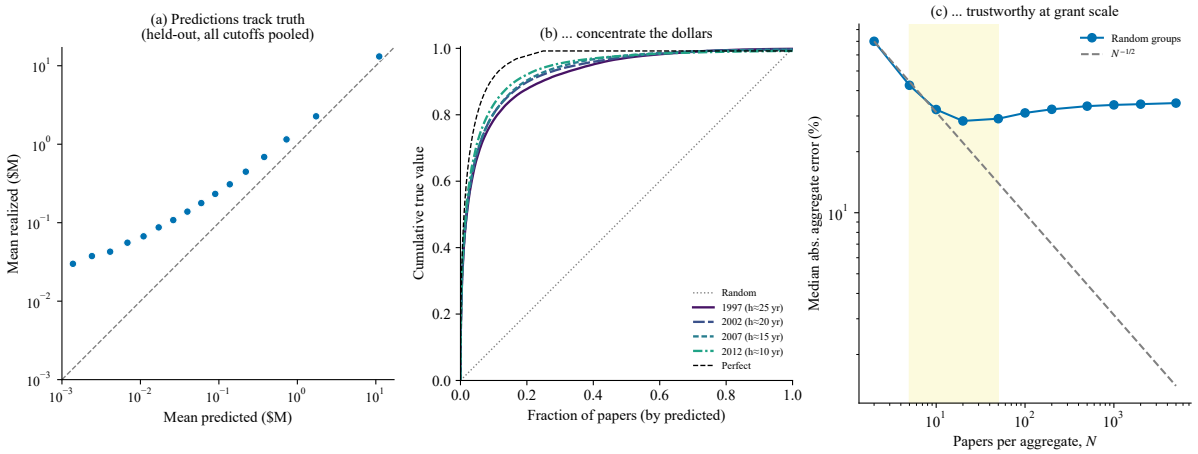
[back to main text]

		Ranking (how much)	Discrimination (whether)		Level bias (% error)		
Pseudo-cutoff	Papers (M)	Spearman ρ	AUC (any spillover)	Top-1% value captured (%)	Aggregate	Young-anchor	Naive
1997	10.8	0.71	0.91	75	−14	−29	−79
2002	17.2	0.74	0.93	79	−3	−11	−69
2007	25.8	0.74	0.94	79	−28	−53	−66
2012	36.7	0.73	0.95	82	−35	−53	−60

Notes: Each row refits the full prediction pipeline using only information available at the pseudo-cutoff year, projects every article forward to the age observable by 2022, and compares predictions to realized spillovers. Spearman ρ ranks articles by predicted versus realized spillover. AUC measures discrimination between articles with and without any realized spillover. Top-1% value captured is the share of total realized dollars found in the articles the model ranks in its top percentile. Aggregate % error is (predicted − realized)/realized on the sum over all articles; young-anchor % error restricts to articles at most five years old at the cutoff. Naive % error uses a forecast that freezes each article's spillover at its cutoff value. *Every metric in this table is sampling-weighted.* The training panel holds all articles with positive spillover and a 5% random sample of the zeros, and each row is weighted by the inverse of its sampling probability; those weights make the metrics describe the article population, not the sample (Section B). The weighting sharpens the model's separation of any-value from zero-value articles and dulls its measured ranking of magnitudes; it leaves the direction and size of the aggregate bias essentially unchanged.

Figure A2: Held-out validation of the spillover projection

[back to main text]



