

Does Peer Review Penalize Scientific Risk Taking? Evidence from NIH Grant Renewals*

Pierre Azoulay
MIT and NBER
100 Main Street
E62-487
Cambridge, MA 02142

Wesley H. Greenblatt
Boston Children's Hospital
and Harvard Medical School
300 Longwood Ave.
Boston, MA 02115

February 11, 2025

Abstract

Scientific projects that carry a high degree of risk may be more likely to lead to breakthroughs yet also face challenges in winning the support necessary to be carried out. We analyze the determinants of renewal for more than 100,000 R01 grants from the National Institutes of Health between 1980 and 2015. We use four distinct proxies to measure risk taking: extreme tail outcomes, disruptiveness, pivoting from an investigator's prior work, and standing out from the crowd in one's field. After carefully controlling for investigator, grant, and institution characteristics, we measure the association between risk taking and grant renewal. Across each of these measures, we find that risky grants are renewed at markedly lower rates than less risky ones. We also provide evidence that the magnitude of the risk penalty is magnified for more novel areas of research and novice investigators, consistent with the academic community's perception that current scientific institutions do not motivate exploratory research adequately.

Keywords: scientific risk, scientific productivity, peer review, NIH, government funding.

JEL Classification: O32, O38, H51

* We thank Alessandro Bonatti, Matt Clancy, Danielle Li, and Scott Stern, as well as seminar participants at the Max Planck Institute for Innovation and Competition, NBER Summer Institute, and NBER Productivity Seminar for useful feedback.

1. Introduction

No funder of science has ever prided itself on supporting “low risk, low reward” projects, yet few ideas elicit as much agreement as the claim that peer review punishes risk taking. These laments appear in the popular press, in U.S. Congressional testimony, errant tweets, or the pages of *Nature*, and are authored by successful investigators and leading science policy makers (Kolata 2009, Kornberg 2007, Nielson 2022, Sutter 2022, Woolston 2014). Prominent scientists have penned editorials satirizing the behavior of risk-averse grant funders, and ponder whether Nobel Prize winners could receive funding for their seminal discoveries in the modern funding climate (Fields 2014, Petsko 2012). While economists would caution that a benevolent social planner might want to tame individual scientists’ risk appetites, the proposition that current institutions do not adequately reward risky research strikes practicing scientists as self-evidently true.

Yet, for all the colorful anecdotes, blue-ribbon panels, and editorials, there is surprisingly little evidence to support the claim that peer review punishes risk taking, and less evidence still that society as a whole would be better served if funders could clamp down on their conservatism. Azoulay et al. (2011) compare two grant funding mechanisms—the Howard Hughes Medical Institute Investigatorships and National Institutes of Health (NIH) R01 grants—and provide evidence that the former encourages more risk taking than the latter, but this does not necessarily imply that peer review at NIH punishes risk taking. Other scholars have examined how novelty shapes grant funding outcomes (Ayoubi et al. 2021, Packalen and Bhattacharya 2020, Veugelers et al. 2022). However, novelty is distinct from risk taking: investigators can pursue novel research which isn’t risky, and risky research which isn’t novel. What is widely agreed upon is that grant funders’ tolerance for risk has direct implications for what research is undertaken in the first place. In a survey of investigators supported by *Fast Grants*, 78% stated they would make significant changes to their research program if they were not constrained by the demands of grant funding agencies, including pursuing more ambitious research programs, pivoting to new topics, and testing hypotheses others see as unlikely to succeed (Collison et al. 2021).

More generally, the choice of what scientists choose to study is fundamental to determining the rate and direction of innovation. An assumption typically made by models of the innovation process is that exploratory projects are less likely to bear fruit than projects that merely seek incremental advances, but with more upside if they are indeed successful (Manso 2011). This is the sense in which science funders always claim to crave “high risk, high reward” projects even though the trade-off between risk and reward in the market for scientific ideas is less obvious than in financial markets (Nielsen and Qiu 2022). While both risky exploratory research and incremental lines of inquiry are valuable from a societal point of view, a first-order concern in contemporary science policy is that funding programs tilt the mix of projects they support too far in the direction of incrementalism. Moreover, the societal costs of undersupplying risky research may be especially prominent for early-stage, basic research, which may lead to substantial knowledge spillovers (Azoulay and Li 2022).

What causes individual scientists to select risky projects? Rather than being solely determined by an individual’s preference for risk taking, this choice is shaped in important ways by the incentives, organization, and supporting institutions that surround a researcher. Tolerance for failure or lack thereof (Tian and Wang 2014), time horizons (Lerner and Wulf 2007), team structure (Wu et al. 2019), and the structure of financial incentives (Graff Zivin and Lyons 2020) shape how much risk innovators choose to bear when selecting projects. The choice of research approach is made more challenging by the difficulty in predicting the ultimate benefits if the project is successful; this is especially the case for early stage and highly uncertain projects (Rosenberg 1995).

In this paper, we study how a specific institution, grant peer review at NIH, shapes scientific risk taking. In grant peer review, scientific proposals are evaluated by a panel of experts, with the resulting scores having a major influence on the likelihood of funding. Scientists who are awarded funding use these resources to conduct the investigations they committed to in their proposal. At the end of the grant cycle, a funded scientist has the option of extending the research program by applying for an additional cycle of support. To study the relationship between peer review and risk taking, one would ideally have access to all proposals (regardless of funding status), reviewer scores, and grant outcomes in the form of publications and citations that can be traced to a particular funding stream.¹ With access limited to publicly available data, we focus on funded applications, and ask whether risk taking is associated with the likelihood that a grant is renewed competitively for an additional cycle. Since we are studying the impact of risk taking conditional on a successful initial grant application, the estimates we present should be construed as a lower bound on the risk taking penalty that would be observed if we had access to the full set of grant applications, included those that were not supported. We also link each grant cycle with the publications they produced, leveraging the requirement that authors acknowledge their NIH grant support in their published work.

R01 grants are the “bread and butter” of academic biomedical research, and a requirement for scientists who aspire to establish and make viable their laboratory over the long run. We study how risk taking is penalized or rewarded in 103,164 NIH-funded R01 and R01-equivalent grant cycles between 1980 and 2015. After controlling for a comprehensive set of covariates for investigator demographics, institution quality, and grant characteristics, we estimate the relationship between risk taking and whether the grant was competitively renewed. Scientific risk taking, much like scientific creativity, is a nebulous concept. There is no consensus on what constitutes risk in science nor how it should be measured (Althaus 2005, Franzoni and Stephan 2023). Since it is unlikely that a single measure can fully represent the phenomenon, we develop four distinct measures of risk taking, each of which captures a different aspect of the risks taken

¹ These data are typically not made available to researchers. NIH employees have access to funded and unfunded grant applications and average scores; even these insiders do not have access to the full distribution of scores at the proposal level, as these scores are destroyed in order to protect reviewer confidentiality.

or borne by scientists, and examine if our findings depend on the particular measure used. All four measures rely on the realized outcomes of the initial grant cycle to characterize the risk profile of subsequent cycles. A maintained assumption in our analysis is therefore that risk taking is imprinted on a particular scientific trajectory at the initial stage, and largely persists across future cycles.

One way to conceptualize risk is through variability in outcomes—projects that produce either groundbreaking breakthroughs or complete failures (Azoulay et al. 2011). Another perspective links risk to disruption, distinguishing research that challenges the status quo from work that consolidates existing knowledge (Funk and Owen-Smith 2017). A third approach considers risk in terms of intellectual movement, where scientists pivot away from their prior work to explore new directions, emphasizing that risk stems not just from the project itself but from who undertakes it. Finally, risk may manifest in intellectual distinctiveness, as researchers choose projects that diverge sharply from those of their peers, making their successes—or failures—particularly visible to the scientific community.

