NBER WORKING PAPER SERIES

UNCERTAINTY AND RISK-TAKING IN SCIENCE: MEANING, MEASUREMENT AND MANAGEMENT

Chiara Franzoni
Paula Stephan

Working Paper 28562
http://www.nber.org/papers/w28562

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2021

The authors wish to thank Erin Baker, Luigi Buzzacchi, Enrico Cagno, Renato Frey, Massimiliano Guerini, Francis Halzen, Chris Hesselbein, Greg Petsko, Henry Sauermann, Viola Schiaffonati, Massimo Tavoni, Paolo Trucco, Reinhilde Veugelers, Roberta Zappasodi for fruitful discussions. The work contained herein has greatly benefited from discussions with scholars and professionals who participated in the Science of Science Funding annual meetings at the NBER Summer Institute and with seminar participants at the Kellogg School of Management, Northwestern. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Uncertainty and Risk-Taking in Science: Meaning, Measurement and Management
Chiara Franzoni and Paula Stephan
NBER Working Paper No. 28562
March 2021
JEL No. H41,I23,I28,J18,O3,O31,O33,O38

ABSTRACT

An underlying rationale for public support of science is that private companies underinvest in research of a risky nature. Yet, risk in science is a poorly understood concept. This paper sets out the foundations for understanding, measuring and managing risk in science. We review insights offered from existing fields that study risk. These contributions, combined with knowledge gained from studies of science, are used to build a conceptual model of risk in science. The model is illustrated with examples drawn from the development of the IceCube Neutrino Observatory. It disentangles different components that determine risk and is used to operationalize an expert-based risk metric, potentially useful in peer review. Moreover, we review emerging empirical work on risk-taking in science, most of which suggests that the current reward structure of science discourages risky research. We develop and outline strategies for hedging and encouraging risk taking. We conclude by proposing a rich agenda for future studies, which is both intellectually challenging and critical for the future of science.

Chiara Franzoni
Department of Management, Economics
and Industrial Engineering
Politecnico di Milano
20133 Milan
Italy
chiara.franzoni@polimi.it

Paula Stephan
Department of Economics
Andrew Young School of Policy Studies
Georgia State University
Box 3992
Atlanta, GA 30302-3992
and NBER
pstephan@gsu.edu

1. Introduction

An underlying rationale for public support of science is that private companies underinvest in research of a risky nature (Arrow 1962; Nelson 1959). Both spillovers and the time horizon necessary to recoup the investment discourage the private sector from investing in risky research. Yet risky research is essential for shifting the knowledge frontier. To the extent this was the case in the past, in the days of Bell Labs and other large private research labs, it is more so today, when private companies, taking a short-run view, engage in less basic, risky research (Arora, Belenzon, and Patacconi 2018; Budish, Roin, and Williams 2015; Fleming et al. 2019). Yet, among researchers today there is concern that risk-taking in science is declining not only among those in the private sector but among those working in the non-profit sector (Edwards et al. 2011; Fedorov, Müller, and Knapp 2010). Many put the onerous on funding agencies, focusing on ways in which science is funded and grants awarded as well as tight budgets faced by agencies (Azoulay, Graff Zivin, and Manso 2011; OECD 2018; Petsko 2011). James Rothman, for example, who shared the Nobel Prize in Physiology or Medicine in 2013, told an interviewer the day after he received the prize, that "he was grateful he started work in the early 1970s when the federal government was willing to take much bigger risks in handing out funding to young scientists." Rothman went on to say "I had five years of failure, really, before I had the first initial sign of success. And I'd like to think that that kind of support existed today, but I think there's less of it. And it's actually becoming a pressing national issue, if not an international issue." (Harris 2013). There is also evidence that the rewards to science not only discourage risk-taking on the part of scientists but increasingly do so (Foster, Rzhetsky, and Evans 2015; Stephan, Veugelers, and Wang 2017; Wang, Veugelers, and Stephan 2017).

This discussion often occurs in the absence of well-defined and developed concepts of what uncertainty and risk-taking mean in science.[1] This paper sets out to address this void. The core contributions of the paper are twofold. First, we lay out a conceptual model of risk in science, which can also be used to obtain expert-based metrics of risk. Second, based on the model and studies of risk taking, we discuss how the current academic environment

[1] Science is not the only domain in which the term *risk* is ambiguous or difficult to define. The problems posed by defining risk have been extensively debated in various fields. For a review see e.g., Hansson (1989) and Aven (2012).

affects risk-taking in science, address ways in which risk can be hedged or encouraged and sketch-out a rich research agenda for future investigation.

The structure of the paper is as follows. Section 2 clarifies concepts and assumptions used in the discussion of risk with the goal of eliminating potential sources of confusion and common misconceptions. Section 3 reviews insights about risk and uncertainty provided by adjacent literatures. Section 4 provides a full conceptual model of risk in science, using the IceCube Neutrino Observatory as an example to illustrate risk components. In Section 5, we explain how the model can be used to obtain an expert-based assessment of risk in science and review alternative approaches to measuring risk. In Section 6 we review studies of risk taking in science and conclude that the current reward structure of science discourages risk-taking. We discuss ways in which risk can be hedged in Section 7 and encouraged in Section 8. We close in Section 9 by outlining directions for future research.

2. Definitions and potential pitfalls

The road to the study of risk in science is paved with misconceptions and prone to misunderstanding. We thus preface this essay by clarifying concepts and assumptions underlying the discussion.

First, the terms risk and uncertainty are ambiguous and prone to generate confusion. Many scholars of economics of science are acquainted with the work of Frank Knight (Knight 1921), who stressed the measurability of the probability of an event in order to distinguish risk from uncertainty. This distinction, however, is rarely found in contemporary economics (Feduzi, Runde, and Zappia 2014). A more recent approach, following Ramsey and De Finetti, stresses the ubiquitous imperfect information that surrounds agents' decisions in real life (Marinacci 2015).[2] Following this approach, it makes sense only to talk about *subjective* probabilities, i.e. degrees of beliefs expressed with more or less confidence. Subjective probabilities are 'measurable' by asking agents their willingness to bet on their beliefs. This view renders the risk-uncertainty distinction void of substantive meaning. Recognizing and subscribing to this view means that we should talk only of *uncertainty* in

[2] In this view, uncertainty is a state of imperfect knowledge that relates to all aspects of a decision, including the states of the world (e.g., what might happen, which consequences might derive), and not just to the probability of each state.

science. However, given that *risk* is the term most commonly used (albeit not in the Knightian sense) by scientists, policy makers and granting agencies to speak about the uncertainty of research, we use the term *risk* as a synonym for *uncertainty* in this essay.

Second, research is not free (Stephan 2012). From the researcher's point of view, time spent on research has an opportunity cost; from a university and/or foundation's point of view, the materials, equipment and space devoted to research come with a cost. These costs should be viewed as *participation costs* and are equivalent to those incurred when purchasing a lottery ticket in which one can gain or lose. In this case, participation costs are those incurred by funders or researchers who 'buy' a chance to discover the outcome of a potential course of action. are fundamentally and conceptually different from the gains and losses that can result from the course of action. Although participation costs can be one element to consider when evaluating the *yield* from doing a research project, they should not be confused with losses.

Third, the value associated with successfully accomplishing proposed research and the prospects of doing so varies across projects. Some projects have a high prospects of achieving their research goals, others have lower prospects of accomplishing their goals. Some projects have a high value if successful; others do not. Individuals or organizations that pursue high value research with low prospects of success can be described as being risk-takers; those that do not can be described as being risk-averse.

The discussion of risk in science often conjures up comparisons of risk associated with financial investments. Such a comparison is beneficial in the sense that the concept of risk in finance, especially after the dramatic financial crises of 2008 and 2020, are well known by the public. However, the comparison has drawbacks in the sense that risk and volatility differ between finance and science in several important respects. First, in finance, the outcomes of investments vary both in the negative and positive spectrum (as many of Bernie Madoff's investors learned all too well). In science, instead, volatility is rarely associated with strictly negative outcomes (losses), unless in cases where research involves harm to a patient, a researcher or significant damage or loss of resources and reputation. More commonly, volatility in science refers to the distribution of research findings that can be thought of as either on the upside or as the status quo, in the sense that, although the research produces no loss, it does not substantially advance the area of study. Second, in finance, risk has both a *speculative* meaning, which focuses on the expected gains associated

with an investment, and a *preventive* meaning, which focuses on avoiding the losses associated with an investment. Science funding organizations that embrace the goal of high risk/high gain, such as the ERC, clearly subscribe to the speculative concept of risk focused on expected gain. The preventive meaning is also present in science, but not with the goal of insuring organizations or people against losses, but rather of insuring organizations against funding research that comes up "empty." One could argue that it is this "empty basket" outcome that review panels insure against, focusing their review efforts on reasons that the research will not yield results rather than focusing on the value of the research if it were to succeed.[3] Third, in finance, volatility is a popular construct because past volatility (or lack thereof) can provide the basis for prediction and calculation of risk associated with an investment. For example, it is common to calculate the alpha associated with an investment, as well as other measures such as the beta, standard deviation, R-squared and Sharpe ratio, based on price fluctuation history. The reason is that past volatility is easy to calculate and can predict future volatility. Instead, in the case of science, past volatility is not necessarily predictive of future volatility.[4] Moreover, even if volatility were predictable, it would be both difficult to observe, because many unsuccessful projects go unpublished (Fanelli 2010; Rosenthal 1979), and to compute, because the value of research findings is not fully understood, especially in the short term.

Before commencing, we address some extreme simplifications and misconceptions that can hijack the discourse on risk and uncertainty in science. One extreme simplification is that "scientific research is always risky." This blunt approximation fails to recognize a considerable degree of nuance, including that, in certain areas of research, risk is limited or mitigated by the character of the research. First, a non-negligible share of research involves virtually zero risks. Consider, by way of example, the Cochrane reviews that consist of collecting, coding and jointly-testing the results of multiple studies of the same medical treatment. Cochrane reviews require rigorous methods and provide valuable results to

[3] A possible reason review panels may focus on the "empty basket" is fear of wasting the limited money they have. This is an empirical question that deserves investigation. Today insurance against some undesirable outcomes is assumed at the institutional level, for example via the actions of the IRB and GPRD on ethical and privacy matters, or via the requirements of funding agencies that the institution certifies access to needed equipment and resources. As a result, while, in the past panels concerned themselves with insuring against some of these problems, today panels arguably focus on insuring against coming up empty handed.

[4] An exception is when the routinization of aspects of research provide scientists with increased confidence in predictions. By way of example, protein structure determination was greatly facilitated by developing an "automated pipeline for protein production and structure determination" which included the development of robots that could grow and screen crystals (Stephan 2012: 93).

scholars and practitioners. Yet, the researcher who chooses to perform a Cochrane review has no doubt that the review can be accomplished and published; the volatility involved in the outcome is practically nonexistent. Second, some research is bound to produce non-zero outcomes, although the value of what is found may be uncertain. For example, many archeological excavations are done after a site has been identified. In these cases, findings are guaranteed, but there is uncertainty with respect to what they will be and their importance. Archival studies, large statistical analyses of galaxies, or research on the collateral effects of approved drugs are other examples of similar situations where risk is very small. Third, some research projects have predictable outcomes, although there is uncertainty concerning the time or costs required to obtain them. An example is the Human Genome Project, which was formally launched in 1990, with the aim of sequencing the 3 billion base pairs of the human genome. At the time the project was started, the set of techniques readily available was sufficient to guarantee eventual success. However, there was uncertainty concerning whether the project could be accomplished in the 15 years that it aspired to. In the end, a working draft of the genome was obtained within ten years, thanks largely to improvements in the technology (Stephan 2012: 88). These three cases suggest that a non-negligible part of research involves projects that are virtually certain to lead to an outcome. In these cases, risk is confined to the uncertainty associated with the value of the outcomes or to the time and resources needed to achieve the research, but not to *whether* an outcome will be forthcoming. In the next sections we expand the discussion by considering the array of components that coalesce to determine risk in science.

Another misconception concerning risk in science is that engaging in a new line of research is always the riskier course of action, while continuing along an existing research path is the play-safe alternative. Consider, for example, the case of James P. Allison. Allison had spent most of his career studying the use of antibodies blocking the immune inhibitory molecule CTLA-4, as a strategy to unleash the immune response to cancer. In 1995, he understood that CTLA-4, a T-cell surface receptor, served to dampen T-cells responses and could be used as a target for cancer immunotherapy (Krummel and Allison 1995). In the same years, other cancer immunotherapy approaches were being investigated, including cancer vaccines and agonist antibodies activating immune stimulatory receptors. However, cancer vaccines showed very limited efficacy, as an influential NIH review pointed-out in 2004 (Rosenberg, Yang, and Restifo 2004), and several agonist antibodies (e.g., targeting CD28, CD40, or 4-1BB) were found to cause serious adverse effects, which in some cases

were life-threating for healthy patients (Suntharalingam et al. 2006). As a result, the clinical community was skeptical of Cancer Immunotherapy in general and pharmaceutical companies were uninterested in further development of these approaches. In spite of this, Allison engaged with a small biotech company to develop a CTLA-4 blocking antibody for clinical use to test in cancer patients, called Ipilimumab. The clinical trial was eventually successful and Ipilimumab became the first immune checkpoint inhibitor drug to receive FDA approval for cancer treatment in 2011 (Wolchok et al. 2013). Immune checkpoint inhibitors have since become one of the most promising frontiers of cancer treatment research. Allison shared the Nobel Prize for Medicine in 2018 (Dobosz and Dzieciątkowski 2019). As the example illustrates, in this case the choice to persist could arguably be described as risky behavior. Thus, undertaking a new line of research is not always the risky behavior and continuing a line of research does not always imply risk avoidance. Kuhn refers to this choice as an "essential tension" (Kuhn 1991), noting that both alternatives, not just the former, are hazardous. Working in a new area of research often requires formulating new theory and using new methods, both of which arguably involve risk. On the other hand, persisting along a line of research can often provide a more predictable path, albeit one with diminishing returns. But it can also lead to a dead-end with no results, perceived as risky behavior by peers. If we want to understand risk, we should look at the uncertainty of prospective results and refrain from taking shortcuts that assume identity between risk and any given observed behavior.

Finally, a further misconception is that some scientists are prone to take risks and others are conservative and risk-averse, as if the attitude of scientists towards risk were a fixed individual characteristic. While recent empirical work has established with reasonable certainty that attitudes towards risk are stable traits, specific to each individual (Frey et al. 2017; Mata et al. 2018), individual risk preference explains only about half of the predicted variance in the measures of risk-taking. The rest depends on erratic behavior and on domain-specific attitudes that individuals show towards various facets of risks encountered in life.[5] Although few if any studies have looked specifically at risk-attitude and risk-taking in scientific research, we would expect that, as for humans in general, the attitude of scientists for embracing risk in research varies at any given point in time, depending on their role or the situation that they face. For example, prior correlational studies suggest that scientists who

[5] Domain-specific factors relate, for example, to preferences for risks exhibited in investments, attitude for seeking thrill and adventure, disinhibition in health risks, et cetera.

have yet to get tenure, may eschew risk; those who have a stable career position may embrace risk (Franzoni and Rossi-Lamastra 2017). Other commentators have proposed that scientists who have obtained success in early-career are less likely to be at risk of losing their reputation and hence more likely to engage in risky research (Henrickson and Altshuler 2012). Regardless of career stage or past success, we expect scientists to share the trait of overconfidence about their ability to achieve success consistently found in individuals' assessment of risk (Camerer and Lovallo 1999; Weinstein 1980).

3. The Meaning of Risk in Adjacent Literatures

We frequently use the word *risk* in everyday life. As such, *risk* at first appears to be a rather intuitive concept. However, when making *risk* a subject of scholarly investigation, we quickly realize that we lack a shared, let alone precise, understanding of the meaning of risk. Scholars of risk, recognizing this, have offered extensive discussion on the topic, and have noted a number of different meanings of risk in different literatures (Althaus 2005; Aven and Renn 2009; Hansson 2002, 2018). Thus, a pre-condition to holding productive discussions is to develop a sound conceptual understanding of risk *in science*, bringing together different pieces of relevant theory which provide insightful concepts and tools. Failure to do so not only undermines our ability to make conscious decisions regarding risk, but also leads to a Tower of Babel in which scientists from different backgrounds implicitly use notions germane to their discipline, but alien to others. In this section we draw insights from four main approaches. All see risk as a manifestation of uncertainty, but each focuses on specific aspects of uncertainty. We explain each in this section. Together, they constitute the building blocks with which we set about defining the meaning of risk in science, presented in the next section.

Risk Analysis. A primary focus in engineering studies of risk regards the occurrence and impact of potential *events*. The intent of this literature is fundamentally utilitarian. The main focus is on representing and quantifying the risks involved in a situation in order to facilitate making decisions. The quantification of risk in Risk Analysis is sometimes called 'Technical assessment' (Renn 1998) or 'Risk metric' (Johansen and Rausand 2014). This is a pragmatic approach that disentangles complex problems into a number of simple pieces, such that each item can be quantified in isolation and then combined. The standard model implies three items (Kaplan and Garrick 1981): the *scenarios* (e.g. what can happen?), the *probability*

associated with each scenario (e.g., how likely is this to happen?) and the *consequences* associated with each scenario (e.g., what loss/gain would this lead to?). More sophisticated models can further disentangle more fine-grained items (Aven 2011). In applied risk analyses, each item is analyzed by one or more experts and the combined outcome is simply calculated as *probability* times *consequence* (Kasperson et al. 1988). This approach for quantifying risk is commonly used in insurance, where the primary focus is on *unwanted events*, i.e. events that deserve special consideration because of their negative impact.[6] Mirroring this view, in the insurance literature it is common to distinguish two families of strategies of risk reduction: reducing the probability of a loss, called '*protection*', and reducing the size of the loss (the consequences), called '*insurance*' (Ehrlich and Becker 1972).