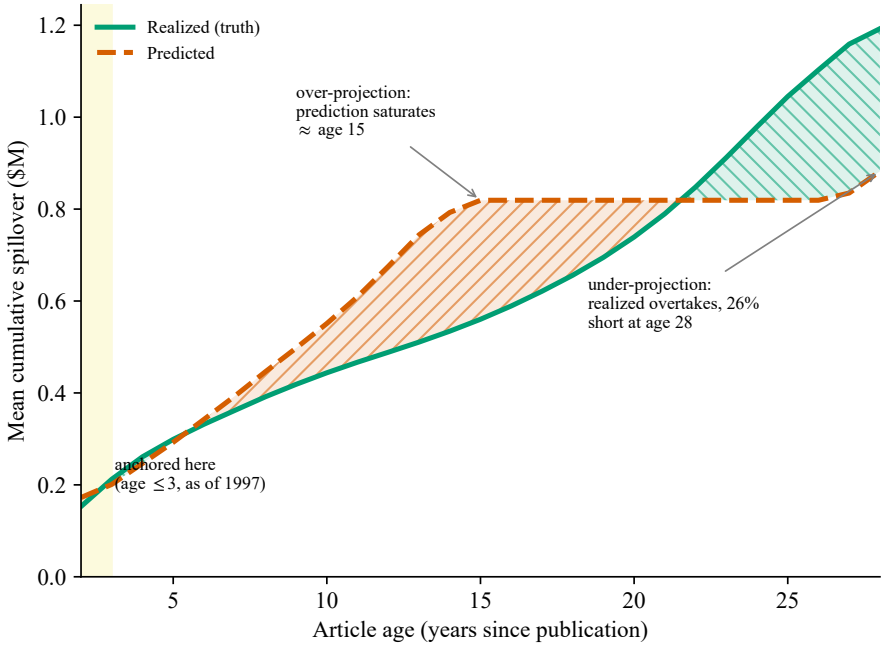
Notes: Held-out diagnostics pooled over the four pseudo-cutoffs of Table A1. The calibration panel plots mean realized against mean predicted spillover by predicted-decile bin, which tracks the 45-degree line across four orders of magnitude. The concentration panel sorts articles by predicted spillover, and the top few percent hold most of the realized value. The aggregation-error panel plots absolute aggregation error against the number of articles per group, alongside the $1/\sqrt{N}$ benchmark: the error falls toward the benchmark as groups grow, so the projection is reliable at the scale of grants and portfolios but not for a single article.

The model discriminates excellently between the articles that will and will not have commercial value: it separates them at an AUC of 0.91–0.95 (the area under the receiver-operating characteristic curve, a discrimination score on which chance is one half and perfect separation is one), and the articles it places in its top percentile hold 75–82% of all the value eventually re-

alized. It ranks magnitudes less well: rank correlations between predicted and realized spillover run 0.71–0.74. We can therefore tell reliably which science matters, but not order the winners among themselves, so we use the model for aggregates over grants and portfolios, never for an individual article. The level is biased low: aggregate spillover falls 3–35% short of realized depending on the cutoff, and 11–53% short for articles young at the cutoff — the regime the terminated grants occupy. The bias is a shape error in the development curve: the model saturates trajectories too early, in the fat tail that dominates aggregates, while the median per-article error is approximately zero. A post-hoc recalibration closes the gap in sample and fails out of it — fitted on the 1997–2007 cutoffs, it moves the held-out 2012 error only from −35% to −30% (Figure A4) — so we do not apply the correction and keep the raw model. Every aggregate level in this paper is therefore a conservative floor.

The backtest validates the global projection but does not separately test its regional split, so we apply no regional correction: the shares of Section F are reported as measured.

