



NIH funding and the pursuit of edge science

Mikko Packalen^a and Jay Bhattacharya^{b,1} 

^aDepartment of Economics, University of Waterloo, Waterloo, ON, Canada N2L 3G1; and ^bCenter for Health Policy/Primary Care and Outcomes Research, Department of Medicine, Stanford University, Stanford, CA 94305

Edited by James M. Poterba, Massachusetts Institute of Technology, Cambridge, MA, and approved April 6, 2020 (received for review June 13, 2019)

The National Institutes of Health (NIH) plays a critical role in funding scientific endeavors in biomedicine. Funding innovative science is an essential element of the NIH's mission, but many have questioned the NIH's ability to fulfill this aim. Based on an analysis of a comprehensive corpus of published biomedical research articles, we measure whether the NIH succeeds in funding work with novel ideas, which we term edge science. We find that edge science is more often NIH funded than less novel science, but with a delay. Papers that build on very recent ideas are NIH funded less often than are papers that build on ideas that have had a chance to mature for at least 7 y. We have three further findings. First, the tendency to fund edge science is mostly limited to basic science. Papers that build on novel clinical ideas are not more often NIH funded than are papers that build on well-established clinical knowledge. Second, novel papers tend to be NIH funded more often because there are more NIH-funded papers in innovative areas of investigation, rather than because the NIH funds innovative papers within research areas. Third, the NIH's tendency to have funded papers that build on the most recent advances has declined over time. In this regard, NIH funding has become more conservative despite initiatives to increase funding for innovative projects. Given our focus on published papers, the results reflect both the funding preferences of the NIH and the composition of the applications it receives.

science | novelty | biomedicine | funding | ideas

With an annual budget of more than \$37 billion, the National Institutes of Health (NIH) funds the work of 300,000 scientists across the globe, and seeks to improve health outcomes by facilitating “fundamental creative discoveries, innovative research strategies, and their applications” (1, 2). As such, the NIH plays a pivotal role in setting the incentives that biomedical scientists have to try out novel ideas in their work, a vital aspect of fruitful scientific investigation (3, 4). The regular exercise of trying out new ideas, which we term edge science, is a crucial contributor to the advance of scientific disciplines (5–8). For scientists, however, edge science is inherently risky, since it can be difficult to predict whether a new idea will produce fruitful results. Working scientists often point to failure as a precursor to success, but there is no guarantee that a particular idea will work (9–12). Public support for edge science, even if failure is likely, can help establish appropriate incentives for novel work in biomedical science. However, when setting its priorities, the NIH considers many factors, including scientific opportunity, disease burden, and availability of private funding (13, 14), and many have questioned the NIH's ability to fund groundbreaking work and innovative science in particular (15–22). Both scientific and political considerations may lead the NIH to underfund the trying out of new ideas. First, because the NIH visibly spends public money, it needs to show discrete manifestations of improvements in health, as well as technological breakthroughs, arising from its supported research. This consideration can lead to a preference to support ideas that have already shown promise, rather than edge science. Second, NIH scientific review panels, for reasons related to their constitution, tend to reward projects that are evidently feasible over novel projects. Despite the importance of NIH support of edge science, there has been little research quantifying its extent. Prior

research has examined the effect of funding decisions on the productivity of scientists, as measured by grants, publications, citation counts, and patents (23–32). Although these outputs are important, they are not quantifiable measures of edge science. In this article, we provide a quantitative assessment of the extent to which NIH policies encourage or impede edge science. Our strategy is to focus on research papers, a key product of NIH funding, rather than grant applications, which are inputs to this end goal. We first determine which research papers try out relatively new ideas, based on a text analysis of a comprehensive corpus of biomedical publications. We then compare the frequency of NIH funding for contributions that represent edge science against the frequency of NIH funding for contributions that represent more conventional science.

Our analysis covers more than 24 million research articles in the MEDLINE database on biomedical research papers. To determine the ideas upon which each paper builds, which we refer to as idea inputs, we employ the Unified Medical Language System (UMLS) metathesaurus, a curated vocabulary of more than 5 million biomedical terms. We determine which UMLS terms appear in the abstract or title of each paper in the MEDLINE database. This reveals a list of idea inputs for each paper. We then determine the vintage of each term based on the year that the term first appears in the MEDLINE database; we refer to this year of first appearance as the cohort of the idea. We classify each term into one of 127 idea types based on the UMLS semantic category for each term. These categories include “Gene and Genome,” “Neoplastic Process,” and “Quantitative Concept.” We determine the research area of each paper (synonymous here with a scientific field) based on the journal in which the paper was published, using the National Library of Medicine's (NLM's) subject category for each journal. There are 125 journal categories, including “Cardiology,” “General Surgery,” and “Molecular Biology.” Delineating idea types and research areas allows us to characterize the novelty–NIH funding link in different settings. While the delineations we

Significance

Innovative research is critical to the advancement of biomedicine. The NIH plays a crucial role in fostering innovation. An empirical assessment of the success of the NIH in this role is an essential step in identifying ways to encourage novel research. This research introduces a measure of the age of the ideas used in published biomedical research papers and comprehensively evaluates the performance of NIH-supported research relative to non-NIH-supported research, using this measure.