Return and volatility. In Finance, risk refers to the uncertainty concerning the *return* on the capital that the entrepreneur or the shareholders bears. This view of risk is well represented by the volatility of returns. The uncertainty is in both the upside -the *profits*- and the downside – the *losses*. The quantification of risk is needed in order to evaluate the assets. The standard approach is to estimate the probability distribution of future returns of the asset and measure the level of dispersion in terms of variance (Markowitz 1952; Tobin 1958), or its square root, the standard deviation, commonly called *volatility.* Because prices of traded assets are widely available and because measures of volatility over time are self-correlated, the volatility of asset prices observed in the past is commonly used to forecast volatility of the asset prices in the future. This is done with econometric methods, which assign diminishing weights to the increasingly distant past.[7] The ability to measure and predict asset volatility and the correlation of volatility has led to the emergence of *portfolio diversification* as a common strategy of risk-coping in finance, that is: investing in a portfolio of assets, such that the average portfolio volatility is stabilized at a level deemed desirable by the investor, conditional on the desired returns.

Probability and Ambiguity. In Probability and Decision Theory, risk and uncertainty refer to *knowledge regarding the likelihood* of contingencies (Marinacci 2015). There are situations in which the cause of uncertainty is rather well-known. For example, when

[6] E.g., fire, theft, injury, et cetera. This view of risk as danger places the uncertainty in the area of losses: from zero -the status quo- downwards.
[7] The most popular are the Autoregressive Conditional Heteroskedasticity (Engle 1982) and the Generalized Autoregressive Conditional Heteroskedasticity (Bollerslev 1986).

throwing a dice, the physical properties of the dice -it having 6 equal-shape faces- determine the possible realizations -6 states with equal probability. In this case, the model that generates uncertainty can be understood and described as an object, and risk can be expressed numerically in terms of *objective probabilities*. Although such situations are not common in real life, objective probability can be applied, with some degree of simplification to other -more common- situations, in which past occurrences, recorded as *frequencies,* are revealing of the model that causes the uncertainty. For example, the mechanism that causes an illness to be fatal may not be fully known. However, it may be reasonable to assume that the mortality rate observed among people with the illness describes some objective property of this mechanism. The relative frequencies can thus be used to shed light on the risk of death and interpreted as objective probability in this case (Hájek 2019). Such calculations have been common during the COVID-19 pandemic. Conversely, when the model that generates uncertainty is unknown and frequencies are unavailable or non-informative, we have a condition of epistemic uncertainty, in which objective probability is inapplicable and we can only resort to *subjective probability*. These situations are common in science. For example, a scientist who prepares an experiment often has no prior observations, because the experiment is new, and has only incomplete theories that suggest potential results. The scientist can of course express her best belief concerning the outcome, based on the scant knowledge that she has and on personal experience. This subjective probability can be thought of in Bayesian terms as the *degree of confidence* in the scientist's belief. Situations like this are also called deep uncertainty or *ambiguity* (Ellsberg 1961).[8]

Recent theoretical works in decision theory have concentrated on modeling decisions under ambiguity (Klibanoff, Marinacci, and Mukerji 2005). Empirical work has shown that humans are not only risk-averse, that is they prefer a sure thing over a gamble of equal expected value but are also ambiguity-averse. That is, they prefer situations in which they face objective probabilities as opposed to situations in which they face subjective probabilities (Ellsberg 1961; Tversky and Fox 1995). For example, they prefer to draw a marble from an urn that they know has half reds and half blacks, than to draw from an urn that one expert says is all reds and another expert says is all blacks. (Berger and Bosetti 2020; Ellsberg 1961). In situations where there is considerable scientific disagreement, such as models that predict the relationship between levels of CO2 abatement and consequent climate

[8] Different scholars have used different terms (see Camerer and Weber 1992: 326).

change, it appears that policy makers prefer options that reduce not only climate change impact, but also the uncertainty/disagreement among the models (Berger, Emmerling, and Tavoni 2016). This leads us to discuss a fourth important stream of contributions.

Human Cognition of Risk. In Social and Cognitive Psychology, risk is examined from a human perspective. Scholars generally agree that the human mind understands risk both through analytical thinking, and through intuition (Epstein 1994; Evans and Stanovich 2013; Sloman 1996). That is, humans are capable of reasoning about risk in a logical and rational way, for example when they consider probabilities. But they also hold instinctive reactions when confronted by risk, for example when they feel danger (Loewenstein et al. 2001; Slovic et al. 2005, 2010).[9] A classical stream of research, known as the *psychometric paradigm*, has conducted extensive empirical research to explain risk perception and to identify the kinds of hazardous situations associated with feelings (emotional intuitions) of risk. The conclusion, which is widely accepted, is that there are two kinds of risks that prompt the strongest emotional reactions in people. These are: "*dread risk*", i.e. the possibility that something uncontrollable, irreversible or catastrophic will occur, and "*unknown risk*", the exposure to new, unforeseen or delayed harms (Fischhoff et al. 1978; Slovic 1987). This can explain, for example, the skepticism and consequently slow progress surrounding several streams of research, such as power generation from nuclear fusion, where it is easy to imagine catastrophic scenarios and uncontrollable unknown events.

A second important focus of psychological research on risk is how people behave and make decisions in conditions of uncertainty. The main contributions in this respect is that human understanding of risk and probability is biased in systematically predictable ways (Tversky and Kahneman 1974). In particular, people systematically undervalue perspective gains, while they exaggerate the magnitude of perspective losses (Kahneman and Tversky 1979). Furthermore, the overvaluing of losses is larger in magnitude than the undervaluing of gains (Kahneman and Tversky 1979), a condition called *loss aversion* (Tversky and Kahneman 1991, 1992). Consequently, when people have to make decisions that involve uncertain gains or losses, their decisions depart systematically from what rational behavior would predict and their behavior is inconsistent and opposite in the spectrum of gains and losses (*reflection effect*) (Kahneman and Tversky 1979). To be more specific, when individuals face decisions that involve prospects of gains, they are *risk-averse*, that is, they

[9] Note that the instinctive reaction is solely associated with potential danger, and thus to a negative view of risk.

would give away gains for more certainty (prefer gambles with greater certainty even if they involve smaller expected gains). When they face decisions that involve prospects of losses, they are *risk-seeking*. That is, they would accept greater losses as long as they are less certain (prefer gambles with less certainty even if they involve a greater expected loss over a smaller but sure loss) (Kahneman and Tversky 1979).

In the following sections, we draw on these four perspectives from distinct disciplines to frame the discussion concerning risk and risk-taking in science in a conceptually-sound way.

4. Components of a Model of Risk in Science

In this section we outline the *components* of a model for representing risk in science. To facilitate our understanding, we take as an example the IceCube Neutrino Observatory at the Amundsen Scott Station at the South Pole. The project was initially proposed in 1987 by Francis Halzen (University of Wisconsin) in a co-authored paper that he presented at a cosmic ray conference in Lodz, Poland (Halzen and Learned 1988) that discussed the possibility of using deep polar ice as a detector.

Neutrinos are subatomic particles of nearly zero mass that have very little interaction with other masses and hence travel undisturbed across matter in outer space. The observation of neutrinos can thus shed light on astrophysical phenomena originating outside our solar system, such as the formation of supermassive black holes; more generally their observation can shed light on the origins of the universe. Although scientists have explored ways to detect neutrinos since the late 1950s, constructing, for example, detectors in mines and lakes, at the time Halzen proposed placing a detector at the South Pole no detection device had successfully observed neutrinos from outside the solar system. The research challenge was thus to build an instrument capable of detecting such neutrinos and determine the direction from which they came and examine "the relevant optical properties of deep Antarctic ice." (Halzen and Learned 1988).

In 1988 Halzen and colleagues were awarded $50,000 from NSF to study the optical quality of ice.[10] The research team at that time knew little about ice or the challenges associated with drilling in ice, which was necessary in order to embed the sensors. This exploratory project evolved into the proof-of concept project AMANDA (Antarctic Muon and Neutrino Detector Array), supported by NSF with additional funding from other foundations and countries. IceCube, which incorporates the AMANDA arrays, received its initial funding from NSF in 2000. At completion of construction in 2010, the project had placed 5,584 digital optical modules in a series of 88 holes drilled into a cubic kilometer of ice, lying 1.5 kilometers below the surface at the South Pole. The ice lying above the sensors shields the sensors from radiation at the earth's surface. The basic principle behind this design is that when a neutrino collides with a nucleon it produces a muon through inverse beta decay. When this occurs, a pale blue light known as *Cherenkov radiation* is emitted, which can be detected. Importantly, the light bounces back in the exact same direction from which the neutrino came. As a result, the position of the cosmic object from which the neutrino originated can be inferred. Some data are sent to the IceCube Project at the University of Wisconsin by satellite. The balance of the data is stored on hard drives and sent once a year to the researchers.

Let us now examine the risks involved in the IceCube Neutrino Observatory. For simplicity, let us assume that the project sought funding all at once and let us take the point of view of a scientist or panel member who is tasked with evaluation of the project. There are sufficiently large numbers of uncertainties involved that it would be difficult to judge the overall risk of the project without a conceptual framework that first identifies and analyses multiple components of risk in isolation and then combines the components into a model of risk. In order to build a suitable framework, we adopt the approach of risk metric used in risk analysis and identify a set of questions concerning risk components. We discuss how the components should be combined in the next section.

A first question to consider is '*what* can be found'. The answer normally entails a range of options. The range can be especially large in exploratory research, the goal of which is to shed light on unknown domains of the natural world (e.g. space exploration, deep ocean exploration). In the case of the IceCube Neutrino Observatory, for example, the goal was to

[10] The funding was obtained in the form of a Small Grant for Exploratory Research (SGER), which did not require external review. (Bowen: 138-139).

explore the cosmos using ice to detect neutrinos and the direction from which the neutrinos came. In empirical research aiming at the test of formal hypotheses, the range is usually narrower.

There are two important things to note concerning what can be found. The first is that uncertainty concerning what can be found cannot be represented by probability. Uncertainty here regards outlining the possible states of the world and captures one essential feature of research: its being an open-ended quest (Nelson 1959). The second is that this component is represented by a range of alternative scenarios. How many scenarios are appropriate is a question of practical relevance that should be addressed with a degree of pragmatism. Research that involves one formal hypothesis may be represented by two scenarios - hypothesis rejected or not-. Other research may require three or more hypotheses. For the sake of clarity, this risk component should outline only the range of alternative *primary outcomes* and leave aside secondary outcomes of research that may be a byproduct of the research agenda. We return to this point later in this section, when discussing the possibility of secondary outcomes (*what else*).

Questions two and three conceive uncertainty in terms of the probability of each scenario happening and focus respectively on methodological and natural risk. Methodological risk—question two-- can be spelled out as: '*how* likely is the proposed approach to work'. Uncertainty here concerns (subjective) probability and is both epistemic and technical. 'Epistemic' because the theory behind the design could be fallacious, or there is uncertainty regarding the scientific knowledge on which a method is grounded. For example, in the IceCube project the strategy for detecting neutrino bursts accompanying the formation of black holes was drawn in part from theoretical work by Shi and Fuller (1998). 'Technical' uncertainty because several details of the execution are typically not yet resolved at the stage of conception. For example, in the IceCube project it was not clear how deep one would have to drill to find bubble-free ice.[11] The original proposal submitted to NSF assumed that one would only have to go to a depth of around 500 meters.[12] This turned out to be off by 1000 meters (Bowen 2017:129-30). There was also uncertainty that "optical attenuation of deep ice in the mid UV range … characteristic of the Cerenkov light emitted

[11] Bubbles cause light to bounce in all directions, which means that if a Cherenkov cone were detected, it would not travel in a straight line.
[12] This assumption was based in part on a conversation with Prof. Edward Zeller (University of Kansas), who thought that "we will obtain good optical clarity below about 150 meters near the pole." (Bowen 2017:130).

by high energy muons, had not been directly measured," (Halzen and Learned 1988, p.2). Uncertainty also existed concerning the difficulty of ice drilling and the time required to drill each hole, as well as potential interferences with signal detection and the frequency of events. In such situations, judgments concerning the risk involved in the method is a matter of subjective beliefs made with little confidence. To quote Halzen: "it's pretty clear we had no idea what we were doing, and so this was real research, right?" Halzen goes on to say that "if we really had [known] what we were doing we would probably not have done it. And, in fact, it turns out that a lot of things we should have known turned out not to be true" (Bowen 2017:147). Situations of deep uncertainty are the norm in projects that demand new and highly-creative methods. In such situations, evaluation often involves assessing whether the team proposing the research is in the best possible position to make it work. Conversely, projects that employ standard or well-known methods do not confront this type of problem and judgments can readily be based on prior experiences or data from past research.

The third question to consider is: '*if* all works, how likely is it that the outcome will be found?', i.e. the odds that the expected primary finding happens within the observation spectrum. This can be thought of as natural risk. Note that this question applies only to observational/experimental research and not to theoretical research. It depends on the relative rarity of the phenomenon in nature, and consequently on the probability of capturing it within the spectrum of observation. We can think of this as the *natural* risk of observing a phenomenon, given the size and conditions of the observations. In our example, assuming that cosmic neutrinos form as theorized, and that the experiment is well executed and capable of detecting neutrinos, the probability that a cosmic event actually happens and is detected poses an element of uncertainty. This uncertainty is caused partly by chance and partly by nature. In many cases, past experiments in adjacent areas or data collected in prior works can provide a basis for computing frequency-based probability. In the case of the IceCube, the probability that a supermassive object collapses was modeled by an equation, in which the key parameters were taken from data collected from previous experiments,

The fourth question to consider concerns '*what else* can be found' and comprises evaluations of secondary outcomes that the project might produce, beyond the stated primary outcome. There are several common situations in research that induce secondary outcomes. First, new instruments, designed with specific goals in mind, are an especially effective source of secondary outcome (Franzoni 2009). Galileo's telescope, for example, intended for navigation, resulted in the discovery of the moons of Jupiter; the radio telescope used by

Jansky, intended to study noise that could interfere in radio transmission, ended-up detecting radio galaxies. Second, because science is an open-ended quest, it is acceptable for a scientist to be flexible and shift the goals driving a project (Nelson 1959). Flexibility can provide backup or recovery plans when a planned task is known to pose problems. For example, when the IceCube team found that bubble-free ice did not exist until a considerably greater depth than they had initially thought, they realized they had a problem. They had, "'goofed up' by placing their instrument in shallow ice." (Bowen 2017:175). Halzen reportedly began to look for "some way to make this disaster look good." (Bowen 2017:188). The answer was supernova. They realized they had "by far the most sensitive supernova detector on the planet." Their disaster "had a mission." (Bowen 2017:189). A third reason why secondary outcomes may be found is that in many scientific projects the need to solve practical and theoretical problems induces learning in ways that often branch-out from the main line of investigation. This learning inevitably suggests directions of additional research that could not have been anticipated at the time of project conception. By way of example, the drilling of ice needed to place the IceCube Neutrino detectors lead to important discoveries concerning the physical properties of deep-ice.

It should be noted that secondary outcomes may exist with or without primary outcomes. Moreover, some types of research are more likely to produce secondary outcomes than others. While it is difficult to think of secondary outcomes arising from a Cochrane review, research that involves exploration of nature and activities never performed before is more likely to have secondary outcomes. More generally, basic research provides, by definition, insurance against coming up empty handed in the sense that non-findings are, in their own way, findings. To quote Bowen, the physicist who chronicled the IceCube project, "There is no such thing as a disaster in basic research. Whatever happens, you learn from it." (Bowen 2017:163). The same cannot be said for research that is extremely path dependent, with few opportunities for secondary discoveries. Protein structure determination is a case in point; it "is either a complete success or complete failure." There are "no intermediate results to publish along the way or to fall back upon if you fail, unlike other fields." (Petsko interview July 3, 2020).

The fifth question to consider is '*how much* is the finding worth,' and concerns the importance of what is found in terms of scientific gains and societal benefits. For example, if cosmic neutrinos detect a cataclysmic astronomical event, what would be the value of this new piece of knowledge? The uncertainty in this case relates to the magnitude of the impact.

There are two important things to note regarding the impact. First, a finding that falsifies a theory may provide informative content, just as (if not more so than) a finding that complies to the theory (Popper 1959). Thus, all scenarios have a non-zero value. Second, the value of a finding in science depends on the context, in the sense that the same finding could be more-or-less valuable depending on the conditions and actions of other scientists in the same field. If the research area is crowded by many competing teams, it is possible that the same finding will be reported independently by multiple researchers and will thus be of lower value to science and society. This is not to say that additional independent reports do not add value by confirming the result, but they do not have the same value as that of the first discovery (Merton 1957; Stephan 1996). Consequently, the ex-ante appreciation of the prospective value of a discovery requires speculation concerning the actions of competing groups.[13]

So far, we have discussed the five questions that help identify risk components from a general point of view. However, if we assume the point of view of the Principle Investigator (PI) there is one additional aspect of risk to consider: *personal* consequences and/or consequences related to the PI's team or lab, germane to the fifth component discussed above. Stated differently, the fifth component concerning *how much* a project is worth can be disentangled into two sub-components: '*how much* is the finding worth for science and society' and '*how much* is the finding worth to the PI'. One or both apply, depending on who is doing the assessment.