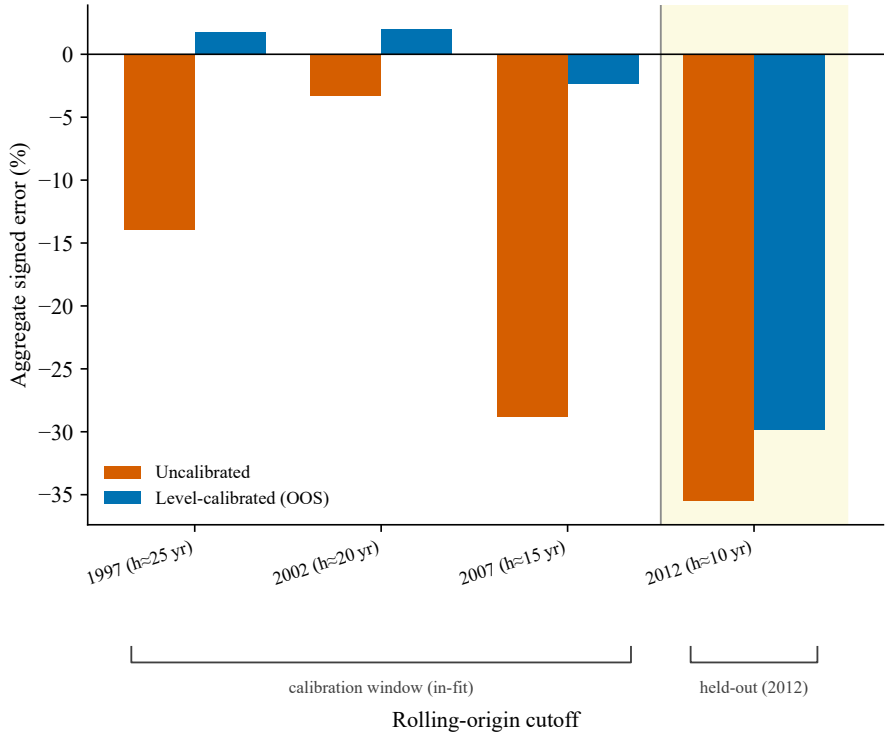
Figure A3: The empirical development curve of article-level spillovers



[back to main text]

Notes: Mean predicted and realized cumulative spillover by article age in the 1997 pseudo-cutoff backtest. The realized profile is the development curve used to distribute 30-year spillovers over calendar years for discounting (Section B): half of the 30-year total accrues only after 18 years. The predicted profile saturates earlier than the realized one — the development-curve shape error behind the aggregate under-prediction in Table A1.

Figure A4: Out-of-sample failure of post-hoc level calibration



[back to main text]

Notes: Aggregate percentage error before and after an age-and-horizon level recalibration fitted on the 1997–2007 pseudo-cutoffs. In-fit cutoffs calibrate to approximately zero by construction; the held-out 2012 cutoff improves only from −35% to −30%. Because the correction does not generalize to the young-article regime the terminated grants occupy, the paper reports raw (uncalibrated) projections and treats aggregate levels as conservative floors.

C. Additional exhibits

[back to main text]

C.1. The headline ratio: the denominator and the accounting variants

[back to main text]

The headline ratio is $8.69 of foregone spillover per dollar of terminated funding (Table 3). One force could move it materially: the grant-to-article linkage has densified over the donor-to-terminated vintage gap, which lifts measured spillover per dollar as a coverage artifact; Section C bounds that drift, and treating it as entirely artifactual lowers the headline to $6.21. The remaining accounting choices move it far less. Re-drawing the funding bands, and restricting the donor pool to its oldest or youngest cohorts, each shift the peer-pool ratio by well under a tenth (Table A2), and the same-investigator donor arm reads about 30% higher. Measured against the budget the terminations actually withhold, rather than the lifetime value of the awards they cancel, the ratio rises to $30.18 per dollar — a different question from the one the headline answers, reported as such (Section C).

The higher-return agency is the NSF. Each terminated NSF dollar forgoes $11.20 of spillover against $7.48 at the NIH, a factor of 1.50. The ordering holds only when both agencies are measured against a comparable denominator: dividing instead by the terminated *award* would deflate the NIH denominator and reverse it, the segment-versus-whole wedge of Section B. The wedge is visible directly in the data. Among the terminated grants for which the award and the lifetime-funding concepts coincide, the two ratios coincide too (10.60 against 10.64, 1,752 grants); among those for which they diverge, the ratios diverge by a factor of 2.37 (469 grants). The wedge is mechanical rather than economic: with a consistent denominator, the NSF ranks first.

