Designing Scientific Grants

Christoph Carnehl* Marco Ottaviani† Justus Preusser‡

April 30, 2024

1 Introduction

Throughout history, financial support from influential patrons has played a crucial role in the dissemination and advancement of knowledge. In the second half of the eighth century, the Abbasid rulers of Baghdad institutionalized support for translators, contributing to a flourishing of science in the Islamic world (Gutas, 1998). In 1610, Galileo Galilei solicited and obtained financial backing from his former pupil Cosimo de' Medici, the Grand Duke of Tuscany, after naming in honor of the Medici family the moons of Jupiter he had just discovered.

In the modern world, scientific grants are a key instrument for stimulating innovation. For example, the National Institutes of Health (NIH) in the US had a budget of around 45 billion dollars in 2022. The Horizon Europe program has a budget of around 95 billion euros for the period 2021-2027.[1]

Scientific grants differ starkly from the traditional market system analyzed by economists. Grants are upfront payments without contractible goals and researchers have significant leeway in directing these resources. The reason for these features lies in the very nature of grant-supported research. For example, it is often impossible to describe ex-ante the value or outcome of research that is inherently open-ended and

*Department of Economics, Bocconi University, *christoph.carnehl@unibocconi.it*.

†Department of Economics, Bocconi University, *marco.ottaviani@unibocconi.it*.

‡Department of Economics, Bocconi University, *justus.preusser@unibocconi.it*.

[1]See https://www.nih.gov/about-nih/what-we-do/budget for the NIH, and https://commission.europa.eu/funding-tenders/find-funding/eu-funding-programmes/horizon-europe_en for the Horizon Europe program.

uncertain. Neither funding agencies nor researchers can anticipate whether years of research will ever lead to a significant discovery, whether the scientific community will realize the significance, or whether the discovery will have an impact beyond the scientific community. Driven by passion, intellectual curiosity, and the hope for scientific reputation, investigators conduct "blue-sky" research to solve open puzzles and challenge conventional wisdom. More than perhaps any other source of funding, grants thus leverage researchers' self-motivation for discovery and recognition within the community.

The uncertainty inherent in research has important and nuanced implications for the design and allocation of grants. In this chapter, we discuss the basic economics of grants, focusing on the relationship between uncertainty, incentives, and optimal grant design.

Compared to other funding instruments, such as patents and research prizes, grants seem to have received considerably less attention in the economics literature. We draw from a diverse set of papers to bring forth some insights into the fundamental trade-offs that researchers face.

Our aim is three-fold. First, we shed light on some of the subtle economic forces at play in grant allocation. For example, we highlight how noise in the evaluation process drives application incentives, leading to unintended consequences of tying the funding budget to the volume of applications, as is often done in practice. Second, we highlight a number of principles of grant design. For example, the analysis will help us better understand how a funding lottery may help grantmakers economize on evaluation costs, but why a lottery may not be helpful for encouraging applicants to improve the quality of their proposals. Third, we outline problems where further empirical and theoretical economic research would help our understanding of grant allocation.

With this chapter, we hope to complement the recent overviews by Price (2019) and Azoulay and Li (2021). Price (2019) explains the details of grant allocation at the National Institutes of Health (NIH) and offers a rebuttal against common critiques of grants. Azoulay and Li (2021) review the empirical economics literature on grants and discuss various aspects of practical grant design.

In the next section we briefly discusses the relative strengths of grants compared to other funding instruments, and we lay out the goals and challenges in grant design. Then, in Section 3, we zoom in on solutions to these challenges.

2 Goals and challenges of grant allocation

2.1 Goals of science funding

Scientific research is central to the economic well-being of modern societies since it creates valuable new knowledge. As explained by Arrow (1962), however, there are several reasons why one may think that markets fail to incentivize an efficient production of knowledge. First, while scientific research consumes large amounts of resources, its output, knowledge, is non-rival since it can be reproduced at (almost) no cost. Second, the production of knowledge involves substantial risk (Franzoni and Stephan, 2023); it is uncertain whether a research endeavor will generate new findings and how valuable these findings will be. Third, knowledge production creates positive externalities by inspiring follow-up research.

For these reasons, there is an economic justification for government interventions that adjust incentives in the market for the production of scientific knowledge. Governments around the world recognize the importance of science policy for the prosperity of societies in general, but also as a strategic investment in national competitiveness and economic security (Dasgupta and David, 1994). Jones (2021) summarizes a large body of evidence that points to the large social gain from subsidizing science and innovation.

Even accepting the need for policies to encourage knowledge production, it is not at all clear what these policies should be. In practice, governments and institutions support research using various instruments with different strengths and weaknesses. For example, grants are awarded before any knowledge is produced and the effective allocation of grants requires a costly ex ante evaluation of proposals. By contrast, the patent system does not generate evaluation costs ex ante. Instead, costly evaluation happens at the patent application stage once the knowledge has been produced. In addition, there may be further costs in the form of market distortions from the temporary monopoly power enjoyed by the patent holder.

To the best of our knowledge, there are no theoretical economic analyses comparing different funding instruments except the stylized treatments by Gallini and Scotchmer (2002) and Wright (1983). In Section 2.2, we delineate the situations where scientific grants seem more appealing than two other ubiquitous instruments—patents and prizes. Then, in Section 2.3, we turn to the challenges in grant design.

2.2 Why support research through grants?

First, grants may be useful for aligning market values with social values. Broadly speaking, a mismatch between the two values may arise from market participants' limited ability or willingness to pay for research output; Azoulay and Li (2021) and Price (2019) give the example of medical treatments, such as Malaria treatments for people in sub-Saharan Africa. A mismatch between social and market values diminishes the appeal of those funding instruments, such as patents, that create incentives by linking the rewards from research to market returns. Bryan and Lemus (2017) provide theoretical support for the idea that patents fall short of guiding researchers to the right topics. In their framework, researchers choose both the intensity and the direction of their research. They show that, while patents can induce the efficient intensity of research, patents do not guarantee directional efficiency. One could conjecture that, here, grantmakers can step in by committing to identifying and supporting socially valuable topics.

Second, grants help address the uncertainty inherent in research. Fundamental research generally involves high uncertainty about the value of discoveries (in addition to the uncertainty of whether something will be discovered at all). Further, often the commercial use of research output reveals itself only in the distant future; see Azoulay et al. (2019b), Bryan and Williams (2021), and Li et al. (2017) for some evidence that grant-funded projects indeed lay the foundation for patented inventions. From theoretical work by Aghion et al. (2008), we can glean some intuition for how this uncertainty impacts funding design. In their framework, researchers are willing to forgo high salaries in exchange for freedom of academic pursuit, consistent with empirical evidence by Stern (2004). Aghion et al. (2008) argue that, in the early stages of research, which is when the commercial value of projects is most uncertain, it is efficient to leave scientists the freedom to decide which questions to pursue, even "if this entails some probability of the scientists wandering off in other directions." Interpreting grants as providing unconditional funding for researchers, their framework provides a rationale for financing fundamental research through grants.

Third, and somewhat related to the previous point, grants may be most suitable when it is difficult to spell out what exactly constitutes a successful research project. In some cases, it is roughly clear what success means. For example, in 1908, the Wolfskehl Prize was announced for the first person to prove Fermat's Last Theorem;

in 1997, the prize was awarded to Andrew Wiles after his proof was deemed complete by the scientific community and published in a journal. However, as argued above, research often takes place at a stage at which the outcomes and the value of undertakings are highly uncertain. In these cases, an upfront payment in the form of a grant may constitute a more compelling incentive to tackle a problem than, say, a prize for attaining a nebulously specified goal.

Fourth, as noted by Azoulay and Li (2021) and Price (2019), grants are useful for targeting specific people and institutions. For example, scientific apprenticeships, expensive laboratory equipment, and other infrastructure are typically financed via grants. Thus grants can direct resources towards underrepresented minorities, or compensate endowment differences across universities to level the playing field.

2.3 Challenges in grant funding

Despite the merits outlined above, grant allocation is not without its own challenges and pitfalls. How can grantmakers identify the most meritorious applications? How can grantmakers be certain that, as Gallini and Scotchmer (2002) put it, grantees do not "take the money and run?" The efficiency of the allocation thus depends on how well grantmakers can screen applications ex-ante, and on whether funded researchers work hard at executing their proposals. Additionally, one would like to achieve all this while economizing on the costs of preparing and reviewing proposals, monitoring awardees, and so forth.

A central premise of modern economic analysis is that the design of institutions is marred by the simultaneous presence of *misaligned interests* and *asymmetric information.*

In the context of a grantmaker deciding whom to fund, it is clear that there are *misaligned interests*: the grantmaker wishes to support the researchers with the "best" project, whereas each researcher may primarily care about receiving funding.

Asymmetric information is typically classified into *hidden information* and *hidden action.*

(1) Hidden information refers to private knowledge of one party that is relevant for the payoffs or strategic incentives of other parties. For example, on the one hand, researchers may have better information than grantmakers about the intrinsic quality of their proposals, but on the other hand researchers may be

uncertain about what exactly grantmakers hope to find in a proposal.

(2) Hidden actions refer to behavior by one party that cannot be observed or controlled by others. For example, after awarding a grant to a researcher, grantmakers may be unable to enforce how exactly the researchers use their grants.

Both forms of asymmetric information features seem particularly pertinent in the context of research because topics and research practices are highly specific and require specialist knowledge to evaluate. The simultaneous presence of misaligned interests and asymmetric information gives rise to incentive problems.

Prospective and retrospective evaluation are the basic instruments that grantmakers can use to mitigate incentive issues. Prospective evaluation (through referee panels, for example) reduces informational asymmetries prior to the allocation. Retrospective evaluation (through publication metrics, for example) rewards researchers for their output, thereby aligning researchers' incentives with those of the grantmaker and mitigating the hidden action problem. Retrospective evaluation, as it turns out, also helps identify meritorious applicants when allocating funds. In the next section, we discuss in detail ways of mitigating incentive issues via prospective or retrospective evaluation.