Research outcomes have many possible consequences for a PI. Here we consider the primary ones: i) career, ii) reputation, and iii) future funding (Stephan 2012). Career consequences depend in part on the position and career stage of the PI. Tenure, for example, shelters PIs from the negative consequences of research that has disappointing results (Franzoni and Rossi-Lamastra 2017). A tenured scholar is arguably less damaged by a non-result than a tenure-track scholar. And a tenure-track scholar arguably benefits more than does a tenured one from a major success. Even among PIs in the same academic rank, the career implications of a non-result may differ. For example, the severity of a non-result for tenure-track scholars is more of a problem for those without other successful projects than for

[13] To illustrate, in 2018 the IceCube team announced that they had detected a cosmic neutrino from a blazar, laying 4-light years away from Earth (Collaboration 2018). Prior to this, the only other identified sources of cosmic neutrinos were limited to the Sun and to a supernova identified in 1987. The fact that the finding was reported only by IceCube, made it more valuable than if the finding were independently reported by other observatories.

those with past successes. Regardless of academic rank, scholars in more demanding institutions may face more severe consequences from failed research.

Outcomes affect one's reputation. Being the lead PI of a successful research project can lead to increased visibility and enhanced reputation. However, if the research results are particularly disappointing, it can also jeopardize the scholar's reputation. This reputational effect is magnified by working in crowded areas, where bibliometric indicators are more granular and thus respond more directly to small changes. Researchers also worry that a proposed idea may appear to be sufficiently on the "wild side" to tarnish their reputation. Halzen had such concerns about the idea of proposing a neutrino observatory in the ice. One reason he chose to propose the idea publicly in Poland, at a relatively small conference, was to shelter himself from such a reputational loss. (Halzen 2010).

Accomplishments (or lack thereof) also have implications for future funding, necessary in most fields and countries for today's scientists.[14] Uncertainty surrounds funding in a variety of dimensions. For example, if the project brings no results or produces results deemed insignificant, funds are unlikely to be forthcoming for further work, while, if it produces results of sufficient importance, future requests are much more likely to be funded. Halzen and colleagues were acutely aware of this and knew that "any type of failure will hamper future funding." (Bowen 2017:175). Risk associated with future funding can discourage scientists from moving in new directions. One eminent scientist who made a major change in his research agenda at age 55 reported that while in making the change he never considered the risk associated with not finding anything, the risk he did consider was "the personal cost: leaving a field in which I was a leader for one in which I was completely unknown, the difficulty in getting funding, the difficulty in getting people to work with me. And indeed, those were the major problems I faced, for years."

Consideration of the value for science and society and for the PI differ in the direction of variability of risk they imply. When we consider consequences from the point of view of science and society, the *how much* has a downward limit of zero, in the sense that it cannot lead to a direct loss. When we consider the point of view of a PI, the *how much* implies a variability that takes the full spectrum from gains to losses. That is, the personal

[14] Although small, "table-top" science exists, research is generally expensive and requires external support. IceCube is at the expensive end of the spectrum, but there are other projects, such as LIGO and the LHC that cost far more.

consequences related to research may imply a direct gain, but also a direct loss in reputation, career status, placement or future funding.

The prior discussion disentangled five components and related questions that must be addressed in a meaningful discussion of risk in science: i) *what,* ii) *how likely,* iii) *if,* iv) *what else* and v) *how much*. They can easily be expanded to include additional components arising in special cases. For example, in the relatively rare case in which research could involve immediate danger to researchers (e.g., developing new explosives), human or animal subjects (e.g., testing new surgical methods), the environment (e.g., research on nuclear power), or it could involve ethical concerns (e.g., cloning of animals), the model can be expanded by adding components regarding *potential losses*.[15] IceCube, for example, confronted this type of risk concerning the safety of the South Pole team. Indeed, one of the drillers was seriously injured while working on the project. Academe has developed provisions to address and cope with such risks. They include ethical and safety protocols and procedures like the IRB that aim at minimizing such risks.

More generally, one could also think of sub-questions, that detail specific aspects. For example, one can ask separately the value of secondary outcomes from the likelihood of finding secondary outcomes. All in all, however, the above questions capture the essential components of risk involved in the majority of scientific research projects. Table 1 summarizes the questions, along with their applicability, focus, and source of uncertainty. This conceptualization is useful for representing research uncertainty and can be used to compare projects. The comparison also reminds us that there are basic differences regarding the risks involved in different types of research. Observational studies and experimental research is subject to natural risk (the *if*), whereas theoretical research is not. In this sense, experimental research faces one more challenge that can lead the research to come-up empty handed. Although basic experimental research may involve a small likelihood of success, due to both methodological and natural uncertainty, in the sense of coming up empty handed, it is not necessarily at high risk of failure compared to applied experimental research. This is because the probability of secondary findings is more favorable than is the probability of secondary findings in applied experimental research.

[15] This can include asking the classical triplets of questions encountered in risk analysis: *what harm can happen*, *how likely is this* and, in that case, *what would be the consequences* (Kaplan and Garrick 1981)

Finally, it should be noted that the framework outlined in Table 1 also highlights that evaluations of risks differ systematically depending on who is doing the evaluation, because the consequences that accrue to science and society are different from those that accrue to the PI.

Table 1. Framework for representing the components of risk in science

Question	Applies	Focus	Source of uncertainty
What can be found?	Always	Primary outcome	Nature State of scientific knowledge
How likely is it that the proposed approach works?	Always	Method Qualifications of the team	Nature State of scientific knowledge Technology
If all works, how likely is the outcome to happen?	Experimental research	Rarity of event	Nature
What else can be found?	Always. More likely in natural explorations, activities never performed before, new scientific instruments, and basic research in general	Secondary outcome	Nature State of scientific knowledge
How much is the finding worth for science and society?	Always	Science Scientific community Societal benefit	Scientific advance in own field Size of field Competing results Scientific advance in other fields Size of other fields
How much is the finding worth for PI?	Principle Investigator	Career stage Professional environment	Personal value of findings Reputation Career status Placement Funding

5. Measuring risk

Model of Expert-based Metric of Risk. The framework set out in Section 4 outlines multiple components of risk in research. In some cases, however, describing risk in its multiple components is not sufficient and there is need for a comprehensive measure of risk. This is the case, for example, in the context of peer review evaluation in which experts are asked to assess the risk of research of proposed research. We thus need a model that

combines components to assess risk. Drawing on models of risk metrics (Budnitz et al. 1998; Haimes 2009; Kaplan and Garrick 1981), we adopt the most simplistic approach and use the question of what and what else to outline the scenarios of possible outcomes (states of the world) —both primary and secondary-- to provide an example of such a model.

The questions *how likely* and *if*, when applicable, shall assess the likelihood of each scenario, intended as subjective probabilities, and expressed in the range 0-1. The questions *how much* shall assess the impact for science and society (u_i); they could also have a separate estimate concerning the value for the scientist. The how much can be quantified with scores, e.g., in the 0-100 range. Table 2 illustrates an operational application of the framework using a simple functional form. It assumes the point of view of a granting agency and so the question *how much* considers only the value for science and society, and not the value for the scientist involved.

Table 2. Example of risk measure for expert-based assessment

Primary/secondary outcome *What / What else*	**Likelihood** *How*	*If*	Total likelihood	**Value for science and society** *How much*	**Expected value for science and society (i-th scenario)**
	Prob. (0-1)	Prob. (0-1)	Prob. (0-1)	Score (0-100)	Score (0-100)
Scenario i Primary outcome (1)	$\mathbf{p_{1i}}$	$\mathbf{q_{1i}}$	$\mathbf{P_{1i} = p_{1i}\, q_{1i}}$	$\mathbf{u_{1i}}$	$\mathbf{V_{i1} = u_{1i}\, P_{1i}}$
Scenario i Secondary outcome (2)	$\mathbf{p_{2i}}$	$\mathbf{q_{2i}}$	$\mathbf{P_{2i} = p_{1i}\, q_{2i}}$	$\mathbf{u_{2i}}$	$\mathbf{V_{2i} = u_{2i}\, P_{2i}}$
Scenario i Total				$\mathbf{U_i = u_{1i} + u_{2i}}$	$\mathbf{V_i = V_{1i} + V_{2i}}$
Scenario (i+1) Primary outcome (1)					
Scenario (i+1) Secondary outcome (2)					
Scenario (i+1) Total				$\mathbf{U_{i+1}}$	$\mathbf{V_{i+1}}$
…					

The main outcomes of the framework are pairs of the value (for science/society) and the likelihood (u_i ; P_i), for each *i-th* scenario of primary and secondary outcome, where the

likelihood P_i is the product of methodological (p_i) and natural (q_i) risk for each scenario.[16] The expected value of the *i-th* scenario (V_i) can be computed as the sum of the expected values of primary (V_{1i}) and secondary outcomes (V_{2i}), given the respective likelihoods. The number of scenarios considered (N) can vary, depending on the case, and on intended uses of the analysis. The pairs can be ordered in descending order of V_i , and eventually plotted and fitted by a risk curve, a standard representation of risk analysis (Kaplan and Garrick 1981). One can, of course, compute the average expected value of all scenarios ($V = \frac{1}{N}\sum_{i=1}^{N} V_i$), but the result would have little meaning. More informative for the purpose of science would be the range of maximum and minimum values for science/society (min U_i ; max U_i) across the scenarios and the range of related expected values (min V_i ; max V_i). High-risk high-gain projects would stand out for having at the same time a high max U_i , a large range between max U_i and min U_i, and a low max V_i . Conversely, low-risk, low-gain projects, like Cochrane reviews, would have a low max V_i, but also a low and not-so-variable max U_i. Likewise, the metric can detect projects that are High-risk (low max V_i), but Low-gain (low and not variable max U_i).

In the case of IceCube Neutrino project, for example, there was large uncertainty concerning the findings, project feasibility (the *how*) and the *if* the sought events would actually occur. The project almost certainly involved secondary outcomes. Thus, it could have led to a large number of scenarios, in which the value of a single potential scenario could be very high, and variable. By way of comparison, Cochrane reviews can lead to a maximum of two/three scenarios (e.g., absolute effect is significant, not significant; mixed/unclear), with no uncertainty in the *how*, no secondary findings and consequently a small and not very variable, value for science/society. The testing of Ipilimumab also involved a limited number of scenarios (e.g., the drug is effective; not effective; mixed), with no uncertainty in the method and no or limited secondary findings. However, the value of discovering an immunotherapy treatment for cancer was very high, given alternative treatments available at the time, and the chances of success (*if*), were small, based on prior experience. Thus, Ipilimumab was at higher risk of coming-up empty-handed (low probability, no secondary finding), although the methodology involved was not risky and the expected value, if successful, was high.

[16] Assumed here as independent from one another. In case the probabilities are not independent, the formula should be edited to include conditional probabilities.

As the example shows, the approach of risk-analysis outlined above can in principle be adopted to evaluate risks involved in research that has yet to be implemented. As such, it potentially could be deployed by funding agencies, in peer-review assessment for evaluating research proposals and could be especially useful in those granting schemas where the high-risk high-gain evaluation is a key criterion. It is important to stress that the application of this method to project evaluation would be an exercise of forecasting, as it involves asking experts to gauge actions and events that could happen in the future and are not pre-determined (Mellers et al. 2015; Tetlock and Gardner 2015). Thus, the opinions would necessarily imply a large margin of uncertainty. In other areas where experts' forecasting is essential, such as climate change and risk analysis, behavioral and decision scientists have elaborated a set of best practices known as 'expert elicitation techniques' to elicit, characterize and treat quantitative expert judgments in the form of subjective probability (Morgan 2014; Raiffa 1997). These techniques are meant to make the best possible use of the information that experts have, including eliciting more than one point of the subjective probability distribution. The approach also takes into account the potential interference of human biases, such as overconfidence and anchoring (Clemen and Winkler 1999; Fischhoff 2015; Fischhoff and Davis 2014; Morgan 2014; Winkler et al. 2019).

In practice, the application of the model we propose is rather laborious. While the approach has much to recommend it, its adoption would undoubtedly increase the burden of review boards that manage large numbers of proposals. At present, we do not know if the potential gains of using a risk metric in grant peer-review would be justified in a cost-benefit logic. This is probably an empirical question that warrants future work. At a minimum the concepts and procedures underlying expert elicitation protocols could be used to help educate panelists regarding aspects of risk.

Text-based correlates of risk. Scholars of the economics of science often wish to measure risk associated with research in large samples and often over long periods of time. The expert-based metric outlined above is clearly too time and resource-consuming in such instances. For these purposes, methods have been developed that can be computed for large samples. These methods do not measure risk directly, but instead infer it from correlates associated with risk that can be obtained from the codified products of research, such as texts of research proposals and published papers.

Many text-based correlates of risk rely on the assumption that research that deviates considerably from past research can lead to breakthroughs but holds risk in that it has yet to be explored by others and can fail. One way to measure the extent to which research deviates from past research is to think of past knowledge as a set of building blocks which can be combined to produce new knowledge. Research that combines past knowledge in well understood ways is referred to as exploitative; research that combines past knowledge in new ways is referred to as explorative. Explorative research arguably is more likely to lead to breakthroughs than exploitative research (Romer 1994; Varian 2009) and to carry higher potential impact (March 1991; Simonton 2003).

Wang, Veugelers and Stephan (2017) draw on this approach to measure *novelty*, the extent to which a published paper draws on references that have not been jointly referred to in previous research. To be more specific, they retrieve references for each paper published in 2001, and construct co-cited journal pairs for each. They then check to see if the pair has previously been made in the last 20 years, identifying pairs (if any) that are novel—that is not previously made and compute the "difficulty" of making each novel pair by examining whether it has "common friends" in the sense of past journal co-citations. The novel score for the paper is the sum of the novel score for each reference combination. The authors find that 89% of the articles contain no novel combinations. Among the 11% that do make a novel combination, they distinguish between highly novel, being in the top 1%, and moderately novel. Importantly, for our perspective, they find that, compared to non-novel papers, the citation distribution associated with highly novel papers has a higher variance and higher mean value, characteristics that we expect in risky research, suggesting that the novel measure correlates with risk.

Uzzi and coauthors (2013) adopt an alternative measure of knowledge recombination, called *atypicality*. For each pair of references found in a paper, they compute a z-score, comparing the probability of making the combination to making the combination by chance. Z-scores for reference pairings greater than zero indicate that the pairing occurred more likely than by chance; those below zero indicate less likely than by chance. For each paper, they then take the lowest 10^{th} percentile z-score of its series of z-scores as an indication of how novel the paper is and the median z-score as an indication of the paper's "conventionality." The authors find that the "vast majority of papers displays a high propensity for conventionality." (Uzzi et al. 2013: 469).

Responses of corresponding authors to the GlobeSci survey (Franzoni, Scellato, and Stephan 2012) concerning a paper published in 2009 are consistent with these findings. Authors were asked to report "with regard to the area of research, this paper is: in a high-risk (of failure), high-reward (if successful) area of research."[17] The variance of citations to papers authors reported as high-risk was greater than the variance of citations to papers reported as low risk as was the mean number of citations. Moreover, the authors' assessment of risk correlated with Uzzi's measure of atypicality (Franzoni, Scellato, and Stephan 2018).[18]

Other approaches to measuring risk rely on words or findings reported in the text to measure the extent to which the current research deviates from past research. Foster et al. (2015), for example, examine the extent to which biochemists introduce novel chemicals and chemical relationships, using abstracts from publications indexed in Medline. The authors distinguish between three types of papers based on the chemical relationships described in the work. Research that makes a *jump* explores previously unexplored chemical relationships - jumping beyond current knowledge. Such research arguably is more likely to fail but, if the research succeeds, it is more likely to make a breakthrough. Research that explores relationships between previously studied entities, and is thus more likely to succeed, is subdivided into research that tests a *new* relationship, not published before, or research that *repeats* an analysis of a previously studied relationship. Again, and important for our perspective, the authors find that the citation distribution associated with *jump* papers and *new* papers has a higher variance and higher mean value than that of *repeat* papers, characteristics that we expect in risky research, suggesting their measures correlate with risk. The authors also find that papers based on *repeat* strategies were six times more likely to be published than that those that used *new* or *jump* strategies during the period 1983-2008 (Foster et al. 2015: 886). Krieger, Li and Papanikolaou (2019) develop a related measure of novelty which compares the chemical structure of new drugs going up for FDA approval to that of previous drug candidates.

---

[17] Approximately 8% said they strongly agreed that it was high risk, high gain. Approximately the same percent said they strongly disagreed; one third neither agreed or disagreed, while almost a quarter said they agreed—but not strongly-- with the statement and a bit more disagreed with it.

[18] The atypicality analysis was only available for a limited subset of respondents. See Franzoni, Scellato and Stephan (2018).

Another measure used as a proxy for the degree to which research deviates from past work is the average age of keywords associated with a publication, measured by the year a keywork first appeared in the scientific literature. The approach works best if applied within a single domain. For example Azoulay et al. (2011) use this method for Medical Subject Headings (MeSH). They also compute a Herfindal index of keyword diversity. Boudreau and coauthors (2014) use the percent of keywords never used before to calculate the novelty of a proposal. A different approach, based on content analysis, measures the degree to which a document employs words that suggest vagueness, probability or unsureness (e.g., uncertainty, weak modals and negative words), as opposed to certainty (Sauermann et al. 2019). Word lists and algorithms exist that can readily provide this metric e.g., for the English language (Loughran and McDonald 2013).