Across each of these measures, we find that risk taking is penalized. When comparing grant cycles in the top and bottom decile of risk taking, grants with greater risk taking have a 9.5% lower renewal rate (20.5% decline) when measuring risk taking using extreme tail outcomes, an 11.1% lower renewal rate (24.4% decline) when measuring risk taking by its disruptiveness, a 6.9% lower renewal rate (15.2% decline) when measuring risk taking by an investigator pivoting from her prior research, and a 4.0% lower renewal rate (9.5% decline) when measuring risk taking by standing out from what other investigators are studying. In contrast to our measures of risk taking, novelty is associated with *higher* grant renewal rates in our data. At the very least, this finding buttresses our claim that novelty and risk taking are different concepts that it is best not to conflate.

Of course, it may well be that the penalty applied to riskier projects is appropriate, given the real-world constraints faced by funders. For public funding at least, the committees and officials doling out funding to individual proposals are accountable to political paymasters, and may feel pressure to present evidence of accumulating and steady success, rather than a series of failed projects occasionally punctuated by an extraordinary outcome (Clancy 2023). We present two separate pieces of evidence consistent with the claim that funders penalize risk even beyond what a benevolent social planner might wish: First, the gradient in the renewal-risk taking relationship appears steeper for novel proposals, relative to less novel ones. Second, for renewed grant cycles, those we measure as taking on more risk lead to elevated rates of citations of the renewed grants' publications in patents, relative to the renewed cycles we measure as taking on less risk.

Our study goes beyond reaffirming the widely held belief that peer review favors conservative research over high-risk, potentially higher-reward projects. By offering systematic evidence of this bias, we take the conversation one step further, entertaining the possibility—as

ridiculous as it may sound to practicing scientists—that the heightened scrutiny placed on risky projects is not entirely misplaced. While society may benefit from imposing a higher bar on risky endeavors—filtering out failures that are not necessary milestones on the way to promising breakthroughs, our findings suggest that the penalty may be excessive, particularly in novel research domains where bold exploration is most likely to pay off. Our results also imply that peer review, far from being an immovable pillar of the scientific process, is a powerful lever for policymakers to shape the kinds of challenges researchers pursue—and it can be reformed to better balance the trade-off between risk and societal benefit, steering innovation toward the breakthroughs that matter most.

Prior studies have explored how aspects of a research proposal influence the evaluation and funding of scientific projects and publications. This includes novelty (Ayoubi et al. 2021, Boudreau et al. 2016, Packalen and Bhattacharya 2020, Teplitskiy et al. 2022, Veugelers et al. 2022), interdisciplinarity (Banal-Estañol et al. 2019, Bromham et al. 2016), and the interplay between novelty and feasibility (Krieger and Nanda 2022, Lane et al. 2022a). Our study builds on this literature in several ways.

First, we introduce an eclectic set of risk-taking measures, each capturing a different dimension of risk, and explicitly distinguish between risk and novelty in the analysis. Second, rather than characterizing an investigator’s prior research trajectory to predict future funding—an approach common in previous studies—we focus on risk measures directly tied to a specific grant proposal and the publications emerging from it. This distinction is important because researchers often manage multiple projects simultaneously, each carrying different levels of risk. An investigator’s past work may not accurately reflect the level of risk in a particular grant proposal, and in some cases, scientists may deliberately avoid proposing high-risk ideas to funders they suspect might penalize bold experimentation, instead pursuing such projects through alternative funding sources.

Our closest intellectual antecedents are Boudreau et al. (2016), who experimentally assign grant proposals to peer evaluators, and Carson et al. (2023), who manipulate the distribution of reviewer scores to assess how peers (rather than funders) value disagreement among reviewers. Like ours, these studies exploit proposal-level variation to uncover the determinants of funding success. Unlike ours, the studies control the treatment of interest in a framed experimental design (List and Metcalfe 2014)—an unambiguous methodological strength. However, these studies do not measure proposal risk, nor can they replicate the high stakes of NIH grant funding competitions. Our approach, by contrast, relies on observational data. While this limits our ability to identify causal effects, it allows us to examine risk-taking as the primary treatment of interest in a real-world setting of substantial policy interest—namely, the entire universe of R01 NIH grants over an extended time period.

The remainder of the paper is structured as follows. Section 2 outlines our argument for the central role of risk taking in scientific research and examines how, or whether, funding peer

review influences the level of risk accepted by investigators, evaluators, and funders. Section 3 introduces the setting, data, and addresses measurement challenges. In Section 4, we describe our empirical strategy and provide descriptive statistics. Section 5 presents the results of our analysis. Finally, Section 6 concludes by discussing the broader implications of our findings.

2. Scientific risk taking and peer review

There is no consensus on what constitutes risk in science (Althaus 2005, Franzoni and Stephan 2023). Following how the term is frequently used in scientific vernacular, we use risk to mean uncertainty surrounding the probability that a scientific project can be successfully brought to fruition. Failure, however, pervades the scientific research process. While some degree of risk is inherent in nearly all innovation projects, a central challenge facing projects with high levels of risk is both that they are less likely to succeed, and that conditional on success, it may be more difficult to predict their ultimate impact (Rosenberg 1995).²

As Franzoni and Stephan (2023) explain, risk in science is a multilayered concept. Different types of risk in science include technical or scientific risk stemming from nature’s answer to the research question being asked, execution risk in being able to successfully carry out the planned experiments, competitive risk from being “scooped” by other investigators studying the same or similar question (Hill and Stein 2024), and risk in how a successfully completed project will be perceived by peers or advance the investigator’s career.

To illustrate how projects may entail different aspects of risk, consider a few examples. The efforts to synthesize insulin using recombinant DNA technology in the late 1970s was perceived to carry limited scientific risk but significant execution and competitive risks. While the foundational work by Cohen and Boyer (Cohen et al. 1973) had already established the basic techniques for gene splicing, challenges in gene sequencing, protein expression, and proper folding presented substantial execution hurdles. Genentech and UCSF led the intense competition to overcome these challenges first and this scientific contest not only transformed diabetes treatment but also heralded the biotechnology era (Hall 1987). In contrast, early efforts to develop cancer immunotherapy entailed very substantial scientific risk, as it was perceived to be very challenging to develop treatments targeting cancer cells but not normal tissue, but they faced arguably lower levels of competitive risk (Littman 2015). Similarly, designing processes to encourage handwashing may carry little technical risk, but entailed career risk for Ignaz Semmelweis, whose empirical observations conflicted with the established scientific and medical opinions of the mid-19th century (Nuland 2004).

Why would individual scientists or society seek out research projects that involve risk? One common justification is the assumption that the probability of success and the impact of

² Of course, the returns or impact of innovative projects are uncertain as well, but we find it conceptually useful to not conflate risk and return in what follows.

scientific results are often inversely related: higher-impact discoveries tend to come with lower chances of success. However, unlike in finance, this relationship is probably not systematic, as there are celebrated cases where both high value and high probability of success coexist. For instance, the Human Genome Project had a very high likelihood of success when it was launched given previous advances in sequencing, and its outcomes were anticipated to be highly impactful for both human health and future scientific progress. Additionally, the value of failure in science weakens the analogy to financial risk and return. Failed experiments often generate insights that increase the expected returns of future research, though individual scientists or firms may struggle to capture that long-term value (Frankel et al. 2023). Despite these nuances—and the limited empirical evidence of the risk/return tradeoff in scientific research—this study will maintain the assumption of a positive correlation between scientific risk and the expected importance of the results.