To conclude this section, it is worth noting that the incentive issues outlined above are partially addressed by the *non-monetary* incentives that organize the conduct of science. Non-monetary incentives include intrinsic intellectual curiosity as well as recognition in the scientific community from publications and establishing priority. Indeed Merton (1957) emphasizes the importance of the priority system for the incentives of researchers. Researchers establish priority by being the first to break ground on important problems. Hill and Stein (2023) find a significant and large negative effect of narrowly failing to establish priority (though they note it is not a winner-take-all contest). Citations reward researchers proportionally to the impact of their work. Thus, the system of non-monetary incentives would seem to partially align researchers' interests with those of grantmakers by encouraging work on important problems. Alas, the system is imperfect. The pressure to publish can have a detrimental impact on the conduct of researchers. When several researchers work on the same problem, there is a race to be the first to publish. Hill and Stein (2024) provide evidence that structural biology projects with higher recognition potential are completed faster and are of lower quality. Further, there is evidence of publication bias against null results (see Andrews and Kasy, 2019, for example). The importance of citation measures in quantifying scientific impact and in promotion decisions can induce researchers to, possibly inefficiently, prioritize research fields that generate higher citation counts over others (see Olszewski, 2020, for example). In summary, even though non-monetary incentives surely play an important role, it is worth examining explicitly the incentive problems tied to the allocation of resources.

3 Grant design

In this section, we discuss remedies to the incentive problems outlined above. We proceed as follows. In Section 3.1, we study the application process, emphasizing the self-selection of applicants. In Section 3.2, we focus on the allocation rule itself. A common theme across Sections 3.1 and 3.2 is on how *prospective* evaluation helps grantmakers allocate efficiently. Next, in Section 3.3 we focus on the benefits of *retrospective* evaluation and how researchers' incentives are shaped by the structure of payments. Finally, in Section 3.4 we consider how the grantmaker choice of what to fund impacts the direction of research.

3.1 The application process

Grant programs thoroughly evaluate proposals via peer review. The NIH, for example, is required by law to do so (Price, 2019). Only the applications deemed to be most promising are funded. For an indication of how competitive the process is, consider that the current fraction of funded applications for "Advanced Grants" at the European Research Council (ERC) is around 10%. Preparing an application is a time-consuming and costly activity,[2] and hence it is natural to expect that candidates are more willing to apply when they are more optimistic about obtaining funding. The costly ordeal of preparing an application thus helps grantmakers screen applicants through self-selection. Conversely, application rates respond to changes in the application costs and the evaluation process.

Adda and Ottaviani (2023) analyze the impact of prospective evaluation on the self-selection of applicants. The baseline version of their model can be understood in terms of demand and supply for grants.

[2]Survey evidence from Hippel and Hippel (2015) indicates an average time of 116 hours spent writing an application for astronomers.

On the demand side, researchers differ in their intrinsic "merit." To give a concrete example, suppose that merit corresponds to the number of additional papers a researcher can publish if awarded the grant. Researchers enjoy a benefit if funded, but bear applications costs.

On the supply side, the grantmaker can only fund a fraction of the applicants, and would thus like to assign grants to the most meritorious applicants. However, the grantmaker only observes a noisy signal about each applicant's merit (by conducting peer review, for example). From the grantmaker's perspective, a higher signal indicates a higher merit of the applicant. Therefore, the grantmaker funds the applicants that obtain a signal above a certain acceptance threshold; the threshold is chosen to exhaust the grantmaker's budget.

Anticipating this allocation rule, only those researchers who are sufficiently optimistic to be evaluated positively will apply. Researchers with high merit know that they are more likely to generate a high signal in the evaluation process. Therefore, there is an application threshold (distinct from the acceptance threshold) such that only researchers with merit above the application threshold will apply.

Building on this characterization, Adda and Ottaviani (2023) shed light on the effect of changes in the grantmaker's budget and evaluation procedure.

First, an increase in the budget of the grantmaker, naturally, increases the incentive to apply, thereby reducing the average merit of applicants. The effect of a budget increase on the success rate—the ratio of funded applicants to the total number of applicants—is more subtle and depends on the distribution of merit across researchers. On the one hand, a budget increase incentivizes more applications (as just noted), pushing the success rate down. On the other hand, a budget increase allows for more applicants to be funded, pushing the success rate up. Thus, the success rate can either increase or decrease. In particular, it decreases if the elasticity of applications with respect to the budget is greater than one.[3] This appears in line with evidence from the 2009 increase in the budget available for research grants in the US due to Obama's Stimulus Package. Grant applications increased more than the budget, thus resulting in a reduction in the fraction of successful applicants (Stephan, 2012, p 145).

[3]Technically, this case arises when the distribution of researchers' merits follows a distribution with thicker tails than exponential. This assumption is in line with an early observation by Lotka (1926) that researchers' productivity in terms of publications follows a power law.

A key insight of Adda and Ottaviani (2023) pertains to the impact of evaluation noise on application incentives. The noise could reflect how carefully evaluators read proposal, and whether a funding lottery is used for applicants at the cusp. A noisier evaluation procedure increases the incentives to apply and thus reduces the self-selection of applicants. Intuitively, as the evaluation becomes noisier, the probability of succeeding in obtaining a grant becomes less responsive to merit, encouraging more low-merit applications. This result suggests that the design of the allocation rule has important consequences for the self-selection of applicants.

The apportionment of budget across different fields is more delicate yet. Research funding organizations grapple with the incentives of panel members to favor research in their own fields of research. Panel members have an incentive to inflate the scores of applications in their field in an attempt to secure more funding for their field. In practice, the overall budget is often apportioned using a mechanical formula. The NIH uses such a system since 1988. The ERC adopted from its inception a similar scheme by apportioning its budget in proportion to the funds requested by applicants in each panel. Proportional budget allocation automatically equalizes the success rate across fields.

At a first glance, proportional apportionment seems fair and balanced. In the framework of Adda and Ottaviani (2023), the system indeed performs well if fields are relatively similar in terms of the noise in the evaluation signal. If, instead, fields are heterogeneous in terms of noise, the performance deteriorates. Noisier fields attract more applications (as argued above) which, proportional apportionment, leads to a proportional increase in the budget; this budget increase in turn induces a further increase in applications. To see how adverse this effect can be, consider as an extreme the case of a field with perfect evaluation (meaning all applicants can predict whether their application will succeed). Applicants who expect to fail, do not apply, and thereby reduce the funding for other applicants from the same field. Whenever the success rate is less than 100%, the process will continue until in equilibrium no applications at all will be submitted from this field.

Beyond this stark illustration, Adda and Ottaviani show that under a general class of allocation rules, including the proportional rule, a reduction of noise in a field leads to a reduction in applications in that field and an increase in applications in all other fields. Adda and Ottaviani also present a broad empirical confirmation of this prediction by exploiting a natural experiment of a change in the application of the proportional apportionment rule at the ERC.

The impact of noise is adverse in terms of efficiency: in fields with low evaluation noise, it is easier to allocate efficiently, and hence, all else equal, these fields should be endowed with more funding. An important challenge for future work is to identify rules that are at once desirable in terms of efficiency and politically feasible.

A second important challenge is to reduce the application and evaluation costs inherent to this system. We briefly discuss *temporary exclusion* with this challenge in mind. By temporary exclusion, we mean that unsuccessful applicants are debarred from reapplying for some period of time. At the ERC, for example, applicants to Starting or Consolidator Grants who are rejected in the early stages of the evaluation process may not reapply to the same grant for two subsequent years (ERC, 2023, p.9).[4] The aim of this restriction is "to allow unsuccessful Principal Investigators the time necessary to develop a stronger proposal."

Azrieli (2024) studies temporary exclusion in a dynamic model similar to the model of Adda and Ottaviani (2023). The threat of temporary exclusion deters researchers from applying if they know that current project is weak and they expect to have a stronger one in the future. This basic logic suggests temporary exclusion indeed reduces excessive applications. The full picture is more nuanced because of a countervailing force: if one's competitors are less likely to apply, then one's own incentive to apply increases. Azrieli (2024) shows that temporary exclusion can at once increase the overall welfare of researchers and decrease the number of applications. In particular, this result holds when the costs of applying are small relative to the gains from obtaining a grant, reflecting situations where the decision to apply is primarily driven by the dynamic considerations described above rather than by application costs.

The effect of temporary exclusion on the grantmaker's payoffs is ambiguous. On the one hand, temporary exclusion increases the expected merit of applicants, and reduces evaluation costs. On the other hand, to exhaust the budget, the evaluator uses a laxer acceptance cutoff, thereby decreasing the expected merit of the funded applicants. Therefore, the effect on the grantmaker's payoffs depends on the overall effect on the expected merit of funded candidates, and on how the grantmaker values project quality relative to evaluation costs.[5]

[4]See `https://erc.europa.eu/apply-grant/starting-grant`.

[5]Azrieli (2024) does not explicitly model the overall payoffs of the grantmaker or evaluation

Further reading In Appendix A.1, we show how to derive the comparative statics result with respect to evaluation noise using techniques pioneered by Lehmann (1988).

3.2 The allocation rule

In the previous subsection, we focused on the application process, taking a stylized perspective on the evaluation and allocation process. We now zoom in on optimal allocation rules through the lens of *mechanism design*.

Mechanism design is a subfield of economic theory that studies how to design institutions—*mechanisms*—to achieve a pre-specified goal. Typically, this goal will not be perfectly achievable. For example, a grantmaker may wish to award grants to the most fund-worthy applicants while economizing on the time spent reviewing proposals. The applicants, who know the nuts and bolts of their proposals, have private information about their fund-worthiness that should guide the efficient allocation. However, if the grantmaker would completely save on evaluation costs by never triggering a review, thereby solely relying on what the applicants claim in their proposals, then the applicants would have incentives to overstate their fund-worthiness. Hence, we do not expect the grantmaker to achieve the first-best outcome of funding the best proposals at zero cost. Notice that this impossibility is a consequence of the simultaneous presence of misaligned interests and asymmetric information (and the fact that there are not enough resources for everyone).

When the first-best is out of reach, mechanism design sheds light on achievable *second*-best outcomes; how does a mechanism optimally balance allocative efficiency with evaluation costs? For example, a recurrent theme in the upcoming sections is that optimal mechanisms often feature some deliberate randomization on the part of the grantmaker. Thus, in this case, the analysis suggests a practically relevant feature of allocation rules and illuminates which forces justify this feature from an efficiency perspective. On a more conceptual level, the analysis provides guidance on how grantmakers can benefit from combining certain instruments (such as a costly review process) with self-reports by the applicants.

Notice that we take the objective of the grantmaker as given; the question is how to implement it. We will be silent on what the grantmaker's notion of fund-worthiness should be, or on how the grantmaker's objective should weigh fund-worthiness against

costs, but we can deduce these informal observations from Azrieli's analysis.

one hour of the reviewers' time.

Before delving into the details, we emphasize that, unless mentioned otherwise, the models discussed below are cast as abstract allocation problems. In particular, we do not claim that the papers from which we draw are focused on the specific problem of allocating grants. The interpretation of the papers' results in the context of grant allocation is our own.