Author contributions: M.P. and J.B. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: All data for the paper are available in the Harvard Dataverse repository (<https://doi.org/10.7910/DVN/QT5OGS>).

¹To whom correspondence may be addressed. Email: jay@stanford.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1910160117/-DCSupplemental>.

First published May 19, 2020.

use have some overlap, the entries in the UMLS ontology provide a reasonable proxy for what type of idea each term represents, and the entries in the NLM ontology provide a reasonable proxy for research areas.

A paper is linked to a particular idea type if it mentions any term of that idea type. At the same time, a paper is linked to a particular research area based on the category of the journal in which it was published. A paper contributes to the scientific discussion regarding the function of that idea type within that research area. Accordingly, we define a contribution as a link from a paper to an (idea type, research area) pair. We count mentions of terms of the same idea type in a paper as one contribution, and we count a paper that mentions ideas from multiple idea types as multiple separate contributions. We adopt a contribution-level approach because it allows us to more fully characterize how a scientific paper contributes to edge science by permitting each paper to advance knowledge on multiple dimensions. On average, a paper published in the 2010s is linked to 7.03 idea types, 1.45 research areas, and 10.28 (idea type, research area) pairs, and is thus counted as 10.28 contributions. With this approach, we can also perform statistical calculations that isolate specific idea type and research areas when comparing NIH funding status and novelty of contributions across different settings. For each contribution, we determine the cohort of the newest idea input of that specific idea type in the paper and refer to this cohort as idea input vintage. We consider contributions that build on at least one relatively recent idea to be novel. By elaborating on a new idea, such work potentially enhances scientists' understanding of the idea's usefulness within a field, and thus represents edge science. In contrast, contributions that build only on well-established ideas represent more traditional science.

Having determined the idea input vintage for each contribution, we then examine how the share of contributions that are in papers with NIH funding varies by idea input vintage. We use paper-level funding acknowledgments information in MEDLINE to construct an indicator variable capturing whether a paper has NIH funding. We limit our analysis to papers for which the (first) author has a US affiliation. This restriction is important because scientists working in the United States are more likely to build on new ideas than the average in biomedicine (33), and because the NIH disproportionately funds US scientists. Hence, if we were to include non-US papers in the analysis, the results would, in part, reflect the relative scientific frontier position of US scientists, as well as the NIH's greater propensity to fund American scientists. By focusing on US papers, our results will reflect a comparison of the vintage of ideas that NIH-funded scientists tend to build on in their work and the vintage of the ideas that comparable scientists without NIH funding tend to build on in their work. We focus on the NIH, as it is by far the largest funder of biomedical work, and as such, plays a large role in setting research priorities in biomedicine. In MEDLINE data, among US research papers published in the 2010s, 46% report NIH funding (of which 98% report extramural support for researchers not directly employed by the NIH), and less than 6% report non-NIH funding, mostly from private foundations and other government agencies. Biomedical research papers are also produced by academic researchers with funding from their university, and by researchers working in pharmaceutical, biotechnology, and other companies, as well as foundations and hospitals.

It bears repeating that we analyze publications rather than NIH grant proposals. An advantage of publications data is that they give information on actually completed projects with NIH funding. This distinction matters because of a potential disconnect between funding proposals and actual research. Regarding the social benefits of science, the work scientists actually do is more important than the work scientists promise to do when applying for funding. The focus on published papers also implies that our results reflect not only the funding preferences of the

NIH and its review committees but also the composition of applications it receives, a distinction that is important when interpreting our results. Our approach is also observational rather than a randomized experiment. The findings point to a need for complementary randomized studies on how changes in NIH funding mechanisms and review structures affect the pursuit of edge science in biomedicine, in terms of novelty of publications and novelty of funded versus unfunded grant applications.

Results

We first examine how the share of NIH-funded contributions varies by idea input vintage. Recall that idea input vintage captures the year that the newest idea that the contribution builds on was introduced in the literature. Fig. 1 shows the overall results for contributions across all idea types and research areas for contributions published between 2010 and 2016. The figure depicts the fraction of contributions that are in papers with NIH funding (the vertical axis) as a function of the idea input vintage (the horizontal axis). The horizontal dashed line in Fig. 1 represents the average share of contributions funded by the NIH over all cohorts.