A scientist’s choice of projects incorporates their subjective assessment of risk with their judgement of the value (to their own career and to society) of the project’s potential impact. Why would scientists ever propose projects whose impact is *ex ante* not commensurate with their perceived probability of success? One possible explanation is that researchers are simply overconfident, holding optimistic beliefs about their projects’ probability of success or the value of their potential results (Gan et al. 2012, Sanchez and Dunning 2023). But a countervailing influence is that of funding institutions. Long gone is the era of the “gentleman scientist” who could, like Charles Darwin, fund his own investigations. Undertaking scientific research almost invariably requires harnessing the support of a financial backer, whether the source of funds is a government agency, a philanthropic foundation, or the private sector. While researcher optimism suggests that science funders might justifiably limit risky projects to align with the goals of a benevolent social planner, the topics scientists ultimately pursue are shaped by the risk tolerance and priorities of their funding sources.

In academic research, peer review is the primary mechanism for allocating research funds. This raises an important question: does peer review bias funders toward or away from risky projects? A reasonable starting point is that funders, by supporting a portfolio of projects with largely independent probabilities of success, are better positioned to absorb risk than individual scientists (Arrow and Lind 1970). Yet, despite the frequent rhetorical emphasis on funders’ tolerance for failure and enthusiasm for “high risk/high reward” projects, there are several reasons to suspect that risk aversion may permeate the peer review process: from the level of individual investigators seeking funds, to the level of peer reviewer evaluating proposals, to review committees, and even to the funding institutions themselves (Azoulay and Li 2022, Franzoni et al. 2022).

Peer review may shape the degree of risk investigators choose to pursue through several mechanisms. Many academic labs rely on external grants to fund their operations. An inability to win external funding may result in a lab’s closure, laying off dependent personnel (Tham et al. 2024), and loss of the investigator’s job and potentially career (Ruben 2017). To the extent

investigators *perceive* grant peer review to punish risk taking, this will lead them to instead propose safer, more incremental projects (Kent 2018, Langer 2012, Nielson 2022). Additionally, the intense competition for funds that is the byproduct of a system in which the vast majority of proposals will not receive support may push investigators to propose projects with relatively quick pay-offs (Alberts et al. 2014, Fang and Casadevall 2015). In short, peer review may shape investigators' incentives for risk taking even in the absence of any actual bias against risk for the evaluators and funders.

The penalization of risk taking may also occur at the level of peer reviewers. Individuals may have a psychological bias against the uncertainty that characterizes risky research compared to more incremental projects (Fox and Tversky 1995, Mueller et al. 2012). In line with this, evaluators of research proposals have been found to give lower scores to proposals with greater novelty (Boudreau et al. 2016) and to emphasize feasibility (Krieger and Nanda 2022, Lane et al. 2022a). The bias against novelty may also lead to delayed recognition of the underlying potential of a scientific work (Wang et al. 2017).

Peer review is conventionally undertaken in a committee. Bias against risk taking can occur due to both committee processes and committee membership. Proposals with greater risk taking may have greater variance in reviewer assessment of their quality (Langer 2012). Committee processes that emphasize consensus and uniformly high evaluations across its members to be funded may penalize risk taking (Carson et al. 2023).³ Similarly, given the wider variation in evaluations of risky projects, processes where each member independently fully evaluates a proposal prior to the committee meeting may be more likely to result in identifying a judge who enthusiastically supports the project compared to processes where one or two committee members have primary responsibility for evaluating each proposal and then present their recommendation to the committee. Exposure to negative assessments of other committee members prior to the final judgment may lead evaluators to revise their own assessment downwards more so than exposure to positive assessments leads them to revise their assessments upwards, biasing against projects with more disagreement (Lane et al. 2022b).

Finally, some aspects of the potential for peer review to penalize risk taking emerge at the level of the funder and innovation ecosystem. Funders can influence risk taking through how they structure committee procedures and membership, set expectations for preliminary evidence, and establish explicit evaluation criteria. Pressure to show political paymasters that administrators' are effective stewards of money can also tilt the selection of projects towards more conventional or incremental approaches (Lorsch 2015).

Of course, many funders think about risk at the level of their entire portfolio of projects, rather than on a grant-by-grant basis. Funders sometimes manage riskier projects using distinct

³ For instance, a process where a proposal's overall score is simply the average of each judge's evaluation will emphasize consensus more than having scores aggregated after dropping the lowest judge score(s) or a process which emphasizes having a few judges with very high levels of support (Gius 2024).

mechanisms or organizational models, with explicitly different evaluation criteria, while applying a more conservative filter in other parts of their project portfolio (Azoulay et al. 2019a).⁴ While there are portions of the NIH portfolio where risk taking is not the goal, such as P01 grants, some clinical trials, or support for research cores where there is more of a focus on operational excellence, we focus on R01 grants, the primary NIH funding mechanisms for biomedical research aimed at extending the scientific frontier and a setting where an excessive penalty for risk taking might plausibly slow scientific progress.

Since beliefs regarding a project’s probability of success and its potential impact are unobservable, under what conditions might it be feasible to show that peer review *excessively* punishes risk taking? Under the maintained assumption that risk and return are positively correlated in the scientific setting, we argue that while a benevolent social planner might in some circumstances wish to discourage risk taking, it would seem perverse for her to do so when a project’s expected impact is especially uncertain.

A useful way to identify cases where penalizing risk-taking is socially inefficient is to consider how informative prior assessments are about a project’s potential value. When prior beliefs are less certain—meaning the initial expectations about a project’s impact are more diffuse—punishing risk-taking is more likely to stifle valuable scientific exploration. This variation in the strength of prior beliefs offers a way to isolate instances where excessive conservatism in funding decisions is particularly costly. When are prior beliefs relatively more uncertain? One case is when a scientist is early in their career and has a limited track record, making it harder to form strong expectations about their projects’ likely outcomes. Another occurs in highly novel research areas, where prior knowledge is sparse, and the value of success is more difficult to predict. Because novel research proposals carry greater option value, funders should, in theory, be more tolerant of risk in emerging fields than in well-established areas of inquiry. Certainly, a finding that risk taking is penalized more strongly for novice investigators or novel research areas appears difficult to reconcile with a reasonable model of optimal scientific risk taking.

Recognizing that the interaction between risk and uncertainty in research returns helps isolate cases of excessive penalization of risk-taking also underscores a crucial distinction: riskiness and novelty are not the same. It is easy to see why risk and novelty are often conflated: truly groundbreaking scientific breakthroughs frequently involve both. A revolutionary idea (high novelty) often carries a high chance of failure (high risk) simply because it ventures into the unknown. Conversely, research that is very safe (low risk) is often incremental and thus less likely to be novel. However, researchers can take risks without being novel, just as they can pursue novelty without taking large risks. Consider, for example, the notion that a bacteria (*H. pylori*)

⁴ The NIH has created a number of distinct funding mechanisms with explicit mandates around high-risk, high-reward research, such as the Pioneer Award and Director’s Transformative R01 Award. In 2021, the NIH funded 106 such awards with an expected accompanying outlay of \$329 million over five years (NIH 2021). This represents only 1.4% by number of awards and 1.8% by funding amount of all R01-equivalent awards that year at the NIH (Lauer 2022).

causes stomach ulcers, the hypothesis proposed by Marshall and Warren (1984) that went against decades of medical dogma blaming stress and acid. Initially, their work seemed risky in terms of acceptance (many peers were skeptical), but the experiment to prove it was straightforward: they identified the bacteria in patients and Marshall even ingested it himself to show it caused gastritis. While bold (and a bit personally risky health-wise!), this was a relatively direct test of a novel hypothesis—not a massive, uncertain project (Blaser 2005). In contrast, Judah Folkman’s early research on the role of angiogenesis in cancer faced extreme skepticism (Stephenson et al. 2013). Because his findings were viewed as highly uncertain and technically challenging to build upon, few scientists were willing to engage with them for over a decade—by which time follow-on work may no longer have been considered novel.