3.2.1 Costly prospective evaluation

How can a grantmaker fund the most promising research while economizing on evaluation costs? Here we study this question using the model of Ben-Porath et al. (2014).

Consider a grantmaker who has one grant to allocate among a group of applicants. Each applicant wants to obtain the grant and privately knows their merit. For the sake of concreteness, suppose each applicant has some estimate about the number of additional publications they could produce if funded. The grantmaker wishes to maximize this number and has no other use for the funds. If the grantmaker were to ask each applicant to self-report their estimate and then fund the top applicant, the applicants would have an incentive to exaggerate. The grantmaker has an additional instrument at their disposal: they can *verify* individual applicants at a cost. Verifying an applicant reveals the applicant's merit. We interpret the act of verifying as conducting an in-depth review by external evaluators; verification costs represent the time burden this act imposes on the evaluators, capturing the idea that the grantmaker wants to economize on evaluation.

To briefly elaborate on the interpretation (see Ben-Porath et al., 2014, p. 3803), what really matters is that verification reveals everything the applicant knows; that is, if the applicant's estimate for the number of additional publications is determined by a set of objective facts about the applicant's research that are ex ante only known to the applicant, then verification reveals those fact. Actually, it is unimportant whether the applicant and grantmaker agree on what the facts imply, but for ease of exposition we will assume that they do agree. Further, there could well be residual uncertainty about the number of additional publications, and neither the applicant nor the principal can further resolve this uncertainty.

In this environment, a mechanism specifies which reports the applicants can make to the grantmaker, which applicants are verified depending on those reports, and how the grant is eventually allocated depending on the reports and verification outcomes. To illustrate, here are three examples of mechanisms. First, the grantmaker could randomly select a winner without requesting any report or verifying any applicant. Second, the grantmaker could outright verify all applicants and then allocate the grant to the most meritorious one. Third, the grantmaker could approach applicants sequentially, asking for self-declarations, possibly verifying some of them, and stopping when an applicant seems "good enough." Given a mechanism, applicants strategically choose which reports to make to maximize the probability of receiving the grant, forming conjectures about what their competitors might do. The scope of possible mechanisms and outcomes is thus quite complex.

Which mechanisms maximize the expected merit of the funded applicant net of verification costs? Ben-Porath et al. (2014, Theorem 1) show that there is an optimal mechanism of the following form:[6] The grantmaker announces a merit threshold and asks all applicants to self-report their merit. If all reports are below the threshold, the grant is allocated uniformly at random. If instead at least one report is above the threshold, then the grantmaker verifies the highest report and, if the report is verified to be true, awards the grant to the corresponding applicant. If an applicant is verified and found to have misreported their merit, then that applicant is *not* allocated the grant.[7] Given these allocation and verification rules, all applicants find it in their best interest to report truthfully; indeed, when misreporting a merit above the threshold, applicants anticipate they will be found out.

Notice that the mechanism allocates the grant through a lottery when no report is sufficiently promising. This feature speaks to the growing interest in lotteries (e.g., Bendiscioli et al., 2022; Heyard et al., 2022). A key argument put forth by Fang and Casadevall (2016) in favor of lotteries is that identifying the best proposals is extremely costly. This concern is roughly reflected in the optimal mechanism. The

[6] The solution described here applies to the case where the applicants are ex-ante symmetric from the perspective of the grantmaker. Ben-Porath et al. (2014) also cover the general asymmetric case.

[7] In this case, as long as the misreporting applicant is denied the grant, there is a lot of flexibility regarding how to proceed with the allocation. The reason for this flexibility is that in the described mechanism the applicants expect that everyone will report truthfully, and therefore they disregard completely the event in which one of the other applicants is verified to have misreported. One way of proceeding would be to allocate the grant randomly among the remaining agents, or not to allocate at all. Another perhaps more reasonable way is to sequentially verify the remaining highest reports until someone is verified to have been truthful, and then award that applicant; in case everyone is verified to have been untruthful, no one is awarded.

grantmaker could verify all reports and guarantee that the best applicant is assigned the grant with certainty. However, the grantmaker chooses not to do so, saving on verification costs by randomly allocating the grant when all reported merits are below the threshold. That said, the mechanism does identify and allocate to the best applicant whenever their merit is above the threshold.[8]

Ben-Porath et al. (2014, Section IV) show that for higher verification costs the optimal mechanism uses a higher threshold; so, there are fewer in-depth reviews and allocation is more frequently via a lottery. They also show that (in the symmetric case) the threshold increases in response to shifts of the distribution of merit in the sense of first-order or second-order stochastic dominance; that is, roughly speaking, if applicants become more meritorious or less heterogeneous, then the threshold increases, meaning the grantmaker more frequently allocates using a lottery.

In practice, grantmakers may not only rely on the outcome of an in-depth review; they also have access to a relatively cheap but noisy signal, perhaps from a preliminary screening of the applicants. Kattwinkel and Knoepfle (2023) study how to optimally combine such a noisy signal with self-reports and verification in a setting with a single applicant (thus, the decision is between funding the applicant and keeping the funds). Kattwinkel and Knoepfle show that, optimally, the grantmaker reveals the results of the preliminary review to the applicant, and then allows the applicant to submit an appeal to be verified. The applicant is funded if the preliminary review is sufficiently positive, or, if upon submitting an appeal the applicant is verified to have sufficiently high merit. An intriguing feature of this mechanism is that the grantmaker does not benefit from obscuring the result of the preliminary review. That is, the grantmaker can use a transparent procedure for evaluating the applicant without sacrificing on economic efficiency. This feature is surprising since, in theory, the principal could try to cross-check the agent's self-report with the preliminary review to deliver stronger incentives for reporting truthfully, potentially letting the grantmaker save on verification costs.

Further reading In Appendix A.2, we sketch some of the formal arguments to derive the comparative statics results discussed above, as well as an optimal mech-

[8] The fact that the optimal mechanism uses a lottery does rely on the assumption that the applicants are ex-ante symmetric (recall footnote 6). In the asymmetric case, there is an optimal mechanism that deterministically allocates to a so-called *favored agent* if no self-report passes the threshold (Ben-Porath et al., 2014, Theorem 1).

anism in the single-applicant case. As for related work, Ben-Porath et al. (2023, 2019) consider, among other things, nearby settings where the grantmaker cannot verify the applicants' claims; instead, the applicants themselves can provide hard evidence about their merit, but must be incentivized to do so. Khalfan (2023) studies a variation where applicants only have a noisy signal about their merit and hence do not know what an in-depth review will reveal. This detail complicates the grantmaker's problem as applicants can now feign ignorance if a review reveals a low merit. Epitropou and Vohra (2019) considers a variation where the applicants arrive sequentially. Li and Libgober (2023) consider a repeated interaction with a single applicant.

3.2.2 Peer selection

Next, we discuss a setting where the applicants for a grant and the reviewers are the same group of people.

On an abstract level, this setting fits the problem of distributing resources within the scientific community. There is a finite amount of resources to allocate. To allocate these efficiently, grantmakers would benefit from knowing how each scientist evaluates their peers. Of course, there is no single grantmaker who allocates all resources, and not all scientists simultaneously apply for a grant, but the model serves as useful abstraction to shed light on an economic tension that we explain momentarily.

On a more concrete level, it has been suggested to *require* applicants to evaluate their competitors' proposals. Merrifield and Saari (2009) suggest such a procedure for allocating telescope with the aim of alleviating the burden that a conventional process would place on external reviewers. A 2014 NSF pilot later used such a procedure to allocate grants; the Gemini Observatory to allocate time on its telescope in Hawaii (Mervis, 2014). The 2016 Neural Information Processing Systems conference, a top-tier machine learning conference, asked authors to volunteer as reviewers to handle the enormous number of submissions (Shah et al., 2018).

Here, we focus on one particular issue that appears relevant when applicants review their competitors: dishonestly exaggerating one's own merits and others' faults. See, for example, the laboratory experiment of Balietti et al. (2016) or the field experiment of Hussam et al. (2022) for evidence that individuals indeed misreport when evaluating their peers (albeit not in the context of scientific grant funding).

We consider a stylized setting, borrowing from work by Alon et al. (2011). Suppose that each applicant has an honest private opinion about who they would nominate as the most deserving recipient of the grant. The grantmaker wishes to allocate to the applicant with the most nominations. At the same time, the grantmaker wishes to use a mechanism under which no applicant can influence their own chance of winning by misreporting their nomination. In the literature, such a rule is called *impartial*.[9]

The key economic tension in this problems stems from the grantmaker's desire to elicit a relative ranking of the applicants. Applicants can improve their rank not only by appraising themselves but also by diminishing the others. Even if self-nominations are prohibited, an applicant may be able to improve their rank by claiming that all others are unworthy of the grant. Notice a key difference to the setting of Ben-Porath et al. (2014). There, the grantmaker is not trying to elicit the applicants' opinions about their peers. Here, these opinions are the key objects of interest.

Here is one example of an impartial mechanism. First, before eliciting any nominations, the grantmaker randomly splits the pool of applicants into two groups—"evaluators" and "candidates". Second, the grantmaker counts only the nominations of the evaluators to pick a winner among the candidates.

Under this mechanism, evaluators know that they are excluded from winning (and hence cannot improve their winning chances by misreporting), and candidates know that their reports are not counted. In practice, instead of randomly splitting applicants into two groups to award one grant, one could also split the grant into two parts and have the evaluators for the first part be the candidates for the second part, and vice versa.

The partition mechanism just described is intuitive and takes to heart a central tenet of peer review—evaluators should have no personal stake in the funding decision. However, there are many other impartial mechanisms, many of which possess desirable properties that the partition mechanism lacks. For example, the partition mechanism performs badly if the applicant with the most approvals happens to have been randomly assigned the role of evaluator.