Measuring both agencies against the same object — the lifetime project funding recorded in Dimensions — puts this counterfactual into agreement with Section A. That section, running on the entire population of funded grants by an independent route and for a different purpose, reaches NSF $20.18 against NIH $8.60 among linked grants ($9.81 against $7.48 counting every grant). Two pipelines, one on 2,221 terminated grants and one on more than a million funded ones, agree on the ordering and roughly on the magnitude. The gap reflects price rather than scientific merit: the median NIH award is $905.4K against $307.1K at the NSF, and although an NSF- and an NIH-funded article have comparable commercial value, the NIH pays about 2.50 times as much grant money per resulting paper. Both ratios are far above one under every accounting variant we report (Table A2), so the qualitative conclusion — the foregone spillover exceeds the budget saved, several times over — does not depend on any single accounting choice. It does depend on the Kogan et al. (2017) value calibration, with which all dollar levels scale proportionally.

Table A2: Sensitivity of the headline ratio to accounting choices

[back to main text]

Variant	Per $
Headline (Arm B, equal split, coherent $; 5.61 at PV@3%)	8.69
Selection bracket	
Arm A (same investigator)	11.31
Credit rule	
Full credit (upper bound; double-counts)	24.57
Donor window	
Oldest donors (2001–2007)	9.21
Youngest donors (2008–2013)	8.18
Youngest three cohorts (2011–2013)	7.66
Denominator	
Per $ of the terminated award	12.33
Per $ withheld	30.18

Notes: Variants of the headline dollars-per-dollar ratio of Table 3. *Selection bracket*: the same-investigator donor arm, which reads high (Section C). *Credit rule*: crediting each article's spillover in full to every grant that supported it, rather than splitting it equally, double-counts articles with more than one funder and is reported as an upper bound only. *Donor window*: restricting the donor pool to its oldest or youngest cohorts. *Denominator*: the last two rows answer a different question rather than testing the same one — dividing by the value of the terminated award, or by the budget not yet outlaid at termination, rather than by the lifetime project funding that bought the science. No transport correction is applied: the cross-vintage gradient an earlier draft corrected for was an artifact of the measurement window and is withdrawn (Section B). All dollar levels scale proportionally with the Kogan et al. (2017) calibration; the per-dollar comparisons and shares do not.

C.2. Who held the terminated grants

The terminated grants were held by established, productive investigators, and were the smaller of the grants those investigators held (Section D). Table A3 reports the investigators' prior publication records; Figure A5 sets the size of the terminated grants against their matched donors.

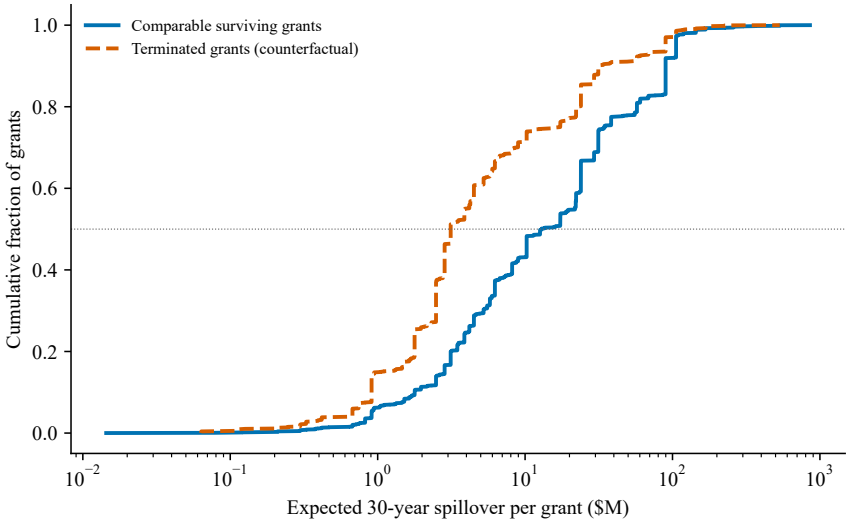
Table A3: Who held the terminated grants

[back to main text]

