Studies in History and Philosophy of Science xxx (xxxx) xxx–xxx

Contents lists available at ScienceDirect

Studies in History and Philosophy of Science

journal homepage: www.elsevier.com/locate/shpsa

Mavericks and lotteries

Shahar Avin

Centre for the Study of Existential Risk, University of Cambridge, Cambridge, UK

HIGHLIGHTS

- A novel analytic synthesis of arguments that support random allocation as an alternative to grant peer review.
- Review of policy implementations of science funding that formally include a random selection element.
- A novel comparison of these policies with the policy-relevant characteristics of the arguments in the literature.

ABSTRACT

In 2013 the Health Research Council of New Zealand began a stream of funding titled ‘Explorer Grants’, and in 2017 changes were introduced to the funding mechanisms of the Volkswagen Foundation ‘Experiment!’ and the New Zealand Science for Technological Innovation challenge ‘Seed Projects’. All three funding streams aim at encouraging novel scientific ideas, and all now employ random selection by lottery as part of the grant selection process. The idea of funding science by lottery emerged independently in several corners of academia, including in philosophy of science. This paper reviews the conceptual and institutional landscape in which this policy proposal emerged, how different academic fields presented and supported arguments for the proposal, and how these have been reflected (or not) in actual policy. The paper presents an analytical synthesis of the arguments presented to date, notes how they support each other and shape policy recommendations in various ways, and where competing arguments highlight the need for further analysis or more data. In addition, it provides lessons for how philosophers of science can engage in shaping science policy, and in particular, highlights the importance of mixing complementary expertise: it takes a (conceptually diverse) village to raise (good) policy.

1. Introduction

As many of the other papers in this collected volume argue, scientific novelty, even in its more extreme ‘maverick’ variety, is an important collective epistemic good. As such, we should spend time thinking about the processes and institutions that encourage or discourage it, and support policies that, all else being equal, increase scientific novelty.[1] One cluster of institutions that affect scientific novelty are the institutions that decide and implement science funding policy: if novel projects are not funded, or if scientists do not apply to work on novel projects, the overall novelty of the scientific community decreases. Criticisms that the current dominant funding method, grant peer review, has an adverse effect on scientific novelty are not new. What is newer is the emergence of academic arguments that support alternative funding mechanisms, specifically with the aim of increasing scientific novelty, though these have also been around for a couple of decades. Newer still are implementations of these alternative mechanisms as actual funding policies around the world.

This paper looks at one such mechanism, the introduction of random selection into the funding process of research projects. §2 outlines the institutional landscape of science funding, and highlights the significance of the emergence of alternatives to the dominant mechanism of grant peer review. §3 recaps the arguments that have been put forward in the academic literature in support of random allocation as an alternative to grant peer review. §4 presents an analytic synthesis of these arguments and shows where they diverge in their assumptions and where evidence is scarce. §5 surveys policy implementations of science funding that formally include a random selection element. §6 compares the surveyed policies against each other and against the policy-relevant characteristics of the arguments in the literature. §7 makes a brief note of the chain of publications, evidenced in the citation record, from some of the arguments in the literature to implemented

E-mail address: sa478@cam.ac.uk.

[1] As a reviewer has keenly pointed out, this is left ambiguous with regards to the optimal level of novelty: is novelty an unalloyed good? From a Bayesian perspective, that seems indeed to be the case: given two statements with the same *a posteriori* credence, the one with lower *a priori* credence is the more valuable one, as it brings us closer to an accurate model. However, society does not care just about the value of information: there is a cost to be paid for changing one's beliefs and actions (even when the change is for the better), and there is a cost to be paid for entertaining, even in passing, low-credence ideas; see, e.g, Polanyi (1962) on the trade-off between novelty and *ex ante* plausibility, two key components of scientific merit. The current paper therefore commits to the weaker claim, that there is value to be gained from increasing the level of novelty in contemporary scientific practice.

https://doi.org/10.1016/j.shpsa.2018.11.006
Received 17 January 2018; Received in revised form 22 March 2018; Accepted 26 November 2018

policy. The aim of the paper is to make two arguments: the first, that there is a coherent justification for encouraging scientific novelty by introducing formal randomness into research funding mechanisms, though much uncertainty remains regarding specific policy details; the second, that the arguments that policy have so far picked up on have taken place outside the philosophy of science, and that now is an opportune moment for philosophers of science to join this interdisciplinary debate.

2. Context of science funding policy

2.1. The landscape of science funding

Investment in Research and Development (R&D) is a substantial global phenomenon. In developed countries, such as the United States of America (USA), Japan, South Korea, and the western member countries of the European Union (EU), spending on R&D is often in the range of 2–3% of the Gross Domestic Product (GDP).[2] This number takes into account both public (government) spending and private (industry) spending. When taking the significant GDPs of these countries into account, we arrive at a global R&D investment of roughly $1.7 trillion.[3]

Spending on R&D is not homogeneous: it can be divided by source of funding, and by type of research. While boundaries between divisions may be blurry, certain distinct categories of R&D funding emerge. The rationale for funding in each category, and the philosophical analysis of desiderata and appropriate mechanisms, may differ between these categories.

The first important distinction to be made within R&D spending is between the sources of funds: public funds, generally originating from tax collection and allocated by the government, and private funds, generated mainly from corporate profit or charitable donations, and expended by for-profit or not-for-profit private organisations. In USA R&D expenditure, the private sector (including charities and universities' own funds) accounts for just under three quarters of total R&D spending, with the dominant private provider by far being industry (65.2% of USA R&D), whereas the public sector (federal and local government) funds just over a quarter of USA R&D.

A second distinction within R&D expenditure is the kind of research or development work taking place. While assignment of individual projects into any category can prove challenging, most within the R&D policy world recognise categories that are similar to the three characters of work defined by the USA National Science Foundation (NSF):

> **Basic Research:** Research that seeks to gain more complete knowledge or understanding of the fundamental aspects of phenomena and of observable facts, without specific applications toward processes or products in mind.
> **Applied Research:** Research aimed at knowledge necessary for determining the means by which a recognised need may be met.
> **Development:** The systematic use of the knowledge or understanding gained from research, directed toward the production of useful materials, devices, systems, or methods, including design and development of prototypes and processes.

(Kennedy, 2012, pp. 4–5, changed styling for clarity).

Somewhat cutting across these categories is the more recently introduced category of 'transformative research', defined by the United States National Science Board as:

> a range of endeavors which promise extraordinary outcomes, such as: revolutionizing entire disciplines; creating entirely new fields; or disrupting accepted theories and perspectives – in other words, those endeavors which have the potential to change the way we address challenges in science, engineering, and innovation (NSB, 2007).

From the perspective of this special issue, transformative research is rich in scientific novelty, and has a certain 'maverick' character. While applied research and development projects can have revolutionary effects, from a policy perspective transformative research is most closely related to basic research, in terms of its high-risk high-reward character, its long term time horizon for impact, and its positive impact outside of the domain in which it was first conceived, all factors which make it unappealing for for-profit organisations (Arrow, 1962). While statistics on transformative research are harder to come by, we can refer to statistics on basic research as a proxy.

Basic research is only a minor component of global R&D. As discussed above, about two thirds of R&D funds in the USA are provided by private industry. Almost all of these funds are directed towards technological development in a few high-tech sectors, and very little (7%) is directed towards basic research, though that amount still makes up 26% of basic research funding. The major supporter of basic research in the USA is the public, via the federal government, which provides 47% of basic research funds.[4] Universities, colleges and charities provide 27% of basic research funds, and the rest comes from local government. The largest single institutions supporting basic research, both in the USA and in the world, are federally funded agencies, of which two are particularly dominant: the National Institutes of Health (NIH), and the NSF.[5] Both NIH and NSF explicitly state they seek and support transformative research, through a range of funding channels (NIH, 2017; NSF, 2017).