Open issues on risk measures It is important to note that limitations and pitfalls exist in the use of these text-based metrics of risk. First, there is the file drawer problem. A large share of failed research remains unpublished (Franco, Malhotra, and Simonovits 2014). Moreover, published articles are formulated to provide an impression of order and fulfilment of the expectations which conceals the trial-and-errors involved (Bourdieu 2001). In this respect, looking at research proposals—which are upstream—addresses part of this concern, although even research proposals do not include intentions and/or ideas that are censored after early-testing. One possible way to reduce the file drawer problem is to look at log-files and lab-notes that scholars deposit in digital repositories when research is in progress or being proposed (Franzoni and Sauermann 2014; Sauermann, Franzoni, and Shafi 2019). Pre-registration of experiments can serve the same purpose.[19]

A second problem relates to measures based on knowledge recombination, like *novelty* and *atypicality*. By construction, they see new research which deviates considerably from past research as riskier. While such measures capture risky research, they do not identify all risky research, such as the risk associated with continuing a line of research that has proven fruitless for some time. For example, they would not have detected risk in the case of the first immunotherapy cancer treatment made by Allison, described earlier in this essay. Third, many of the available metrics rely on citations, which have well-known drawbacks and are but one measure of the extent to which research is successful. Furthermore, some of the riskiest research may not have a measurable effect for many years.

[19] E.g. https://www.cos.io/initiatives/prereg. Accessed November 23, 2020.

Last, but certainly not least, these measures provide a single assessment of risk that relates to a full project or completed research, and do not discern risk components within projects, the importance of which we have stressed in Section 3. Developing complementary measures or metrics that rely on alternative assumptions could thus prove valuable. It is possible that textual analysis and other AI-based metrics will offer new opportunities in this direction.

6. The Relationship of Rewards to Risk Taking; Extent to Which Funding Agencies are Risk Averse

This paper began with the concern that risk taking in science is on the decline. Many put the onus on funding agencies and changes in the way in which science is funded and grants awarded. Others stress that the rewards to doing science not only discourage risk-taking, but increasingly do so. In this section we examine the extent to which the rewards to science discourage risk taking and whether there is evidence to suggest that these trends, to the extent they exist, are increasing. We also examine the extent to which funding agencies are risk averse.

Reputation plays an important role in science although by no means is it the only reward to doing science (Merton, Stephan and Levin).[20] But reputation is key. It affects hiring and promotion opportunities, as well as funding decisions and plays a major role in the acquisition of the position and resources necessary to engage in a research. Measures of reputation in recent years rely increasingly on bibiometrics, where citation counts and their derivatives (e.g., the H-Index, the Journal Impact Factor), play a prominent role.

The heavy emphasis on reputation, particularly when measured with bibliometric indicators, arguably discourages risk taking on the part of scientists. This is not obvious; examples of scientists who have taken a risky course receiving a Nobel Prize, for example, are readily available [21] and there is research suggesting that prestigious prizes can encourage risk taking (Rzhetsky et al. 2015). But overall, the citation-premium for doing risky-research,

[20] Rewards also include the satisfaction derived from puzzle solving, and financial gain that often accompanies a successful research career (Stephan 2012; Stephan and Levin 1992). Cohen, Sauermann and Stephan (2020) also show that scientists are strongly motivated by an interest in contributing to society.
[21] Jim Allison is but one case in point.

compared to that of doing not-so-risky research, is arguably insufficient to encourage risk taking. To illustrate, Foster and colleagues (2015), using the method described earlier, find that "jump" papers reporting highly innovative chemical combinations receive 52% more citations on average than "repeat" papers reporting known combinations, while "new" papers reporting moderately-innovative combinations enjoy 30% more citations than those reporting known combinations. Their results suggest that taking the risk associated with "jump" and "new" research makes it more likely to achieve high impact, but the additional rewards are small and arguably do not compensate for the possibly of failing. Stephan (2019) has called this the "quad" effect, referring to the fact that competitive female figure skaters attempt fewer quadruple jumps, arguably because the incremental score they can earn for completing a quad, compared to successfully completing a triple jump, is insufficient to compensate for the risk of failing to complete the quad jump. For male figure skaters, scoring is different: the incremental score is larger and provides sufficient incentive to attempt the quad. The work of Uzzi et al. (2013) is consistent with the findings of Foster et al. (2015), and shows that "The highest-impact science is primarily grounded in exceptional conventional combinations of prior work yet simultaneously features an intrusion of unusual combinations." Stated differently, a little bit of risk adds spice to the research; but conventionality is the dominant characteristic of highly cited papers.[22]

Wang, Veugelers and Stephan (2019) find that in the short run highly novel papers are less likely to be top-cited (1%) than moderately novel or non-novel papers but over time highly novel papers are significantly more likely to be top-cited. The authors also find that highly novel papers are less likely to be published in High Impact Journals. Causality of these findings cannot, of course, be determined but the results are consistent with the idea that in the short run the rewards to science are biased against risk-taking. This can discourage risky research, given that universities make crucial career decisions, such as hiring, third-year review and tenure, using evaluations based on relatively short-time windows. While the tenure-track system has long been common in the US and Canada, it has only recently been introduced in several continental European countries, such as Germany, Sweden and Italy, suggesting that the rewards for risk-taking have shrunk globally.

[22] Papers characterized as having high medium conventionality coupled with a high tail "novelty" have a hit rate in the top 5 percent 9.2 times out of 100.

The role that reputation plays in discouraging risk taking arguably is growing not only because of an extension of the tenure-track system to an increasing number of countries but, and related, because citation counts are becoming increasingly important and readily available. Sixty-five years ago, there was no ready way to measure citations to published work. As late as the early 1990s the only way to count citations was to laboriously look in the volumes published by the Institute of Scientific Information. Today citation counts are readily available and arguably affect the direction of the PI's research. Consider, for example, what the computational chemist Richard Catlow said upon becoming Foreign Secretary of the Royal Society: "I was lucky. When I began my scientific career in the 1970s, I had no real sense of how my work was cited. […] If I had been citation-driven, I might have abandoned a field that is now central […]. By the 1990s, when citation data became prominent, I was already a full professor." [23] The work of Foster et al. (2015) is consistent with Catlow's view, finding that papers which focus on already established relationships have been growing over time, consistent with the idea that high risk in research is on the decline.

Although we lack systematic studies of the relationship of risk taking to funding success, the work that does exist suggests that reviewers and panels are risk averse. In an experiment conducted at the Harvard Medical School, Boudreau and coauthors (2014), for example, find that more novel research proposals, as measured by the percent of keywords not previously used, receive more negative evaluations during peer-review. The result is driven by proposals with particularly high levels of novelty. Their preferred explanation for this finding rests on bounded rationality of reviewers. To quote the authors: "experts extrapolating beyond the knowledge frontier to comprehend novel proposals are prone to systematic errors, misconstruing novel work. This implies that, rather than receiving unbiased assessments (with zero mean errors), novel proposals are discounted relative to their true merit, quality and potential." (Boudreau et al. 2014: 2779). Veugelers and coauthors (2019) find that applicants to the ERC Starting Grant program, with a history of highly novel publications, are significantly less likely to receive funding than those without such a history. The major effect comes during stage one, when panel members screen a large number of applicants based entirely on a five-page summary of the proposed research and a CV. The finding suggests that reviewers rely on bibliometrics, which, as we have seen, are biased

[23] https://media.nature.com/original/magazine-assets/d41586-017-08289-z/d41586-017-08289-z.pdf. Accessed August 6, 2020.

against risk taking in the short run, in making decisions. Lanöe (2019), using a measure of novelty, finds evidence that funding decisions made by French National Research Agency are biased against risk-taking. Wagner and Alexander (2013) evaluate the SGER NSF program designed to support high risk, high reward research that ran from 1990 to 2006. Funding decisions were made entirely by program officers with no external review. The authors find that program officers routinely used but a small percent of available funds. The findings suggest that either officers were averse to funding risky research, despite the number of funded proposals that had transformative results or, that risk taking was not rewarded within NSF. Conversely, Sauermann, Franzoni and Shafi (2019), using speculative words as a measure of risk, find no evidence to suggest that funding made by citizens who pledge money on the research crowdfunding platform Experiment.com either favor or disfavor risky research.

7. Hedging Risk

Financial instruments such as futures are designed to hedge risk in markets where volatility exists, as does the ability to assemble a portfolio of assets with various degrees of risks. Insurance reduces the size of losses in cases of a negative event. A question of interest is the extent to which strategies exist in science for hedging risk, either at the individual or institutional level. If so, understanding such risk hedging strategies has the potential to help policy makers and administrators advance risk taking in science.

Individual level At the individual level, a strategy for hedging risk that is not publicly advertised, but sufficiently common to inspire a cartoon,[24] is to pursue the early stage of research with funding drawn from an earlier grant. In some instances, research begun on the back-burner may eventually be proposed for funding, having been "de-risked" by the back-burner treatment. Such strategies are only available to scientists who are sufficiently senior to have established a funding pipeline. Scientists also attempt to manage risk by seeking funding with a sufficient time horizon to allow them to recoup from possible failures that may occur along the way.[25] Scientists also hedge their bets by pursuing research with potential secondary objectives, as we have seen in the case of IceCube.

---

[24] http://phdcomics.com/comics/archive.php?comicid=1431. Accessed October 27, 2020.
[25] HHMI funding is highly prized not only for its status and size but also for the fact that the funding is for 7 years.

Another important strategy for hedging risk at the individual level consists in outsourcing parts of the research process to others in exchange for payment. For example, ten years ago it could take a postdoctoral fellow a year to try to create a transgenic mouse in a PI's lab. At the end of the year there was a reasonable possibility that the postdoctoral fellow would come up empty-handed and the lab would be "mouseless." The availability of outsourcing the creation of such a mouse to a company has de-risked the activity. While the company may experience some failure along the way, by preserving steps of the process it is able to guarantee a mouse in a reasonable period and, if it fails completely, the PI will get a refund.[26] Similar strategies of risk transfer via outsourcing are possible in several other domains including experimental psychology.

University level At the institutional level, hiring policies help universities insure against investing in scientists who are not productive. "Soft money" positions, for example, come with no salary guarantee, but instead salary is funded (or almost fully funded) from grants for which the researcher is responsible. Such arrangements put faculty under considerable pressure to produce results. The bio-physicist Stephen Quake called this situation "funding or famine."[27] If their research is not deemed fundable or comes up empty handed, the university can cut its losses and hire another individual into the position. It is notable that soft money positions have been on the rise in recent years. In the US, for example, the majority of basic medical faculty are hired in soft money positions and are responsible for bringing in most of their own salary (Stephan 2012). Soft-money positions also are common outside of medical institutions. Stephan documents that during the years when the NIH budget doubled, the majority of new hires were made into soft money positions (Stephan 2007). Soft money positions not only transfer risk to the faculty; they also discourage risk taking on the part of the faculty given the importance of continued funding. A second way universities insure against risk is hiring faculty into "tenure-track" positions and/or implementing third-year review during the tenure-track period. Such practices mean that the university can cut its losses if it views performance to be inadequate. Finally,

---

[26] Cyagen, for example, which at the date of this writing had delivered over 50,000 animal models, states that "we will fully refund the client's service fee if animals with the specified genotype are not generated (except for genetic modifications severely affecting viability, morbidity, or fertility.) https://www.cyagen.com/us/en/service/transgenic-mice.html?gclid=EAIaIQobChMI5aCYqo_46gIVDvDACh2PpgnmEAAYAiAAEgK7VfD_BwE. Accessed August 6, 2020.

[27] https://opinionator.blogs.nytimes.com/2009/02/10/guest-column-letting-scientists-off-the-leash/

universities also recruit faculty with strong funding streams rather than hire individuals without funds[28], again a strategy for minimizing risk.

Universalities and institutions also insure against the risk that research could result in serious harm or damage by means of policies that mandate pre-approval for certain types of potentially harmful research. For example, most universities require pre-approval by the Institutional Review Board (IRB) for all research that involve human subjects. Moreover, they often engage external organizations to train and certify faculty in compliance regarding the conduct of research, thereby minimizing their own risk in so doing.

Granting institution level Funding institutions which award money to support research have several ways of managing risk. Perhaps the most straightforward is to favor proposals that include preliminary findings, thereby insuring that the research is feasible.[29] Another is to adopt a portfolio approach. There are in principle several ways to do so. One widely used in practice is to have separate calls with separate budgets for risky and 'regular' research, limiting the budget of the former. In principle, the approach assumes the ability of panel members to assess risk involved in a proposal and select accordingly. In practice, panel committees are given little guidance on how to discern risk and even less on how to select accordingly.

A third strategy granting institutions use to hedge risk is to fund in stages. In the first stage, a pool of projects is selected for initial short-term funding. In the second stage, an interim evaluation is conducted. Funds for less promising projects are curtailed; funds for more promising ones are continued. Although this strategy is used by DARPA,[30] it is infrequently used among most funding agencies. A stage-funding approach is, by way of contrast, common in the Venture Capital industry for funding entrepreneurial projects, where it is called the 'spray and pray' strategy (Lerner and Nanda 2020). Although the two domains are fundamentally different, the parallel provides some interesting insights. The interim evaluation is especially useful when the initial estimate is unreliable, but can be quickly updated with initial funding (Vilkkumaa et al. 2015). It is thus especially suitable for research

[28] During the NIH doubling, universities recruited senior faculty with more than one grant. Post-doctoral fellows in the US are reportedly more likely to get a faculty position if they have secured a K-99 grant from NIH. More generally, getting ERC funding as a starter is seen as a path to obtaining faculty status.
[29] When the National Institute of General Medical Sciences at NIH was funding protein structure projects, the mantra was "no crystal, no grant," code for the requirement of preliminary results
[30] https://fas.org/sgp/crs/natsec/R45088.pdf Accessed August 30, 2020.

that can make substantial steps forwards in a relatively short period of time and does not require large fixed-costs to be started (Ewens, Nanda, and Rhodes-Kropf 2018).[31] Some research fields meet these conditions, but the conditions are more the exception than the norm in the natural sciences, where the share of research that requires expensive equipment and substantial effort is large (Stephan 2010).

A fourth strategy granting institutions use to hedge risk relies on portfolio diversification. It applies specifically to challenge grants, that is grants designed to address a specific challenge or set of pre-defined goals, such as the full sequencing of the human genome, or the development of a vaccine for Covid-19. In these special cases, the desired outcomes are known, but the way to achieve them is unknown and/or there is uncertainty concerning which approach is more likely to be successful, or more efficient or quicker. Risk in this case can be managed by funding a pool of projects that take diverse routes to reach the same outcome and thus the risks are non-correlated. By way of example, Operation Warp Speed, launched in the spring 2020 by the US government to advance a COVID-19 vaccine, aimed at having at least one approved vaccine available by the end of 2020, a record-time given that normal vaccine development takes about 10 years and has a success rate of 6% (Mullard 2020). The approach of the US administration was to select a certain number of candidates to maximize the odds that at least one would make it to the finish line. In May 2020, the task force identified 14 vaccine candidates out of more than 100 that existed at a pre-clinical stage. Of the 14 candidates, 8 were projected to reach early stage small clinical trials, 3-to-5 were projected to reach large scale clinical trials.[32] The selection also aimed at vaccine candidates based on different technologies and vaccinal strategies in an attempt to accelerate the time frame and increase the odds of having a successful vaccine in the near future. By way of example, the AZD1222 developed by University of Oxford and AstraZeneca, is an adenovirus-based vaccine. Its strategy is to train the immune system to recognize the spike protein typical of the SARS-CoV-2 surface, by carrying DNA for the spike antigen in host cells though a vector. In the case of AZD1222 the vector is a genetically-engineered adenovirus of chimps. The Moderna's mRNA-1273 aims at the same training strategy, but with a nucleotide-based vaccine. In this case, a synthetic lipid

[31] In the VC industry, this has largely favored IT and digital companies.
[32] https://www.hhs.gov/about/news/2020/05/15/trump-administration-announces-framework-and-leadership-for-operation-warp-speed.html. Accessed November 2020.

nanoparticle is engineered to carry mRNA templates. Each approach may have drawbacks: for example, some patients have pre-existing immunity to adenovirus and the antigens encoded by mRNA may not confer sufficient protection against pathogens. Indeed, no vaccine based on either adenovirus or nucleotides had ever been approved in the USA or Europe (Mullard 2020) at the time of this writing. A different and possibly more dependable strategy is the protein subunit approach taken by Sanofi-GlaxoSmithKline. In this case the candidate is the spike antigen itself, combined with an immunogenic adjuvant, to trigger an immune response.

A question of considerable interest is whether a portfolio diversification strategy that resembles financial portfolio management can be used to manage risk in science. This would consist of choosing a mix of research projects with levels of expected value and volatility that lead to a desired average future value with a desired risk-exposure. As before, the feasibility of this approach depends on being able to forecast the outcomes of research projects with reasonable accuracy. Moreover, and to complicate things, in this case one would need not only to assess projects in isolation, but to also estimate the covariance among the outcomes of different projects. A 'bare bones' approach is to ask reviewers to classify proposals by levels of risk involved, ask if any two pairs of projects have prospective outcomes which are correlated,[33] then fund projects in each level of risk, keeping the correlation below a maximum threshold. It is possible that future experimental studies on protocols for risk assessment in peer review will test this approach.