Principal investigators of the terminated grants	*N* (resolved)	Prior publications	Career (years)	Publications per year
Research grants	1,742	54	22	2.4
Career-development	71	24	12	2.0
Fellowships	129	5	4	1.0
All resolved	1,942	47	21	2.3

Notes: The principal investigators of the terminated grants, measured strictly as of the grant year from each investigator's publication record, so that nothing published after the award enters. The principal investigator is resolved to a publication record for 87% of terminated grants; the unresolved remainder are disproportionately fellowship holders, whose thin records are genuine rather than a resolution failure. A surviving-grant benchmark is available only for donors whose investigator also resolves, and resolution correlates with productivity strongly enough that the resolvable slice of the donor pool returns about twice the spillover per dollar of the full pool — so a terminated-versus-surviving contrast drawn here would be measuring who is findable in a bibliographic database. The same-vintage comparison in Section D, against contemporaries resolved identically on both sides, does not have that defect. All entries are medians across the resolved investigators in each row.

Figure A5: The terminations removed the smaller grants



[back to main text]

Notes: Cumulative distributions of the expected 30-year spillover *of a grant*, in millions of 2010 dollars, with each grant placed at the mean of its agency × field × funding-band cell on both sides: the terminated grants (matched-donor counterfactual) and the comparable surviving grants that serve as their donors. The horizontal axis is a level per grant, not a rate: the terminated distribution sits to the left because the terminated grants are *smaller*. The per-dollar comparison, once size is divided out, is the 38th percentile reported in Section D, alongside the peer comparison of Table 4. *Source*: authors' calculations on the matched treated-and-donor grant panel.

C.3. The field partition

Table A4 reports the full field partition of the foregone spillover; Section E discusses the three patent-proximate divisions that dominate it.

Table A4: Foregone spillover by field: full partition

[back to main text]

Field (ANZSRC FOR-2020 division)	Terminated grants	Foregone ($M)	Share (%)
Biomedical and clinical sciences	426	7,779	29.8
Information and computing sciences	212	6,170	23.6
Biological sciences	218	4,430	17.0
Education	742	2,220	8.5
Health sciences	157	1,451	5.6
Engineering	36	1,030	3.9
Chemical sciences	15	919	3.5
Commerce, management and services	66	582	2.2
Psychology	59	339	1.3
Built environment and design	12	250	1.0
Creative arts and writing	23	153	0.6
Human society	104	148	0.6
Physical sciences	3	128	0.5
Environmental sciences	14	116	0.4
Language, communication and culture	43	98	0.4
Economics	22	73	0.3
Philosophy and religious studies	15	70	0.3
Law and legal studies	26	59	0.2
Agricultural, veterinary and food sciences	4	34	0.1
Earth sciences	9	31	0.1
History, heritage and archaeology	12	18	0.1
Mathematical sciences	2	3	0.0
Field not recorded	1	13	0.0
Total	2,221	26,112	100.0

Notes: The foregone 30-year spillover of Table 3 partitioned by research field (ANZSRC 2020 two-digit divisions, from the Dimensions classification of each grant). We assign grants spanning several divisions to their primary division, and give the single grant Dimensions leaves unclassified its own row rather than dropping it, so the partition matches Table 3 up to rounding. Rows are rounded to the nearest $M and so need not sum to the displayed total; the grant counts and shares reconcile exactly.

C.4. Break-even, by route and by field

The break-even calculations of Section B are reported in full here: by the route the value travels (Table A5) and by research field (Table A6).

Table A5: Payback time of a research dollar, overall and by route

	% of funding recovered	Break-even (years)	Half-life (years)
Panel A. Overall (NIH+NSF, cohorts 1990–2002, each observed ≥20 years)			
Undiscounted	339.7%	9.8	
Discounted at 3%	232.6%	11.3	
Linked grants only[a]	400.8%	8.6	
Highest-coverage cohorts[b]		10.7	
Panel B. By route (undiscounted; half-life is on the paper's own clock)			
Direct	163.7%	13.1	10.9
Science chain	103.8%	19.6	18.6
Patent chain	47.9%	not reached	19.3
Mixed	24.4%	not reached	23.7
All four		*9.8*	*16.8*
Panel C. The budget window (cohorts 1990–92, observed a full 30 years)			
Realized inside 3 years	2.4%		
Realized inside 5 years	7.4%		
Realized in 30 years	557.2%	7.9	20.5