2.2. Desiderata for a public science funding mechanism

Before proceeding to discuss the merits and shortcomings of different science funding mechanisms, we should consider the properties of a good funding mechanism. The following list of desiderata is adapted from Chubin (1994); Cole, Cole, and Rubin (1977); Martino (1992), and is based on surveys of practicing scientists[6]

> **Effectiveness:** The mechanism should be effective at identifying high quality research. The results of effective allocation would be that high quality research is supported, leading to scientific progress and new knowledge.
> **Efficiency:** All parties involved would prefer, *ceteris paribus*, for the process to be as efficient as possible, meaning for it to require as little time and resources while providing an adequate level of effectiveness. This applies both to the time and resources on the reviewing side (administration, internal and external reviewers) and the applicants' side.
> **Accountability:** Funders want the allocation process to maintain accountability, to make sure scientists are spending the funds in ways that would further scientific research, to make sure they remain within the bounds of the project outlined in their proposal, that due process is followed in the allocation process, and that laws and regulations are adhered to, e.g. regarding treatment of human and animal subjects.
> **Fairness:** The allocation mechanism should distribute funds fairly, meaning it should not unjustifiably discriminate against any individual or group.

2.3. The process of funding by peer review

Both of the leading institutions in funding of global basic research,

[2] All statistics in this section are from NSB (2016), for fiscal year 2013, in 2016 PPP US dollars.

[3] Total national R&D expenditure: US$456.1 billion (largest), China $336.5 billion (second largest); All EU nations combined $342.4 billion.

[4] Note, though, that this amount only makes up 31% of the federal R&D budget, as much goes into supporting applied R&D in industries that cannot easily be privatised, such as defence and energy.

[5] Expenditure on basic research: NIH $14.7 billion, NSF $4.4 billion.

[6] Chubin's list of desiderata also includes Responsiveness, Rationality and Reliability, but I see these as sub-components of effectiveness.

NIH and NSF, allocate funding by a scheme of project choice called *peer review*, where research proposals originating from practising scientists are reviewed and ranked by other scientists working in the same or adjacent fields of research (their 'peers'). The following summary of the operation of peer review is based on consideration of the application and review process at several governmental science funding agencies, including NIH (2013a,b), NSF (2013a,b), the Australian National Health and Medical Research Council (Graves, Barnett, & Clarke, 2011), and the Austrian Science Fund (Dinges, 2005).

The process of resource allocation for basic research by peer review is the dominant contemporary form of resource allocation for scientific projects, the 'gold standard' of science funding. Some aspects of the process are strongly conserved across nations and institutions:

Investigator freedom: Project proposals originate from the investigators, not dictated by the funding body or a central organising committee. The extent to which investigators are free to design projects is limited under various guideline constraints, but there are many opportunities for significant levels of freedom.

Individual projects: As proposals originate from the investigators, they arrive at the funding body as discrete, compartmentalised funding opportunities. The funding bodies have the role of choosing among them, but they do not, to any significant extent, coordinate between different investigators to form overarching research programmes.

Information provision: As proposals originate from the investigators, they must inform the funding body about the contents and merits of their proposed projects. This is often done using a detailed written research plan, accompanied by various supporting documents.

Peer assessment: Funding bodies seek the expert opinion of one or more scientists in evaluating the merit of the proposed projects. While there are guidelines for component categories of evaluation, the decisions are still significantly subjective, i.e., not algorithmic or box-ticking.

Integration of assessments: Often assessment is sought from more than one source, e.g. multiple reviewers or a mix of internal and external reviewers. The different assessments are always combined in some way to form a single judgement per proposal, which is then compared to the judgements of other proposals.

Ranking and cutoff: There are never enough resources to fund all projects proposed. As such, comparisons of integrated assessments are used to decide which projects will get funded and which will not.

Other aspects of the process exhibit more variability, such as the identities of reviewers, the method of integrating assessments, and the guidelines for merit evaluation. Nonetheless, the practice of science funding by peer review is very strongly entrenched, as it has been around since World War Two (Agar, 2012; Greenberg, 1999, 2003). Given the emphasis of grant peer review on the ability and responsibility of scientific experts to select research projects for funding, the proposal to select projects for funding at random seems very odd in contrast.

2.4. Criticisms of grant peer review

Chubin and Hackett (1990) present a critical examination of grant peer review, based on evidence from surveys of practising scientists. Their stated aim is to overcome the nearly-mythical standing of peer review as a pillar of modern science, and to highlight the fact that very little has been done to subject peer review to methodical analysis, despite known tensions and probable shortcomings. This paucity of evidence regarding the effectiveness of peer review persists to this day, as was reported in a recent comprehensive literature review (Guthrie, Ghiga, & Wooding, 2017).

The main argument of Chubin and Hackett is that peer review serves a function for multiple stakeholders, each having slightly different expectations from the process and its products. These different desiderata are often in tension with each other, and so the process of peer review often fails to fully satisfy any of the desiderata, to the chagrin of stakeholders. Chubin and Hackett note, using survey data, an increase in the concern scientists and other stakeholders report regarding peer review. However, they note that as total success rates of applicants in grant peer review declined, because the increase in the scientific cohort size outpaced growth of allocated funds, pressure increased within the scientific community, which led to increased scrutiny of the allocation mechanism, though it was not any feature of peer review *per se* that caused the increase in pressure. Nonetheless, the increased attention to peer review brought to the fore explicit statements about what different parties considered the proper function of peer review, and what were the perceived shortcomings in fulfilling this function.

In the two decades since Chubin's surveys such pressures have only increased, as discussed in the next section, leading to the current experiments with alternative funding models such as funding by lottery.

3. Arguments for random allocation

3.1. Greenberg

Greenberg (1998) presented the first argument in an academic journal (to my knowledge) for funding science by lottery. In a short piece in *The Lancet*, Greenberg, a science journalist and author of several books on the politics of science, enumerates many of the common complaints against peer review. These include time and resource costs, peer review being close to random anyway, and lack of evidence supporting the claim that peer review is a good way to pick meritorious projects, including some deeper concerns about the possibility of conducting meaningful studies on the counterfactuals involved. So far, these have appeared time and again, including in the published literature (Cole et al., 1977; Martino, 1992).

However, Greenberg follows the criticism with (what was then) a novel proposal:

> So, as a first step towards either verifying peer review or moving on to a better system, the powers that be should slice off some respectable percentage of the research funds – say, 15–20% over 5 years – and set them aside. These funds would be awarded by lottery to applicant scientists whose qualifications and projects have been certified as respectable, ratings easily determined at a small fraction of the cost of peer review.
>
> Details aside, the basic principle is clear; instead of dodging the fact that chance plays a big part in awarding money, the system will sanctify chance as the determining factor. After a few years, let's look back and evaluate the science that came out of this system.
>
> If it's no worse than what we're getting now, let's chuck peer review, and thereby save a lot of needless effort and money (Greenberg, 1998, p. 686).

Some of the key features in Greenberg's proposal will keep recurring, so it is useful to break these down:

Pilot of random allocation: The change to the existing system should start with a small pilot that will employ a formal random element, and results will be compared to the results of peer review.

Randomness of peer review: The low reliability of peer review as a measure of merit means it will not significantly outperform a lottery.[7]

Efficiency advantage: A lottery will be cheaper and faster than existing peer review.

Pre-lottery screening: Not all scientists and not all projects should be admitted to the lottery. There should be some quality check,

[7] Greenberg, and many later authors, tie this low reliability to low success percentages, such that reviewers are asked to differentiate between the top applicants.

though this is expected to be much less time- and cost-intensive compared to peer review of proposals.

The entire piece can be summarised by one of its opening sentences, in Greenberg's characteristic style:

> [W]hen it comes to providing research money, there's got to be a better way than the cumbersome, snail-paced, expensive, and unproven peer review derby long in effect in the USA and elsewhere.
>
> Why not try a lottery among qualified researchers?

3.2. Brezis

Nearly a decade later, Brezis (2007) presented an economic model-based proposal for introducing a random element into R&D project selection. Brezis focused on criticisms of the effectiveness of peer review due to bias, and especially on concerns regarding a conservative bias in peer review. This bias tends to direct funds to low-risk projects, underfunding highly innovative projects that could lead to significant progress, but are also more likely to fail (the kind of novel, or 'maverick', research this issue is focused on).