Fischer and Klimm (2014) show that the grantmaker can do better by splitting applicants into more than two groups (though the procedure for aggregating their nominations becomes more intricate). They propose a mechanism that guarantees the

[9]Some papers of the literature instead use the terms *strategyproofness* or *dominant-strategy incentive-compatibility* instead of *impartiality*. These notions all coincide under the assumption that each applicant seeks to maximize their individual chance of winning and is otherwise indifferent to who wins.

selection of an applicant who receives at least half as many approvals as the applicant with the most approvals; the fraction of one half is, in fact, the best theoretical guarantee (Alon et al., 2011; Fischer and Klimm, 2014). Niemeyer and Preusser (2023) provide conditions under which the partition mechanism is approximately optimal among all impartial mechanisms when there are many applicants, albeit under a different notion of optimality.[10]

Further reading A large number of papers build on the framework of Alon et al. (2011). For example, Kurokawa et al. (2015) impose the additional constraint that each applicant can only review a limited number of other proposals. De Clippel et al. (2008) and Holzman and Moulin (2013) take a different approach, studying which other desirable properties are compatible with impartiality. Olckers and Walsh (2023) survey the literature on impartial mechanisms. In related settings, Bloch et al. (2023) and Kattwinkel et al. (2022) consider mechanisms that are not necessarily impartial but which nevertheless provide the applicants with incentives for truthful reports; for example, they allow for mechanisms where the grantmaker cross-checks applicants' reports to detect misreports.[11]

3.2.3 Incentivizing investment

Up to this point, we treated each applicants merit as exogenous; applicants do not control the quality of their proposals. However, we would think that in practice researchers adjust their proposals in response to, say, the fierceness of their competition or the grantmaker's evaluation criteria.

In this section, we discuss how the grantmaker's mechanism incentivizes productive investment, based on work by Augias and Perez-Richet (2023). The headline insight is that, here, the grantmaker does *not* benefit from allocating randomly. Moreover, the optimal evaluation is noisy for low-merit proposals, but exact for high-merit

[10] Alon et al. (2011) and Fischer and Klimm (2014) evaluate a rule using a worst-case criterion: in the worst case, what is the ratio between the nominations received by the chosen applicant and the most nominations received by any individual applicant? Niemeyer and Preusser (2023) instead evaluate a rule by computing the number of approvals of the chosen applicant, taking expectations with respect to the grantmaker's belief about the profile of nominations.

[11] In technical terms, Bloch et al. (2023) and Kattwinkel et al. (2022) consider *Bayesian incentive-compatible* mechanisms when the applicants' private types are correlated across applicants or with an external signal. See also Kattwinkel (2020). As noted in footnote 9, impartiality coincides with the stronger notion of dominant-strategy incentive-compatibility.

ones.

In the model, there is a unit mass of grants that can be assigned to a unit mass of applicants. Each applicant is endowed with an initial merit level, which has some distribution in the applicant population. Crucially, applicants can invest in their merit, incurring a cost that increases in the distance between the initial and the final merit. Applicants care about the probability of receiving funding minus the investment costs. The grantmaker observes each applicant's final merit (but neither initial merit nor investment) and then decides whom to give a grant. In the baseline model, the grantmaker has enough resources to fund everyone, but would only like to fund those whose final merit exceeds some threshold; let us call this the grantmaker's preferred threshold.

Augias and Perez-Richet (2023, Theorem 1) provide a condition on the distribution of initial merit under which a deterministic threshold mechanism is optimal.[12] This mechanism deterministically funds the applicants whose final merit exceeds the mechanism's acceptance threshold; all others are rejected. The acceptance threshold is higher than the grantmaker's preferred threshold. That is, the grantmaker commits to rejecting some applicants who the grantmaker would actually like to fund. This commitment to being tough generates powerful investment incentives.

The condition on the distribution of initial merit entails that the grantmaker's preferred threshold lies in the upper tail of the distribution. We interpret this condition to mean that competition is tough: the grantmaker only wants to approve exceptional applicants.

The optimal mechanism has an interesting alternative interpretation (Augias and Perez-Richet, 2023, Section 5.3): instead of committing to being tough, the grantmaker allocates optimally based on a *noisy* evaluation of final merit. One way of implementing the optimal mechanism via a noisy evaluation is as follows: if an applicant's final merit fails to pass the acceptance threshold, the evaluation only reveals that the merit failed to pass; else, if the final merit passes the threshold, the evaluation perfectly reveals the final merit. The coarse reporting of failure is important for circumventing the lack of commitment to being tough; if the grantmaker learned

[12] Interestingly, a deterministic threshold mechanism is also optimal (under the condition on the merit distribution) if the grantmaker is concerned with welfare (Augias and Perez-Richet, 2023, Proposition 5). Here, welfare means the final merit of the funded applicants, plus the private payoffs of the funded applicants from winning, minus total investment costs. However, the value of the threshold changes relative to the case where the grantmaker only cares about final merit.

that an applicant narrowly failed to clear the acceptance threshold, a non-committed grantmaker would be tempted to approve the applicant, thereby upsetting the investment incentives. These observations give us a sense of how the grantmaker benefits from delegating the evaluation to referees who are tougher than the grantmaker. The alternative interpretation of the optimal mechanism also speaks to the optimal level of evaluation noise: to stimulate productive investments into merit, evaluation should be exact at the top but coarse at the bottom.

Let us briefly elaborate on why it is quite subtle that there is an optimal deterministic mechanism. It is key to understand the induced investment incentives. Given a deterministic threshold mechanism, applicants with initial merit narrowly below the threshold will invest to win a grant. However, applicants with initial merit above the threshold or too far below the threshold have no incentive to invest. The grantmaker could try to generate better investment incentives by approving *randomly* for final merit levels near the threshold. If an applicant's initial merit is above the threshold, they are now incentivized to invest so as to escape the lottery. Of the applicants with low initial merit who previously did not invest, some are now incentivized to invest a little to become eligible for the lottery. The downside of approving randomly, of course, is that some applicants that the grantmaker would like to approve will be rejected. Further, some applicants who previously invested to push themselves above the threshold are discouraged from investing as they are now only rewarded randomly.

Further reading Since the optimal mechanism is implementable via noisy evaluation, there is a tenuous connection to work on contest models. Specifically, Morgan et al. (2022) study how investment incentives depend on the noisiness of the evaluation (when the grantmaker is non-committed). An important caveat is that Morgan et al. (2022) focus on the total investment across all applicants, including those that go unfunded (as is typical in the contest literature).[13] This focus potentially limits the applicability to grant allocation since grantmakers may care more about the effort of the applicants that are funded than those who are unfunded, especially if the projects of unfunded researchers are never realized.

[13]We refer to Fu and Wu (2019) and Vojnović (2015) for further reading on contest models.

3.2.4 Falsification

Applicants may not only undertake productive investments in their proposals. They may instead spend considerable effort on falsification, letting their proposals appear stronger than they actually are. This non-productive effort is socially wasteful, hampers the grantmaker in identifying good applicants, and hurts the political sustainability of the grant-awarding institution. We briefly discuss mechanisms that anticipate such falsification incentives, using work by Perez-Richet and Skreta (2022).[14]

A high-level insight is that optimal mechanisms involve *productive falsification*. The rough idea is as follows. Even if everyone can craft a brilliant proposal, doing so will be easier for those with better research ideas. The fact that applicants will give in to the temptation of falsifying their proposals thus gives the grantmaker an (imperfect) instrument for inferring their underlying merit. Optimal mechanisms exploit this falsification incentive productively.

Another insight is that, depending on the magnitude of the falsification costs, the optimal mechanism again involves a lottery. Intuitively, a lottery deters excessive falsification by lowering the reward from falsification to a high level. An added welfare benefit is that fewer resources are devoted to wasteful falsification, which seems in line with common arguments in favor of lotteries (for example, Gross and Bergstrom, 2019).

Further reading Even though the grantmaker may benefit from productive falsification, Perez-Richet and Skreta (2024) note that mechanisms prone to falsification may be politically unsustainable and generate unfair advantages for applicants with higher gaming abilities. Perez-Richet and Skreta (2024) study optimal mechanisms that are immune to falsification attempts. Li and Qiu (2024) study a related setup with multiple applicants and multiple grants. Gross and Bergstrom (2019), roughly, study a variation of the above setup where the grantmaker lacks commitment power. Gross and Bergstrom argue that lotteries may improve welfare. Myers (2022) considers a similar setup as Gross and Bergstrom (2019) and argues that lotteries may fail to improve welfare if falsification effort also generates positive externalities (for example, if attempts to embellish one's proposal generates genuinely valuable ideas or clarifies the exposition).

[14] Here we use Proposition 1 of Perez-Richet and Skreta (2022, p. 1118) to reinterpret the model that they introduce in their Section 2.

3.3 Retrospective evaluation and post-award management

While the prospective evaluation of applicants surely plays an important role, grantmakers and funded reseachers continue interacting after the grant is conferred. Post-award management and the retrospective evaluation of researchers expand the toolbox that grantmakers have at their disposal. Goldstein and Kearney (2020) document the extent of post-award management of grants at ARPA-E; see also Azoulay et al. (2019a).

As suggested earlier, a concern in the post-award stage may be that grantmakers cannot monitor how researchers exactly use their funds (the hidden action problem). We next illustrate, using the simple setup of Maurer and Scotchmer (2004, Section 5), how retrospective evaluation helps mitigate the hidden action problem. As an added benefit, retrospective evaluation turns out to also help the grantmaker in the *pre*-award stage.

3.3.1 A simple model of post-award management

The model of Maurer and Scotchmer (2004) unfolds over multiple periods. A researcher in each period needs a grant to carry out their project. When awarded a grant, the researcher obtains an immediate private benefit corresponding to, say, career advances or reputation. Working on the project is costly for the researcher, but benefits society through knowledge production. The researcher can also choose to shirk, for example by diverting the funds to other activities. Whenever the research does not work on the project, no knowledge is produced.

Here we have a classic hidden action problem: the grantmaker wants the researcher to work, while the researcher has an incentive not to work as the costs of research are borne privately.

The model also features a hidden information component. Researchers differ in how likely they are to develop fund-worthy research ideas in future periods. This productivity is privately known to the researcher but not to the grantmaker.

Suppose that the grantmaker can retrospectively evaluate the funded researchers at zero cost. At the end of each period the grantmaker learns whether the funded researcher worked on their project. In particular, the grantmaker can withhold future funding from researchers who shirked in the past.

This policy of retrospective evaluation is effective whenever the prospect of receiving future funding is sufficiently valuable to researchers. It turns out that this is not the case for all researchers. Only those researchers who are sufficiently optimistic about having fund-worthy ideas in the future are willing to incur the costs of working today. For the other researchers, the threat of losing access to funds in the future has little bite since they consider it unlikely to have fund-worthy ideas in the future.

Thus, the optimal grant-allocation policy in Maurer and Scotchmer (2004) has the following structure. The grantmaker assigns a grant to every first-time applicant. In the following periods, the grantmaker evaluates each researcher's past performance and only awards a grant to researchers who delivered results in the past. It follows that researchers with high productivity in generating ideas will work whenever they hold a grant and receive a grant whenever they apply for one. The other researchers will receive a grant only once and not deliver results. Thus, retrospective evaluation and repeated grant allocation jointly allow the grantmaker to channel funds to productive researchers and to provide them incentives to use the grant effectively.