Further complicating the relationship between novelty and risk, scientists’ careers depend on how their contributions are perceived by peers. As Merton (1973) notes, employment and promotion decisions reward originality. A researcher who consistently works on problems deemed insufficiently “cutting edge” may face significant career risk, even if their chosen projects are technically safe. Likewise, even a risk-averse funding agency has a mandate to pursue new ideas, which may lead it to favor “safe” forms of novelty—projects that appear innovative without being truly high-risk.

We view the study of scientific novelty as essential in its own right. However, research in this area should (i) carefully distinguish between novelty and risk-taking, (ii) develop empirical measures that do not conflate the two, and (iii) treat any correlation between them as an object of empirical interest rather than a taken-for-granted assumption. The present study endeavors to abide by these principles.

3. Setting, data, and measurement

The National Institutes of Health

The NIH plays a prominent role in funding biomedical research. In 2012, it comprised 27% of all biomedical research funding in the United States and nearly two-thirds of all public and philanthropic research funding (Moses et al. 2015). In 2021, the NIH spent \$32 billion to support extramural research projects (Lauer 2022). While the NIH funds clinical research (including clinical trials, health services research, and behavioral studies), the bulk of their funding supports basic research aimed at advancing the frontier of scientific understanding. With research project grants having an average size of over \$580,000 in 2021, NIH grants are a major source of funding for the researchers they support. Not only does this funding play an important role in supporting the purchase of materials and the salaries of lab personnel, many academic researchers are in “soft money” positions, funding their own salary out of awarded grants. The ability to successfully raise grant funding is an important metric upon which academic employment and promotion decisions are made.

Understanding the process the NIH uses to evaluate and fund grant applications can provide useful context for our empirical analysis. Briefly, the NIH is comprised of 27 separate institutes or centers (IC) which are typically organized around a set of diseases (e.g., National Cancer Institute) or organ systems (e.g., National Heart, Lung and Blood Institute). Each IC is responsible for funding research that is potentially relevant to its mission. Scientific evaluation of grant proposals occurs primarily in standing review committees known as study sections (e.g., “Cellular and Molecular Immunology” or “Atherosclerosis and Vascular Inflammation”). Each study section may review and score grants from multiple ICs, which are then funded by the ICs in order of their study section score until the IC exhausts its funds.⁵ The amount of resources devoted to reviewing grant proposals—both from the NIH with staff to manage the process and from study section members—is substantial. In 2021, the NIH reviewed nearly thirty-eight thousand applications for R01-equivalent grants (Lauer 2022).

The NIH is an attractive setting for studying how peer review might respond to risk taking. First, professional norms lead most academic research to be published in peer-reviewed journals (Dasgupta and David 1994). When combined with rich bibliometrics in biomedicine, this provides a robust “paper trail” we can follow and characterize to measure the degree of grant risk taking. Second, as NIH research is a source of significant improvements in health and positive economic spillovers, understanding how the NIH may more effectively utilize its funds to support breakthrough innovation is of considerable policy interest (Azoulay et al. 2019b, Fleming et al. 2019).

Data

As we do not observe grant applications, we focus on whether an awarded grant was renewed. We rely on the NIH’s Consolidated Grant Applicant File to identify all those researchers who received an NIH grant, 1980-2015. We start our analysis in 1980 as we are not able to reliably link publications to the NIH grants which supported them before this date, and end in 2015 to allow sufficient time for follow up to observe research outcomes and renewal status of grants awarded in that year. We limit our analysis to R01-equivalent grants.⁶ These investigator-initiated grants are a major funding mechanism to support the work of principal investigators who have achieved career independence and run their own laboratory. We also exclude grant awards from special emphasis study sections, which may not have an option for grant renewal, and grants without associated publications for which many of our measures of risk taking would be undefined.

⁵ Institute directors do have discretion to fund applicants out of order considering factors such as fit with the institute’s mission, the overall portfolio of projects being supported, and their own evaluation of application quality.

⁶ We consider R01-equivalent grants to be those with an NIH activity code of R01 or R37. We exclude several grant mechanisms considered by the NIH to be R01-equivalent but which may have distinct dynamics around risk taking. For instance, we exclude DP1, DP2 and DP5 which are part of the NIH’s effort to fund high risk, high reward projects; R56 and RL1 which may not be applied for by the investigator; U01 in which NIH staff actively support the research; and R35 and RF1 which provide longer-term support to investigators.

We focus our analysis on the grant cycle, which corresponds to a single competitive review of a funding proposal. A cycle starts the year a grant proposal is funded and ends either at the conclusion of grant funding or when a proposal for grant renewal is competitively reviewed and funded. As grants may have multiple principal investigators and investigator characteristics are explicitly considered during renewal decisions, we analyze our sample at the investigator-grant-cycle level.⁷ This gives a final sample of 103,164 investigator-grant-cycles from 63,101 grants and 37,222 investigators.

Our publication data is from *PubMed*, an online resource from the National Library of Medicine to support access to the biomedical research literature that indexes over 30,000 journals and over 34 million articles. We limit our analysis to original research publications; this excludes other types of publications such as review articles, letters, and editorials. For each grantee, we identify their career publications in *PubMed* using Author-ity (Torvik and Smalheiser 2021, Torvik and Smalheiser 2009, Torvik et al. 2005). Author-ity disambiguates authors using the insight that authors of the same name sharing coauthors, affiliations, keywords, and other such characteristics are likely, in fact, to be the same person. Author-ity has been shown to be highly accurate (Lerchenmueller and Sorenson 2016).

We use NIH RePORTER to link publications attributable to each NIH grant. This data relies on the grant having been listed in the publication’s acknowledgement section as a funding source, which is required as a condition of accepting NIH funding. We consider publications attributable to a grant cycle if they are published up to one year after the last year of the grant cycle to allow for publication delays.⁸ As investigators may have different degrees of risk taking across different projects and funding sources, in all cases, our measures of risk taking are calculated solely using publications during a grant cycle that acknowledge funding from the grant rather than also including publications which do not acknowledge funding from the grant.

To measure research topic and methodology, we capitalize on Medical Subject Headings (MeSH). MeSH is a hierarchical controlled vocabulary created and maintained by the National Library of Medicine used for indexing and searching the biomedical literature. In the 2019 edition of MeSH, there are over 29,000 terms, with an average of about 12 MeSH keywords associated with each *PubMed* publication. Crucially, MeSH are assigned by professional indexers. While it is difficult to assess accuracy as there is no single “correct” way to index articles, MeSH have generally been found to be consistent across indexers (Coletti and Bleich 2001, Funk and Reid 1983). In addition to MeSH associated with grant publications, we use the National Library of

⁷ 138 (0.13%) grant cycles appear in our sample more than once as they have multiple principal investigators.

⁸ Among publications acknowledging a focal NIH grant in our sample, 20.8% are published more than one year after the end of grant funding. These publications and the results they are based on will typically not have been available to study sections when they review a grant for competitive renewal and so do not inform our outcome or measures of interest.

Medicine Medical Text Indexer to assign MeSH to grant abstracts. Grant abstract text is available from the NIH RePORTER.

We supplement publication data with citations from Clarivate’s *Web of Science*. All grants are adjusted to 2020 dollars using the Biomedical Research and Development Price Index.

Measuring risk taking

Given the different sources of risk in science, no single measure can fully capture the dynamics of risk-taking across the diverse range of projects in our data. We therefore employ four distinct measures, each reflecting a different aspect of scientific risk. Since these measures have varying strengths and weaknesses, considering them together provides a more comprehensive perspective than any single measure in isolation. Importantly, our analysis remains non-parametric, examining the full distribution of risk-taking levels. This ensures that our findings are not driven solely by conventional levels of risk-taking but also account for whether penalties for risk arise primarily in the tails of the distribution.