The results in Fig. 1 suggest an inverted U shape to the relationship between NIH funding and the novelty of idea inputs. The share of research supported by the NIH is the highest for papers that build on ideas that are relatively new but not too recent; there appears to be a substantial time lag for funding work on a new idea. The NIH funding rate is lower for contributions that build on the most recent ideas (post-2005 cohorts) than it is for contributions that build on a bit more mature ideas (1990 to 2005 cohorts). The share of NIH support is low also for contributions that only reference well-established knowledge (i.e., ideas introduced to the literature a long time ago; pre-1970 cohorts). Specifically, contributions that build on very recent ideas are less often NIH funded than contributions that build on ideas that have had a chance to mature for at least 7 to 10 y. The magnitude of these idea input vintage-related NIH funding differences is also considerable: the NIH funding rate is over 60% for contributions that make use of 10- to 25-y-old ideas, whereas

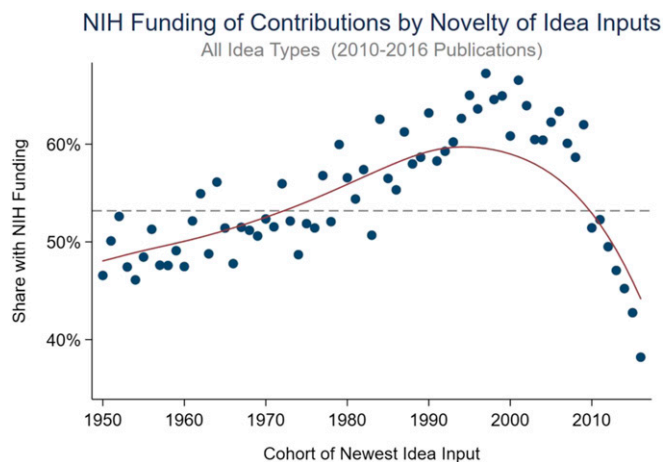


Fig. 1. Share of NIH funding by novelty of idea inputs (2010 to 2016): all idea types. Calculated based on 992,633 biomedical research papers published during 2010 to 2016. The horizontal axis captures the idea input vintage (the cohort of a contribution is the year when the newest idea input used in the contribution was introduced to the literature). Later (earlier) cohort years represent more novel (more conventional) science. The vertical axis captures the rate of NIH funding. The markers capture the mean NIH funding rate for each idea input vintage. The solid line represents a non-parametric regression line estimate. The dashed line represents the average funding rate across all cohorts.

it is only about 50% for contributions that either build on only well-established knowledge or build on some very recent ideas.

Fig. 2 shows the same relationship for a cohort of contributions published during 2 earlier decades, namely, the 1990s (Fig. 2, *Top Left*) and the 2000s (Fig. 2, *Top Right*), as well as for two subperiods of the current decade, namely, 2010 to 2014 (Fig. 2, *Bottom Left*) and 2015 to 2016 (Fig. 2, *Bottom Right*). The results in Fig. 2 show that even in earlier decades, the NIH funded novel science at higher rates than it funded traditional science. However, unlike in the current decade, in the 1990s, the funding rate for work on very recent ideas was not yet lower than the funding rate for work on well-established ideas. The tendency to fund the most novel work at a lower rate appears to be a recent phenomenon. *SI Appendix, Table S7* shows that the secular decline in funding rates for work that builds on very recent ideas is also statistically significant ($P < 0.001$).

We next examine variation in the novelty–NIH funding link across idea types. For this analysis, we first classify each of the 127 idea types to one of three broad idea type groups: “Basic Science,” “Clinical,” and “Miscellaneous.” In Fig. 3, the top left, top right, and bottom left panels show the funding rates by idea input vintage when included contributions are limited to each of these three broad idea type groups. For ease of comparison, the bottom left panel of Fig. 3 shows the overall result again. The results show that among contributions that build on basic science ideas, the NIH funds edge science considerably more often than it funds more traditional science, although the pattern again follows an inverted U-shape. The U-shape pattern is present also for clinical ideas, but the rising part of the link is very weak. *SI Appendix, Table S1* summarizes results for individual idea types. The positive novelty–NIH funding link is generally present for basic science ideas, but not for clinical ideas. For example, work that references genes or proteins is much more likely to be NIH funded when it mentions newer gene or protein ideas relative to work that only mentions gene or protein ideas first introduced to the literature a long time ago. In contrast, contributions that

reference new ideas in the neoplastic process category have a lower share with NIH funding than contributions that rely on older ideas in that category, and contributions that mention a new drug have a much lower share of NIH funding than contributions that only mention older drugs.

The NIH thus tends to fund novel published work more, but only for some idea types, not for others. There is also substantial variation in the average NIH funding rate across idea types, as can be seen from Fig. 1 and *SI Appendix, Table S3*. The aggregate novelty–NIH link might thus be driven by the NIH’s tendency to fund certain idea types where a lot of progress is taking place, rather than by the NIH’s tendency to disproportionately fund novel work within idea types. Perhaps, for example, the overall result is mainly driven by the fact that the NIH has funded at a high rate any research that makes use of either novel or well-established advances in genomics.

To investigate this possibility, we calculate NIH funding rate by idea input vintage, while holding the average funding rate the same across all (idea type, research area) pairs. In this calculation, we thus allow the funding rate to vary only across cohorts within each (idea type, research area) pair. That is, our procedure holds fixed variation in NIH funding rates due to the idea type and research area of a contribution, while continuing to reflect variation in funding due to differences in the age of idea inputs. Fig. 4 reports the results of this calculation.