Brezis' paper presents three quantitative models to explore efficient allocation of R&D funds. In the first model, all reviewers are the same, and applications can be clearly delineated into two pools: innovations (gradual improvements on existing knowledge) and inventions (radically novel technologies). The value of each project is composed of three components: its distance from existing technologies (the greater the better), its cleverness, and its degree of inventiveness (which only inventions have). Reviewers are assumed to be able to correctly assess distance and cleverness, but are unable to assess inventiveness, such that a flat (and, on average, too low) inventiveness score is assigned to inventions. Under such assumptions, reviewers are likely to underscore inventions, resulting in lost effectiveness. The proposed solution is to separate the inventions from the funding pool (focalisation) and select from amongst them by lottery (randomisation); the remaining innovations are ranked and funded as usual. This method, that relies on expert assessment to select a subsection of proposals to enter into a lottery, is labelled by Brezis as 'focal randomisation'.

In Brezis' second model, reviewers are no longer assumed to be equal. They vary in their degree of diligence, such that more diligent reviewers spend more time on their reviews and arrive at more accurate assessments of distance. In the third model reviewers are also allowed to vary in their creativeness, such that more creative reviewers are better able to assess the inventiveness of proposals. With this increased and variable ability to assess the value of proposals, Brezis suggests that the high-value inventions will be occasionally highly ranked, but only by some (more diligent and creative) reviewers. These will show up in the ranked list of proposals as mid-ranking with high variability between reviewers, whereas high quality innovations will be consistently scored highly, and low-quality proposals will be consistently scored low. The information from the reviewers, both in terms of score averages and in terms of variability, can be used to transfer those proposals about which reviewers disagree into a lottery, while also funding the unanimously high-scored top-ranking proposals. This version of focal randomisation does not rely on reviewers' ability to tell apart in advance which proposals are innovations and which are inventions, but rather leverages the variance in reviewer scores as a proxy for this distinction.

Brezis shows, with worked numerical examples, that focal randomisation provides better outcomes than the alternative of straightforwardly funding according to aggregated reviewers' scores. This result is robust across the models, and stems from the ability of random selection to pick highly innovative projects, which is something reviewers (at least in the models) struggle to do.

On top of making a clear policy recommendation, two further features should be noted about Brezis' argument:

Bounded reviewer ability: Reviewers in Brezis' models are not able to ascertain the true value of proposed projects. This lack of reliability is further analysed by breaking down both features of proposal merit (cleverness, distance, inventiveness) and reviewer characteristics (creativity, diligence).

Link between lack of epistemic access and randomisation: The reason focal randomisation is superior to straightforward peer review in Brezis' models is that reviewers make inaccurate evaluations of proposals' values, and these inaccuracies are systemic – some projects, namely inventions, simultaneously have high intrinsic value and low epistemic access, leading to a conservative bias. Focal randomisation leverages information that is available, such as score variability, to identify those regions of low epistemic access, and directs them to a lottery. This results in *higher* effectiveness, in terms of the value of projects chosen, compared to peer review.

Consensus precludes randomisation: As an extension of the above point, when reviewers agree that a proposal should be accepted or rejected, this indicates good epistemic access to the true value of the proposal, and obviates the need for randomisation.

While she focuses on conservative bias, Brezis concludes that focal randomisation could also help ameliorate further biases, expanding the virtues of random allocation beyond effectiveness and into fairness:

> It could also be that referees choose projects in which they are not completely disinterested. They could act not in the public interest exclusively, but might have self-interest, and might, for some subjective reason, dislike a project. […] This 'public choice' perspective would strengthen the importance of introducing randomization, into which no elements of sympathy, approval or power enter. Randomization on the projects where referees disagree not only increases diversity, but is a way of avoiding the tendency to accept projects of 'club' insiders (Brezis, 2007, p. 15).

3.3. Barnett, Graves, Clarke, Herbert and Blakely

Barnett, Graves, Clarke, Herbert and Blakely are public health and biomedical researchers who set out to collect up-to-date and reliable evidence on the shortcomings of grant peer review, mainly its lack of *reliability* or high randomness (Graves et al., 2011) and its low *efficiency* or high costs (Herbert, Barnett, Clarke, & Graves, 2013).

Graves et al. (2011) presents the most thorough measurement published to date of the variability of grant peer review scores.[8] The authors used the raw peer review scores assigned by individual panel members to 2705 grant proposals. All proposals were submitted to the National Health and Medical Research Council of Australia (NHMRC) in 2009. The scores were given by reviewers sitting on panels of seven, nine, or eleven members, and the average score of the panel was used to decide whether a project was funded or not, based on its rank relative to other proposals.

The authors used a bootstrap method to generate a counterfactual population of possible review scores for each proposal, yielding a mean and a variance in that mean, and a confidence interval around the mean. This confidence interval, labelled by the authors the 'score interval', was then compared to the funding cut-off line: proposals whose score interval was consistently above or consistently below the funding line were considered 'efficiently classified' by the review system, whereas proposals whose score interval straddled the funding line were considered as problematic, or 'variably/randomly classified'.

The results showed that overall, 61% of proposals were never funded (score interval was consistently below the funding line), 9% were always funded (score interval consistently above the funding line), and 29% were sometimes funded (score interval straddling the funding line).

[8] An earlier review paper by Cicchetti (1991) covers various measurements with smaller sample sizes.

In the authors' opinion, the discrepancy between the observed levels of variability, and the importance of funding decisions to individuals' careers, is cause for concern. The authors claim the results show "a high degree of randomness", with "relatively poor reliability in scoring" (p. 3). The authors follow with a list of possible improvements to the peer review system. One of their suggestions is to investigate the use of a (limited) lottery, similar to Brezis' focal randomisation:

> Another avenue for investigation would be to assess the formal inclusion of randomness. There may be merit in allowing panels to classify grants into three categories: certain funding, certain rejection, or funding based on a random draw for proposals that are difficult to discriminate (Graves et al., 2011, p. 4).

In addition to the variability of grant peer review, the group also evaluated its cost. The cost of the grant peer review system can be broken down into three components:

1. The cost of writing the applications (both successful and unsuccessful), incurred by the applicants.
2. The cost of evaluating the proposals and deciding on which application to fund, incurred by internal and external reviewers.
3. The administrative costs of the process, incurred by the funding body.

According to Graves et al. (2011), in the funding exercise discussed above, the largest of these costs was, by far, the cost incurred by the applicants, totalling 85% of the total cost of the exercise (p. 3). The authors used full costing of the review process and administration budget, but only a small sample of applicant reports. To complete their data, a more comprehensive survey was conducted amongst the researchers who submitted applications to NHMRC in March 2012. The results of this survey, discussed below, are reported in Herbert et al. (2013).

The authors received responses from 285 scientists who submitted in total 632 proposals. These provide a representative sample of the 3570 proposals sent to NHMRC in March 2012, and display the same success rate of 21%. Based on the survey results the authors estimated, with a high degree of confidence, that 550 working years went into writing the proposals for the March 2012 funding round. When monetised based on the researchers' salaries, this is equivalent to 14% of the funding budget of NHMRC. New proposals took on average 38 days to prepare, and resubmissions took on average 28 days. The average length of a proposal was 80–120 pages.

Using survey data, the authors also tried to detect a correlation between extra time spent on a proposal and the proposal's likelihood of being funded. Surprisingly, no such correlation was found, and given the power of the study this suggests that, on average, 10 extra days spent on a proposal is likely to *at most* increase the likelihood of success by 2.8% (p. 3).

Based on their findings, the authors hypothesise the existence of a curve which associates the accuracy of the peer review system in evaluating the merit of a proposal to the amount of information provided by each applicant (Fig. 1). The hypothetical graph has certain interesting features:

- The graph hypothesises the existence of an 'ideal', which is the amount of information required for the optimal level of accuracy. In the paper this level of accuracy appears close to, though not equal to, 100%.
- In the area left of the 'ideal', i.e. where the information provided is less than the ideal amount, the graph displays diminishing returns, such that equal increases in information provided results in less increase in accuracy the more information has already been provided.
- In the area right of the 'ideal', the graph displays an 'overshoot' effect, with accuracy decreasing as information increases. In the text, this is explained by the claim that the reviewers are being overburdened with too much information.