A key consideration in our measurement strategy is that we construct our risk measures using information from publications emerging from the initial grant cycle. Ideally, risk would be assessed directly from the content of the renewal proposal rather than inferred from past outcomes. This necessarily embeds an assumption: that a project’s risk profile at the renewal stage largely reflects the level of risk taken in the initial funding period. While investigators may maintain a consistent approach to risk-taking across multiple grant cycles—whether due to inertia, specialization, or the constraints of their research agenda—this remains an inherent limitation of our method.

Access to full proposal texts, including those of unfunded applications, would allow for a more direct and precise assessment of risk at the time of renewal. However, given the constraints of publicly available data, our strategy represents a meaningful step forward in understanding how risk-taking influences funding outcomes. We aim to be transparent about these limitations so that readers can judge for themselves whether our approach is reasonable. In our view, the compromises we have made are a necessary price for progress in studying this important and otherwise difficult-to-analyze question. With these general considerations in place, we now turn to the construction of the four risk-taking metrics.

First, risk-taking may lead to greater variance in outcomes, increasing both the likelihood of a breakthrough and the likelihood of failure (Azoulay et al. 2011). We capture this dynamic with a measure we term “extreme tail outcomes,” defined as the difference between the highest and lowest vintage-adjusted citation percentiles among publications from the initial grant cycle. We focus on long-run vintage-adjusted citation percentiles for several reasons. Relying solely on short-run citation data would introduce considerable noise, as the impact of scientific contributions—especially those employing unconventional approaches—often takes time to be fully recognized (Wang et al. 2017). Additionally, using percentiles enables meaningful comparisons

across grants awarded in different time periods, ensuring that differences in citation practices over time do not distort our measure.⁹

Research can be either consolidative, reinforcing earlier findings, or disruptive, challenging the status quo. Our second measure leverages this distinction with the disruption index proposed by Funk and Owen-Smith (2017), which has received increased attention (and validation) from the science of science community (Bornmann et al. 2020; Wu et al. 2019). This index classifies a paper as highly disruptive if future research cites it while largely ignoring its references, whereas highly consolidative research is characterized by future studies citing both the focal paper and its references. Our measure of “disruption” for a grant cycle is the highest percentile of the disruption index for the publications emerging out of the initial round of funding (calculated relative to all publications appearing in *PubMed* in the same year).

As a third measure, we examine the intellectual distance between a grant cycle’s publications and the investigator’s prior work. This approach recognizes that some aspects of risk stem not just from the project itself but also from who is conducting the research. Scientists who integrate their existing expertise with new knowledge—whether by addressing novel topics, adopting new techniques, or working with unfamiliar model organisms—are taking on risk. Such shifts may yield valuable discoveries but can also devalue prior expertise and invite skepticism from peer evaluators (Arts and Fleming 2018; Hill et al. 2021). We term this measure “pivoting.” To compute it, we calculate the fraction of MeSH term pairs in a grant cycle’s publications that are new relative to those in the investigator’s prior work during the preceding five years.

Our final measure of risk taking captures intellectual distance from what other NIH-funded investigators are studying within the same scientific domain. This measure reflects the risks of pursuing contrarian research—projects that diverge from the mainstream of one’s peers—thereby increasing career risk (Fang and Casadevall 2015). We term this measure “standing out” and operationalize it as the fraction of MeSH term pairs in a grant cycle’s publications that appear exclusively in the focal investigator’s work, meaning no other funded researchers in the same IC-study section-year use these term pairs in their grant proposal abstracts.^{10,11}

⁹ A limitation of this approach is that we infer failure from exceptionally low citation rates. While this is a reasonable proxy, truly failed projects are more likely to remain unpublished in a file drawer rather than being published in a minor outlet.

¹⁰ The relevant sample size varies across each of these measures. Extreme tail outcomes is only defined for those grant cycles resulting in at least two publications indexed in Web of Science. Disruption is only defined for those grant cycles resulting in at least one publication that is cited at least once. Pivoting is defined only if there is at least one publication with MeSH terms in both the current grant cycle and the preceding five years. Standing out is defined only if there is one publication with MeSH terms in the current cycle and sufficient funded grant proposal abstracts in the same IC-study section-year are available.

¹¹ We compare an investigator’s grant cycle publication MeSH to that for all available awarded grant proposal abstracts, including those that are not R01-equivalent grants, in the same IC-study section-year. Among the 103,164 investigator-grant-cycles in our sample, we have abstracts for other grants in the same IC-study section-year for 83,451. The mean number of abstracts in the comparison group for each investigator-grant-cycle is 99. We limit our measure to those investigator-grant-cycles with at least 20 other

Measuring novelty

To empirically distinguish novelty from risk taking, we construct a separate measure of novelty. Prior research has taken various approaches, including analyzing unusual or intellectually distant combinations of references (Shibayama et al. 2021, Uzzi et al. 2013, Wang et al. 2017), uncommon combinations of patent classes (Fleming 2001, Verhoeven et al. 2016), citation network centrality (Shibayama and Wang 2020), and the age of keywords or keyword combinations (Arts et al. 2024, Misha and Torvik 2016, Packalen and Bhattacharya 2020). Other studies have used domain-specific measures, such as novelty in chemical structures (Krieger et al. 2022, Rzhetsky et al. 2015). Following Boudreau et al. (2016), we define novelty based on new combinations of MeSH terms. Specifically, our measure captures the fraction of MeSH term pairs in a grant cycle’s publications that were first used in any *PubMed*-indexed publication within the preceding three years.¹²

4. Research design and a descriptive look at the data

We observe the universe of NIH grants actually awarded, but unfortunately, do not observe grant applications which are not ultimately funded. Because of this, we empirically focus on NIH R01-equivalent grants that were awarded and measure whether or not the grant was renewed. This approach lets us reliably observe all investigators who were at risk of having their grant renewed. In this population, a grant may not be renewed either because the investigator applied for but did not receive a competitive renewal, or because they never applied for a renewal. Given the career importance of grant funding in maintaining a lab, we expect most investigators will apply for renewal.¹³ Furthermore, if an investigator with risky research is dissuaded from applying for grant renewal due to their perception of a low likelihood of renewal and the time involved in applying, the implications for our research question are very similar to having applied and not been selected. As our sample evaluates the degree of risk taking among scientists all of whom have an R01-equivalent grant, this suggests these scientists will have similarly benefited from the act of applying for and receiving an R01-equivalent grant (Ayoubi et al. 2019, Jacob and Lefgren 2011).

Using this research design, the primary threat to unbiased estimation of the impact of risk taking on grant renewal is omitted variable bias, i.e., unmeasured factors correlated with both our proxies for risk taking and determinants of a grant’s likelihood of renewal. To mitigate this concern,

abstracts in the same IC-study section-year. We prefer this approach to alternatives, as the comparison group is thicker for those with sufficient available data, facilitating our focus on the tail-measures of intellectual distance from other investigators. Similar results are obtained using other thresholds for the number of abstracts in the same IC-study section-year.

¹² 7.9% of MeSH pairs in our sample have three or fewer years between when they were first used in *PubMed* and their use in a grant cycle publication. Our findings are robust to using thresholds other than three years, including using MeSH pairs new to *PubMed* in the past year.