Further reading In Appendix A.3 we spell out the Maurer-Scotchmer model in more detail.

3.3.2 The pre-award benefit of post-award management

In our discussion of the model of Maurer and Scotchmer (2004), we suggested that post-award management also helps screen productive researchers prior to the allocation. In this section, we develop this idea further, using the model of Mylovanov and Zapechelnyuk (2017) to study how retrospective evaluation helps the grantmaker choose whom to fund.

A grantmaker allocates one grant among a group of applicants. Each applicant privately knows their own individual merit. After allocating the grant, the grantmaker learns the funded applicant's merit and can impose a limited penalty. Our interpretation is that the grantmaker is able to infer the applicant's true merit by observing, say, the number of their publications. To penalize the applicant, the grantmaker can debar them from future calls or terminate the grant prematurely. However, penalization is imperfect; the winner may enjoy some career gain from winning the grant, which the grantmaker cannot undo ex post. Notice that premature termination has bite when researchers rely on the grant to carry out their work, or when they are intrinsically motivated about working on their projects. This feature contrasts with the Maurer-Scotchmer model, in which researchers only care about getting funding.

How can the grantmaker use the threat of ex-post penalization to screen the applicants? Mylovanov and Zapechelnyuk (2017, Theorem 1) derive the following optimal mechanism, taking the form of a *binary shortlisting procedure*.[15] First, the grantmaker announces a threshold. Second, applicants are asked to declare whether their merit lies above the threshold (and they will find it in their best interest to do so truthfully). Third, applicants above the threshold are shortlisted, and applicants below the threshold are shortlisted with some probability strictly less than one. Fourth, a winner is selected uniformly at random from the shortlist (if the shortlist is empty, a winner is selected uniformly at random from the full applicant pool). Finally, if the winner is ex-post found to have misreported their quality, the penalty is triggered.

We highlight two qualitative insights from this result.

First, the optimal mechanism involves a lottery. To gain an intuition, consider applicants with merit below the threshold. Applicants are more likely to be shortlisted (and hence win the grant) if they claim to pass the threshold. Therefore, if a low-merit applicant would never win the grant by reporting truthfully, they would have an incentive to misreport, even under the threat of ex-post penalization. By introducing the lottery, the grantmaker incentivizes these low-merit types to report truthfully. Thus, the lottery serves as an incentive device for ensuring truthful reports.

The second insight is that ex-post penalties indeed provide incentives ex ante. Applicants are more likely to be short-listed (and hence win the grant) if they pass the threshold. Thus the optimal mechanism is more likely to allocate the grant to a researcher with high merit than a mechanism that allocates grants completely at random.

Further reading Li (2020) studies a nearby setting where the grantmaker does not learn the winner's type for free (as in Mylovanov and Zapechelnyuk (2017)) but must pay a cost to do so.

[15] Here, we focus on the optimal mechanism in the case with sufficiently many agents. See Theorem 1 in Mylovanov and Zapechelnyuk (2017) for the general solution.

3.3.3 Post-award incentives

We next ask how the structure of grants shapes the incentives of funded researchers. As we have seen in Section 3.3.1, the nature of grants as upfront payments creates a hidden action problem: conditional on being funded, a researcher may have an incentive to use the funds to serve their own interests rather than those of the grantmaker or society. The model in Maurer and Scotchmer (2004) emphasizes that grantmakers can re-align interests by monitoring researchers and, possibly, cutting them off from future funding. However, slashing funding might not be a sufficiently powerful instrument. Researchers can apply for grants from different grantmakers or might not require external funding frequently. Further, it is somewhat stringent to assume that grantmakers can perfectly monitor whether researchers work.

In classical hidden action models (Holmström, 1979), incentives are typically provided through performance-dependent payments. Grants do not usually use outcome-dependent rewards, but we can still gain some insights from this literature. In the following, we will first give an overview of relevant insights from the literature assuming that performance-dependent rewards are available to the funder. Then, we will briefly discuss how such performance-dependent rewards can be interpreted in the context of research grants.

Consider a researcher who has been awarded a grant, and there is no uncertainty about the researcher's merit. Once the grant has been allocated, the researcher must decide how to allocate their resources; for example, whether to spend time and effort on the projects in their proposal. While the grantmaker wants the researcher to devote time to those projects, the researcher has private incentives not to; for example, because the researcher prefers working on alternative projects or traveling to conferences. In general, a grantmaker cannot directly observe or verify how a researcher allocates their time. Therefore, the grantmaker must use verifiable outputs that depend on the researcher's choices to provide incentives, such as publications related to the grant projects. Due to the uncertain nature of research, publication outcomes are not a perfect measure of the researcher's choices, but they are likely correlated. Suppose that, by allocating more time to grant-related projects, a researcher can increase the likelihood of having publications at the end of the grant period. Then, the grantmaker can design a reward scheme for the researcher that depends on publication outcomes. Good publication outcomes serve as a signal of the researcher's having allocated their time as the grantmaker desired and are therefore rewarded. If a researcher has worse publication outcomes, the reward should be lower to ensure that the researcher is discouraged from allocating their time differently.

This basic hidden action model has been extended in many directions. In the context of scientific research, we focus on two relevant modeling ingredients. First, grants usually have long time horizons and researchers decide how to allocate their resources repeatedly. Second, over time researchers learn about their projects—for example, the feasibility of producing any (publishable) results.

For simplicity, suppose that the grantmaker wants to encourage a researcher to work on a project for a specified duration. The successful completion is observable (in the form of journal publications, for example). However, whether and when the researcher successfully completes the project is uncertain, even when working on the project as desired. Therefore, the grantmaker cannot infer whether the absence of a publication at each point is due to bad luck of the researcher or due to the researcher's not having worked on the project. How can the grantmaker incentivize the researcher to work on the project?

Let us consider the model of Green and Taylor (2016). The grantmaker desires that the researcher successfully completes the funded project and chooses both the duration of the project and how the researcher is rewarded. The researcher then chooses at each point in time whether or not to work on the project. When working on the project, the researcher foregoes the value of using the funds differently. Naturally, the project can only be completed when the researcher works on the project. To capture the uncertainty in research, the researcher may fail to complete the project even when working on it. Green and Taylor (2016) show that the optimal grant design in this case is as follows. The grant provides funding until a (finite) deadline. The reward for completing the project is time-dependent; the earlier the researcher completes the project, the higher the reward. Intuitively, the researcher has to be compensated for two reasons. First, when working on the project, the researcher foregoes the benefits of allocating the funds differently. Second, upon completing the project before the deadline, the researcher foregoes the private enjoyment from the funds for the remaining time—this is because the researcher could always decide to use the funds differently and claim that the absence of success was due to bad luck.

To capture the idea that a researcher learns about a project while working on it, consider the following model studied by Moroni (2022). It is uncertain whether the project is feasible. If the project is infeasible, it cannot be completed. Thus, the researcher learns about the project's feasibility while working. Specifically, when working without success, the researcher becomes more pessimistic about the project, and hence also more pessimistic about the prospect of winning whatever reward the grantmaker promised for successful completion. The more pessimistic the researcher, the higher the reward necessary to incentivize effort. Therefore, one might expect the reward to increase over time in the presence of learning. However, this is not the case. Suppose that the researcher would expect a higher reward tomorrow than today. Then, the researcher would rather delay working on the project today (which is unobserved by the grantmaker) and work on the project tomorrow. Because the researcher did not work today, the researcher did not become more pessimistic about the project. Therefore, the researcher must be compensated for earlier successes. The optimal design features a project deadline and a constant reward.[16]

Halac et al. (2016) show how to incorporate hidden information into this model. Suppose that the researcher has better information about the project's feasibility than the grantmaker. In this case, the grantmaker can offer a menu of different grant designs so that the researcher reveals their private information via self-selection. Green and Taylor (2016) show how the optimal grant design changes if completing the project requires an intermediate step that is unverifiable to the grantmaker. In this case, the grantmaker optimally incentivizes the researcher by staging the project into two steps with a soft first and a hard second deadline.

Manso (2011) considers a different variant of the dynamic hidden action model focusing on the researcher's choice between two projects. One project is safe while the other project is riskier; in particular, it is unknown how likely it is to succeed on the risky project and, in expectation, a success is less likely than on the safe project. Manso (2011) shows that to encourage work on the riskier project, the grant design must exhibit tolerance for early failure (reflected, for example, in a longer duration) and rewards for long-term success. If early successes are rewarded relatively highly, the researcher has an incentive to work on safe, less creative problems rather than on novel but risky problems.

On the empirical side, Myers and Tham (2023), using survey-data of US profes-

[16]The constant reward relies on the agent's not discounting the future. If the agent would discount the future, the optimal reward schedule would be declining over time as the incentive to delay the work is slightly mitigated.

sors, provide evidence that a longer grant duration indeed encourages riskier projects, but only for researchers with long-term job security.[17] Further, in response to larger grants researchers shift resources towards existing projects rather than opening new ones. Myers and Tham (2023) also discuss estimates for the rate at which researchers are willing to substitute duration for size.[18]

The results in the preceding discussion of hidden action models rely on the ability of the grantmaker to design and commit to performance-dependent rewards. Whenever the use of such rewards is infeasible or undesirable, the results provide a benchmark on what would be achievable with these additional instruments. We consider it an interesting question to think about how a grantmaker can design instruments to resemble performance-dependent payments. Maurer and Scotchmer (2004) provide one example of repeatedly assigning grants of fixed value. One could imagine that the grantmaker can condition the access to or the design of future grants on past performance to approximate the optimal performance-dependent rewards. Gross and Bergstrom (2023) interpret the reward in a moral hazard model normatively as how the scientific community *should* evaluate scientific contributions.

Another strand of the literature on dynamic hidden action models focuses on the use of costly inspections for incentivizing researchers to work. Ball and Knoepfle (2023) consider a dynamic model in which a researcher decides whether to work on a project. The grantmaker cannot use time-dependent payments in their model but chooses the timing of costly inspections that reveal the researcher's past choices. Similar to the papers discussed above, the researcher increases the probability that a project is completed by working and project completion terminates the grant. How can the grantmaker time inspections in such a way that the researcher will work while economizing on the cost of inspections? Ball and Knoepfle (2023) show that the grantmaker optimally inspects the researcher in fixed time intervals.[19]

[17] Here, when we say a grant has a size of $\$x$ and a duration of n years, we mean that the grantee can freely spend up to $\$x$ within n years. After n years, all unused funds expire.