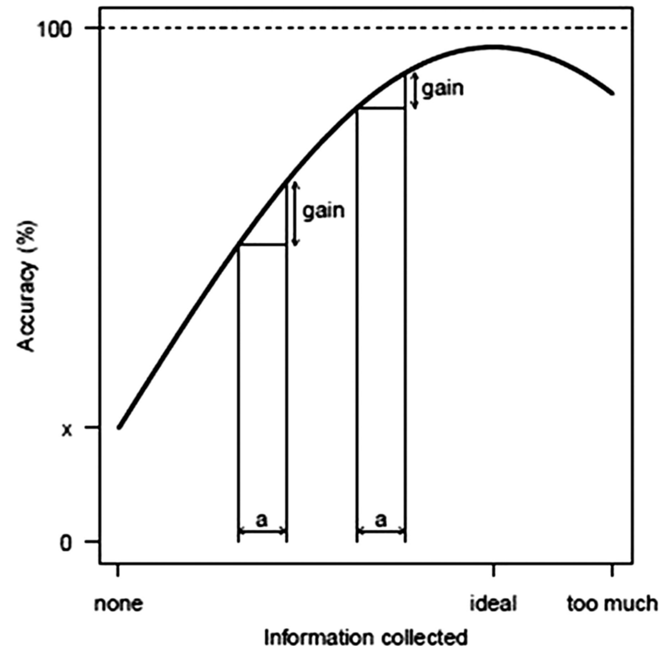

Fig. 1. The accuracy of peer review assessment as a function of information provided. Reproduced from Herbert et al. (2013, Fig. 2, p. 5), published under CC−BY−NC licence: http://creativecommons.org/licenses/by-nc/3.0/legalcode.

The authors rely on their result, that no statistically significant correlation was found between extra time spent on a proposal and its likelihood of success, to argue that the current amount of information provided is more than the ideal. However, one does not follow the other because increased accuracy does not imply higher merit for a proposal. Nonetheless, the authors' description of reviewers having to read 50-100 proposals of 80–120 pages does suggest an unnecessary cognitive burden. Based on their hypothetical curve, the authors' suggestions for reducing the amount of information gathered implies a lower accuracy for the peer review system. The authors believe this lowered accuracy is justified, on cost/benefit grounds, even though in their model a high level of accuracy is possible.

In addition to reinforcing many of Brezis' key points, the key contribution of the group can be seen as providing quantifiable measures for *bounded reviewer ability* and the potential *efficiency advantage* of a lottery. They draw a clear connection between *cost* and *epistemic access*.

3.4. Gillies

Gillies (2014) is the first published paper by a philosopher of science to call for random allocation of research funds. The argument is developed from a critique of the peer review process written by Sir James Black, a Nobel Laureate and discoverer of two blockbuster drugs. Black argues that the slowdown in the rate of discovery of new drugs at the beginning of the 21st century was not due to "the culling of low-hanging fruit" (Gillies, 2014, p. 2), but rather due to a failure in the system of peer review, a system that, according to Black, destroyed scientific creativity. The anti-creative prejudice of peer review, Black and Gillies argue, is believed throughout the research community, leading researchers with creative ideas to keep silent about them and instead seek funding for 'safer' projects. This effect is labelled by Black as 'undesirable feedback'.

Gillies argues that the failure of peer review in supporting creative research stems from 'researcher narcissism' (Gillies, 2014, p. 8). Under this condition, individual researchers believe their chosen approach to the topic is the best possible approach. The first reason that 'researcher narcissism' emerges is psychological. Each individual researcher spends a lot of time choosing their approach, and in the process of choosing they come up with reasons for justifying their eventual decision. These reasons are likely to be shared by other researchers surrounding the individual, as individuals choose a research community that will support their approach. This echoing of reasons further entrenches the individuals' belief in the justification of the superiority of their chosen approach. The second reason for the emergence of 'researcher narcissism' is personal benefit. Individual researchers will enjoy more successful careers if others, including funding bodies, awarding bodies, and students, choose to endorse or follow the same approach as the individual researcher. Conversely, if the field rejects the approach of the researcher, negative career consequences are likely to follow. This incentive to have one's own approach succeed will often lead individuals to believe that it will indeed succeed, or at least behave as if this was their belief.

While it is clear that in a heterogeneous research community not all researchers can be justified in believing their chosen approach is the best one, Gillies goes further to argue that *virtually all* researchers are mistaken in their 'narcissist' beliefs. This is "because no one really knows in advance what research projects are going to succeed. This is because research is, by definition, the exploration of things, which are as yet unknown" (Gillies, 2014, p. 7). While not explicit in the 2014 piece, this position echoes Gillies' Kuhnian reading of science, and his numerous detailed case studies that show researchers under-valuing innovative research from competing paradigms or research schools, an effect that is amplified by peer review (Gillies, 1992, 2005, 2007, 2008, 2009, 2012).

As an alternative to grant peer review, Gillies proposes selecting which researchers to fund using random selection. Proposals will still be generated by individual researchers, as they are under peer review, to check that they fall within the field of the grant. In addition, applying individuals will be evaluated for competence in the field of research suggested, for example, by checking for qualifications such as a research PhD, experience as a research assistant in other projects, or by the applicant's track record of research and publication. These checks, Gillies argues, would be relatively straightforward and much cheaper than peer review, while still addressing the need to filter out 'cranks'.[9] Once proposals have been screened, a random selection is performed from amongst the remaining proposals.

While Gillies considers the likely immediate response to be that random selection will be less effective than peer review in supporting good research, he argues that the converse is true, due to the widespread systemic bias resulting from 'researcher narcissism' as described above. Furthermore, it is envisaged that once peer review is removed applicants will have no reason to opt for 'safe' proposals, and the level of innovation in proposals will increase, thus leading to more creative and beneficial research. Gillies proposes that random choice may first be introduced in a restricted set of institutions, and research can then be carried out to monitor the performance of random choice.

Largely we see Gillies repeating and reinforcing the key points we've seen above, such as *pilot of random allocation*, *efficiency advantage* and *pre-lottery screening*, as well as a more detailed account of the causes of *bounded researcher ability*. Gillies also puts forward a more radical version of *lack of epistemic access*. In addition to the above, two important features should be noted about Gillies' work (especially when taking into account his earlier papers and books):

Detailed historical examples: Cases where we are led to believe, despite the difficulties of counterfactual historical reasoning, that random allocation would have led to better outcomes than peer review (e.g. zur Hausen, pharmaceutical R&D), or that peer review would have led to worse outcomes than whatever system existed at the time (e.g. Frege, Semmelweis, Wittgenstein).

Systemic effects of peer review: Beyond the immediate effect on applicants in a funding round, Gillies draws out the systemic effect on the balance between competing research programmes or paradigms, and on the kinds of proposals individuals consider fundable. Both of these systemic effects lead to a systemic reduction in diversity and creativity across the research community.

3.5. Avin

Another philosopher of science who has argued for science funding by lottery is Avin (2015, 2017, 2018). Avin, like Brezis, utilises formal models to explore the effect of different allocation mechanisms on collective generation of epistemic progress. Unlike Brezis, though, Avin focuses on long-term dynamic effects that emerge as a society of researchers explores a scientific topic of interest. In that, Avin's work can be seen as belonging to the growing literature on formal social epistemology.

Avin (2015, 2017) presents a modification of the epistemic landscape model of Weisberg and Muldoon (2009). Adapted from its original context in evolutionary biology, Weisberg and Muldoon's model depicts a population of investigators exploring a hilly landscape, where different coordinates of the landscape depict different approaches to the study of a scientific topic of interest; the height at each coordinate in the landscape corresponds to the significance of the results that would accrue from pursuing the corresponding approach. The population of investigators is seeded on the landscape at random, with each agent having limited information about the landscape – only the significance of immediately adjacent positions is known. In each simulation turn, the agents follow simple rules, based on their individual characteristics, in the aim of exploring more of the landscape and attaining significant results. The performance of the entire community, rather than of individual agents, is then measured as the individual characteristics of the agents (the rules they follow) are varied.

In Avin's model, the characteristics of individual agents remain the same throughout the simulation. Avin, instead, varies the population itself, through the funding mechanism in a process akin to selection, as well as the topology of the landscape, as approaches gain or lose significance in response to investigators' actions. In Avin's simulation, investigators need to *propose* to work on an approach; at every turn all proposals are aggregated, and a subset is selected based on different selection mechanisms, modelled to represent idealised versions of funding mechanisms. The selected investigators are placed on the landscape for a limited period, representing a time-limited grant, and pursue their research; once the grant period is over, investigators submit their results, and then re-apply for funding to pursue research on the same, or one of the adjacent, research approaches. The selection mechanisms modelled include the community's best estimate of the merit of the proposal, based on past experience (akin to grant peer review), a triage of proposals somewhat akin to Brezis' focal randomisation, a completely random lottery, and automatic renewal such that no new entrants are accepted.