Fig. 4 still indicates an inverted U-shaped relationship between the novelty of idea inputs and NIH funding, but the results are quantitatively different. The highest NIH funding rate for novel work is now only a little above the funding rate for work that builds on well-established ideas, and the funding rate for work that builds on the most recent ideas is now markedly below the average funding rate across all cohorts. Together with the unadjusted results shown in Fig. 1, this result suggests that while the NIH has been successful in funding innovative science, this success has been mostly due to its differential support for certain idea types where progress has been faster in recent decades, in

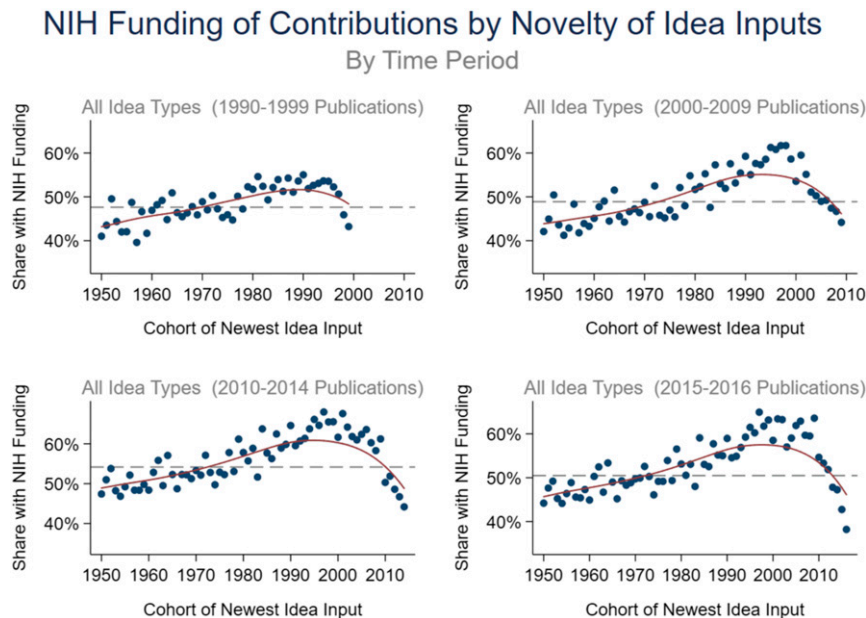


Fig. 2. Share of NIH funding by novelty of idea inputs: for 1990s (*Top Left*), 2000s (*Top Right*), 2010 to 2014 (*Bottom Left*), and 2015 to 2016 (*Bottom Right*). Results in these panels are calculated based on 763,079, 1,155,199, 735,460, and 257,173 biomedical research papers published during 1990 to 1999, 2000 to 2009, 2010 to 2014, and 2015 to 2016, respectively. In each panel, the horizontal axis captures the idea input vintage, with later (earlier) cohort years representing more novel (more conventional) science, and the vertical axis captures the rate of NIH funding. The markers capture the mean NIH funding rate for each idea input vintage. The solid line represents a nonparametric regression line estimate. The dashed line represents the average funding rate across all cohorts.

NIH Funding of Contributions by Novelty of Idea Inputs By Idea Type Group (2010–2016 Publications)

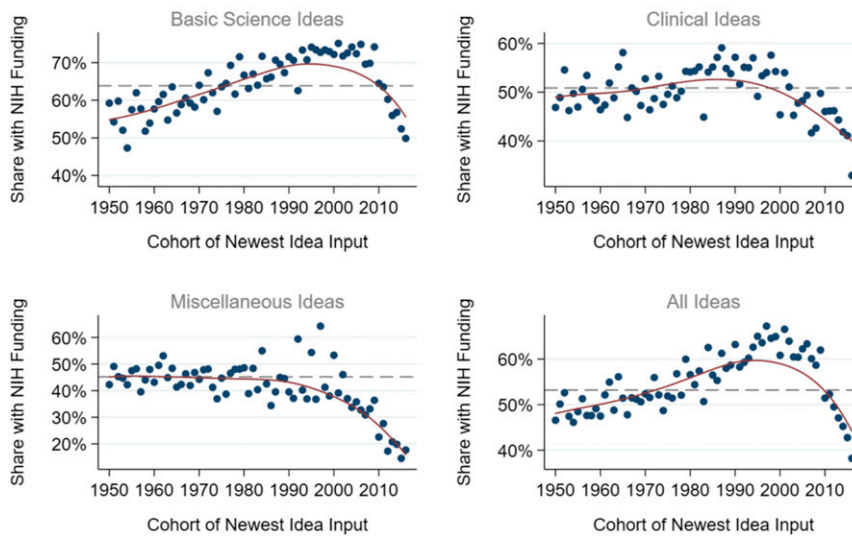


Fig. 3. Share of NIH funding by novelty of idea inputs by idea type group: basic science (*Top Left*), clinical (*Top Right*), miscellaneous (*Bottom Left*), and all idea types (*Bottom Right*). Calculated based on 992,633 biomedical research papers published during 2010 to 2016. In each panel, the horizontal axis captures the idea input vintage, with later (earlier) cohort years representing more novel (more conventional) science, and the vertical axis captures the rate of NIH funding. The markers capture the mean NIH funding rate for every idea input vintage. The solid line represents a nonparametric regression line estimate. The dashed line represents the average funding rate across all cohorts.