8. Encouraging Risk

Seed-funding Universities, for example, can promote risky research with the potential of high payoff by providing seed funding to faculty. The California Institute of Technology, by way of example, had such a program whereby faculty could submit a short proposal to the Vice Provost for Research and get a decision in a matter of days. Funds ranged from $25,000 to $250,000 a year for a period of two years. The idea was to give faculty the wherewithal to engage in early-stage risky research that, given the risk aversion of granting agencies, was deemed not yet ready for submission. If the initial findings looked promising, and produced enough preliminary data, the faculty would then submit a full grant

[33] Prospective outcomes can be correlated, for example, if projects follow a similar approach and methodology.

proposal. Universities also provide "bridge funding" to keep labs afloat between grants or if there is a lapse in funding. For example, soon after IceCube was funded, George W. Bush became president and declared "no-new-starts" which restricted NSF from funding any new capital projects in the coming year. The University of Wisconsin responded by loaning IceCube $4.5 million to keep it afloat (Bowen 2017:288).

Block-funding Several scholars have stressed the role of block-funding—the practice of assigning resources to scientists with no strings attached and without the need to commit to a project—as a way to encourage risk taking. (Heinze 2017; Laudel 2017). Block funds can be provided by universities, the employing institution or the national research systems to all staff, irrespective of achievement, or can require some minimum level of research activity to be funded. The practice is not without its critics, who point out that block funding allows scientists to sit on resources and does not hold scientists accountable in case of misuse, thus requiring more monitoring than does competitive funding. However, such a system arguably encourages longer-term research trajectories and shelters scientists from the negative consequences of early failure, in contrast to competitive research-funding. Wang, Lee and Walsh (2018) compare the novelty of research funded under the two systems in Japan. They find that research performed under block funding was more novel than that performed under competitive funding for low-status investigators (e.g., junior and female investigators), but the reverse was true for high-status scholars, where competitive funding was associated with more novelty. The latter finding is consistent with work of Veugelers et al. (2019), which finds that junior applicants to the ERC are penalized for having a history of novel research, but senior applicants are not penalized.

Grants for a Longer Duration of Time The Howard Hughes Medical Institute (HHMI), as noted above, funds successful applicants for seven years, rather than for three to five years, as is common for most other funding organizations. Furthermore, it does not demand early results nor does it penalize researchers for early failure. Azoulay et al (2011) compare the research output of HHMI investigators to a group of similarly accomplished NIH investigators using propensity scoring. They find that HHMI investigators use more novel keywords and produce more hits and more flops, compared to the NIH investigators. Although it is not clear whether the results depend upon the longer duration of grants and the

practice of HHMI to not demand early results nor penalize researchers for early failure or other variables,[34] the results suggest that these practices encourage risk taking.

***Special High-Risk High-Gain funding* initiatives** Competitive funding can also be directed specifically to High Risk High Gain science in ways that encourage risk-taking. We have noted before that some institutions have research programs especially targeted to risky science. Examples include the grants of the *ERC*, the *Director's Awards for High-Risk, High-Reward Research program* of NIH, the, *IDEAS Factory* of the Engineering and Physical Sciences Research Council of the UK and the Early-Concept Grants for Exploratory Research (EAGER) program of NSF that replaced SGER grants.

Pro-Risk Peer Review Design Targeting funds to support high risk high gain science assumes that high-risk high-gain research can be identified and supported by reviewers. The lack of understanding and suitable methods to assess risk, however, has been a critical obstacle to the implementation of such a strategy. As an officer of a private foundation said when speaking about the topic: "One of the challenges is we don't have any measures of riskiness. We don't know when a project is high risk, high reward." (Michelson 2020:142). Furthermore, many practices used in traditional peer review may discourage funding risky research rather than promote the funding of risky research.[35] One such practice is that of requiring strong panel consensus to grant. Assuming that risky research is more uncertain and sparks more disagreement than non-risky research (Linton 2016), the requirement of consensus may work against risky endeavors.

At present we have scant knowledge concerning which practices of peer review design encourage or discourage the selection of risky projects. Amid a lack of scientific understanding, some granting institutions are using creative ways to design peer-review assessment that encourage risk taking. For example, the Audacious Project hosted by the TED organization places emphasis on goals and type of reviewers involved. They choose to focus on research with strong potential global impact, and to employ reviewers who are "not typical experts in the relevant fields" (Price 2019:317). Other organizations place emphasis on rules of deliberation that do not stress consensus. One example is the provision of a

[34] Despite the authors' efforts to match the HHMI sample with comparable NIH investigators, selection is still a concern.
[35] For a humorous account, see Petsko, 2012. https://genomebiology.biomedcentral.com/articles/10.1186/gb-2012-13-5-155. Accessed August 6, 2020.

golden-ticket to each panel member that provides immediate selection, regardless of the opinion of other panelists (Sinkjaer 2018). This system has been used by the Villum Foundation; a similar system is used by the Melinda and Bill Gates Foundation. HHMI takes a different perspective and selects people, instead of projects, encouraging researchers to ask "tough questions in science, even at the risk of failure."[36] Leslie Voshall, a highly productive mosquito researcher at Rockefeller University, for example, is on record saying that her application to HHMI, which was funded by the institute, involved doing something "bold and new" and was supported with no preliminary data.[37] The Chan Zuckerberg Biohub program follows a somewhat similar philosophy, awarding fellowships to researchers for projects that are based on "bold ideas that lack preliminary evidence" (Maxmen 2017). The Open Philanthropy Project takes a somewhat similar approach, supporting projects that have "high odds of failure" and often have been turned down by other organizations, such as NIH, on the grounds that they are too risky.[38]

Although these approaches are interesting, at the present time we lack both a scientific understanding of peer review design and empirical evidence concerning the supposed efficacy of the various approaches. Future research is certainly needed in this area.

9. Conclusions

We began this paper by asserting that a scholarly understanding of risk taking in research was underdeveloped, yet critical given the key role that risk plays in advancing the knowledge frontier. We set out to address this void by reviewing insights offered from other fields that study risk. These contributions, combined with knowledge gained from studies of science, led us to propose a conceptual model of risk in science that we hope can frame and accelerate future research on risk and inform the related policy debate. The model we developed disentangles different components that determine risk. It can also be used to operationalize an expert-based risk metric, potentially useful in future peer reviewers' evaluations. We also reviewed various text-based metrics of risk currently employed in statistical analyses of large samples of research. Most of the studies that employ these metrics

[36] https://www.hhmi.org/programs/biomedical-research/investigator-program. Accessed August 6, 2020.
[37] https://podcasts.apple.com/us/podcast/the-inner-scientist/id1419667345. Accessed October 27, 2020.
[38] https://www.nature.com/articles/d41586-017-08795-0. Accessed October 27, 2020.

suggest that the current reward structure of science discourages risk taking. There is also evidence to suggest that this trend towards "play-it-safe science' is increasing. This led to a review and discussion of strategies for hedging and for encouraging risk taking.

Although this paper advances our understanding of risk-taking in science, there is much that we do not know. The research agenda going forward is challenging but rich. We encourage others to take up this important subject. Key issues to be addressed include, but are not limited to, the following.

First, there is a need to develop alternative measures of risk in science. We have proposed a new metric, which awaits testing. The pluses are that it elicits different components of risk. But the cost in terms of implementation could be quite high. Text-based correlates of risk are much easier to implement, but, as we have outlined in Section 5, come with serious pitfalls and limitations. Some of these can and should be addressed in future research. Second, evidence suggests that selection by panels, which takes place in many granting agencies, is biased against risky research, even in cases where risky research is a stated priority. Although several funding organizations have started to experiment with creative ways to select risky research, the reality is that none of these approaches are based on a scientific understanding of which designs lead to more accurate assessments and are more appropriate when supporting high-risk science is a priority. Moving forward, there is a strong need for theorical and empirical investigations to build a scientific understanding of peer review design. We hope that the notion of risk in science that we have offered will pave the way to more scientific inquiries. Some of the questions that need investigation include the following. Which experts or pool of experts are more accurate in assessing risky research? Do experts eschew risk individually? Do they do so when they meet to discuss proposals in consensus meetings? Do bolder peer review approaches lead to more polarized views, sparking more disagreement? Can alternative deliberation rules, such as golden tickets, promote risk? More generally, peer review is an understudied subject that invites research into whether the way in which it is currently organized, and consensus formed works against risky science.

Third, is the need to consider how workload associated with large funding initiatives affects panel decisions regarding risk. Likewise, there is the question of whether small scale promotes risk taking. Is HHMI's apparent willingness to fund researchers who take risk related to the small number of awards it makes and therefore its ability to look more closely

at the applicant and their research? If so, could such a model be altered to situations where scale is an issue? One possibility, for example, might be to allot a certain share of the budget to risky proposals, then ask reviewers to identify risky proposals and choose among them, using some random method such as a lottery, rather than try to score each individually on risk. A related question is what should be an appropriate target budget for risk? 10%, 25%, 50%? Clearly there is no one answer to this question. The answer depends upon the funding agency and the source of its funds. But in almost all cases the answer is greater than zero and among public funders—whose mission is to fund research the private sector eschews-- it is considerably larger.

Fourth, risk and rewards are correlated in financial markets because there is a market equilibrium that works through asset price adjustments. To what extent does science operate in a parallel universe, inducing individuals to enter research areas where the reward—but also the failure rate—is high? More generally, the issue of risk taking in science seems ripe for study by financial economists. A promising and understudied line of inquiry is whether portfolio approaches similar to those adopted in finance can be adopted by granting agencies.

Fifth, there are virtually no studies of the attitudes of scientists towards risk, even though numerous scholars study risk in other groups and conditions. It would be interesting to study the risk attitude of scientists in different subfields and see how these relate to the pace at which subfields advance. A related question is whether PhD training and early experience of scientists exert a long-term effect on attitudes towards risk.

Sixth, career conditions and the incentive systems appear to affect scientists' willingness to take on a risky research agenda. But much more work needs to be done in this area. Would less emphasis on bibliometrics in hiring, promotion and funding decisions have the expected effect of encouraging risk taking? Would fewer soft money positions encourage risk-taking? Would publishing failed research enhance risk-taking?

Seventh, we need to know more about how risk correlates across projects. We lack knowledge, for example, concerning how outcomes of research projects that share similar research approaches are correlated. For example, if two projects on protein structure determination use the same methodology, is the success/failure rate of projects not correlated? Or, if a group of Alzheimer's studies focus on the same root cause, are findings not likely to either miss the mark or coalescence on a treatment? This question also arises at the more macro level. By way of example, mice constitute 90% of all animal models used in

medical research (Stephan 2012). Yet only one out of nine drugs that work on animals work on people. Why? Is the mouse a poor model? Have researchers, in an effort to standardize experiments, controlled mouse environments to such an extent that the mouse is no longer a useful model? Can the use of other animal models improve the success rates of drugs? The question raises the possibility that while the progress of science is hindered by being risk averse at the micro level, it is harmed by assuming too much risk at the macro level in terms of highly correlated research models. This is an area that is ripe for research.

To conclude, we set out to provide a conceptual foundation to inform and advance discourse concerning risk in science. As this discussion of key issues shows, a rich agenda for future studies, which is both intellectually challenging and critical for the future of science, awaits. We call on scholars of science and adjacent disciplines to join in this effort.

Bibliography

Althaus, Catherine E. 2005. "A Disciplinary Perspective on the Epistemological Status of Risk." *Risk Analysis* 25(3):567–88.

Arora, Ashish, Sharon Belenzon, and Andrea Patacconi. 2018. "The Decline of Science in Corporate R&D." *Strategic Management Journal* 39:3–32.

Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources for Invention." Pp. 609–26 in *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press.

Aven, Terje. 2011. "On Some Recent Definitions and Analysis Frameworks for Risk, Vulnerability, and Resilience." *Risk Analysis* 31(4):515–22.

Aven, Terje. 2012. "The Risk Concept-Historical and Recent Development Trends." *Reliability Engineering and System Safety* 99(0951):33–44.

Aven, Terje and Ortwin Renn. 2009. "On Risk Defined as an Event Where the Outcome Is Uncertain." *Journal of Risk Research* 12(1):1–11.

Azoulay, Pierre, J. Graff Zivin, and Gustavo Manso. 2011. "Incentives and Creativity: Evidence from the Howard Hughes Medical Investigator Program." *The Rand Journal of Economics* 42(3):527–54.

Berger, Loïc and Valentina Bosetti. 2020. “Are Policymakers Ambiguity Averse?” *The Economic Journal* 130(626):331–55.

Berger, Loïc, Johannes Emmerling, and Massimo Tavoni. 2016. “Managing Catastrophic Climate Risks Under Model Uncertainty Aversion.” *Management Science* 63(3):587–900.

Bollerslev, Tim. 1986. “Generalized Autoregressive Conditional Heteroskedasticity.” *Journal of Econometrics* 31(3):307–27.

Boudreau, Kevin, Eva Guinan, Karim R. Lakhani, and Christoph Riedl. 2014. “Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance and Resource Allocation in Science.” *Management Science* 62(10):2765–2783.

Bourdieu, Pierre. 2001. *Science de La Science et Réflexivité*. Ital. vers. Paris: Raisons D’Agir.

Bowen, Mark. 2017. *The Telescope in the Ice. Inventing a New Astronomy at the South Pole*. New York: San Martin Press.

Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2015. “Do Firms Underinvest in Long-Term Research ? Evidence from Cancer Clinical Trials.” *American Economic Review* 105(7):2044–85.

Budnitz, Robert J., George Apostolakis, David M. Boore, Lloyd S. Cluff, Kevin J. Coppersmith, C. Allin Cornell, and Peter A. Morris. 1998. “Use of Technical Expert Panels: Applications to Probabilistic Seismic Hazard Analysis.” *Risk Analysis* 18(4):463–69.

Camerer, Colin and Dan Lovallo. 1999. “Overconfidence and Excess Entry: An Experimental Approach.” 89(1):306–18.

Camerer, Colin and Martin Weber. 1992. “Recent Developments in Modeling Preferences: Uncertainty and Ambiguity.” *Journal of Risk and Uncertainty* 5:325–70.

Clemen, Robert T. and Robert L. Winkler. 1999. “Combining Probability Distributions From Experts in Risk Analysis.” *Risk Analysis* 19(2):187–203.

Cohen, Wesley M., Henry Sauermann, and Paula Stephan. 2020. “Not in the Job Description: The Commercial Activities of Academic Scientists and Engineers.” *Management*

Science 66(9):4108–4117.

Collaboration, IceCube. 2018. “Neutrino Emission from the Direction of the Blazar TXS 0506+056 Prior to the IceCube-170922A Alert.” *Science* 361(6398):147–51.

Dobosz, Paula and Tomasz Dzieciątkowski. 2019. “The Intriguing History of Cancer Immunotherapy.” *Fronteer Immunology* 10:2965.

Edwards, Aled M., Ruth Isserlin, Gary D. Bader, Stephen V. Frye, Timothy M. Willson, and Frank H. Yu. 2011. “Too Many Roads Not Taken.” *Nature* 470(7333):163–65.

Ehrlich, Isaac and Gary S. Becker. 1972. “Market Insurance , Self-Insurance, and Self-Protection.” *Journal of Political Economy* 80(4):623–48.

Ellsberg, Daniel. 1961. “Risk, Ambiguity, and the Savage Axions.” *Quarterly Journal of Economics* 75:643–69.

Engle, Robert F. 1982. “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of U.K. Inflation.” *Econometrica* 50(4):987–1008.

Epstein, Seymour. 1994. “Integration of the Cognitive and the Psychodynamic Unconscious.” *American Psychologist* 49(8):709–24.

Evans, Jonathan St B. T. and Keith E. Stanovich. 2013. “Dual-Process Theories of Higher Cognition: Advancing the Debate.” *Perspectives on Psychological Science* 8(3):223–41.

Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf. 2018. “Cost of Experimentation and the Evolution of Venture Capital.” *Journal of Financial Economics* 128(3):422–42.

Fanelli, Daniele. 2010. “Do Pressures to Publish Increase Scientists’ Bias? An Empirical Support from US States Data.” *PLoS ONE* 5(4):e10271.

Fedorov, Oleg, Susanne Müller, and Stefan Knapp. 2010. “The (Un)Targeted Cancer Kinome.” *Nature Chemical Biology* 6:166.

Feduzi, Alberto, Jochen Runde, and Carlo Zappia. 2014. “De Finetti on Uncertainty.” *Cambridge Journal of Economics* 38(1):1–21.

Fischhoff, Baruch. 2015. “The Realities of Risk-Cost-Benefit Analysis.” *Science* 350(6260).

Fischhoff, Baruch and Alex L. Davis. 2014. “Communicating Scientific Uncertainty.” *Proceedings of the National Academy of Sciences* 111(3):13664–71.

Fischhoff, Baruch, Paul Slovic, Sarah Lichtenstein, Stephen Read, and Barbara Combs. 1978. “How Safe Is Safe Enough? A Psychometric Study of Attitudes Technological Risk and Benefit.” *The Perception of Risk* Policy Sci(9):127–52.

Fleming, L., H. Greene, G. Li, M. Marx, and D. Yao. 2019. “Government-Funded Research Increasingly Fuels Innovation.” *Science* 364(6446):1139–41.

Foster, Jacob G., Andrey Rzhetsky, and James A. Evans. 2015. “Tradition and Innovation in Scientists’ Research Strategies.” *American Sociological Review* 80(5):875–908.

Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. “Publication Bias in the Social Sciences: Unlocking the File Drawer.” *Science* 345(6203):1502–5.

Franzoni, Chiara. 2009. “Do Scientists Get Fundamental Research Ideas by Solving Practical Problems?” *Industrial and Corporate Change* 18(4):671–99.

Franzoni, Chiara and Cristina Rossi-Lamastra. 2017. “Academic Tenure, Risk-Taking and the Diversification of Scientific Research.” *Industry and Innovation* 24(7):691–712.