In addition to modelling the population makeup as dynamic, Avin also models the landscape as dynamic. When investigators complete their grants, three effects are triggered:

1. The significance of the approach is significantly reduced, to reflect the one-off nature of discovery.
2. The approach, and all nearby approaches, lose some significance due to reduced novelty.
3. A new hill is added to the landscape in a random position to indicate new research avenues.

When the population dynamics and landscape dynamics are taken together, Avin's result indicate that funding by lottery can significantly outperform funding by peer review.

[9] The responsibility of science funding mechanisms to filter out 'cranks' has been highlighted by Polanyi (1962) in a seminal paper which also presents a strong defence of grant peer review against the foil of project selection by government bureaucrats.

Avin (2018) goes beyond the model, to consider the policy factors that are required to bridge the model results and actual policy recommendations. Two factors of particular interest are the cost and fairness of a lottery. On both aspects, a lottery compares favourably to peer review, as effort required by applicants is significantly reduced, and bias is eliminated (even if that bias is unknown). Avin notes, however, that for these benefits to accrue, the lottery's entry criteria need to be minimal, which could undermine accountability, and so a pre-screening solution is proposed that can capture some fairness and efficiency gains, as well as the effectiveness gains indicated by the model, while maintaining a high degree of accountability. Avin concludes with a list of certain domains where, according to the model or other policy considerations, a lottery should not be applied, for example in applied research with an urgent deadline, in very well-established narrow fields where the ‘terrain’ of open questions is thoroughly mapped out, and in cases of ‘Big Science’, where very large sums are awarded to individual projects employing hundreds or thousands of researchers.

3.6. Fang and Casadevall

Fang and Casadevall (2016) provide the most recent published paper calling for funding by lottery, and in some sense the most ambitious, as it calls for the introduction of random allocation at the NIH, the world's largest funder of basic research. Fang and Casadevall are biomedical researchers, and their proposal is partly based on their previous research (Fang, Bowen, &Casadevall, 2016) that showed reviewers' scores of the top 20% of applications are a very weak predictor of grant productivity (in terms of numbers of publications and citations). The focus on the top 20% matters, as Fang and Casadevall argue that peer review is *no longer* effective, due to falling success rates. They accept that reviewers are able to broadly distinguish between meritorious and non-meritorious proposals, but reject (with the evidence above) reviewers' ability to make fine grained distinctions amongst the meritorious.[10]Note that contra Brezis, Fang and Casadevall thus argue that the greatest uncertainty is near the top of the pile, not in the middle.

Beyond low reliability in distinguishing amongst meritorious projects, Fang and Casadevall list all the objections to peer review that we are by now familiar with: conservatism, bias (they note specifically concerns about gender and race bias), and large time investments. Of these, they consider bias to be the main reason to shift to a lottery, once the lack of reliability of peer review is acknowledged:

> Although the proposed system could bring some cost savings, we emphasize that the primary advantage of a modified lottery would be to make the system fairer by eliminating sources of bias. The proposed system should improve research workforce diversity, as any female or underrepresented minority applicant who submits a meritorious application will have an equal chance of being awarded funding (Fang & Casadevall, 2016, p. 5).

Fang and Casadevall provide extensive details about their proposed lottery system. In addition to *pre-lottery screening*, they propose:

Automatic re-entrance to future lotteries: This would apply to those proposals that have been deemed meritorious in the prescreening, but failed to win the lottery. This has two claimed advantages: it frees up the time of applicants who would otherwise be writing new proposals every year, and it allows institution administrators to plan ahead given the number of proposals from their institution that are waiting for their lucky turn.

Limit to one proposal per applicant: This is to avoid abuse of the system, and to give all researchers who have produced meritorious proposals equal chances.

Detailed feedback only for rejected applications: This is to allow authors of failed applications to revise and resubmit in future prescreening rounds.

Separate lottery for newcomers: This is to guarantee a good representation amongst the funding portfolio for new entrants, increasing diversity.

Funders can hand-pick winners: Programme officers (who manage NIH funding) could use various payment mechanisms at their disposal to provide funds for researchers who end up particularly unlucky, or if a field is at risk of drying up due to a series of unsuccessful draws.

Fang and Casadevall conclude with ten benefits of funding by their modified lottery over current NIH practices. Most of these are familiar from the above works, but the last is worth mentioning: as the lottery will make visible the number of meritorious projects who are going unfunded from year to year, it will help signal the amount of untapped research potential to funders, politicians, and the public, and may help garner more resources for research.

4. Analytic synthesis

We have seen several arguments put forward to support science funding by random allocation. These have emerged from diverse corners of academia, and utilise a wide range of tools to support the arguments: closed-form models (Brezis), novel data collection (Barnett et al.), historical examples (Gillies), agent based simulations (Avin), and surveys of empirical findings (Fang and Casadevall). Here I aim to put together a single coherent synthesis that explains the basic thrust of the argument supporting the policy, as well as points of disagreement that explain the variations we see between different implementations of this policy (and further variations we may see in the future).

P1. Selection by peer evaluation is the natural policy ‘foil’, as it is the dominant form of funding.

P2. Peer evaluation is used to provide a reliable and accountable measurement of proposal merit, such that public funds could be spent on the best science.

P3. The cost of peer evaluation is composed of the cost of preparing proposals, of reviewing them, and of administering the process. Of these, proposal preparation is the largest cost. The costs increase as the level of detail asked for in the proposals increases.

P4. Qua measurement, at least in some domains, peer review is subject to random errors, or at least does not offer a good cost-benefit trade-off at the level of accuracy sought:

- P4a. General argument from cluelessness: research is inherently about venturing into the unknown, precluding any ability to evaluate project merit *ex ante.*
- P4b. Specific argument from cluelessness: most research is incremental and can therefore be reliably evaluated *ex ante*, but some research projects are highly innovative, which precludes an accurate evaluation of their merit. However, these are some of the most valuable research projects. Reviewer disagreement may be an indicator that a proposal is highly innovative.
- P4c. Lack-of-accuracy argument: like all measurements, peer review is subject to random errors, which limit its ability to make fine-grained distinctions between proposals even if coarse grained distinctions are possible, e.g. between meritorious and non-meritorious projects.
- P4d. Cost addendum to lack-of-accuracy argument: like most measurement procedures, there are diminishing returns to accuracy from information provided, such that much more information is required from applicants to make fine grained distinctions than is required for coarse grained distinctions.

[10] A recent study, that replicated the NIH funding process, further supports the claim that reviewers do not agree amongst themselves on the relative merit of high-quality proposals, and therefore calls for a consideration of a modified lottery proposal (Pier et al., 2018).

P5. A main putative shortcoming of funding by lottery is its lack of reliability, but if it comes close to or matches the reliability of peer review in some domains, then other features of the lottery will make it a more favourable selection mechanism.
P6. In addition to random errors, peer review, at least in some domains, is also subject to systemic errors, or biases. These affect both reliability and fairness:
P6a. Researcher narcissism: all reviewers are biased with regards to their own research agenda and competing research agendas. These biases compound when a field is composed of competing schools of thought, or when an existing paradigm is challenged by a new one.
P6b. On top of (P4b), which indicates lack of accuracy when evaluating radically novel proposals, there may also be bias against novel research, as part of or in addition to (P6a).
P6c. Apart from bias for/against ideas and projects, there can be biases for/against individuals, individual characteristics (such as race and gender), institutions and network membership.
P7. With regards to bias for/against ideas, applicants are incentivised to learn and adapt to these biases, creating a feedback loop.
P8. A lottery is an inherently unbiased selection mechanism.
P9. A selection mechanism needs to maintain accountability. Specifically, awarding grants to proposals that are considered, by consensus, to be low quality should be avoided (though this may clash with P4a and P6a).
P10. There is currently only little, and somewhat controversial, data on the accuracy and bias of peer review. Such data is hard to come by.
P11. There is currently very little data on how funding by lottery works/will work in practice.
C1. In areas where the reliability of peer review is low (defined by which of P4 and P6 sub-premises one accepts), a lottery should be trialled as an alternative funding mechanism with potentially better effectiveness and efficiency.
C2. Unless one is very concerned about P4a and/or P6a, proposals and/or applicants should be pre-screened to maintain (coarse-grained) accountability while still reducing costs and (some) bias.
C3. At present, funding lotteries should be run as trial policies with a small subset of all available funds, and outcomes should be compared to the outcomes of funding by peer review.