¹³ Applications for competitive grant renewal, using the NIH’s more expansive definition of R01-equivalency (see footnote 6) were funded at a rate of 34.5% in 2015 (Lauer 2016), close to the 35.9% renewal rate in our sample in the same year, and consistent with the assumption that most grantees do apply for renewal.

we incorporate a rich set of controls for investigator, institution, and grant cycle characteristics.¹⁴

We also consider controlling for publication and citation outcomes during the initial grant cycle, but the inclusion of these controls warrants careful consideration. Study sections evaluating renewal applications assess an investigator’s progress, often reflected in publications acknowledging the grant’s support. Omitting research productivity and impact measures could introduce omitted variable bias, as investigator quality may correlate with both productivity and risk-taking. However, since the relationship between risk-taking and research output (quantity or quality) may vary in magnitude and direction, the bias introduced by exclusion is difficult to predict. At the same time, publication outcomes lie on the causal pathway between risk-taking and grant renewal, making their inclusion problematic. Controlling for them effectively conditions on a collider (Elwert and Winship 2014), potentially biasing estimates of risk-taking’s effect on renewal. To address this tradeoff, we present results both with and without these productivity controls. Including them generally attenuates effect sizes toward zero, but our qualitative insights remain robust. Importantly, we do not claim to estimate the causal effect of risk-taking on funding success; rather, we document robust conditional correlations.

Methodological considerations

Our primary estimating equation relates the dependent variable of interest, $renewal_{ij,t(j)}$, for investigator i , grant j , and grant cycle $t(j)$ to risk taking assessed for the preceding cycle of the same grant j , $t(j)-1$. Formally, we estimate variations of:

$$E[renewal_{ij,t(j)} = 1 \mid X_{ijt}] = \beta_0 + \beta_1 risk_taking_{ij,t(j)-1} + X_{ij,t(j)}$$

where $renewal$ is an indicator variable equal to 1 if grant j was renewed for the next cycle $t(j)$, and 0 otherwise. We include high dimensional vectors of controls for investigator and grant cycle characteristics $X_{ij,t(j)}$. This vector includes sets of indicator variables for investigator age, degree and gender; log investigator career citations, log prior investigator NIH funding, an indicator variable for no prior investigator NIH funding, and log institutional NIH total funding as measures of investigator and institutional quality; and sets of indicator variables for year of evaluation, initial cycle length, and IC-study section fixed effects. Select specifications also include controls based on measures of research quality and quantity during cycle $t(j)$, including log of mean journal impact factor, sets of indicator variables for the number of publications during the grant cycle (separately for those acknowledging and not acknowledging funding from the grant), and indicator

¹⁴ The NIH instructs reviewers to score each application based on five criteria: significance, innovation, approach, investigator, and environment. While the first three assess the importance, novelty, and feasibility of the proposed project itself, the last two evaluate investigator qualifications and ability of the institution to support the work, respectively.

variables for having publications in highly selective journals.^{15,16} We use robust standard errors with two-way clustering at the investigator and grant level throughout. For β_1 to have a causal interpretation, the measures of $risk_taking_{ij,t(j)-1}$ must be uncorrelated with the error term of the corresponding regression equation, after conditioning on the observable covariates included in X . In spite of the high dimensionality of X , this remains a very strong assumption, reinforcing the earlier caveat that our estimates correspond to robust conditional statistical associations rather than causal effects.

Our primary statistical model is a linear probability model (LPM). While LPM can produce fitted values outside the interval $[0,1]$ —which makes it unsuitable for predictive purposes—our objective is to estimate the marginal effects from the conditional expectation function. In this context, the LPM is typically more appropriate (Angrist and Pischke 2008: pp. 94-106). Specifically, LPM offers a straightforward estimation of marginal effects and is computationally efficient when dealing with a high-dimensional vector of covariates and fixed effects, which is crucial for our analysis.

Descriptive Statistics

Tables 1a and 1b present descriptive statistics for investigators and grant cycles, respectively. Reflecting the intense competition for R01-equivalent grants, the sample consists of highly accomplished scientists. On average, investigators begin the grant cycle 18.7 years after completing their terminal degree, with 56 prior publications and 2,234 career citations. Approximately 5% of investigators in the sample have achieved elite status, defined as having won a Nobel Prize, Lasker Award, Howard Hughes Medical Institute Investigatorship, or membership in the Institute of Medicine or National Academy of Sciences.

The majority of an investigator’s research output is not directly funded by the focal R01-equivalent grant. Among publications during the cycle period, 39.9% acknowledge the focal grant, while 59.6% of non-grant-funded publications still receive NIH support, either through other NIH grants to the investigator or coauthors. Overall, 45.2% of grant cycles were renewed, with renewal rates higher for subsequent cycles (51.5%) compared to the first cycle (40.4%).¹⁷ Renewed grant cycles were associated with younger investigators (17.6 vs. 19.6 years of career experience), greater research productivity in cycle $t(j)$ (8.1 vs. 7.0 publications and 663 vs. 422 long-run citations, limiting the comparison to articles directly acknowledging the focal grant), and a greater percentage of the investigator’s research activity funded by the focal grant (46% vs. 38.2% of

¹⁵ We use journal impact factor as a proxy for publication quality in grant cycle publications, as this metric is visible to peer review committees at the time of evaluation, whereas early citations provide a noisy signal of quality. Our results remain robust when instead measuring grant productivity using citations accrued by the time of renewal or, alternatively, long-run citations.

¹⁶ For this measure we include those publications which were published in *Nature*, *Cell*, *Science*, and the *New England Journal of Medicine*.

¹⁷ In contrast, the success rate of applicants for a new R01-equivalent grant is lower. In 2015, 16.1% of applications for a new R01-equivalent grant were funded (Lauer 2016).

publications). Figure 1 provides an histogram of the distribution of the number of publications per grant cycle, separately for renewed and terminated grants.

Table 1c presents descriptive statistics for the four measures of risk taking, as well as novelty. Although the mean level of risk taking is generally higher for terminated grants than for renewed ones, the corresponding magnitudes are modest. Figure 2 displays histograms for each of the four measures by grant cycle renewal status, providing initial evidence that renewal rates vary across different levels of risk taking. Finally, Table 1d presents correlations between measures of risk taking and novelty. The relatively low correlation between these measures highlights that each captures a distinct approach to risk taking, with novelty offering a contrast to the risk-taking metrics.

5. Results

Risk taking and grant renewal

Table 2 reports estimates of the impact of risk taking on the likelihood of grant renewal for each of the four measures. Overall, we observe a clear inverse relationship: high levels of risk are associated with lower renewal rates, while low levels of risk correlate with higher renewal rates across all measures. Not only are these effects statistically significant, but their magnitudes are also considerable. For example, when comparing grant cycles in the bottom versus top decile of the “extreme tail outcomes” metric in Table 2A, column 2, those in the bottom decile—reflecting low levels of risk taking—experience a 5.1% higher renewal rate (an 11.0% increase relative to intermediate deciles), whereas those in the top decile have a 4.4% lower renewal rate (a 9.5% decrease relative to intermediate deciles). Similar patterns emerge across the other risk-taking measures for the top and bottom deciles: disruption (8.9% higher [19.6% increase] vs. 2.2% lower [4.8% decrease]) in Table 2B, pivoting (4.1% higher [9.0% increase] vs. 2.8% lower [6.2% decrease]) in Table 2C, and standing out (2.7% higher [6.4% increase] vs. 1.3% lower [3.1% decrease]) in Table 2D.¹⁸ These findings remain even after controlling for publication quantity and quality during the grant cycle, though in some cases, the magnitudes of the risk penalty shrink somewhat. Importantly, we find no evidence suggesting that the penalty for risk taking diminishes in later cycles (i.e., on the second or third competitive renewal for the same grant).