[18] See also Tham (2023) and Tham et al. (2024) for the impact of funding delays on career outcomes.

[19] Ball and Knoepfle (2023) also study the case in which working reduces the probability of a *breakdown*, that is, the failure of the project. In this case, random inspections are optimal.

3.4 Grant supply and the direction of research

The previous sections shed light on *how* the grantmaker should structure payments and design the allocation rule to achieve a particular goal. We now shift our focus to the choice of *what* to fund. How does a grantmaker's funding a specific topic impact researchers' incentives and the direction of research?

In practice, grantmakers frequently offer grants that target specific topics, so-called mission-oriented grants. Examples include *Requests for Applications (RFAs)* at the NIH (Myers, 2020), and the *SWITCHES* program at ARPA-E (Azoulay and Li, 2021). Further, in the context of investigator-initiated proposals, grantmakers can selectively fund proposals that align most closely with their goals, thereby indirectly steering the direction of research. Investigator-initiated proposals perhaps run a higher risk than mission-oriented grants of attracting proposals that do not closely fit the grantmaker's interests. However, investigator-initiated proposals can more easily leverage researchers' knowledge about the feasibility and prospects of research topics.

Besides the availability of funding and personal intellectual curiosity, researchers' preferences over questions are shaped by feasibility—how difficult is progress given current knowledge?— and the reward bestowed by the community for establishing priority. Importantly, the incentives from feasibility and priority may fail to induce an efficient allocation of researchers to questions due to *congestion* and *competition* effects.

When a new researcher joins a field, the rate of discovery in this field increases. *Congestion effects* arise when this increase is decreasing in the number of incumbent researchers. From a social planner's perspective, it may be preferable for the new researcher to reallocate to a field with fewer incumbent researchers. In theoretical work, Hopenhayn and Squintani (2021) develop this idea in a model of corporate R&D, showing how such congestion effects lead to socially inefficient entry on "hot" areas. As they note, "[t]he source of market failure in [the] model is the absence of property rights on problems to be solved[.]" While theirs is a model of patents and corporate R&D, a similar line of reasoning may apply to academic research since researchers cannot explicitly acquire property rights to open questions. Thus, there may be an economic value to communal norms that informally bestow such property rights. For example, Ramakrishnan (2018) and Strasser (2019) describe how such norms prevailed in the early days of structural biology. Nowadays, structural biologists are more secretive and races to establish priority are prevalent (Hill and Stein, 2023, 2024).

Competition effects arise when researchers attempt to outrace one another to establish priority. When the reward for establishing priority on a problem is high, many researchers begin to work on the problem, leading to adverse incentives on research quality. Hill and Stein (2024) model these entry and research incentives and provide evidence of a negative effect on research quality.[20] The evidence suggests that projects with "higher ex-ante potential generate more competition, are completed faster, and are lower quality."

How can grants improve the allocation of researchers to research questions and curtail congestion and competition effects? Let us speculate on some channels. First, grantmakers can create "hot" areas by offering many or particularly attractive grants, effectively making certain topics more valuable to investigate and thereby pulling researchers away from congested areas. Naturally, implementing this would require detailed knowledge of what is feasible and on what researchers would otherwise work. Second, if unfunded researchers cannot investigate certain questions (due to the costs of specialized equipment, for example), grantmakers effectively place an upper bound on the number of researchers working on a question, reducing both competition and congestion effects. Third, grants may act as a commitment device: a researcher with a grant credibly signals to the community that they have the resources to investigate their area.

How effective are grants at guiding the direction of research? Recent empirical work by Myers (2020) quantifies the costs of steering the direction using RFAs at the NIH. The evidence suggests that researchers do not simply follow the money, as it were, but are much more likely to apply for RFAs that are topically similar to their prior research. Overall, Myers (2020) suggests it is costly to incentivize specific

[20] Their model does not assume that establishing priority is a winner-takes-all-contest, but only that being second yields a lower reward than being first; this aligns with empirical evidence from structural biology (Hill and Stein, 2023). The model of Hill and Stein (2024) builds on earlier theoretical work by Bobtcheff et al. (2017), who cast competition between two researchers as a winner-takes-all contest without endogenous entry. Bobtcheff et al. (2017) show that "when breakthroughs become more frequent, researchers are under increasing competitive pressure and have decreasing incentives to wait and let their ideas mature." That is, the adverse incentives to publish low-quality work to establish priority are especially pertinent when technological progress allows the research community to produce insights quickly.

topics. In equilibrium, “RFAs must make more funds available to attract the same number of applications as the investigator-initiated mechanism.”

The costs and benefits of investigating a question are difficult to quantify since they depend on the current knowledge frontier. Carnehl and Schneider (2023) develop a model for understanding the evolution of scientific knowledge and the researchers’ incentives for pursuing questions. The model conceptualizes knowledge as a set of questions to which the answer is known. The answer to any particular question has a spillover effect on nearby questions through improved conjectures about their answers. Having precise information about answers is valuable since it guides decision-making in practical problems.

In this model, the interests of researchers and society are misaligned for two reasons. First, researchers do not fully internalize that uncovering an answer provides a guiding light for future researchers. Second, researchers bear the costs of uncovering answers; these costs are lower for questions similar to ones with a known answer. The misaligned interests lead to inefficiently low novelty in research. Hence this model provides a formal framework for investigating the effects of policy tools (such as RFAs) when researchers have a stronger preference than society for questions close to existing knowledge.

Interpreting grants as reductions in the cost of research, grants can induce researchers to investigate more novel questions, even when grants are not topic-specific. In this model, however, such grants cannot incentivize “moonshots”—research on extremely novel questions that guide future researchers to choose more efficiently. To encourage researchers to work on moonshot questions, one may have to resort to other tools, such as research prizes or mission-oriented grants.

Appendices

Appendix A Some formal details

Here we illustrate some of the formal ideas from the main text.

A.1 The application process

Adda and Ottaviani (2023) use techniques pioneered by Lehmann (1988) to study the impact of noise. Here, we spell out the conditions that characterize equilibria, and then use those conditions to study the impact of noise.

Let θ denote the merit of an applicant. As noted, in equilibrium there is an *application threshold* $\hat{\theta}$ such that exactly the applicants with merit above $\hat{\theta}$ apply. Further, there is an *acceptance threshold* $\hat{x}$ such that the grantmaker funds an applicant if and only if the evaluation produces a signal above $\hat{x}$. Let $F_\theta(x)$ denote the probability that the evaluation of an applicant with merit θ produces a signal below x, and let F_θ^{-1} denote the associated quantile function.[21] Thus, an applicant with merit θ expects to be funded with probability $1 - F_\theta(\hat{x})$. The application threshold $\hat{\theta}$ is such that an applicant with merit $\hat{\theta}$ is indifferent between applying and abstaining; if the application costs are given by $c \in (0, 1)$ and the benefits from being funded are normalized to 1, the thresholds solve the equation $1 - F_{\hat{\theta}}(\hat{x}) = c$; equivalently

$$F_{\hat{\theta}}^{-1}(1 - c) = \hat{x}.$$

An additional constraint on $\hat{x}$ and $\hat{\theta}$ is that the grantmaker's budget is exhausted. Specifically, if G denotes the distribution of merit and B denotes the grantmaker's budget, then

$$B = \int_{\hat{\theta}}^{\infty} (1 - F_\theta(\hat{x}))\; dG(\theta)$$

holds. Here, $1 - F_\theta(\hat{x})$ is the probability that a researcher with merit θ is funded when applying. Since exactly the researchers with merit above $\hat{\theta}$ apply, integrating $1 - F_\theta(\hat{x})$ over $[\hat{\theta}, \infty)$ yields the total volume of funded researchers. This volume must equal the grantmaker's budget B.

Let us now consider the following thought experiment. The signal distribution changes from F to $\tilde{F}$, and the acceptance threshold $\hat{x}$ changes to a point $\hat{y}$ that preserves $\hat{\theta}$ as the application threshold; that is, $\tilde{F}_{\hat{\theta}}^{-1}(1 - c) = \hat{y}$. Now consider an applicant with merit θ above the application threshold $\hat{\theta}$; by construction of the

[21] Here, F_θ is assumed to possess a continuously differentiable and strictly positive density. Moreover, the density is strictly supermodular in the signal and merit, meaning the monotone likelihood ratio property holds; higher signals indicate high merit.

application threshold, this applicant will apply both under F and $\bar{F}$. Further, this applicant's funding probability is lower under $\tilde{F}$ than under F if and only if

$$1 - \tilde{F}_\theta(\hat{y}) < 1 - F_\theta(\hat{x})$$

Plugging in for $\hat{x}$ and $\hat{y}$, we obtain the inequality

$$F_\theta(F_{\hat{\theta}}^{-1}(1-c)) < \tilde{F}_\theta(\tilde{F}_{\hat{\theta}}^{-1}(1-c)).$$

If $\tilde{F}$ is strictly less accurate than F in the sense of Lehmann, then this inequality indeed holds. Therefore, if signals become strictly less accurate and the application threshold is held fixed, then all applicants who previously applied now enjoy a *strictly lower* funding probability (except the applicant at the application threshold). It follows that if the application threshold is held fixed, then the budget will be underspent. Therefore, in equilibrium, the application threshold must decrease in response to a decrease in accuracy; consequently, applications increase.

A.2 Costly prospective evaluation

A.2.1 The single-applicant case

In the framework of Ben-Porath et al. (2014), let us focus on the case with a single applicant where it is significantly easier to derive the (essentially unique) optimal mechanism.

With a single applicant, the grantmaker's decision is whether to fund that applicant or keep the funds. Let t denote applicant's privately known merit (the applicant's "type"). The grantmaker's payoff from funding the type t applicant equals t. The grantmaker's payoff from keeping the funds is some fixed number x; hence t may be greater or less than x. The grantmaker's verification costs are c. The applicant only cares about their probability of winning the grant.