Versions of the argument above play out in the justification for the funding by-lottery policies surveyed in §5. We should, however, also notice the absence of certain arguments, to do with power and control, which I will only sketch here.

4.1. Control of funding

There is an ongoing debate over who should control scientific research, the state or the scientific community (Bernal, 1939; Polanyi, 1962; Greenberg, 1999, 2003).[11] The main argument for control by the scientific community over state funds (albeit with appropriate oversight and accountability mechanisms) is that the scientific community is best positioned to judge which research avenues are most promising, which researchers are most competent, and which proposals are most plausible and meritorious. Radical scepticism about the ability of the scientific community, as suggested by some who support funding by lottery, threatens the status quo, and a shift of power to the state.

In particular, it is important to differentiate two readings of the lottery proposal. On a per-project reading, the lottery implies that any project is as good as any other – in which case state officials, or other interested parties, may argue for picking their favourites. On the social level, however, such exercise of control would be precisely antithetical to the aims of the lottery, which is to eliminate bias – it is the distribution generated by the lottery, rather than the individual projects chosen, that makes the lottery favourable. If, however, it would be difficult to defend this distinction and sustain the lottery in the long run from interventions by interested parties, then the lottery proposal may well end up backfiring.

4.2. Personal benefits from reviewing proposals

We often talk about the cost for reviewers, in terms of time wasted in reading tall stacks of proposals and endless meetings with bad coffee. However, the researchers who sit on reviewing panels, often quite senior and accomplished in their field, also gain certain benefits from the process. Gillies (2014) hints at one of these: as long as the scientific community is divided over the best framework, theory or approach in a certain domain (which is often the case), reviewers get to exercise some power over the battle, towards the position they most strongly identify with. This, in fact, is *expected* of them, as it is part of the exercise of their expertise. On a more humdrum level, the slow process of scientific publication means that reviewers gain privileged access to scientific progress and plans made by their peers.

As one reviewer notes, the loss of such benefits to proposal reviewers is in no way part of the normative criticism of the lottery proposal; rather, it is a factor that may explain why some actors might be less willing to implement such proposals. Since there is a high degree of overlap between the senior scientists who sit on grant review panels, and the scientific advisors who are likely to consult on matters of science policy, including on the overall mechanism of science funding, there may be personal incentives to support the status quo and reject proposals that would diminish reviewers' benefits. Note, however, that most policy proposals of funding by lottery involve pre-screening by experts, such that reviewers retain at least some of the benefits discussed.

5. Current examples of random allocation of research funds

The above summarises the arguments in the literature for formally introducing a random selection element into science funding. As mentioned in the framing of the paper, however, this is no longer a merely theoretical debate: there are now at least three science funders who employ random selection as part of their funding mechanisms, surveyed below.[12]

5.1. New Zealand Health Research Council – Explorer Grants

In 2013, the New Zealand Health Research Council (HRC) launched a new funding stream, titled 'Explorer Grants', to promote transformative research HRC (2017). Explorer Grants are available in any health research discipline and are worth NZ$150,000 for a term of up to 24 months. Applicants are required to provide a brief project proposal (circa 10 pages for the entire application in 2017). The proposals are anonymised and presented to a panel for initial evaluation of transformative potential and viability. Applications that pass the initial screening are submitted to a lottery, and the lottery winners are funded. Unlucky proposals which fail to win the lottery can be resubmitted in

[11] In recent decades attention has largely shifted to questions of corporate control, but for basic research the main locus of funding, and therefore debate over control, is at the state level.

[12] The Foundational Questions Institute (FQXi) used to run a mini-grants program via lottery, which received a mention as the only such policy at the time in a policy review of alternatives to grant peer review (Guthrie et al., 2013). These grants are of a much smaller scale than the funding streams surveyed below, and at the time of writing were not accepting new applications, and so are not included in this survey.

subsequent rounds.

These are the numbers of applicants and grants awarded to date:

Year	# Applications	# Awards
2013	116	3
2014	24	4
2015	45	4
2016	38	9
2017	34	11

From the above, there have been 31 grants awarded to date (all involving random selection), for a total amount of NZ$4.65M (US$3.25M in current exchange rates).

5.2. *Volkswagen Foundation – experiment!*

The Volkswagen Foundation (VolkswagenStiftung) is the largest private research funding foundation in Germany, with an annual funding volume of around 150 million euros. It has been running the 'Experiment!' funding stream since 2012, with the objective of:

> an exploration of fundamentally new research topics disregarding a high project risk and the vagueness of a successful outcome. The funding is meant for an exploratory phase, which is limited with respect to duration and finance, in order to demonstrate preliminary evidence for the concept's potential. In case of disappointment the scientific explanation of obstacles is a desired result (VolkswagenStiftung, 2017).

Grants are awarded for up to EUR 120,000 and 1.5 years, and cover "experiments and theory in science, engineering, behavioral and life sciences (VolkswagenStiftung, 2017).

The 2017-2020 funding rounds are expected to award 30–40 grants per year, doubling the amount from the 2013–2016 period. The first years of the program saw around 500 applications per year.

Applicants are required to submit brief proposals (circa 5 pages). An initial screening takes place by Foundation staff to guarantee proposals meet the program criteria; this stage filters down from 500 applications to 120–140 proposals. In a second stage, the applications that passed the initial screening are presented to an international and interdisciplinary jury, which selects the 15–20 most promising applications (akin to peer review, though no written evaluations are produced); during this process, each jury member is given a single 'funding joker' which they can use to propose consideration of an application that does not reach consensus.

In the 2017–2020 rounds, a second mechanism will be introduced in the second selection stage, whereby in addition to the 15–20 proposals selected by the jury, an additional 15–20 proposals will be selected at random from the qualifying proposals. Applicants will not know whether their proposals have been selected by the jury or have been lucky in the lottery. The lottery selection is considered to be in 'trial phase', and will be re-evaluated in 2020.

Given the details above, it is estimated that during 2017–2020 there will be 70 proposals selected by lottery, for a value of up to EUR 8.4M (US$9.9M in current exchange rates).

5.3. *New Zealand science for Technological Innovation challenge – Seed Projects*

The Science for Technological Innovation challenge (SfTI) is one of New Zealand's 11 National Science Challenges. It was launched in 2015 as "a 10-year, multi-million-dollar investment aimed at growing a future high-tech New Zealand economy." (Science for Technological Innovation, 2017) In 2016 and 2017 SfTI offered funding for 'Seed Projects', which are intended to bring new people with fresh ideas into the SfTI Challenge. They align with the SfTI Challenge Research Themes and involve high-risk research with potentially high rewards (Science for Technological Innovation, 2017).

Funded Seed Projects receive up to NZ$100,000 per year, and last for up to three years (2016 round) or two years (2017 round). A portion of the funding pool is reserved for Vision Mātauranga, which "aims to unlock the science and innovation potential of Māori knowledge, resources and people for the benefit of all New Zealanders." (Science for Technological Innovation, 2017) In 2016, projects were selected by a panel, in a system akin to peer review. In 2017, however, the panel only provided initial screening for matching the eligibility criteria, both of the challenge in general and of Vision Mātauranga, and were entered into two corresponding ballots. A total of 18 projects were selected from the ballots, out of 79 applications, for a total worth of NZ$3M (US $2.1M in current exchange rates). There are currently no plans to fund more Seed Projects as part of the SfTI challenge.