We further investigate the full distribution of risk-taking levels and their association with grant renewal, as depicted in Figures 3, 4, 5, and 6. In these specifications, separate indicator variables are included for each ventile of the risk measure. Consistent with Table 2, we find higher

¹⁸ For extreme tail outcomes, pivoting, and standing out, percentiles are calculated relative to other grant cycles in the same evaluation year. For disruption, we emphasize work that is disruptive in an absolute sense and so use the percentile of the most disruptive paper, calculated relative to all publications in *PubMed* appearing in the same year.

renewal rates for lower levels of risk taking and lower renewal rates for higher levels. Notably, several measures show a nearly monotonic decrease in renewal rates as risk taking increases.¹⁹

Additionally, we examine how the penalty for risk taking varies across NIH Institutes and Centers. While both the magnitude of the penalty and the types of risk taking most penalized differ across Institutes, the pattern persists across the entire NIH. We also examine temporal changes in the penalty for risk taking. Given the increasing competition for R01-equivalent grants and evolving NIH policies, caution is needed when interpreting trends over time. Nevertheless, we find that the magnitude of the risk-taking penalty has remained relatively stable (Appendix Figure A2).

Novelty and risk taking

We next examine whether novelty confounds the relationship between risk taking and grant renewal. We estimate models similar to those shown in Figures 3 through 6, but with the addition of indicator variables for each ventile of the novelty measure. The results of this analysis are presented in Figure 7. Two key findings emerge from this analysis. First, in contrast to our measures of risk taking, we observe that novelty is rewarded. The magnitude of this reward is comparable to the penalty for risk taking. Second, the inclusion of novelty does not substantially reduce the observed penalty for risk taking. This provides empirical evidence that novelty and risk-taking measures may operate independently, supporting the notion that novelty, at least in some contexts, is not a valid proxy for risk taking. Appendix Figure A4 replicates this analysis by assessing the effect of novelty on grant renewal without incorporating measures of risk taking. As in Figure 7, we find that novelty is associated with a higher probability of grant renewal.

Is the risk penalty excessive?

As previously discussed, we propose that variation in the informativeness of prior beliefs about the potential value of research pursued in a grant may help identify cases where the penalty for risk taking is likely excessive from a social standpoint. To test this empirically, we examine the interaction between novelty and risk taking on the likelihood of grant renewal. We also analyze how career stage interacts with risk taking, as less established or younger scientists may lack a sufficient track record for evaluators to form an informative prior about the potential success of their investigations, compared to their senior counterparts.

Following NIH classifications, we focus on new investigators—those in their first R01-equivalent grant cycle—and early stage investigators—those within 10 years of their terminal

¹⁹ An exception is with disruption, where there is a non-linearity in the association between disruption and grant renewal: grant cycles with the very highest degree of disruption have higher renewal rates than those with elevated but less extreme levels of disruption.

research degree or clinical training.²⁰ Additionally, we assess the risk-taking penalty for members of the National Academy of Sciences, whose rarefied status likely provides evaluators with more informative prior beliefs about the value and feasibility of their proposed research.²¹

Figure 8 summarizes these findings, with Appendix Tables A1 and A2 providing the full estimates. Several trends emerge. First, across most measures of risk taking, more novel research, new investigators, and early stage investigators face a higher penalty for risk taking, while members of the National Academy of Sciences experience a lower penalty.²² Second, at higher levels of novelty, the penalty for risk taking increases for pivoting and standing out, but remains statistically indistinguishable from less novel research for the extreme tail outcomes and disruption measures (Appendix Table A2). Finally, some of these results are imprecisely estimated, particularly with regard to status, due to the small number of elite investigators.

We provide additional evidence suggesting that the penalty for risk-taking may be excessive. Specifically, we ask whether, among grant cycles that were renewed, those characterized by higher levels of risk-taking yield better outcomes than their lower-risk counterparts (Table 3, as well as Appendix Tables A3 and A4). To assess this, we examine the relationship between risk-taking and research relevance, measured by citations to renewed grant cycles’ publications in patents (Marx and Fuegi 2020, 2022). Across all four measures, we find a positive association between risk-taking and impact on private-sector patenting—although for the “extreme tails” metric, this effect reverses when controlling for initial grant cycle productivity (column 5).²³ In line with Becker (1957), these results suggest that NIH peer review systematically disadvantages risky proposals, imposing a higher bar for funding despite their potential for significant private-sector impact.

²⁰ As we do not observe residency or fellowship training, we assume for determining early stage investigator status an average length of postgraduate clinical training of five years for researchers holding an M.D. or M.D./Ph.D.

²¹ This measure also includes members of the National Academy of Medicine. We exclude any members who are also Howard Hughes Medical Institute Investigators who may have distinct dynamics and constraints for risk taking (Azoulay et al. 2011).

²² New and early stage investigators may need to show some degree of pivoting from the research done during the training phase of their career to establish independence and a distinct research identity from their mentors.

²³ The Appendix (Table A4, Panel A) provides similar evidence after limiting the patents being considered to those linked to FDA-approved drugs (Durvasula et al. 2023). Because this outcome is more skewed, the corresponding analysis subsamples are smaller than in Table 3. We also examine the association between risk taking in the current grant cycle and publications and patents in the next grant cycle for those grant cycles that were renewed (Table A4, Panel B). We find these analyses less compelling since they essentially regress publication or citation outcomes in cycle $t(j)$ on a transformation of publication/citation outcomes in cycle $t(j)-1$. Yet, as above, the results imply that risk taking is associated with superior outcomes in the future on the sample of renewed grants. An exception is observed for the pivoting measure, which is associated with fewer publications in the subsequent grant cycle, consistent with a decline in productivity when shifting to new research topics.

Robustness checks and ancillary results

We conduct several robustness checks to assess the sensitivity of our estimates to alternative assumptions and subsamples. First, we test alternative measures of risk taking and novelty (Appendix Table A5). This includes the difference between the maximum and median, median and minimum, winsorized maximum and minimum, and variance in vintage-adjusted citation percentiles for extreme tail outcomes (Appendix Figures A5 and A6); the mean rather than the maximum disruption index and a variation of the index following Bornmann et al. (2020) for disruption (Appendix Figures A7 and A8); fraction new single MeSH rather than MeSH pairs and cosine dissimilarity score for single MeSH terms after term frequency-inverse document frequency weighting for pivoting (Appendix Figures A9 and A10); fraction unique single MeSH rather than MeSH pairs and cosine dissimilarity score for single MeSH terms after term frequency-inverse document frequency weighting for standing out (Appendix Figures A11 and A12), and MeSH pair age as well as two measures—the fraction of bigrams which are new and the average number of new bigrams per publication—constructed from textual analysis of publication title and abstract compared to the corpus of publications in *OpenAlex* following Arts et al. (2024) for novelty (Appendix Figures A3 and A4). Emphasizing how pivoting and standing out are distinct from novelty, the findings with these measures are robust to exclusion of MeSH terms and MeSH term pairs which were first used in *PubMed* within the past decade. As risk taking measures may be undefined for some grant cycles, the sample used for analysis varies across each risk taking measure (see footnote 10 for details). Our results are robust to using a subsample for which all four risk taking measures are defined for all grant cycles (Appendix Table A6).

The penalty for pivoting might be smaller for those grants for which the investigator proposed—and peer review approved and funded—to pivot in the initially submitted planned research studies. To examine this, we identify the fraction of MeSH terms shared between the grant proposal abstract and an investigator’s prior work and split the sample at the median for this measure. We find overall a similar magnitude penalty for pivoting across both groups (Appendix Table A7). In other words, grants which proposed a pivot and then investigators actually pivoted in the work carried out are still renewed at lower rates.

Author-ity has been shown to be highly accurate (Lerchenmueller and Sorenson 2016). It is likely the case, however, that its accuracy of publication disambiguation is higher for rare names within *PubMed*. To ensure our results are not driven by investigators with poorly disambiguated publications in Author-ity, we split the sample at the median name frequency based on frequency within the entire corpus of *PubMed* and find similar results across both groups (Appendix Table A8).