the sense that the ideas used as building blocks in research tend to be relatively novel (such as genomics). Conversely, the results imply that on average, within research areas and idea types, the NIH funds have not gone to the most novel projects.

Discussion

Scientific progress depends on the openness of a scientific community to try out new ideas systematically and continuously (5–8, 34–36). When new ideas are first born, they are often raw and poorly understood and only develop into transformative ideas (if at all) through revision and debate on their merits. One of the primary goals of the NIH is to support innovative biomedical research, including work that tries out new ideas. However, identifying novel projects that are likely to be successful is difficult. Facilitating novel work is also often in tension with other factors, including public health needs, that the NIH considers when setting its priorities. The NIH faces pressure to deliver short-term successes that can be at odds with a systematic commitment to edge science. A particular temptation is to focus funding priorities based on the likelihood of producing high-impact publications (as measured by citations) without regard for whether the work represents novel or more traditional science (15–22). But a focus on high-impact science alone runs the risk of undermining more explorative work that often fails, in the sense that it is not rewarded with many citations, but which lays the foundation for breakthroughs that arrive later as the ideas mature. From this perspective, public funding agencies should strive to fund both edge science and high-impact science. The NIH's aim and initiatives to increase funding for innovative science signal that it holds this view.

Our analysis of published biomedical research articles finds qualified support for the idea that the NIH supports innovative work. Despite this, our findings suggest that there is substantial room for the NIH to do more to promote edge science. We find that in the current funding environment, there is a time lag for NIH funding of contributions that rely on new ideas. The NIH funds work that builds on very recent ideas at a lower rate than

work that builds on ideas that have had a chance to mature for at least 7 y. Moreover, contributions that rely on very recent ideas are NIH funded less often even than work on mature ideas. This pattern has not always been true of NIH-funded work. In the 1990s, contributions that tried out very recent advances were not at a disadvantage. In this regard, the NIH has thus become more conservative over time, despite a variety of policies that the NIH

NIH Funding of Contributions by Novelty of Idea Inputs If Funding Rates Were Constant Across Idea Types and Research Areas

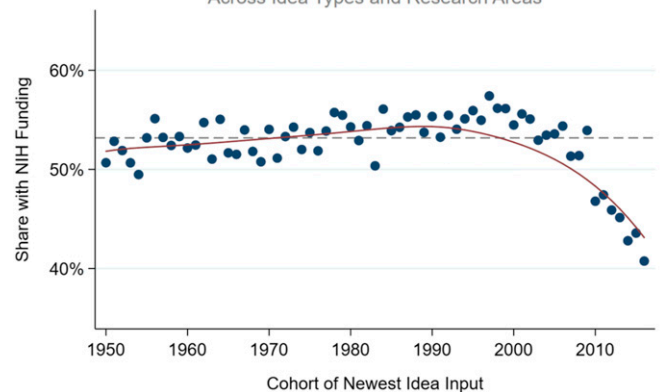


Fig. 4. Share of NIH funding by novelty of idea inputs: if funding rates were the same across idea types and research areas. Calculated based on 992,633 biomedical research papers published during 2010 to 2016. The horizontal axis captures the idea input vintage, with later (earlier) cohort years representing more novel (more conventional) science. The vertical axis captures the rate of NIH funding. The markers capture the mean NIH funding rate for each idea input vintage. The solid line represents a nonparametric regression line estimate. The dashed line represents the average funding rate across all cohorts. In this analysis, funding rates are adjusted so that all variation in funding rates is due to variation across cohorts within each (idea type, research area) pair.

has implemented in the past 2 decades to reward innovative and high-risk project proposals (37–40).

Because not all biomedical research projects seek NIH funding, our findings reflect the funding preferences of the NIH and its review committees regarding research proposals it actually receives. As a consequence, our findings may reflect, in part, funding opportunities for novel ideas by non-NIH funders. For example, it may be that the pursuit of novel work opens up generous non-NIH funding opportunities, and that for this reason, contributions that build on very recent ideas are less often NIH funded than contributions that build on a bit more mature ideas. It may also be that non-NIH funding opportunities have improved relative to NIH funding opportunities in recent decades, explaining the apparent decline in how often very novel papers are NIH funded. It should be noted that the NIH's stated aim and initiatives regarding innovative science suggest that the NIH itself does not see sufficient non-NIH funding opportunities for novel work. Furthermore, the sheer size of the NIH relative to others makes it likely that whatever the NIH does sets the agenda in an important way. If the NIH were found not to fund innovative science, it would be hard to imagine the other funding agencies could compensate for this to an extent that would result in a favorable overall funding environment for edge science.