6. Comparison between policies and arguments

How do the policies surveyed above compare to the proposals that emerge from the arguments for funding by lottery? To what extent do they follow arguments from cluelessness, lack of accuracy, or cost? How do they deal with fairness and accountability issues? These are not always explicit, but we can look at policy elements that may serve as proxies for these concerns. For proxies of accountability, we can see if institutional affiliation is required, if pre-lottery screening is required, and whether the funder follows up with awardees to check progress and quality of outputs. For proxies of concerns regarding effectiveness of selection, we can look at the kind of proposals accepted (only transformative research?), the composition of reviewers per proposal (specific peers, broad interdisciplinary panels, or light touch filtering by the institution's staff), the extent to which a random lottery is used, whether reviewer scores affect lottery chances (e.g. as a weighted lottery or in a focal randomisation process), and whether applicants that fail the lottery are allowed to resubmit. For proxies of fairness concerns, we can look at whether proposals are anonymised, and whether fairness is explicitly mentioned as a benefit of lottery selection. For proxies of efficiency, we can look at whether proposals are required to be short, and whether detailed plans and budgets are mandatory.

A comparison of the policies across these characteristics is presented in the following table:

Characteristic	Explorer Grants	Experiment!	Seed Projects
Just transformative research	Yes	Yes	Yes
Need institutional affiliation	Yes	Yes	Yes
Pre-lottery screening	Yes	Yes	Yes
Screening done by	Experts	Staff	Experts
Random selection	All	Half	All
Reviews affect lottery	No	No	No
Short proposals	Yes	Yes	Unknown
Detailed budget required	No	No	No
Explicit fairness reasoning	No	No	Yes
Anonymous proposals	Yes	Yes	No
Allow resubmission	Yes	Unknown	Yes
Post-funding monitoring	Yes	Yes	Yes

One clear difference between the implemented policies and the proposals in the literature is that the policies only apply for a small subset of R&D, namely transformative research; this contrasts with the arguments which each see themselves as applying to all R&D, or at least all basic research.[13] By limiting their scope, the implemented policies

[13] There are also several related philosophical puzzles here: why is it that noticing the lack of a certain collective property (novelty) in a class of objects (research outputs) does not lead to new action, but carving out (and officially naming) a sub-class of the original class (transformative research), that is identified by certain values of that property (high novelty), does enable such policy change? Can research projects be carved into transformative and non-transformative *in practice*? Can they be so separated *ex ante*? Why is it desirable to have *any* non-transformative research? A survey of transformative research,

shy away from directly attacking the established mechanism of grant peer review, which is something all arguments in the literature do. Indeed, all implemented lottery mechanisms are operated by funders who also run other funding streams, where the selection mechanism for other, non-transformative, funding streams is grant peer review. It may be that data on the results of these trial policies will be required before traditional funding streams will consider shifting from grant peer review to random selection — it is easier to try something new than to change an existing policy.

From the above it is clear that the implemented policies lean more on the side of Fang and Casadevall, in seeing lotteries as a way to make funding of transformative research more efficient (shorter proposals) and to remove explicit biases (anonymised proposals or explicit statement that lotteries contribute to fairness). They do not follow Brezis' prescription of utilising reviewer disagreement to mark proposals for random selection, nor Graves et al.‘s proposal of triage, though the pre-filtering for transformative research may be seen as providing this function. They also do not follow Gillies' more extreme version of concern regarding ‘research narcissism’, as they require both institutional affiliation and passing a pre-filtering panel, which leaves space for bias to sneak back in, though the interdisciplinary makeup of the panel may correct for this bias to some extent. The Experiment! program is explicitly described as a pilot, and is run alongside panel selection, which matches with the general observation that evidence on funding selection is lacking, and more data is required to compare between alternative mechanisms. Another data gathering exercise is being run alongside the Explorer Grants program (Barnett, Graves, Clarke, & Blakely, 2015). For all three streams, it is too early to tell whether they had the kinds of effects that the academic arguments for lottery suggest, but preliminary results are expected in the next few years; more comprehensive comparisons, however, may require much longer, as the value of some research projects is only revealed long after publication, in the case of so-called ‘sleeping beauty’ papers (Ke, Ferrara, Radicchi, & Flammini, 2015).

7. A note on impact

How did we get from the world of 1998, that had one *Lancet* article on funding by lottery, to the world of 2017, with three funding streams implementing (some version of) that policy? The full story will need to be uncovered by historians or sociologists of science policy, but it is helpful to track the way these arguments have related to each other in the literature.

While concerns about various shortcomings of grant peer review have existed for decades, there has been little systematic evaluation of the process (Guthrie et al., 2017). Given this background, integrative interdisciplinary reviews of alternatives to peer review have had a significant role to play in shaping alternative policies. Ioannidis (2011), published in the interdisciplinary journal *Nature*, presents a harsh criticism of current funding practices, and surveys a range of possible measures for improvement, including funding by lottery (citing Graves et al.‘s work on the randomness in current selection methods). Guthrie, Guerin, Wu, Ismail, and Wooding (2013) presents a comprehensive policy-oriented survey of the literature on alternatives to grant peer review. The report was produced by RAND Europe, a think tank often contracted by governments and large funders with a long history of informing policy decisions. In the section of funding by lottery, the report references Ioannidis's review, as well as Graves et al.‘s results and Brezis's model. This RAND Europe report was cited by a New Zealand government review of the Health Research Council's Explorer Grants program (HRC, 2014).

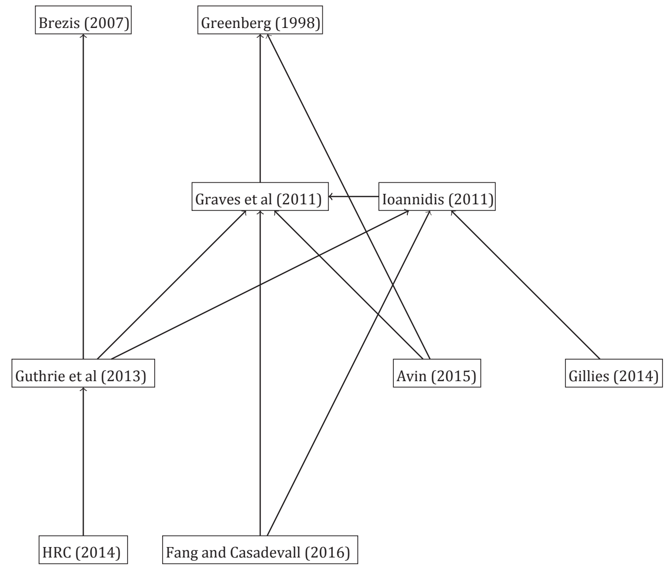

Fig. 2. Partial citation graph for works discussing science funding by random allocation.

The full citation graph for the works discussed in this paper is presented in Fig. 2.[14] Without drawing too many conclusions from a very limited sample size, it seems that philosophy of science could be doing better in engaging the audiences that are relevant for increasing the allocation of resources to help promote scientific novelty.

8. Conclusion

The world in 2018 contains three research funding streams that allocate funds randomly amongst qualifying proposals. Though their total funding amount is negligible when compared to the global basic R&D funding budget, the near-total dominance of funding by peer review as the gold standard of science funding since World War Two marks these policies as surprising changes to the landscape. This surprise is also partly attributable to the non-intuitive idea of funding by lottery, despite arguments being put forward to attempt such policies starting two decades ago. Looking at all the different arguments for randomised funding as a whole, we can see a common structure emerging that emphasises the faults and biases in peer review as a failure to accurately measure (or predict) future research quality, in particular when evaluating novel research. Given this argument, a novel mechanism is proposed – randomisation of selection post initial filtering – that either cuts superfluous investment in useless evaluation (the cost saving argument for lottery) or eliminates harmful bias in selection (the pro-novelty and pro-diversity argument for lottery). While disagreements remain between versions of the argument, and enough uncertainties remain to support different specific implementations, it seems

(*footnote continued*)
its philosophy, limits and criticisms, is beyond the scope of this paper. It should be noted, however, that the way science policy institutions conceptualise science has a direct effect on the policies they are willing to entertain. It should also be noted that the origin of transformative research as a science policy category is linked to Kuhn's picture of science as divided into Normal and Revolutionary. It shouldn't be news that ideas from philosophy of science travel far, or that their criticisms travel much more slowly (if at all). It might be more surprising, at least to some, that large sums of money depend on such concepts, or that the time it takes for them to play out is measured in decades.

[14] Where authors have published more than once on the topic of funding by lottery, I have picked the earliest publication. Later publications (including this one) show more citations across disciplines, including from non-philosophy papers to philosophy of science papers on this topic, though there are yet no citation paths from policy reports such as HRC (2014) to works in philosophy of science.

justifiable to run these policies as trial versions to learn more about the outcomes. This is now happening, and provides an opportunity for engagement for social epistemologists, for philosophers interested in contributing to good policy, and for the (highly valuable, from a researcher narcissism perspective) intersection of these two groups.