One reason investigators may choose not to renew their grant is that their primary research focus lies outside of the scope of what the NIH typically supports. For instance, some organic chemists, electrical engineers, or sociologists may have an isolated project at the interface of biomedicine within the context of a broader research agenda elsewhere. To ensure our results are

not driven by investigators with only a weak attachment to NIH funding, we calculated the fraction of career publications with any NIH funding and split the sample at the median of this measure (Appendix Table A9). We find a similar penalty for risk taking across both those with above and below the median career attachment to NIH funding.

Finally, we use a bounding technique proposed by Oster (2019) to gauge the sensitivity of our results to omitted variable bias. The intuition behind this approach is that the stability of the coefficient for risk taking when varying the set of control variables included in the model, scaled by the movement in R^2 , provides information about the potential impact of unobserved covariates. To generate these bounds, the analyst must assume proportionality between the covariances of the outcomes with observed and unobserved covariates and posit a maximum value for R^2 if the regression could include all observed and unobserved covariates. Oster’s technique generates δ , which can be interpreted as the degree of selection on unobservables relative to observables necessary to reduce the magnitude of the effect of the regressor of interest to zero. Appendix Table A10 reports the results of this exercise. For all four risk taking measures, δ is above one, the threshold recommended by Oster to suggest robustness to the influence of unobservable covariates.²⁴

6. Discussion

Using a comprehensive dataset of awarded NIH R01-equivalent grants and multiple measures of risk-taking, we find that investigators who take greater risks experience lower renewal rates than their lower-risk counterparts. This provides systematic evidence that NIH peer review penalizes risk-taking. We also show that novelty functions differently from our risk-taking measures, cautioning against the practice of using the former as a proxy for the latter. Given the NIH’s central role in U.S. biomedical research funding, these findings have significant policy implications for peer review and innovation funding.

A key strength of our study is its examination of risk-taking across the full distribution of observed values, capturing effects at the extremes. Not all scientific innovation contributes equally to progress, and the most radical advances often emerge from the highest-risk projects (Rzhetsky et al. 2015, Veugelers and Wang 2019, Wang et al. 2017). That peer review punishes risk-taking is particularly striking given mixed evidence on its ability to predict scientific productivity and impact (Cole et al. 1981, Danthi et al. 2014, Li and Agha 2015, Pier et al. 2018).²⁵

²⁴ Some values for δ are negative. This implies that the correlation between unobservables and risk taking would need to be in the opposite direction as the correlation between observables and risk taking to make the corresponding results disappear.

²⁵ This parallels the dynamics seen for innovation in other settings, such as entrepreneurship. For instance, even among VC-backed startups, a handful companies that are outliers in their success account for the majority of VC profits and impact on the larger economy. Likewise, while investors may be able to identify

Our results cannot be given a causal interpretation if unobserved covariates drive both risk-taking and grant renewal. However, given the rich set of controls included in the statistical models, we believe this concern may be less significant than it appears. In fact, our estimates may understate the true penalty for risk-taking. Some of our measures depend on the existence of published work, meaning we may miss projects that failed to generate publications, potentially underestimating risk-related attrition. Additionally, because we observe only funded grant applications, any initial selection bias against risky projects would make our sample less venturesome than the full applicant pool. Investigators aware of the NIH’s risk aversion may also self-select out of the funding process altogether, particularly for their high-risk projects (Langer 2012). Finally, young researchers training under NIH-funded mentors may internalize a preference for safer projects, reinforcing long-term conservatism in research agendas.

Each of our risk-taking measures captures different facets of risky research, each with its own strengths and imperfections. An investigator branching into a new research area, for instance, may be engaging in a promising recombination of ideas—or straying too far from their expertise, reducing their chances of success (Arts & Fleming 2018). Similarly, pursuing underexplored topics might signal a potential breakthrough or simply an unpromising research avenue. Given these differences, one could reasonably quibble with any single measure. Yet despite their idiosyncratic flaws and the fact that they are constructed in distinct ways—sometimes even drawing on different underlying data—the results remain remarkably consistent across measures. This consistency is reassuring: if each metric is flawed in its own way, much like Tolstoy’s unhappy families, they are at least flawed differently. And yet, they tell the same story.

Both high-risk, high-reward and low-risk, incremental research play vital roles in scientific progress. However, there are strong reasons for the NIH to increase support for high-risk research. Established research fields naturally lend themselves to incremental work, while emerging areas require greater exploration and risk-taking. Instead, we find the opposite: the penalty for risk-taking is largest in more novel areas, suggesting a potential welfare loss. NIH risk aversion may also disproportionately affect research areas with limited private-sector incentives, where public support is most essential (Aghion et al. 2008). Governments are uniquely positioned to bear risks too large for private industry, especially given declining corporate investment in science (Arora et al. 2017, Fleming et al. 2019). NIH’s size and influence mean its hesitancy to fund risky fundamental research is unlikely to be offset by other funding sources, potentially slowing technological progress (Babina et al. 2023). The same logic applies to NIH-funded clinical research, where industry has limited incentives to invest in high-risk studies with uncertain commercial returns (Greenblatt et al. 2023). While private philanthropies provide alternative funding, their priorities may not align with the greatest public needs (Murray 2012).²⁶

those companies with little growth potential, they have only a very limited ability to discern those within their portfolio most likely to generate outsized returns (Kerr et al. 2014).

²⁶ Researchers who do not receive NIH funding may be able to continue their research with support from other funders. For instance, Chuban and Hackett (1990) find that many investigators are able to pursue

Recognizing these challenges, the NIH has created dedicated high-risk, high-reward funding mechanisms, such as the Pioneer Awards, Director’s Transformative R01 Awards, and New Innovator Awards (NIH 2019). In 2021, these programs funded 106 awards totaling \$329 million over five years (NIH 2021). However, such mechanisms represent only a small fraction of NIH’s overall funding portfolio (see footnote 4). Several private science funders—including the Howard Hughes Medical Institute, Open Philanthropy Project, Fast Grants, Chan Zuckerberg Initiative, and Wellcome Trust—have also launched initiatives to support high-risk research.

While our study focuses on peer review in public funding, its implications extend to other contexts where high-risk innovation is evaluated. Rather than treating risk tolerance as an afterthought or relying on vague exhortations to evaluators, funders should proactively design processes that encourage the desired level of risk-taking. Dedicated funding mechanisms are an important step, but they must also address inherent biases against risk. For instance, as high-risk projects may have greater variation in evaluator assessments of quality, moving from a consensus-based to a champion-based approach or emphasizing enthusiastic support rather than mean score may facilitate the selection of riskier projects (Gius 2024). The composition of committees and panels should also be scrutinized. Expert evaluators may have more pronounced biases against novelty in their own field (Wang et al. 2017). Increasing diversity of expertise and traditions and considering an evaluator’s prior risk-taking behavior when selecting reviewers may help mitigate risk penalization (Nicholson and Ioannidis 2012).

Finally, a fundamental challenge in motivating risk-taking in science is that, *ex post*, a failed but risky research project and a low-quality idea may generate similar outcomes—no publication, patent, or product. While fame and fortune can both act as powerful motivators for pursuing risky science, investigators may still hesitate if failed risky projects lead to job loss or lab closure. Without compromising high standards and accountability, motivating risk-taking may require simultaneous adjustments to multiple features of the funding environment, including feedback, time horizon, and job security (Manso 2011, Azoulay et al. 2011).

their proposed research even if they do not receive the focal grant. Similarly, Jacob and Lefgren (2011) find the marginal impact of receiving an NIH grant to be only about one additional publication over five years, in part because researchers can substitute to other funding sources. Yet, NIH’s penalization of risk taking may influence the type of proposals investigators submit to these alternative sources. The preeminent position of NIH within the biomedical research funding ecosystem implies that its risk aversion could “contaminate” the behavior of other funders.