Notes: Cumulative spillover realized through 2022 — never projected. Cohort curves are aggregate ratios (total spillover ÷ total funding), not means of per-grant ratios. Funding in constant 2010 USD; spillover credited to a grant at an equal split across the grant(s) supporting each paper. The half-life in Panel B is on the *paper's* clock (years after publication); the one in Panel C is on the *grant's* (years after the award starts), and the gap between them is the publication lag. [a] The denominator is the funding of grants with at least one traceable paper, rather than all funding; the gap between the two is linkage coverage, not economics. [b] Cohorts 2000–02, the highest-coverage cohorts we can still observe for 20 years; every estimate here is an upper bound on the payback *time*, because a grant whose papers we cannot see contributes funding but no spillover. **The cross-cohort comparison is not identified and no trend is claimed**: linkage coverage doubles across cohorts and the patent-value pool itself grew four- to six-fold, and holding that pool fixed reverses the sign of the apparent gradient.

[back to main text]

Table A6: Break-even by field (NIH+NSF, cohorts 1990–2002, undiscounted)

		All four routes		Direct route only[a]	
Field (ANZSRC division)	Grants	Cum. value (% of cost)	Break-even (yrs)	Cum. value (% of cost)	Break-even (yrs)
Biomedical and Clinical Sciences	50,262	320%	9.6	171%	12.6
Biological Sciences	45,731	480%	7.4	234%	9.5
Information and Computing Sciences	22,647	378%	11.3	132%	16.0
Engineering	21,947	770%	6.8	280%	8.2
Earth Sciences	17,039	42%	not reached	17%	not reached
Chemical Sciences	11,015	1123%	4.6	509%	5.2
Education	9,749	32%	not reached	14%	not reached
Physical Sciences	7,083	112%	19.0	44%	not reached
Health Sciences	6,533	65%	not reached	34%	not reached
Mathematical Sciences	6,134	85%	not reached	32%	not reached
Psychology	3,495	81%	not reached	41%	not reached
Agricultural, Veterinary and Food Sciences	3,333	311%	10.8	141%	14.7
Human Society	2,683	33%	not reached	18%	not reached
Economics	2,073	65%	not reached	36%	not reached
Commerce, Management, Tourism and Services	1,934	41%	not reached	17%	not reached
Philosophy and Religious Studies	1,784	120%	18.3	54%	not reached
History, Heritage and Archaeology	1,723	89%	not reached	40%	not reached
Language, Communication and Culture	1,705	34%	not reached	16%	not reached
Built Environment and Design	1,670	250%	12.7	110%	18.5
Law and Legal Studies	1,106	44%	not reached	20%	not reached
Environmental Sciences	1,097	155%	16.9	76%	not reached
Creative Arts and Writing	506	86%	not reached	39%	not reached
Unclassified	27,687	278%	8.9	143%	12.9

Notes: Grant-side ANZSRC Field-of-Research divisions — the grants' native taxonomy, so no crosswalk is imposed. *Unclassified* is the 22% of grants Dimensions leaves uncoded; it is not a field. [a] The conservative variant: value reaching the paper only because a patent cites it directly, discarding every indirect route. **This table measures commercial value and nothing else.** A field that does not break even here has not been shown to be without value; it has been shown to deliver little of its value *through the patent channel*, which is the only channel this paper prices — and which is precisely the wrong instrument for reading the worth of, say, history or law.

Table A6 reports the break-even calculation of Section B field by field, and it is easily misread. The fields closest to the patent frontier repay their funding quickly — chemical sciences in under five years, engineering in under seven — while several do not repay it at all inside the twenty years we can observe. Among those are education, human society, law, language, and economics — and, more surprisingly, mathematics and health sciences.