Franzoni, Chiara and Henry Sauermann. 2014. “Crowd Science: The Organization of Scientific Research in Open Collaborative Projects.” *Research Policy* 43(1):1–20.

Franzoni, Chiara, Giuseppe Scellato, and Paula Stephan. 2012. “Foreign-Born Scientists: Mobility Patterns for 16 Countries.” *Nature Biotechnology* 30(12):1250–53.

Franzoni, Chiara, Giuseppe Scellato, and Paula Stephan. 2018. “Context Factors and the Performance of Mobile Individuals in Research Teams.” *Journal of Management Studies* 55(1):27–59.

Frey, Renato, Andreas Pedroni, Rui Mata, Jörg Rieskamp, and Ralph Hertwig. 2017. “Risk Preference Shares the Psychometric Structure of Major Psychological Traits.” *Science Advances* 3(10):1–13.

Gilbert, Nigel G. and Michael Mulkay. 2003. *Opening Pandora’s Box. A Sociological Analysis of Scientists’ Discourse*. Cambridge: Cambridge University Press.

Haimes, Yacov Y. 2009. “On the Complex Definition of Risk: A Systems-Based Approach.”

Risk Analysis 29(12):1647–54.

Hájek, Alan. 2019. “Interpretations of Probability.” *The Stanford Encyclopedia of Philosophy. Edward N. Zalta (Ed.)*. Retrieved (https://plato.stanford.edu/archives/fall2019/entries/probability-interpret/).

Halzen, Francis. 2010. “The IceCube Project.”

Halzen, Francis and J. G. Learned. 1988. “High-Energy Neutrino Detection in Deep Polar Ice.” in *Proceedings of the 5th International Symposium on Very High-Energy Cosmic Ray Interactions*.

Hansson, Sven Ove. 1989. “Dimensions of Risk.” *Risk Analysis* 9(107–112).

Hansson, Sven Ove. 2002. “Philosophical Perspectives on Risk.” *Keynote Address at the Conference Research in Ethics and Engineering (25 April 2002)*.

Hansson, Sven Ove. 2018. “Risk.” in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. Stanford University.

Harris, Richard. 2013. “Scientists Win Nobel For Work On How Cells Communicate.” Retrieved (https://www.npr.org/templates/story/story.php?storyId=230192033).

Heinze, Thomas. 2017. “How to Sponsor Ground-Breaking Research: A Comparison of Funding Schemes.” *Science and Public Policy* 35(5):302–18.

Henrickson, Sarah and David Altshuler. 2012. “Risk and Return for the Clinician-Investigator.” *Science Translational Medicine* 4(135):135cm6 LP-135cm6.

Johansen, Inger Lise and Marvin Rausand. 2014. “Foundations and Choice of Risk Metrics.” 62:386–99.

Kahneman, Daniel and Amos Tversky. 1979. “Prospect Theory: An Analysis of Decision under Risk.” *Econometrica* 47(2):263–92.

Kaplan, Stanley and B. John Garrick. 1981. “On The Quantitative Definition of Risk.” *Risk Analysis* 1(1):11–27.

Kasperson, Roger E., Ortwin Renn, Paul Slovic, Halina S. Brown, Jacque Emel, Robert Goble, Jeanne X. Kasperson, and Samuel Ratick. 1988. “A Social Amplification of Risk: A Conceptual Framework.” *Risk Analysis* 8(2):177–87.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. “A Smooth Model of Decision Making under Ambiguity.” *Econometrica* 73(6):1849–92.

Knight, Frank. 1921. *Risk, Uncertainty, and Profit*. Boston, MA: Houghton Mifflin.

Krieger, Joshua L., Danielle Li, and Dimitris Papanikolaou. 2019. *Nber Working Paper Series Missing Novelty in Drug Development*.

Krummel, Matthew F. and James P. Allison. 1995. “CD28 and CTLA-4 Have Opposing Effects on the Response of T Cells to Stimulation M F Krummel 1, J P Allison.” *Journal of Experimental Medicine* 182(2):459–65.

Lanoë, Marianne. 2019. “The Evaluation of Competitive Research Funding: An Application to French Programs.” L’Universite de Bordeaux.

Laudel, Grit. 2017. “How Do National Career Systems Promote or Hinder the Emergence of New Research Lines?” *Minerva* 55(3):341–69.

Lerner, Josh and Ramana Nanda. 2020. “Venture Capital’s Role in Financing Innovation: What We Know and How Much We Still Need to Learn.” *Journal of Economic Perspectives* 34(3):237–61.

Linton, J. D. 2016. “Improving the Peer Review Process: Capturing More Information and Enabling High-Risk/High-Return Research.” *Research Policy* 45(9):1936–38.

Loewenstein, George F., Chrisopher K. Hsee, Elke U. Weber, and Ned Welch. 2001. “Risk as Feelings.” *Psychological Bulletin* 127(2):267–86.

Loughran, Tim and Bill McDonald. 2013. “IPO First-Day Returns, Offer Price Revisions, Volatility, and Form S-1 Language.” *Journal of Financial Economics* 109(2):307–26.

March, James G. 1991. “Exploration and Exploitation in Organizational Learning.” *Organization Science* 2(1):71–87.

Marinacci, Massimo. 2015. “Model Uncertainty.” *Journal of the European Economic Association* 13(6):998–1076.

Markowitz, Harry M. 1952. “Portfolio Selection.” *Journal of Finance* 7(1):77–91.

Mata, Rui, Renato Frey, D. Richter, J. Schupp, and Ralph Hertwig. 2018. “Risk Preference: A View from Psychology.” *Journal of Economic Perspectives* 32(2):155–72.

Maxmen, Amy. 2017. "'Riskiest Ideas' Win $50 Million from Chan Zuckerberg Biohub." *Nature* 542:280–281.

Mellers, Barbara, Eric Stone, Terry Murray, Angela Minster, Nick Rohrbaugh, Michael Bishop, Eva Chen, Joshua Baker, Yuan Hou, Michael Horowitz, Lyle Ungar, and Philip Tetlock. 2015. "Identifying and Cultivating Superforecasters as a Method of Improving Probabilistic Predictions." *Perspectives on Psychological Science* 10(3):267–81.

Merton, Robert K. 1957. "Priorities in Scientific Discovery: A Chapter in the Sociology of Science." *American Sociological Review* 22(6):635–59.

Michelson, Evan S. 2020. *Philanthropy and the Future of Science and Technology*. Abigdon, Oxon (UK) and New York (NY): Routledge.

Morgan, M. Granger. 2014. "Use (and Abuse) of Expert Elicitation in Support of Decision Making for Public Policy." *Proceedings of the National Academy of Sciences* 111(20):7176–84.

Mullard, Asher. 2020. "COVID-19 Vaccine Development Pipeline Gears Up." *Lancet* 395(10239):1751–52.

Nelson, Richard R. 1959. "The Simple Economics of Basic Scientific Research." *Journal of Political Economy* 67(3):297–306.

OECD. 2018. "Effective Operation of Competitive Research Funding Systems." *OECD Science, Technology and Industry Policy Papers* (57).

Petsko, Gregory. 2011. "Risky Business." *Genome Biology* 12(119).

Popper, Karl. 1959. *The Logic of Scientific Discovery*. London: Hutchinson & Co.

Price, Michael. 2019. "TED Offers Funding Path, Not Just Stage, for 'audacious' Ideas." *Science* 364(6438):317.

Raiffa, Howard. 1997. *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. McGraw-Hill College.

Renn, Ortwin. 1998. "Three Decades of Risk Research: Accomplishments and New Challenges." *Journal of Risk Research* 1(1):49–71.

Romer, Paul. 1994. "The Origins of Endogenous Growth." *Journal of Economic Perspectives*

8(1):3–22.

Rosenberg, Steven A., James C. Yang, and Nicholas P. Restifo. 2004. “Cancer Immunotherapy: Moving beyond Current Vaccines.” *Nature Medicine* 10(9):909–15.

Rosenthal, R. 1979. “The File Drawer Problem and Tolerance for Null Results.” *Psychological Bulletin* 86:638–641.

Rzhetsky, Andrey, Jacob G. Foster, Ian T. Foster, and James A. Evans. 2015. “Choosing Experiments to Accelerate Collective Discovery.” *Proceedings of the National Academy of Sciences* 112(47):14569–74.

Sauermann, Henry, Chiara Franzoni, and Kourosh Shafi. 2019. “Crowdfunding Scientific Research: Descriptive Insights and Correlates of Funding Success.” *PLoS ONE* 14(1):e0208384.

Simonton, Dean Keith. 2003. “Scientific Creativity as Constrained Stochastic Behavior: The Integration of Product, Person, and Process Perspectives.” *Psychological Bulletin* 129(4):475–94.

Sinkjaer, Thomas. 2018. “Fund Ideas, Not Pedigree, to Find Fresh Insight.” *Nature* 555:143.

Sloman, Steven A. 1996. “The Empirical Case for Two Systems of Reasoning.” *Psychological Bulletin* 119(1):3–22.

Slovic, Paul. 1987. “Perception of Risk.” *Science* 236:280–85.

Slovic, Paul, Melissa L. Finucane, Ellen Peters, and Donald G. MacGregor. 2010. “Risk as Analysis and Risk as Feelings: Some Thoughts about Affect, Reason, Risk and Rationality.” Pp. 21–36 in *The Feeling of Risk*, edited by P. Slovic. Abingdon, Oxon: Earthscan.

Slovic, Paul, Ellen Peters, Melissa L. Finucane, and Donald G. MacGregor. 2005. “Affect, Risk, and Decision Making.” *Health Psychology* 24(4 SUPPL.):35–40.

Stephan, Paula. 2007. “Early Careers for Biomedical Scientists: Doubling (and Troubling) Outcomes Presentation Harvard University.” in *Presentation Harvard University*. Boston, MA.

Stephan, Paula. 2012. *How Economics Shapes Science*. Cambridge, MA: Harvard University

Press.

Stephan, Paula. 2019. “Practices and Attitudes Regarding Risky Research.” Stanford University: Metascience Symposium.

Stephan, Paula E. 1996. “The Economics of Science.” *Journal of Economic Literature* 34(3):1199–1235.

Stephan, Paula E. 2010. “The Economics of Science.” Pp. 217–73 in *Handbook of Economics of Innovation*, edited by B. H. . Hall and N. Rosenberg. Amsterdam, NL: North Holland.

Stephan, Paula and Sharon Levin. 1992. *Striking the Mother Lode in Science: The Importance of Age, Place, and Time*. Oxford, UK: Oxford University Press.

Stephan, Paula, Reinhilde Veugelers, and Jian Wang. 2017. “Blinkered by Bibliometrics.” *Nature* 544(7651):411–12.

Suntharalingam, Ganesh, Meghan R. Perry, Stephen Ward, Stephen J. Brett, Andrew Castello-Cortes, Michael D. Brunner, and Nicki Panoskaltsis. 2006. “Cytokine Storm in a Phase 1 Trial of the Anti-CD28 Monoclonal Antibody TGN1412.” *New England Journal of Medicine* 355:1018–28.

Tetlock, Philippe and Dan Gardner. 2015. *Super-Forecasting. The Art and Science of Prediction*. London: Penguin Random House.

Tobin, James. 1958. “Liquidity Preference as Behavior towards Risk.” *Review of Economic Studies* 25(2):65–86.

Tversky, Amos and Craig R. Fox. 1995. “Weighing Risk and Uncertainty Tversky Fox 1995.Pdf.” *Psychological Review* 102(2):269.

Tversky, Amos and Daniel Kahneman. 1974. “Judgment under Uncertainty: Heuristics and Biases.” *Science* 185(4157):1124–31.

Tversky, Amos and Daniel Kahneman. 1991. “Loss Aversion in Riskless Choice: A Reference-Dependent Model.” *Quarterly Journal of Economics* 106(4):1039–61.

Tversky, Amos and Daniel Kahneman. 1992. “Advances in Prospect Theory: Cumulative Representation of Uncertainty.” *Journal of Risk and Uncertainty* 5(4):297–323.

Uzzi, Brian, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. “Atypical

Combinations and Scientific Impact." *Science* 342(6157):468–72.

Varian, Hall R. 2009. "The Economics of Combinatorial Innovation. Guglielmo Marconi Lecture June 16, 2009."

Veugelers, Reinhilde, Paula Stephan, and Jian Wang. 2019. "Excess Risk-Aversion at ERC." *Presented at the ASSOF Meeting*.

Vilkkumaa, Eeva, Ahti Salo, Juuso Liesiö, and Afzal Siddiqui. 2015. "Fostering Breakthrough Technologies - How Do Optimal Funding Decisions Depend on Evaluation Accuracy?" *Technological Forecasting and Social Change* 96:173–90.

Wang, Jian, You Na Lee, and John P. Walsh. 2018. "Funding Model and Creativity in Science: Competitive versus Block Funding and Status Contingency Effects." *Research Policy* 47(6):1070–83.

Wang, Jian, Reinhilde Veugelers, and Paula Stephan. 2017. "Bias against Novelty in Science: A Cautionary Tale for Users of Bibliometric Indicators." *Research Policy* 46(8):1416–36.

Weinstein, Neil D. 1980. "Unrealisitc Optimism about Future Life Events." *Journal of Personality and Social Psychology* 39(5):806–20.

Winkler, Robert L., Yael Grushka-Cockayne, C. Lichtendahl Jr, and Richmond R. Jose. 2019. "Probability Forecasts and Their Combination : A Research Perspective." *Decision Analysis* 16(4):239–333.

Wolchok, Jedd D., F. Stephen Hodi, Jeffrey S. Weber, James P. Allison, Walter J. Urba, Caroline Robert, Steven J. O'Day, Axel Hoos, Rachel Humphrey, David M. Berman, Nils Lonberg, and Alan J. Korman. 2013. "Development of Ipilimumab: A Novel Immunotherapeutic Approach for the Treatment of Advanced Melanoma." *Annals of the New York Academy of Sciences* 1291(1):1–13.

NBER WORKING PAPER SERIES

FUNDING RISKY RESEARCH

Chiara Franzoni
Paula Stephan
Reinhilde Veugelers

Working Paper 28905
http://www.nber.org/papers/w28905

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2021

This paper was written for the NBER Entrepreneurship and Innovation Policy and the Economy (EIPE) meeting held April 27, 2021. The authors wish to thank the organizers, Josh Lerner and Scott Stern, for thoughtful comments on an earlier version. Comments from participants at the meeting are also highly appreciated. The authors also thank Katalin Karikó for helpful responses to our emails and her willingness to review our discussion of mRNA research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Funding Risky Research
Chiara Franzoni, Paula Stephan, and Reinhilde Veugelers
NBER Working Paper No. 28905
June 2021
JEL No. I23,O31,O38

ABSTRACT

The speed with which Covid-19 vaccines were developed and their high-performance underlines how much society depends on the pace of scientific research and how effective science can be. This is especially the case for vaccines based on the new designer mRNA technology. We draw on this exceptional moment for science to reflect on whether the government funding system is sufficiently supportive of research needed for key breakthroughs, and whether the system of funding encourages sufficient risk-taking to induce scientists to explore transformative research paths. We begin with a discussion of the challenges faced by scientists who did pioneering-research related to mRNA-based drugs in getting support for research. We describe measures developed to distinguish risky from non-risky research and their citation footprint. We review empirical work suggesting that funding is biased against risky research and provide a framework for thinking about why principal investigators, panelists and funding agencies may eschew risky research. We close with a discussion of interventions that government agencies and universities could follow if they wish to avoid a bias against risk.

Chiara Franzoni
Department of Management, Economics
and Industrial Engineering
Politecnico di Milano
20133 Milan
Italy
chiara.franzoni@polimi.it

Reinhilde Veugelers
KU Leuven
MSI & ECOOM
Naamsestraat 69
3000 Leuven
Belgium
and Bruegel and PIIE and CEPR
reinhilde.veugelers@kuleuven.be

Paula Stephan
Department of Economics
Andrew Young School of Policy Studies
Georgia State University
Box 3992
Atlanta, GA 30302-3992
and NBER
pstephan@gsu.edu

1. Introduction

The Covid-19 pandemic has underlined how much society depends on the pace of scientific research and how effective science can be. The speed with which Covid-19 vaccines were developed and their high performance surpassed even the most optimistic expectations. This is especially the case for those based on the new designer mRNA technology which enabled the identification of a vaccine with high efficacy in less than three months after sequencing of the virus and holds huge promise for the future of vaccines and medicine more broadly.[1]

We draw on this exceptional moment for science to reflect on an important and pressing theme. Is the government funding system sufficiently supportive of the science needed for key breakthroughs, such as mRNA-based drugs? If the science needed for such breakthroughs requires transformative research, particularly in its early phases, does our system of science funding encourage sufficient risk-taking to induce scientists to explore transformative research paths?

In this contribution, we discuss risk-taking and the funding of risky science. We start in Section 2 by describing problems faced by Katalin Karikó, a scientist who did pioneering-research related to mRNA-based drugs. In section 3 we briefly describe measures developed to distinguish risky from non-risky research and the extent to which the citation footprint of risky research differs from that of non-risky research. We then review empirical work concerning the funding of research, which suggests that funding is biased against risky science. Section 4 provides a framework for thinking about why funding agencies may eschew funding risky research. We focus first on how factors within the research system, such as pressure to show results in a short time period and the widespread use of bibliometrics, contribute to risk aversion. We then focus on three key players affecting research funding and the role the three play in determining the amount of risky research that is undertaken: (1) principal investigators, (2) panelists and (3) funding agencies. Section 5 closes with a discussion of interventions that government agencies and universities could do if they wish to avoid a bias against risk.