With these considerations in mind, our finding that the NIH has become more conservative is consistent with at least three potential explanations. First, it may be that review committees have become more cautious in terms of funding the most novel work, or that reviewers today are unable to discern which research projects are edge science. This interpretation is supported by our finding that, after controlling for variations in funding rates across idea types and research areas, there is no strong link between novelty and NIH funding, and by reports that biomedical researchers themselves have become more conservative in their research choices (41). Second, ideas generated in the 2000s may have been considered less fruitful by the scientific community than ideas generated during the 1990s. Concerns about stagnation in technological progress in medicine (42) are supported by empirical evidence showing that health advances have become less common and harder to achieve in recent decades (43). Third, while in our understanding many NIH grantees do not feel obligated to adhere to their research plans in light of new results [geneticist Mario Capecchi is a famous example (44)], it may be possible that some NIH grantees feel locked into their funded research aims, which have aged by the time of funding. Lock-in may render funded scientists unable to work on emerging ideas. An intriguing direction for future work is to examine to what extent such lock-in occurs and whether the NIH should seek to eliminate it.

Our evidence also points to considerable heterogeneity in the NIH's tendency to fund novel work. Contributions that build on a novel basic science idea, especially a recently discovered gene or protein, are more often NIH funded than are contributions that build on comparable older ideas. In contrast, this novelty–NIH funding link is not present for clinical ideas. Our analysis does not directly reveal the mechanisms behind this heterogeneity. While it is well known that the NIH has provided generous support for research linked to genomics in recent decades (45), this preference does not explain the positive link between novelty and frequency of NIH support among contributions that build on advances in genomics. Here, a limitation of our analysis is that we do not explicitly consider the role of non-NIH funders, although our findings suggest that they play an important role in promoting edge science. For example, given our finding that NIH supported work in pharmaceutical research tends toward older ideas rather than edge science, it must necessarily be the case then that non-NIH funders such as pharmaceutical companies play a key role in supporting work on new drugs (and hospitals likely fund novel clinical work more broadly). This heterogeneity may be

driven by differences in the appropriability of benefits from novel research, as patents are more useful for pharmaceutical research than for other settings. This division of labor may also be optimal. Since the NIH relies on public funding, it has a unique responsibility to solve market failure problems in science. Ideally, the NIH should support both novel and impactful scientific work that would have trouble receiving funding from private sources.

Our results suggest that the NIH's overall tendency to fund edge science papers is driven mostly by its funding at higher rates rapidly advancing areas of investigation such as genomics and its applications, rather than funding the most novel work within each area of investigation. To the extent that rewarding novelty within idea types and research areas is infeasible, it may be that the only possible approach for directing funds to novel work is to provide ample funding to entire areas of investigation that are thought to be advancing at a faster rate than others. Of course, a top-down approach to research resource allocation also has risks (45–50). The NIH may end up directing the funds to areas that turn out to be less fruitful than expected, and inertia in funding decisions will likely keep some areas well-funded long past their eventual stagnation. It is also questionable whether anyone can know in advance which discipline or organism holds the key to the next important development in biomedicine. These risks of a less diversified funding portfolio must be weighed against its benefits, including the ability to direct funds to novel work. A move to a more egalitarian funding model, one that does not favor some areas of investigation as much as current practice, would likely lead to a significant decrease in the novelty of NIH-supported work unless the NIH addresses the conservatism of scientific review committees.

Our text-based approach for identifying innovative research suggests a fruitful agenda for quantitative analyses of funding agencies. Constructively, measuring new idea adoption can help the NIH to better understand the impact that its review practices and funding decisions have on the scientific enterprise, and can help the NIH design policies to counteract the long-term trend that we identify in the funding of work that builds on new ideas, and thus speed scientific advance. Researchers could also apply the approach to NIH grant application data to gain insights into what an increase in the NIH budget would yield in terms of the pursuit of edge science. While we have focused on work that is novel in the sense that it builds on new ideas, the analysis could also be extended to capture also work that tries out novel combinations of old ideas. Finally, the approach can be used to evaluate the novelty of work funded by agencies, such as the National Science Foundation, that operate on scientific fields outside of biomedicine.