Acknowledgements

I would like to thank the editors and anonymous referees for valuable feedback, as well as the participants at the Risk & the Culture of Science workshop in Cambridge. This publication was made possible through the support of a grant from Templeton World Charity Foundation. The opinions expressed in this publications are those of the author(s) and do not necessarily reflect the views of Templeton World Charity Foundation.

References

Agar, J. (2012). *Science in the 20th century and beyond. Polity*.

Arrow, K. (1962). Economic welfare and the allocation of resources for invention. *The rate and direction of inventive activity: Economic and social factors. Nber* (pp. 609–626). .

Avin, S. (2015). Funding science by lottery. *Recent Developments in the Philosophy of Science: EPSA13 Helsinki* (pp. 111–126). Springer.

Avin, S. (2017). Centralised funding and epistemic exploration. *The British Journal for the Philosophy of Science*. https://doi.org/10.1093/bjps/axx059.

Avin, S. (2018). Policy considerations for random allocation of research funds. *RT. A Journal on Research Policy and Evaluation, 6*(1)https://riviste.unimi.it/index.php/roars/article/view/8626.

Barnett, A. G., Graves, N., Clarke, P., & Blakely, T. (2015). What is the impact of research funding on research productivity? http://eprints.qut.edu.au/83127/.

Bernal, J. D. (1939). *The social function of science.* Cambridge, MA: M.I.T. Press.

Brezis, E. S. (2007). Focal randomisation: An optimal mechanism for the evaluation of R&D projects. *Science and Public Policy, 34*(10), 691–698.

Chubin, D. E. (1994). Grants peer review in theory and practice. *Evaluation Review, 18*(1), 20–30.

Chubin, D., & Hackett, E. (1990). *Peerless science: Peer review and US science policy.* Albany: State University of New York Press.

Cicchetti, D. V. (1991). The reliability of peer review for manuscript and grant submissions: A cross-disciplinary investigation. *Behavioral and Brain Sciences, 14*(01), 119–135.

Cole, S., Cole, J. R., & Rubin, L. (1977). *Peer review and the support of science.* WH Freeman.

Dinges, M. (2005). *The Austrian Science Fund: Ex post evaluation and performance of FWF funded research projects.* Vienna: Institute of Technology and Regional Policy.

Fang, F. C., Bowen, A., & Casadevall, A. (2016). Research: Nih peer review percentile scores are poorly predictive of grant productivity. *eLife, 5*, e13323. https://doi.org/10.7554/eLife.13323.

Fang, F., & Casadevall, A. (2016). Research funding: the case for a modified lottery. *mBio, 7* e00422-16.

Gillies, D. (1992). The Fregean revolution in logic. In D. Gillies (Ed.). *Revolutions in mathematics* (pp. 265–305). Oxford: Oxford University Press.

Gillies, D. (2005). Hempelian and Kuhnian approaches in the philosophy of medicine: the Semmelweis case. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences, 36*(1), 159–181.

Gillies, D. (2007). *Lessons from the history and philosophy of science regarding the research assessment exercise, Vol. 61*, Royal Institute of Philosophy supplement37–73.

Gillies, D. (2008). *How should research be organised?* London: College Publications.

Gillies, D. (2009). How should research be organised? An alternative to the UK Research Assessment Exercise. In L. McHenry (Ed.). *Science and the pursuit of wisdom. Studies in the philosophy of Nicholas Maxwell* (pp. 147–168). Ontos Verlag.

Gillies, D. (2012). Economics and research assessment systems. *Economic Thought, 1*, 23–47.

Gillies, D. (2014). Selecting applications for funding: why random choice is better than peer review. *RT. A Journal on research policy and evaluation, 2*(1)http://riviste.unimi.it/index.php/roars/article/view/3834.

Graves, N., Barnett, A. G., & Clarke, P. (2011). Funding grant proposals for scientific research: retrospective analysis of scores by members of grant review panel. *BMJ, 343*.

Greenberg, D. S. (1998). Chance and grants. *The Lancet, 351*(9103), 686.

Greenberg, D. S. (1999). *The politics of pure science.* Chicago: University of Chicago Press.

Greenberg, D. S. (2003). *Science, money, and politics: Political triumph and ethical erosion.* Chicago: University of Chicago Press.

Guthrie, S., Ghiga, I., & Wooding, S. (2017). What do we know about grant peer review in the health sciences? *F1000Research, 6*(1335).

Guthrie, S., Guerin, B., Wu, H., Ismail, S., & Wooding, S. (2013). *Alternatives to peer review in research project funding.* RAND Corporation.

Herbert, D. L., Barnett, A. G., Clarke, P., & Graves, N. (2013). On the time spent preparing grant proposals: an observational study of Australian researchers. *BMJ Open, 3*(5).

HRC (April 2014). Report to the Health Research Council of New Zealand Board: Evaluation of the Explorer Grant Fund. https://www.parliament.nz/resource/en-nz/51SCHE_EVI_00DBSCH_ANR_68456_1_A492798/3c5855c2b7c0dcfc7bd685e16806b80552ed7ca4.

HRC (August 2017). Explorer grant application guidelines. EX218 https://gateway.hrc.govt.nz/funding/downloads/2018_Explorer_grant_application_guidelines_(EX218).pdf.

Ioannidis, J. P. A. (2011). More time for research: Fund people not projects. *Nature, 477*(7366), 529–531. URL https://doi.org/10.1038/477529a.

Ke, Q., Ferrara, E., Radicchi, F., & Flammini, A. (2015). Defining and identifying sleeping beauties in science. *Proceedings of the National Academy of Sciences, 112*(24), 7426–7431. http://www.pnas.org/content/112/24/7426.

Kennedy, J. V. (2012). The sources and uses of U.S. science funding. *New Atlantis, 36*, 3–22. URL http://www.thenewatlantis.com/publications/the-sources-and-uses-of-us-science-funding.

Martino, J. P. (1992). *Science funding: politics and porkbarrel.* New Brunswick, NJ: Transaction Publishers.

NIH (2013a). NIH grants policy statement. http://grants.nih.gov/grants/policy/nihgps_2013/, Accessed date: 9 November 2013.

NIH (2013b). NIH grants process overview. http://grants.nih.gov/grants/grants_process.htm, Accessed date: 18 November 2013.

NIH (2017). NIH director's transformative research awards. https://commonfund.nih.gov/tra.

NSB (2016). Science & Engineering Indicators 2016. https://www.nsf.gov/statistics/2016/nsb20161/#.

NSB (May 2007). Enhancing support of transformative research at the National Science Foundation. https://www.nsf.gov/nsb/documents/2007/tr_report.pdf.

NSF (2013a). Grant Proposal Guide. http://www.nsf.gov/publications/pub_summ.jsp?ods_key=gpg, Accessed date: 9 November 2013.

NSF (2013b). US NSF - Merit Review. http://www.nsf.gov/bfa/dias/policy/merit_review/, Accessed date: 19 November 2013.

NSF (2017). Introduction to transformative research. https://www.nsf.gov/about/transformative_research/.

Pier, E. L., Brauer, M., Filut, A., Kaatz, A., Raclaw, J., Nathan, M. J., et al. (March 20, 2018). Low agreement among reviewers evaluating the same nih grant applications. *Proceedings of the National Academy of Sciences, 115*(12), 2952–2957. https://doi.org/10.1073/pnas.1714379115 published ahead of print March 5, 2018.

Polanyi, M. (1962). The republic of science: Its political and economic theory. *Minerva, 1*, 54–73.

Science for Technological Innovation (2017). 2017 SEED PROJECT FUND- ING Q & As. http://www.sftichallenge.govt.nz/sites/default/files/2017-06/2017%20SEED%20funding%20Q%20and%20As_0.pdf, Accessed date: 7 September 2017.

VolkswagenStiftung (November 2017). Experiment! in search of bold research ideas. information for applicants 100. https://www.volkswagenstiftung.de/fileadmin/downloads/merkblaetter/MB_100_e.pdf.

Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy in Science, 76*(2), 225–252. http://www.jstor.org/stable/10.1086/644786.