[1] https://www.nature.com/articles/d41586-021-00019-w. Accessed March 4, 2021.

2. mRNA-design: difficulties encountered advancing a risky agenda

We begin by describing the development of designer messenger RNA (mRNA), the breakthrough technology used by Pfizer-BioNTech and Moderna to develop the first two vaccines against Covid 19 to obtain FDA approval in the US and EMA approval in the EU.

The mRNA is a protein-coding-single-stranded molecule, produced by cells during the transcription process, when the genes encoded in DNA are copied into the molecule of RNA. The discovery of mRNA was reported in *Nature* in May of 1961, a result of scientists' search concerning the synthesis of proteins coded in DNA. The path to synthesize mRNA in a test tube was made possible in 1984 when Paul Krieg and Doug Melton, scientists at Harvard, identified that, by using SP6 RNA polymerase, functional mRNA can be produced in vitro. By 1990, a group of scientists demonstrated that injection of synthetic, in vitro transcribed mRNA into animals led to expression of the encoded protein (Wolff et al. 1990). Soon, the scientific world realized that this system could potentially be used to turn human bodies into medicine-making factories and treat a variety of diseases, ranging from infection, to cancer to rare diseases and possibly mend such things as damaged heart tissue (Sahin, Karikó, and Türeci 2014). But, at the time, mRNA was not the only conceivable way to introduce protein expression into cells: other nucleic acid-based technologies were under investigation. Moreover, there remained two critical problems that needed to be addressed. In vitro-transcribed mRNA, when delivered to animals, could either be destroyed by the body as the body fielded an immune response before reaching its target, or worse yet, cause serious side effects (Sahin et al. 2014). No one knew how to make mRNA effective in humans despite years of interest on the part of scientists.

Katalin Karikó was determined, by all accounts, on finding a way to make synthetic mRNA applicable to treat human diseases. Born and educated in Hungary, she came to the US in 1985, first as a postdoctoral fellow to Temple University, then to USUHS. In 1989, she moved to a faculty position at the Medical School of the University of Pennsylvania.[2] She submitted more than 20 grants, initially for smaller sums, to the University of Pennsylvania and the American Heart Association, then for larger sums to NIH. [3] As hard as she tried, she repeatedly failed to get funding for her research. "Every night I was working:

[2] For a summary of Karikó's early career see https://www.wired.co.uk/article/mrna-coronavirus-vaccine-pfizer-biontech. Accessed May 27, 2021.
[3] Emails from Kariko to coauthors, March 17, 2021 and March 18, 2021.

grant, grant, grant," recounts Karikó. "And it came back always no, no, no." [4] Her inability to support her research on grants eventually resulted in her being taken off her faculty position by the university. In 1995 she accepted a non-faculty position, that she describes as "more like a post-doc position" at the University of Pennsylvania, without any prospect of advancing.[5]

Two years later, Drew Weissman, an MD PhD immunologist, moved from NIH (where he had worked with Anthony Fauci) to the University of Pennsylvania. The same year, Karikó and Weissman met at the school's photocopy machine. While chatting informally, they recognized that they shared an interest in developing a synthetic mRNA vaccine against HIV. They realized the potential of combining their biochemistry, molecular biology and immunology expertise and decided to begin working together. At the time, Karikó was focused on mRNA-based therapy for treating cerebral diseases and strokes. With Weissman, Karikó switched focus to mRNA-based vaccines. Weissman supported the early-stage work partly on one of his existing NIH grants, which had no direct connection to mRNA research. [6] Their breakthrough occurred when they recognized that uridine was the nucleoside in the mRNA that provoked the human immune system. They discovered that, when replacing uridine with pseudouridine, another naturally occurring nucleoside in the mRNA, it could enter into cells without alerting the RNA sensors. Their research was eventually published in *Immunity* in 2005, after being rejected by several leading journals.[7] Karikó was the first author, Weissman the senior author. It eventually became a highly cited paper, receiving to date more than 1000 Google-Scholar citations, although it took until 2015 to reach its first 500 citations. They disclosed their joint work to the University of Pennsylvania, which filed and obtained patents. Karikó and Weissman were listed as co-inventors. These patents, in line with the Bayh-Dole act, acknowledge NIH grants, including the grant that had no direct connection to mRNA research. The patents were licensed exclusively to CellScript by the University of Pennsylvania. CellScript sublicensed the University's patent to the German-based firm BioNTech, incorporated in 2008, and US-based Moderna, incorporated in 2010. The subsequent development of the

[4] https://www.statnews.com/2020/11/10/the-story-of-mrna-how-a-once-dismissed-idea-became-a-leading-technology-in-the-covid-vaccine-race/. Accessed March 4, 2021

[5] Correspondence with Kariko and https://www.ae-info.org/ae/Member/Karikó_Katalin for Karikó's CV. Accessed May 27, 2021.

[6] The grant was for the Role of gp-340 in HIV Infection and Transmissions.

[7] Gina Kolata, "Kati Kariko Helped Shield the World from the Coronavirus," New York Times, April 8, 2021. https://www.nytimes.com/2021/04/08/health/coronavirus-mrna-kariko.html. Accessed May 27, 2021.

mRNA-based drugs was conducted by these companies with equity investments, but also with a large involvement of public money.[8]

Even after the 2005 discovery, the two found funding for mRNA research difficult to obtain. According to Weissman "We both started writing grants. We didn't get most of them. People were not interested in mRNA. The people who reviewed the grants said, "mRNA will not be a good therapeutic, so don't bother." [9] Weismann, however, continued to receive funding from NIH, some, but not all of which, was for mRNA research; Karikó continued to have difficulty getting funding.[10] A 2007 R01 application that Karikó submitted to NINDS at NIH, for example, was not discussed at the study section meeting, having been judged by reviewers to be in the lower half of the applications. The proposal focused on the anti-inflammatory effects of neurotropics in ischemia stroke. Two reviewers described the proposed work as "novel," the third described the proposal as suffering from a "relative lack of novelty." Other comments from reviewers included statements such as: "Preliminary data should be provided to support that the proposed experiments can be carried out," "insufficient preliminary data." The work related to one of the aims was described as "very preliminary and, there is high likelihood, that these experiments, especially in vivo, will not work." [11] An application submitted in 2012 with Drew Weismann, with the goal of developing "a new therapeutic approach to treat ischemic brain injury by delivering specific mRNAs" was scored, but neither it nor the resubmission received a sufficiently strong score to be funded. Concerns included: "preliminary data presented are insufficient to suggest that this approach is worthy of in-depth evaluation in a stroke model" and that the first aim of the study was "largely descriptive." [12]

In 2006 Karikó and Weissman founded the company RNARx, with the intention of using mRNA to treat anemia. Karikó was the CEO of the company from 2006 to 2013. In

---

[8] E.g. in 2013 DARPA awarded Moderna a grant for up to $25M for developing mRNA-based therapeutics. https://investors.modernatx.com/news-releases/news-release-details/darpa-awards-moderna-therapeutics-grant-25-million-develop. Accessed May 27, 2021.

[9] Gina Kolata, "Kati Kariko Helped Shield the World from the Coronavirus," New York Times, April 8, 2021. https://www.nytimes.com/2021/04/08/health/coronavirus-mrna-kariko.html. Accessed May 27, 2021.

[10] Drew Weissman appears as the principal investigator on a total of 10 projects funded by the National Institutes of Health (NIH) between 1998 and 2021. Retrieved from https://reporter.nih.gov/ in March 2021.

[11] Reviews provided to the authors by Karikó March 18, 2021.

[12] The application was for the continuation of an R01 that Karikó "inherited" from Frank Welsh when he retired.

her role, she applied-for and received one STTR grant from NIH.[13] In 2013, Karikó became Senior Vice President at BioNTech.

We have no way of knowing what would have played out in terms of research outcomes if Karikó's early applications for funding had not been turned down or if she had gotten research support from the University of Pennsylvanian in the early period. Perhaps mRNA-based vaccines would have been available for Swine Flu in 2009. But, without the casual meeting with Weissman at the photo copy machine, she could also have given-up researching a way to make designer mRNA technology effective for drug development in humans. What we do know is that her early proposals were not funded and that the University of Pennsylvania moved her out of her soft-money faculty position. This could reflect a failure to address the problem of the immune system response, which later on was facilitated by her collaboration with Weissman. It could also reflect risk aversion on the part of review panels, that considered the area too risky to be fundable at the time, especially since Karikó had, at that time, few publications and citations, few preliminary results and no prior record of funding. More generally, the example tells us that the early-funding of designer mRNA research, now considered a promise of future medicine, was difficult.

Karikó is not the only scientist to hear "no, no, no." Similar anecdotal evidence is not difficult to find. A researcher at a top research institution in the US, in speaking of NASA and NSF, said: "programs are not very adventurous." And "what I experienced was that I couldn't get any new idea or anything I was really excited about funded by NSF. It never worked…the feedback is 'well this is too new: we don't know whether it's going to work'." (Franzoni and Stephan 2021). James Rothman, the day after he shared the Nobel Prize in Medicine or Physiology in 2013 told an interviewer that "he was grateful he started work in the early 1970s when the federal government was willing to take much bigger risks in handing out funding to young scientists." Rothman went on to say "I had five years of failure, really, before I had the first initial sign of success. And I'd like to think that that kind of support existed today, but I think there's less of it. And it's actually becoming a pressing national issue, if not an international issue." (Harris 2013).

[13] Katalin Karikó was the principal investigator of an STTR award project funded by the NIH between 2007 and 2011: https://www.sbir.gov/sbirsearch/detail/294077; https://grantome.com/grant/NIH/R42-HL087688-02. Accessed March 4, 2021.

3. Risk aversion in science funding: A review of empirical evidence

Concerns that the selection of grant proposals is overly conservative has been growing in recent years. Commentators on science policy have long lamented that science funders are too conservative and risk averse and skimp on supporting breakthrough research (e.g., Laudel 2017; Mazzucato 2015; Viner, Powell, and Green 2004). Funding agencies are accused of placing too much emphasis on the downside of avoiding failure and too little emphasis on the upside potential of supporting truly courageous ideas (Azoulay, Graff Zivin, and Manso 2012; Nicholson and Ioannidis 2012). But do we have more than anecdotal evidence on risk bias by science funding agencies? Do we have any science of science funding insights on this?

Although research in the area is limited, a handful of recent empirical works have begun to address the topic. The results support the view of risk aversion in funding. Before we review this evidence, it is important to note that risk remains an ill-defined concept (Althaus 2005; Aven 2011; Franzoni and Stephan 2021; Hansson 2018). Moreover, it is difficult to measure. Here we follow Franzoni and Stephan (2021) and use the term *risk* in its *speculative* meaning, in the sense that risk refers to uncertainty concerning the outcomes of research, where the outcomes vary predominantly in the spectrum of gains and potentially lead to exceptional results, but also to no results.[14] We preface the literature review with ways to measuring risky research.

Measures of risky research

Empirically identifying risky research is challenging (Franzoni and Stephan 2021). Most researchers who study risk depend on partial measures that look at the degree to which research results deviate from past results and/or look at the building blocks upon which the research is based. Foster and colleagues (2015) adopt the first approach and distinguish between three types of papers based on the chemical relationships described in the work. Research that makes a *jump* explores previously unexplored chemical relationships -jumping beyond current knowledge-. Such research arguably is more likely to fail but, if the research succeeds, is more likely to make a breakthrough. Research that explores relationships between previously studied entities is subdivided into research that tests a *new* relationship,

[14] We do not use the term *risk* in the *preventive* meaning, to mean the possibility of a negative event (e.g., a loss or harm), a use that is common in the Risk Analysis literature.

not published before, or research that *repeats* an analysis of a previously studied relationship. Foster and colleagues (2015) find that *jump* papers (reporting highly innovative chemical combinations) receive 52% more citations on average than "repeat" papers (reporting known combinations), while *new* papers (reporting moderately-innovative combinations) enjoy 30% more citations than those reporting known combinations. Their findings suggest that taking the risk associated with *jump* and *new* research makes it more likely to achieve high impact. But they also find this research is more likely to "fail." The authors thus find that the citation distribution associated with *jump* papers and *new* papers has a higher mean than that of *repeat* papers. They also find that the distribution has a higher variance, both characteristics that we expect in risky research, suggesting their measure correlates with risk. The additional rewards associated with *jump* papers are, however, relatively small and may not compensate sufficiently for the possibility of failing, suggesting higher expected returns of a safer research path.[15]

Wang et al. (2017) view scientific research as a combinatorial process and measure novelty in science by examining whether a published paper makes first time ever combinations of scientific knowledge components as proxied by referenced journals, accounting for the difficulty of making such combinations. Almost all new combinations made by novel papers cross subject categories. While recognizing that novelty is but one dimension of risk, they show that novel papers have patterns consistent with risky research of a higher mean and higher variance in citation performance; they also have a higher probability of becoming a highly cited paper, but at the same time a higher probability to be a no/low cited paper. Wang et al. (2017) also find strong evidence that novel research takes more time to become top-cited (Figure 1) and that it is published in journals having a lower impact as measured by the Journal Impact Factor. These findings suggest that bibliometric indicators based on citation counts and Journal Impact Factors with a short citation window, may be biased against risky, novel research. They also show that citations to novel papers are more likely to come from a broader set of disciplines and from disciplines that are more distant from their "home" field (Figure 2), suggestive that novel research has a tendency to both be best appreciated and to spark applications well beyond the disciplinary boundaries.

[15] They also find that papers based on *repeat* strategies were six times more likely to be published than those that used *new* or *jump* strategies during the period 1983-2008.

Figure 1 Delayed Recognition: Novel papers take more time to be cited

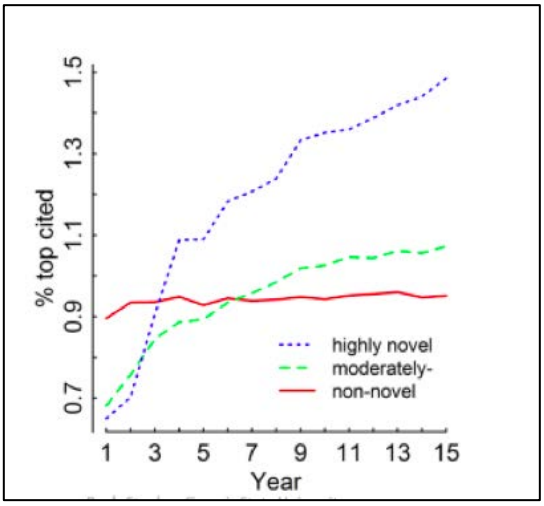

Figure 2 Trans-disciplinary impact over time

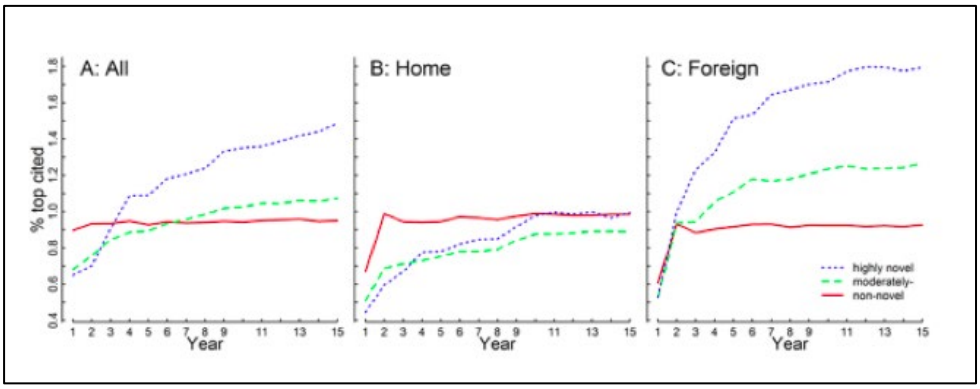

Source: Wang et al. (2017).

Assessing bias in funding risky research

Using their novelty measure, Veugelers, Stephan and Wang (2021) examine whether the ERC, the most important funding agency of the European Commission, set up in 2007 with the explicit aim to fund "high gain/high risk" research, is biased against novelty. They find that applicants to the ERC Starting Grant program with a history of highly novel publications are significantly less likely to receive funding than those without such a history. The major penalty for novelty comes during the first stage of selection, when panel members screen a large number of applications based on a short summary of the proposed research and a CV listing the candidate's main publications, for which they are most likely to use readily

available bibliometric indicators. The finding is thus consistent with the use of bibliometric indicators in grant selection as a source of bias against risky research.

In an experiment conducted at the Harvard Medical School, Boudreau and coauthors (2014) find that more novel research proposals, as measured by the percent of keywords not previously used, receive more negative evaluations during peer-review. This result is driven by proposals with particularly high levels of novelty. Their preferred explanation for this finding rests on bounded rationality of reviewers. To quote the authors: "experts extrapolating beyond the knowledge frontier to comprehend novel proposals are prone to systematic errors, misconstruing novel work. This implies that, rather than receiving unbiased assessments (with zero mean errors), novel proposals are discounted relative to their true merit, quality and potential." (Boudreau et al. 2014: 2779). Lanöe (2019), using a measure of novelty, finds evidence that funding decisions made by French National Research Agency are biased against risk-taking. Wagner and Alexander (2013) evaluate the SGER NSF program designed to support high risk, high reward research that ran from 1990 to 2006. Funding decisions were made entirely by program officers with no external review. The authors find that program officers routinely used but a small percent of available funds. The authors interpret the findings as suggesting that either officers were averse to funding risky research, despite the number of funded proposals that had transformative results, or that risk taking was not rewarded within NSF. Packalen and Bhattacharya (2018), using the vintage of ideas embodied in a paper as a proxy for novelty, find that NIH's propensity to fund projects that build on the most recent vintage of advances has declined over the last several decades.