Materials and Methods

We describe the main features of our approach in the introduction; we describe further details here and in the *SI Appendix*. Coverage of the MEDLINE database spans years 1946 to 2016 (as of February 2018). In determining the ideas on which a paper builds and the vintage of each idea, we use the title and abstract text (the latter is available only for articles published since 1975). Our analysis of NIH funding starts in 1990 because we limit the analysis to papers with a US first author and because coverage of author affiliation starts in 1988. We infer the country of affiliation based on the first author because only information on the first author is available for articles published before 2014. *SI Appendix, Table S3* lists the UMLS categories that we use as a proxy for idea type and their division to the three broad idea type groups. *SI Appendix, Table S4* lists examples of idea inputs uncovered by our approach. *SI Appendix, Table S5* lists the NLM's journal-level subject categories, which we use as a proxy for research area. We counts a paper that is linked to K idea types and J research areas as $K \times J$ contributions. Results from analyses with an alternative delineation of research areas and from analyses with a paper-level approach are shown in the *SI Appendix* and confirm the conclusions we report here. The depository for data and code for the statistical analyses of this paper is available in Harvard Dataverse (51).

ACKNOWLEDGMENTS. We thank Bruce Weinberg, Joel Blit, Jeremy Goldhaber-Fiebert, Neeraj Sood, David Studdert, Partha Bhattacharyya, Neesha Joseph, and Walter Schaeffer for helpful discussions. We also thank seminar participants at the Institute for Fiscal Studies (London); the University of Illinois at Chicago Institute of Government and Public Affairs; Stanford Medical School; Ca'Foscari University of Venice; Johns Hopkins University; American Economics Association Annual Conference; Asia Pacific Innovation

Conference (Beijing); The Ohio State University; the University of Southern California; the Latin American, Caribbean Economics Association-Latin American Meeting of the Econometric Society Annual Meeting; India Conference on Innovation, Intellectual Property, and Competition at the Indian School of Business (Hyderabad); and the National Bureau of Economic Research working group on Innovation in an Aging Society for helpful feedback. We thank the National Institute of Aging for funding for this research through grant P01-AG039347.