A low number in this table does not show that a field is without value; it shows that little of the field's value travels through the patent channel, the one channel this paper prices. The patent channel is a reasonable measure of the commercial value of chemistry. It is a poor measure of the value of history, law, or education research, whose returns run through human capital, institutions, public health, and practice, none of which cite a patent. For such fields the measure is silent, and we report the field split so that these silences are explicit.

C.5. Funded versus unfunded science

The funded-versus-unfunded contrast of Section C is confounded by the coverage of grant acknowledgements, not identified. Table A7 reports the contrast and Figure A6 the confound.

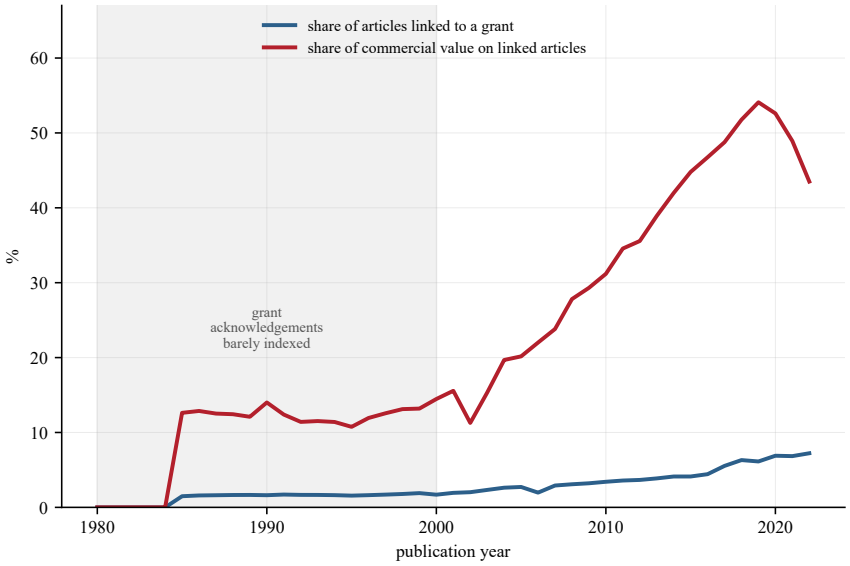
Table A7: Realized commercial value of science, funded versus unfunded

[back to main text]

		N articles	Mean value ($M)	% with positive value	Funded/unfunded ratio
All articles	funded	12.3M	0.799	82.7	8.4×
	unfunded	440M	0.095	15.5	
High-coverage stratum	funded	2.03M	1.158	96.7	3.2×
	unfunded	5.13M	0.367	60.1	

Notes: Realized commercial value of science on grant-linked versus unlinked articles. *Unfunded* means not linked to any grant we observe, so the raw contrast is confounded by linkage coverage, which collapses before 2000 and outside acknowledgement-indexing funders. We therefore report a high-coverage stratum — US-authored articles published 2005–2018 in dense-coverage fields — alongside all articles, and disclose the confound directly in Figure A6. The gap falls from 8.4 to 3.2 times from the all-article to the stratum panel, and even the stratum figure is a correlational lower bound: it removes neither the coverage confound nor selection into funding, so it is not a causal effect of funding. N is the population-scaled article count — each sampled article is scaled by its sampling weight back to the full population (Section B) — not a raw count of articles. No test is implied; the ratio is a raw ratio of means.

Figure A6: Funded versus unfunded is confounded by coverage, not identified



[back to main text]

Notes: Grant-acknowledgement indexing, not a change in funding, drives most of the pre-2000 rise. The figure plots the share of articles linked to a grant (a count share) and the share of commercial value on linked articles (a value share), by publication year; acknowledgements are barely indexed before 2000. This is the confound behind the funded-versus-unfunded contrast of Table A7. Commercial value accrues with a lag, so the most recent publication years are right-truncated: the decline in the value share after 2019 reflects value not yet realized, not a fall in the grant-linked share.