One can show that, when deriving an optimal mechanism, it is without loss to focus on mechanisms of the following form:[22] the applicant directly reports a type $\hat{t}$; then the applicant is verified with some probability $v(\hat{t})$; if the applicant's report is verified, the applicant is funded with certainty if the report was truthful, and with

[22] This step combines a version of the Revelation Principle with an optimality argument.

probability 0 if the applicant misreported; if the applicant's report is not verified, the applicant is funded with some probability $q(\hat{t})$. Moreover, these probabilities are such that the applicant always finds it in their best interest to report truthfully; this is the case if and only if for all t and $\hat{t}$,

$$v(t) \cdot 1 + (1 - v(t)) \cdot q(t) \geq (1 - v(\hat{t})) \cdot q(\hat{t}) \tag{IC}$$

holds. The left side is the overall allocation probability from reporting t truthfully; the right side is from misreporting $\hat{t}$. In the literature, this is constraint is known as *incentive compatibility (IC).*

It is convenient to define the overall allocation probability as $p(t) = v(t) \cdot 1 + (1 - v(t)) \cdot q(t)$. Thus (IC) can also be written as $p(t) \geq p(\hat{t}) - v(\hat{t})$. Since this inequality holds for all t and $\hat{t}$, it also holds if we flip their roles, yielding $p(\hat{t}) \geq p(t) - v(t)$. We can now rearrange to obtain

$$v(t) \geq p(t) - p(\hat{t}). \tag{IC*}$$

This inequality says if type t enjoys a higher winning probability than $\hat{t}$ (in which case the right side of (IC*) is positive), then type t must be verified with sufficient frequency to deter $\hat{t}$ from misreporting to t.

The grantmaker's expected payoff equals

$$x + \mathbb{E}[p(t) \cdot (t - x) - v(t) \cdot c],$$

where the expectation is taken with respect to the distribution of t.[23]

We would now like to maximize the grantmaker's payoff subject to (IC*). To see how this is done, let $\underline{p}$ denote the smallest value of $p(t)$ across all t; that is, $\underline{p} = \inf_{t \in \mathbb{R}} p(t)$. Since the applicant reports to maximize their funding probability, the constraint ensuring that the applicant report truthfully, (IC*), is hardest to satisfy for the worst-off applicant types who enjoy $\underline{p}$. Indeed, (IC*) holds if and only if $v(t) \geq p(t) - \underline{p}$ holds for all t. The grantmaker wishes to minimize on verification

[23] That is, in the background there is a cummulative distribution function with finite mean, representing the uncertain distribution of the applicant's merit from the grantmaker's perspective.

costs, and therefore sets $v(t)$ as low as possible, meaning

$$v(t) = p(t) - \underline{p}.$$

If we plug this equation into the grantmaker's expected payoffs, we obtain

$$x + \mathbb{E}[(p(t) - \underline{p})(t - x - c) + \underline{p}(t - x)].$$

This modified expected payoff has the following interpretation: when the grantmaker wishes to raise $p(t)$ above $\underline{p}$, the grantmaker must simultaneously raise the verification probability $v(t)$ by the same amount since, else, the worst-off merit type would misreport. Thus, the gain from raising $p(t)$ is, effectively, the "net merit type" $t - x - c$. It follows that optimally $p(t)$ equals 1 whenever the net merit type is positive (that is, $t - x - c \geq 0$); when the net merit type is negative, $p(t)$ optimally equals $\underline{p}$. The optimal value of $\underline{p}$ depends on c and the distribution of t. For example, setting $\underline{p} = 0$ is optimal if the prior mean of t is less than x (that is, $\mathbb{E}[t] \leq x$), capturing a scenario where the grantmaker is a priori unwilling to fund the applicant.

A.2.2 Comparative statics

Let us now return to the model with multiple applicants and provide some intuition for the comparative statics discussed in the text.

Let t denote a generic level of merit of a generic applicants. The merit levels are distributed identically and independently across agents according to a cummulative distribution function F that admits a density. Let c denote the verification costs. Departing from our discussion of the single-applicant case, suppose merit is always positive while the payoff from keeping the funds equals 0.

Ben-Porath et al. (2014) show that the optimal threshold equals $t^* - c$, where t^* solves the equation

$$\mathbb{E}[t] = \mathbb{E}[\max(t - c, t^* - c)] \tag{A.1}$$

where both expectations are taken with respect to F. To interpret this equation, consider the following hypothetical scenario (which roughly follows Section III of Ben-Porath et al., 2014). Fixing one of the applicants—say, applicant 1—, suppose the best payoff the grantmaker can obtain by not consulting applicant 1 equals some number x. Informally speaking, the value x acts as an outside option that is roughly analogous to the option of keeping the funds in the single-applicant case. The grantmaker now has at least two ways of combining the outside option with applicant 1.

(1) First, the grantmaker could simply allocate to applicant 1 (without consulting them), yielding an expected payoff of $\mathbb{E}[t]$.

(2) Second, the grantmaker could elicit applicant 1's merit, allocating to applicant 1 if the net merit $t - c$ lies above x, and else taking the outside option x. The expected payoff would now be $\mathbb{E}[\max(t - c, x)]$. Note that the logic of focusing on the net merit $t - c$ rather than t is the same as in the single-applicant case—if applicant 1's report is to inform the allocation, then the grantmaker must verify the report.

Equation (A.1) says if x happens to equal $t^* - c$, then the grantmaker is indifferent between these two ways of combining the outside option with applicant 1.

Using (A.1), we can deduce the comparatice statics with respect to c and F described earlier. Indeed, (A.1) rearranges to (see Ben-Porath et al., 2014, equation (3)):

$$\int_{\underline{t}}^{t^*} F(s)\, ds = c,$$

where $\underline{t}$ denotes the lowest possible type. If c increases, then clearly t^* must increase to maintain equality. A FOSD- or SOSD-increase in F decreases the left side,[24] so t^* increases accordingly.

A.3 A simple model of post-award management

In the model of Maurer and Scotchmer (2004), researchers privately know the probability λ of obtaining a fund-worthy idea in each period. If a researcher works on their idea, they incur an effort cost κ. The researcher's immediate private benefit from being awarded is t. Researchers discount future payoffs at rate r. Under the rule where researchers are funded if and only if they worked in all past period, a

[24]Recall, a cdf F is FOSD-higher than a cdf G if $F(s) \leq G(s)$ holds for all $s \in \mathbb{R}$, and F is SOSD-higher than G if $\int_{-\infty}^{t} F(s)\, ds \leq \int_{-\infty}^{t} G(s)\, ds$ holds for all $t \in \mathbb{R}$.

researcher of type λ has an incentive to work in the first period if

$$v \leq v - \kappa + \frac{\lambda}{r}(v - \kappa).$$

By shirking (on the left side), the researcher enjoys the private benefit t at zero cost, but loses out on all future net benefits. By working, they pay the cost κ; additionally, they enjoy future benefits from funding equal to the per-idea rent $v - \kappa$ times the probability λ of obtaining an idea in any given period, in perpetuity from next period (and, thus, divided by the discount rate r). This no-shirking condition can be rearranged to

$$\lambda \geq \frac{r\kappa}{v - \kappa}.$$

This inequality shows that only sufficiently productive researchers (high λ) have the incentives to work hard for continuous funding.

References

Adda, Jérôme and Marco Ottaviani (2023). “Grantmaking, Grading on a Curve, and the Paradox of Relative Evaluation in Nonmarkets”. In: *The Quarterly Journal of Economics*, qjad046.

Aghion, Philippe, Mathias Dewatripont, and Jeremy C Stein (2008). “Academic freedom, private-sector focus, and the process of innovation”. In: *The RAND Journal of Economics* 39.3, pp. 617–635.

Alon, Noga, Felix Fischer, Ariel Procaccia, and Moshe Tennenholtz (2011). “Sum of us: Strategyproof selection from the selectors”. In: *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 101–110.

Andrews, Isaiah and Maximilian Kasy (2019). “Identification of and correction for publication bias”. In: *American Economic Review* 109.8, pp. 2766–2794.

Arrow, Kenneth J. (1962). “Economic Welfare and the Allocation of Resources for Invention”. In: *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton: Princeton University Press, pp. 609–626. ISBN: 978-1-349-15486-9. DOI: 10.1007/978-1-349-15486-9_13. URL: https://doi.org/10.1007/978-1-349-15486-9_13.

Augias, Victor and Eduardo Perez-Richet (2023). *Non-Market Allocation Mechanisms: Optimal Design and Investment Incentives*. arXiv: 2303.11805 [econ.TH].

Azoulay, Pierre, Erica Fuchs, Anna P Goldstein, and Michael Kearney (2019a). "Funding breakthrough research: promises and challenges of the "ARPA Model"". In: *Innovation policy and the economy* 19.1, pp. 69–96.

Azoulay, Pierre, Joshua S Graff Zivin, Danielle Li, and Bhaven N Sampat (2019b). "Public R&D investments and private-sector patenting: evidence from NIH funding rules". In: *The Review of economic studies* 86.1, pp. 117–152.

Azoulay, Pierre and Danielle Li (2021). "Scientific Grant Funding". In: *Innovation and Public Policy*. Chicago: University of Chicago Press, pp. 117–150. ISBN: 9780226805597. DOI: doi:10.7208/chicago/9780226805597-008. URL: https://doi.org/10.7208/chicago/9780226805597-008.

Azrieli, Yaron (2024). "Temporary exclusion in repeated contests". In: *arXiv preprint arXiv:2401.06257*.

Balietti, Stefano, Robert L Goldstone, and Dirk Helbing (2016). "Peer review and competition in the Art Exhibition Game". In: *Proceedings of the National Academy of Sciences* 113.30, pp. 8414–8419.

Ball, Ian and Jan Knoepfle (2023). "Should the Timing of Inspections be Predictable?" In: *arXiv preprint arXiv:2304.01385*.

Ben-Porath, Elchanan, Eddie Dekel, and B Lipman (2023). *Sequential mechanisms for evidence acquisition*. Tech. rep. working paper.

Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman (2014). "Optimal allocation with costly verification". In: *American Economic Review* 104.12, pp. 3779–3813.

— (2019). "Mechanisms with evidence: Commitment and robustness". In: *Econometrica* 87.2, pp. 529–566.

Bendiscioli, Sandra, Teo Firpo, Albert Bravo-Biosca, Eszter Czibor, Michele Garfinkel, Tom Stafford, James Wilsdon, and Helen Buckley Woods (June 2022). "The experimental research funder's handbook (Revised edition, June 2022, ISBN 978-1-7397102-0-0)". In: DOI: 10.6084/m9.figshare.19459328.v3. URL: https://rori.figshare.com/articles/report/The_experimental_research_funder_s_handbook_final_version_/19459328.

Bloch, Francis, Bhaskar Dutta, and Marcin Dziubiński (2023). "Selecting a winner with external referees". In: *Journal of Economic Theory*, p. 105687.

Bobtcheff, Catherine, Jérôme Bolte, and Thomas Mariotti (2017). "Researcher's dilemma". In: *The Review of Economic Studies* 84.3, pp. 969–1014.

Bryan, Kevin A and Jorge Lemus (2017). "The direction of innovation". In: *Journal of Economic Theory* 172, pp. 247–272.