1. National Institutes of Health, Budget. <https://www.nih.gov/about-nih/what-we-do/budget>. Accessed 5 September 2019.
2. National Institutes of Health, Mission and goals. <https://www.nih.gov/about-nih/what-we-do/mission-goals>. Accessed 5 September 2019.
3. F. S. Collins, Exceptional opportunities in medical science: A view from the national Institutes of health. *JAMA* **313**, 131–132 (2015).
4. B. N. Sampat, Mission-oriented biomedical research at the NIH. *Res. Policy* **41**, 1729–1741 (2012).
5. A. Marshall, *Principles of Economics* (Macmillan and Co., London, 1920).
6. A. P. Usher, *A History of Mechanical Inventions* (McGraw-Hill, New York, 1929).
7. T. S. Kuhn, *The Structure of Scientific Revolutions* (Chicago University Press, Chicago, 1962).
8. T. S. Kuhn, "Objectivity, value judgment and theory choice" in *The Essential Tension*, T. S. Kuhn, Ed. (University of Chicago Press, Chicago, 1977), pp. 320–339.
9. M. Livio, *Brilliant Blunders: From Darwin to Einstein – Colossal Mistakes by Great Scientists That Changed Our Understanding of Life and Universe* (Simon & Schuster, 2013).
10. S. Firestein, *Failure: Why Science is So Successful* (Oxford University Press, 2015).
11. R. Popovian, "Dedicated scientists driven to discover cures." *Morning Consult* (2016). <https://morningconsult.com/opinions/dedicated-scientists-driven-to-discover-cures/?source=acsh.org>. Accessed 5 December 2017.
12. M. Zaringhalam, "Failure in science is frequent and inevitable—and we should talk more about it." *Scientific American Blog* (2016). <https://blogs.scientificamerican.com/guest-blog/failure-in-science-is-frequent-and-inevitable-and-we-should-talk-more-about-it/>. Accessed 5 December 2017.
13. Institute of Medicine, *Scientific Opportunities and Public Needs: Improving Priority Setting and Public Input at the National Institutes of Health* (The National Academies Press, Washington, DC, 1998).
14. National Institutes of Health, NIH-wide strategic plan: Fiscal years 2016–2020. <https://www.nih.gov/sites/default/files/about-nih/strategic-plan-fy2016-2020-508.pdf>. Accessed 5 September 2019.
15. D. Hegde, B. N. Sampat, Can private money buy public science? Disease group lobbying and federal funding for biomedical research. *Manage. Sci.* **61**, 2281–2298 (2015).
16. J. S. Langer, Enabling scientific innovation. *Science* **338**, 171 (2012).
17. J. M. Nicholson, J. P. A. Ioannidis, Research grants: Conform and be funded. *Nature* **492**, 34–36 (2012).
18. R. M. Cook-Deegan, Does NIH need a DARPA? *Issues Sci. Technol.* **13**, 25–28 (1996).
19. T. R. Cech, Fostering innovation and discovery in biomedical research. *JAMA* **294**, 1390–1393 (2005).
20. G. Kolata, Grant system leads cancer researchers to play it safe. *New York Times*, 28 June 2009. <https://www.nytimes.com/2009/06/28/health/research/28cancer.html>. Accessed 27 April 2020.
21. B. Albers, M. W. Kirschner, S. Tilghman, H. Varmus, Rescuing US biomedical research from its systemic flaws. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 5773–5777 (2014).
22. D. Geman, S. Geman, Opinion: Science in the age of selfies. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 9384–9387 (2016).
23. M. Murray, P. Aghion, M. Dewatripont, J. Koev, S. Stern, Of mice and academics: Examining the effect of openness on innovation. *Am. Econ. J. Econ. Policy* **8**, 212–252 (2016).
24. B. A. Jacob, L. Lefgren, The impact of NIH postdoctoral training grants on scientific productivity. *Res. Policy* **40**, 864–874 (2011).
25. B. A. Jacob, L. Lefgren, The impact of research grant funding on scientific productivity. *J. Public Econ.* **95**, 1168–1177 (2011).
26. P. Azoulay, J. S. Graff Zivin, G. Manso, Incentives and creativity: Evidence from the academic life sciences. *RAND J. Econ.* **42**, 527–554 (2011).
27. P. Azoulay, J. S. Graff Zivin, D. Li, B. N. Sampat, Public R&D investments and private-sector patenting: Evidence from NIH funding rules (National Bureau of Economic Research, Cambridge, MA). Available at <https://www.nber.org/papers/w20889>. Accessed 12 May 2020.
28. M. E. Blume-Kohout, Does targeted, disease-specific public research funding influence pharmaceutical innovation? *J. Policy Anal. Manage.* **31**, 641–660 (2012).
29. M. Blume-Kohout, K. B. Kumar, N. Sood, University R&D funding strategies in a changing federal funding environment. *Sci. Public Policy* **42**, 355–368 (2015).
30. D. Li, Expertise versus bias in evaluation: Evidence from the NIH. *Am. Econ. J. Appl. Econ.* **9**, 60–92 (2017).
31. D. Li, L. Agha, Research funding. Big names or big ideas: Do peer-review panels select the best science proposals? *Science* **348**, 434–438 (2015).
32. K. J. Boudreau, E. C. Guinan, K. R. Lakhani, C. Riedl, Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science. *Manage. Sci.* **62**, 2765–2783 (2016).
33. M. Packalen, Edge factors: Scientific frontier positions of nations. *Scientometrics* **188**, 787–808 (2019).
34. D. Besancenot, R. Vranceanu, Fear of novelty: A model of strategic discovery with strategic uncertainty. *Econ. Inq.* **53**, 1132–1139 (2015).
35. M. Packalen, J. Bhattacharya, Age and the trying out of new ideas. *J. Hum. Cap.* **13**, 341–373 (2019).
36. M. Packalen, J. Bhattacharya, Neophilia ranking of scientific journals. *Scientometrics* **110**, 43–64 (2017).
37. M. Larkin, Can Zerhouni create a bold, risk-taking NIH? To succeed, reforms will have to change NIH's organisation and culture. *Lancet* **362**, 1382–1383 (2003).
38. F. Collins, Opportunities and challenges for the NIH—an interview with Francis Collins. Interview by Robert Steinbrook. *N. Engl. J. Med.* **361**, 1321–1323 (2009).
39. J. Avorn, A. S. Kesselheim, The NIH translational research center might trade public risk for private reward. *Nat. Med.* **17**, 1176 (2011).
40. C. Woodward, National Institutes of Health seek to speed up therapeutic innovations. *CMAJ* **183**, E91–E92 (2011).
41. A. Rzhetsky, J. G. Foster, I. T. Foster, J. A. Evans, Choosing experiments to accelerate collective discovery. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 14569–14574 (2015).
42. J. Le Fanu, "Science's dead end." *Prospect Magazine*, 21 July 2010. <https://www.prospectmagazine.co.uk/magazine/sciences-dead-end>. Accessed 27 April 2020.
43. N. Bloom, C. I. Jones, J. Van Reenen, M. Webb, *Are Ideas Getting Harder to Find?* (Manuscript, 2018).
44. M. R. Capecchi, *Gene Targeting 1977-Present* (Nobel Lecture. The Nobel Foundation, 2007).
45. M. J. Joyner, N. Paneth, J. P. A. Ioannidis, What happens when underperforming big ideas in research become entrenched? *JAMA* **316**, 1355–1356 (2016).
46. M. Rosbash, A threat to medical innovation. *Science* **333**, 136 (2011).
47. B. Albers, The end of "small science"? *Science* **337**, 1583 (2012).
48. G. A. Petsko, Big science, little science. *EMBO Rep.* **10**, 1282 (2009).
49. N. Vermeulen, J. N. Parker, B. Penders, Big, small or mezzo? Lessons from science studies for the ongoing debate about 'big' versus 'little' research projects. *EMBO Rep.* **11**, 420–423 (2010).
50. M. Peifer, The argument for diversifying the NIH grant portfolio. *Mol. Biol. Cell* **28**, 2935–2940 (2017).
51. M. Packalen, Replication data for: NIH funding and the pursuit of edge science. Harvard Dataverse. <https://doi.org/10.7910/DVN/QTSOGS>. Deposited 7 May 2020.