Bryan, Kevin A and Heidi L Williams (2021). "Innovation: market failures and public policies". In: *Handbook of industrial organization*. Vol. 5. 1. Elsevier, pp. 281–388.

Carnehl, Christoph and Johannes Schneider (2023). "A quest for knowledge". In: *arXiv preprint arXiv:2102.13434*.

Dasgupta, Partha and Paul A David (1994). "Toward a new economics of science". In: *Research policy* 23.5, pp. 487–521.

De Clippel, Geoffroy, Herve Moulin, and Nicolaus Tideman (2008). "Impartial division of a dollar". In: *Journal of Economic Theory* 139.1, pp. 176–191.

Epitropou, Markos and Rakesh Vohra (2019). "Optimal on-line allocation rules with verification". In: *Algorithmic Game Theory: 12th International Symposium, SAGT 2019, Athens, Greece, September 30–October 3, 2019, Proceedings 12*. Springer, pp. 3–17.

ERC (2023). "Information for applicants: Starting and Consolidator Grant 2024 calls". URL: https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/horizon/guidance/information-for-applicants_he-erc-stg-cog_en.pdf.

Fang, Ferric C and Arturo Casadevall (2016). *Research funding: The case for a modified lottery.*

Fischer, Felix and Max Klimm (2014). "Optimal impartial selection". In: *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 803–820.

Franzoni, Chiara and Paula Stephan (2023). "Uncertainty and risk-taking in science: Meaning, measurement and management in peer review of research proposals". In: *Research Policy* 52.3, p. 104706. ISSN: 0048-7333. DOI: https://doi.org/10.1016/j.respol.2022.104706. URL: https://www.sciencedirect.com/science/article/pii/S004873332200227X.

Fu, Qiang and Zenan Wu (July 2019). *Contests: Theory and Topics*. DOI: 10.1093/acrefore/9780190625979.013.440. URL: https://oxfordre.com/economics/view/10.1093/acrefore/9780190625979.001.0001/acrefore-9780190625979-e-440.

Gallini, Nancy and Suzanne Scotchmer (2002). “Intellectual property: when is it the best incentive system?” In: *Innovation policy and the economy* 2, pp. 51–77.

Goldstein, Anna P and Michael Kearney (2020). “Know when to fold ‘em: An empirical description of risk management in public research funding”. In: *Research Policy* 49.1, p. 103873.

Green, Brett and Curtis R Taylor (2016). “Breakthroughs, deadlines, and self-reported progress: Contracting for multistage projects”. In: *American Economic Review* 106.12, pp. 3660–3699.

Gross, Kevin and Carl T Bergstrom (2019). “Contest models highlight inherent inefficiencies of scientific funding competitions”. In: *PLoS biology* 17.1, e3000065.

— (2023). “Rationalizing risk aversion in science”. In: *arXiv preprint arXiv:2306.13816*.

Gutas, Dimitri (1998). *Greek Thought, Arabic Culture: The Graeco-Arabic Translation Movement in Baghdad and Early ‘Abbasid Society (2nd–4th/8th–10th c.)* Routledge.

Halac, Marina, Navin Kartik, and Qingmin Liu (2016). “Optimal contracts for experimentation”. In: *The Review of Economic Studies* 83.3, pp. 1040–1091.

Heyard, Rachel, Manuela Ott, Georgia Salanti, and Matthias Egger (2022). “Rethinking the funding line at the Swiss national science foundation: Bayesian ranking and lottery”. In: *Statistics and Public Policy* 9.1, pp. 110–121.

Hill, Ryan and Carolyn Stein (2023). “Scooped! Estimating rewards for priority in science”.

— (2024). “Race to the bottom: Competition and quality in science”.

Hippel, Ted von and Courtney von Hippel (2015). “To apply or not to apply: A survey analysis of grant writing costs and benefits”. In: *PloS one* 10.3, e0118494.

Holmström, Bengt (1979). “Moral hazard and observability”. In: *The Bell journal of economics*, pp. 74–91.

Holzman, Ron and Hervé Moulin (2013). “Impartial nominations for a prize”. In: *Econometrica* 81.1, pp. 173–196.

Hopenhayn, Hugo and Francesco Squintani (2021). “On the direction of innovation”. In: *Journal of Political Economy* 129.7, pp. 1991–2022.

Hussam, Reshmaan, Natalia Rigol, and Benjamin N Roth (2022). “Targeting high ability entrepreneurs using community information: Mechanism design in the field”. In: *American Economic Review* 112.3, pp. 861–898.

Jones, Benjamin F (2021). “Science and innovation: The under-fueled engine of prosperity”. In.

Kattwinkel, Deniz (2020). “Allocation with Correlated Information: Too good to be true”. In: *Proceedings of the 21st ACM Conference on Economics and Computation*, pp. 109–110.

Kattwinkel, Deniz and Jan Knoepfle (2023). “Costless information and costly verification: A case for transparency”. In: *Journal of Political Economy* 131.2, pp. 504–548.

Kattwinkel, Deniz, Axel Niemeyer, Justus Preusser, and Alexander Winter (2022). “Mechanisms without transfers for fully biased agents”. In: *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 332–333.

Khalfan, Nawaaz (2023). *Optimal Allocation with Noisy Inspection.* Tech. rep. Working Paper.

Kurokawa, David, Omer Lev, Jamie Morgenstern, and Ariel D Procaccia (2015). “Impartial Peer Review.” In: *IJCAI*. Vol. 15. Citeseer, pp. 582–588.

Lehmann, Erich Leo (1988). “Comparing Location Experiments”. In: *The Annals of Statistics* 16.2, pp. 521–533.

Li, Danielle, Pierre Azoulay, and Bhaven N Sampat (2017). “The applied value of public investments in biomedical research”. In: *Science* 356.6333, pp. 78–81.

Li, Yingkai and Xiaoyun Qiu (2024). *Screening Signal-Manipulating Agents via Contests*. arXiv: 2302.09168 [econ.TH].

Li, Yunan (2020). “Mechanism design with costly verification and limited punishments”. In: *Journal of Economic Theory* 186, p. 105000.

Li, Zihao and Jonathan Libgober (2023). “The Dynamics of Verification when Searching for Quality”.

Lotka, Alfred J (1926). “The frequency distribution of scientific productivity”. In: *Journal of the Washington academy of sciences* 16.12, pp. 317–323.

Manso, Gustavo (2011). “Motivating innovation”. In: *The Journal of Finance* 66.5, pp. 1823–1860.

Maurer, Stephen M and Suzanne Scotchmer (2004). “Procuring knowledge”. In: *Intellectual property and entrepreneurship*. Emerald Group Publishing Limited, pp. 1–31.

Merrifield, Michael R and Donald G Saari (2009). “Telescope time without tears: a distributed approach to peer review”. In: *Astronomy & Geophysics* 50.4, pp. 4–16.

Merton, Robert K. (1957). “Priorities in Scientific Discovery: A Chapter in the Sociology of Science”. In: *American Sociological Review* 22.6, pp. 635–659. ISSN: 00031224. URL: http://www.jstor.org/stable/2089193 (visited on 02/26/2024).

Mervis, Jeffrey (2014). “A radical change in peer review”. In: *Science* 345.6194, pp. 248–249. DOI: 10.1126/science.345.6194.248. eprint: https://www.science.org/doi/pdf/10.1126/science.345.6194.248. URL: https://www.science.org/doi/abs/10.1126/science.345.6194.248.

Morgan, John, Justin Tumlinson, and Felix Várdy (2022). “The limits of meritocracy”. In: *Journal of Economic Theory* 201, p. 105414.

Moroni, Sofia (2022). “Experimentation in organizations”. In: *Theoretical Economics* 17.3, pp. 1403–1450.

Myers, Kyle (2020). “The elasticity of science”. In: *American Economic Journal: Applied Economics* 12.4, pp. 103–134.

— (2022). “The Potential Benefits of Costly Applications in Grant Contests”. URL: https://www.kylemyers.org/working.

Myers, Kyle and Wei Yang Tham (2023). “Money, Time, and Grant Design”. In: *arXiv preprint arXiv:2312.06479*.

Mylovanov, Tymofiy and Andriy Zapechelnyuk (2017). “Optimal allocation with ex post verification and limited penalties”. In: *American Economic Review* 107.9, pp. 2666–2694.

Niemeyer, Axel and Justus Preusser (2023). “Simple Allocation with Correlated Types”.

Olckers, Matthew and Toby Walsh (2023). *Manipulation and Peer Mechanisms: A Survey*. arXiv: 2210.01984 [cs.AI].

Olszewski, Wojciech (2020). “A theory of citations”. In: *Research in Economics* 74.3, pp. 193–212. ISSN: 1090-9443. DOI: https://doi.org/10.1016/j.rie.2020.06.001. URL: https://www.sciencedirect.com/science/article/pii/S109094432030140X.

Perez-Richet, Eduardo and Vasiliki Skreta (2022). “Test design under falsification”. In: *Econometrica* 90.3, pp. 1109–1142.

— (2024). “Fraud-proof non-market allocation mechanisms”.

Price, W Nicholson (2019). “Grants”. In: *Berkeley Technology Law Journal* 34.1, pp. 1–66.

Ramakrishnan, Venki (2018). *Gene Machine: the race to decipher the secrets of the ribosome*. Basic Books.

Shah, Nihar B, Behzad Tabibian, Krikamol Muandet, Isabelle Guyon, and Ulrike Von Luxburg (2018). "Design and analysis of the NIPS 2016 review process". In: *Journal of machine learning research* 19.49, pp. 1–34.

Stephan, Paula (2012). *How economics shapes science*. Harvard University Press.

Stern, Scott (2004). "Do scientists pay to be scientists?" In: *Management science* 50.6, pp. 835–853.

Strasser, Bruno J (2019). *Collecting experiments: Making big data biology*. University of Chicago Press.

Tham, Wei Yang (2023). "Science, interrupted: Funding delays reduce research activity but having more grants helps". In: *Plos one* 18.4, e0280576.

Tham, Wei Yang, Joseph Staudt, Elisabeth Ruth Perlman, and Stephanie D Cheng (2024). "Scientific Talent Leaks Out of Funding Gaps". In: *arXiv preprint arXiv:2402.07235*.

Vojnović, Milan (2015). *Contest theory: Incentive mechanisms and ranking methods*. Cambridge University Press.

Wright, Brian D (1983). "The economics of invention incentives: Patents, prizes, and research contracts". In: *The American Economic Review* 73.4, pp. 691–707.