Although the evidence discussed above is preliminary, it suggests that risky research is disfavored in the competition for funding. This seems the case not only when the funding is directed to 'standard' science, but even when a deliberate goal of the funding agency is to support high-risk, high gain, as in the case of the ERC.

Assuming that the preliminary evidence is correct, why do funding agencies eschew supporting risky research and thus possibly miss the opportunities of funding breakthroughs? Is this a conscious or unconscious choice, implanted in their modes of operating? And what can be done to encourage risk taking among funders? These are the questions we explore further in the next sections.

4. Why would funding agencies eschew supporting risky research?

To the best of our knowledge there is no research that directly addresses why funding agencies may be light on risky research. Moreover, there may be multiple factors that concur to play a role. Given the lack of research that informs on clear causes, we can only formulate a set of hypotheses. In this section we provide an overview of hypotheses that deserve more scrutiny in the future. We start with hypotheses that arise from outside the funding agencies and relate to the broader research system. We then discuss how these translate into a set of incentives and opportunities that could induce funding agencies to eschew supporting risky research and instead fund "safe" research at different levels of analysis. We consider three such levels: 1) The principle investigators who write and submit grant proposals, 2) the panels in charge of selecting research to fund (composed by panelists and research officers), and 3) the funding agencies. Figure 3 provides a summary of the hypotheses.

4.1 Research system

Calls for more accountability when using public funds and the trend towards more and regular evaluations of policy programs put increasing pressure on publicly-funded science institutions to show results, especially those aligned with political cycles. Shorter windows for results already bias against basic research programs in general; witness the heated discussions on the share of overall public R&D budgets for bottom-up basic research programs like ERC, NSF and NIH, compared to more directed, applied and closer to market programs. But even within basic research programs, the pressure to show results quickly may discourage publicly funded agencies from funding more risky research. A major factor is the length of time it takes risky research to have an impact: novel breakthroughs typically take a long time to materialize. As shown in Figure 1, Wang et al. (2017) find that novel research requires a longer time window than non-novel research to achieve impact. For the first three years after publication, the probability that a highly novel paper is among the top 1% of cited papers is below that of non-novel papers; but beyond these three years, highly novel papers start accumulating a lead and 15 years after publication, novel papers are nearly twice as likely to be in the top 1% of highly cited papers. This longer time window for big impact means that it takes the community a longer time to learn about new approaches and switch from established research paths to adopting new ones. Funding agencies, in an effort to maintain or expand funding, may feel that they cannot afford to wait for risky research to

show its impact and opt instead to fund safer research that has measurable effects in the near term, even if these effects are less likely to become breakthroughs.

An important factor for discouraging risky research may be the general lack of tolerance for failure and the meager rewards for those that take uncertain paths within the science system. As hiring, promotion and funding decisions are important conditions to engage in research and as reputation is a key reward to doing science, a lower inclination of researchers to engage in risky research can be traced back to biases against risk in the science system in general.

Universities routinely make crucial career decisions, such as hiring, mid-year review, tenure and promotion. When these career evaluations are made using bibliometric indicators with relatively short-time windows (like Journal Impact Factors or short windows for calculating citations), to measure research "quality", they can discourage risky research, as these measures appear to be biased against risk-taking (Stephan, Veugelers, and Wang 2017).

Career status and progression is not only an important reward for scientists themselves, it is also an element that goes into the track record and reputation that funding agencies and their panels consider as part of the applicant's profile. Any bias against risk in career decisions may thus have indirect effects on funding decisions, which may in turn affect career progression negatively when a candidate's productivity in acquiring external funding is a crucial factor in determining career progression, as the story of Karikó illustrated.

The large number of researchers in the US on "soft money" positions, i.e. in positions where salary is funded from grants that the researcher is responsible for obtaining, encourages the submission of proposals with little risk. If their research is not deemed fundable or comes up empty handed, the university can cut its losses and hire another individual into the position. It is notable that soft money positions have been on the rise in recent years. In the US, for example, the majority of basic medical faculty are hired in soft money positions and are responsible for bringing in most of their own salary (Stephan 2012). Soft-money positions also are common outside of medical institutions. Stephan documents that during the years when the NIH budget doubled, the majority of new hires were made into soft money positions (Stephan 2007). Soft money positions not only transfer risk to the faculty; they also discourage risk taking on the part of the faculty given the importance of continued funding.

Beyond affecting career progression and funding decisions, track record matters more generally as it affects a scientist's reputation and recognition. Although by no means the only reward to doing science, peer recognition and reputation are key drivers of scientists' choices (Merton 1957; Stephan and Levin 1992).[16] With peer recognition and reputation biased against risky research, scientists may be less prone to choose risky research paths. This is however not obvious; examples of scientists who have taken a risky course receiving a Nobel Prize, for example, are readily available [17] and there is research suggesting that prestigious prizes can encourage risk taking (Rzhetsky et al. 2015)

4.2 Principal investigators

The "lack of risky proposals" hypothesis

When Story Landis was Director of NINDS at NIH, she noticed that the amount of support the institute provided for what it classified as "basic-basic" research was declining, compared to what it was spending on "basic-disease" and applied research. Finding this disconcerting, the Institute set out to investigate why. Somewhat to their surprise, they found that the amount of dollars researchers requested to do "basic-basic" research had declined by 21%. Landis, when asked why she thought the decline was occurring replied: "My concern is that the decrease in the number of basic-basic applications reflects the perception that NINDS is only interested in disease-oriented research."[18] Basic research is not, of course, the same thing as risky research but the two are arguably close cousins. The example may suggest that researchers, anticipating that risky proposals have a difficult time at review, may simply refrain from conceiving and submitting risky research proposals. This is difficult to test, given a lack of data on proposals. But it is a plausible hypothesis. Some evidence consistent with a lack in supply of risky proposals is reported in Veugelers, Stephan and Wang (2021). They find that non-funded junior ERC applicants who fail in the second stage have significantly lower likelihood of producing novel papers after being rejected, compared to the successful ones. The evidence is consistent with rejected applicants learning that risk is not rewarded. Faced with the pressures to (re-) apply for funding, they adjust their research

[16] Rewards also include the satisfaction derived from puzzle solving, and financial gain that often accompanies a successful research career (Stephan 2012; Stephan and Levin 1992). Cohen, Sauermann and Stephan (2020) also show that scientists are strongly motivated by an interest in contributing to society.
[17] Jim Allison, who shared the Nobel Prize for immunotherapy for cancer in 2018 is but one case in point.
[18] https://blog.ninds.nih.gov/2014/03/27/back-to-basics/. Accessed March 30, 2021.

portfolio away from risky research, something which the successful applicants are "freed" from doing.

Overall, the reward-premium awarded by the science system for doing risky-research, compared to that of doing not-so-risky research, appears insufficient to encourage risk taking. The findings from Foster and colleagues (2015) reported supra, suggest that taking the risk associated with *jump* and *new* research makes it more likely to achieve high impact. But the additional rewards in terms of the extra citations they find are relatively small and may not compensate sufficiently for the possibility of failing in terms of not getting published and its negative impact on the researcher's careers. Their results thus suggest that returns may be higher for following a safer research path. Stephan (2019) has called this the "*Quad effect*", referring to the fact that competitive female figure skaters attempt fewer quadruple jumps, arguably because the incremental score they can earn for completing a quad, compared to successfully completing a triple jump, is insufficient to compensate for the risk of failing to complete the quad jump. For male figure skaters, scoring is different: the incremental score is larger and arguably provides sufficient incentive to attempt the quad. The work of Uzzi et al. (2013) is consistent with the findings of Foster et al. (2015), and shows that "[t]he highest-impact science is primarily grounded in exceptional conventional combinations of prior work yet simultaneously features an intrusion of unusual combinations", suggesting that a risky approach, when embedded in a more standard conventional approach can better escape the citation premium bias. Stated differently, a little bit of risk adds spice to the research; but conventionality is the dominant characteristic of highly cited papers.[19]

The "Loss aversion by principal investigators" hypothesis

The preferences of scientists for the level of risk involved in the projects they wish to pursue may not only reflect biases against risk in the reward structure of science, as discussed supra, but also loss aversion on the part of scientists. Behavioral psychology has shown that humans are generally loss-averse. They over-estimate the magnitude of perspective losses and under-estimate the magnitude of perspective gains (Kahneman and Tversky 1979; Tversky and Kahneman 1991). It is not implausible to expect that scientists are no exception to the rule.

[19] Papers characterized as having high medium conventionality coupled with a high tail "novelty" have a hit rate in the top 5 percent 9.2 times out of 100.

4.3 Research agencies

Lack of a portfolio approach hypothesis

A common approach taken by investors to stabilize the volatility of outcomes is to include in the same portfolio stocks that have uncorrelated outcomes, or that have outcomes that are negatively correlated, i.e. when one loses, the other gains. In finance, where the investors normally want to maximize the overall return of the portfolio and are risk-averse, a portfolio approach enables purchasing more risky stocks than when the investor buys stocks "one by one" without reference to what is in her portfolio. The "one-by-one" practice is frowned upon in the investment literature, given that the choice of a stock whose outcomes are highly correlated to those already in the portfolio may expose the investor to extreme gains, but also extreme losses, foregoing any advantages from using the portfolio to diversify away the risk.

The same logic holds to some extent for funding agencies. Agencies generally review proposals one by one, rank them in descending order of overall aggregated score, and then distribute funds according to the score until the budget is exhausted.[20] This "one by one" approach may arguably restrict the level of risk agencies take. To the extent that they are risk averse, the "one by one" approach only aggravates the risk-taking problem.

The "interdisciplinary bias" hypothesis

Review panels are often designed by funding agencies to be discipline-based. This, for example, is generally the case at NSF, and ERC. The latter for instance operates with 25-panels which are mostly discipline-focused. It follows that investigators who want to propose research involving multiple disciplines often must make hard choices concerning the most appropriate panel to consider their proposal. It also means that their proposal may face obstacles at review that discipline-focused proposals do not face. Bromham, Dinnage and Hua (2016) studied more than 18 thousand proposals submitted to the Australian Research Council Discovery Program. They found that the probability of receiving funding decreased

[20] At some agencies, such as NSF, program officers have some leeway in making decisions, but this is not common.

as the degree of interdisciplinarity of the proposal increased.[21] Banal-Estanol and colleagues (2019) studied the success rate of teams of co-investigators that sought funding at the UK Engineering and Physical Sciences Research Council. They showed that team-members with interdisciplinary backgrounds (i.e. who had balanced shares of publications in different fields) were penalized, even if those with an interdisciplinary background who were eventually funded were more successful ex-post.

A penalty directed at interdisciplinary research may work against funding risky science, because, as noted supra, papers of high novelty are often interdisciplinary. Moreover, Wang et al. (2017) also find that novel work that is highly cited is more likely to garner citations from outside, not from within its own field, suggesting that the research is appreciated more by others than by colleagues. Monodisciplinary panels may thus more likely be biased against risks associated with novel interdisciplinary research.

Peer review protocols conceal uncertainty hypothesis

Peer review opinions, especially for risky proposals, involve forecasting research outcomes in conditions of uncertainty (Knight 1921; Nelson 1959). However, protocols commonly used to elicit experts' opinions arguably provide little room for uncertainty, usually requiring reviewers to provide a single score on a numeric ordinal scale to represent a criterion. For example, the ERC requires a single score to rate the "ground-breaking nature and potential impact of the research project". Given the uncertainty of future outcomes, the request of a single-point estimate score can conceptually be thought of as the *median* value of the possible outcome distribution envisaged by the reviewer. Whereas in peer review of 'standard' science, the provision of a single point-estimate may provide a necessary time-saving compromise, in evaluations of risky research, the outcomes of interest can be expected to be in the tails and a single-point estimate may have little meaning. Furthermore, uncertainty regarding the outcomes is *the* key piece of information in this case (Morgan 2014; Morgan and Henrion 1990). It seems plausible that similar practices that demand

[21] The study uses Interdisciplinary Distance (IDD), a measure that takes into account the fields indicated as pertinent to the proposal by the principle investigator and the distance between the fields, based on the relative frequency with which the fields co-occur throughout the entire sample.

experts to express a score that conceals, rather than represents uncertainty, may induce poor judgments.[22]

Practices that stress reviewers' agreement may disfavor risky science hypothesis

It is customary in grant peer review to collect several expert opinions about each proposal before taking a decision. The underling idea is that *aggregation*[23] of a larger number of opinions improves accuracy (Kaplan, Lacetera, and Kaplan 2008; Snell 2015), because random errors likely cancel each-other out when averaging results (Larrick and Soll 2006). This would especially be important for risky proposals, which are more difficult to evaluate, hence more exposed to imprecisions and misjudgments.

The efficacy of this approach relies on two assumptions. First, that a large number of independent reviewers are available. Second, that the mechanisms for aggregating multiple views are unbiased towards risk. In practice, however, both assumptions are problematic. The costs of the review process (Ismail, Farrands, and Wooding 2008) and the unwillingness of reviewers constrain the number of opinions that can be collected (e.g., the NIH advises 4 reviews for each proposal, ERC panels solicit between 3 and 10 external reviews, but only for proposals that are short-listed to go to the second-stage of evaluation). Moreover, reviewers may not be independent and instead have correlated errors (biases), because they share the same background knowledge or beliefs (Clemen and Winkler 1999; Ottaviani and Sorensen 2015).

Another critical point when moving from multiple opinions to a single aggregated opinion and to a final *deliberation* (e.g., binary choice to fund or not)[24] relates to whether methods and rules used in the decision are unbiased towards risk. Prior studies of peer review have evidenced low levels of agreement among reviewers even in the evaluation of "standard'" proposals (Pier et al. 2018). Risky proposals probably spark even greater disagreement, given the larger uncertainty involved. Current practices at NIH and ERC use

[22] The scholars of expert elicitation have elaborated and tested a number of techniques which are commonly used in drug-approval, risk analysis, climate-change forecasting, and other areas where uncertainty is key and expert opinions are the only way to collect information (Morgan and Henrion 1990).
[23] The term *aggregation* (Bolger and Rowe 2015; List 2012) means the combination of multiple opinions. Aggregation can be computed with rules or algorithms (e.g., average, quantile average, ..) or can be done behaviorally, with a *consensus meeting* (Hora et al. 2013; Martini and Sprenger 2018).
[24] *Deliberation*, i.e. the binary choice to fund or not, can be directly dependent on the aggregation method or involve additional rules (e.g., aggregation with arithmetic average and deliberation in descending order of aggregated score until budget saturation.)

behavioral aggregation, i.e. *consensus meetings*, during which multiple views are confronted and disagreement resolved with discussion (Lamont 2009). However, behavioral aggregation is exposed to groupthink (Cooke 1991; Lamont 2009) and may lead people to herd away from the truth, following influential opinions (Banerjee 1992; Mengel 2019). Consistent with the findings from Della Vigna and Pope (2018) that academics overestimate the accuracy of beliefs of highly-cited scholars, this may lead to herding on their beliefs. Furthermore, the requirement of consensus may arguably induce a bias against risky research. Assuming that risky proposals lead to outcomes in the 'tails' of the distribution, i.e., either "hits" of "flops" (Azoulay, Graff Zivin, and Manso 2011), it is plausible that the related opinions would also be polarized. If this is the case, methods of aggregation and deliberation that do imply consensus may be systematically biased against risk-taking (Linton 2016). Alternative methods that do not imply consensus exist or are conceivable, such as gold cards, or lotteries among those above a given threshold (Fang and Casadevall 2016; Gross and Bergstrom 2019; Roumbanis 2019), but their limited use has not to date enabled analyses.

4.4 Panelists and research officers

The "insurance agent" hypothesis

Many agencies and panels are acutely aware that the future of their program depends upon supporting researchers who do not come up "empty-handed." They may look at the opportunity cost of funding more risky research and compare it with benefits from funding safer research. These concerns may be magnified by the size of the grant. It is one thing to place $200,000 on a project that may come up empty handed. It is entirely another to place $2M.

Such concerns can lead panels to place considerable emphasis on "what can go wrong", rather than "what can go right" during the review process. One of the "what can go wrong" concerns is that the proposed research cannot be accomplished. This concern undoubtedly fuels the heavy emphasis at many funding agencies on strong "preliminary findings" or, at some agencies, contingency plans, as part of the proposal. In so doing, the implicit requirement is that research be de-risked before it is funded. In this way, the panel supports research with little chance of failure, funding sure bets rather than research that is not a sure-bet but may have considerable up-side potential.

The "bibliometric screening and workload" hypothesis

Scientists and agencies collectively invest a huge amount of time in peer review. For example, the NIH evaluates approximately 80,000 applications annually, engaging over 2,000 reviewers per years and has more than 150 standing Study Sections.[25] The ERC averages 15 members on each of its 25 separate panels. The average panel member for the Starting Grants looks at 137 proposals per call; for Advanced Grants, 83 proposals.[26]

Given the heavy workload, it is not surprising that reviewers and panel members may seek ways to rapidly screen proposals, especially on a first pass. One of the easiest ways to do so is to focus on the publishing record of the scientist proposing the research, by examining readily available measures of citations to papers and other bibliometric indicators on platforms such as Google Scholar and Scopus. Such was not always the case: As late as the early 1990s the only way to count citations was to laboriously look in the volumes published by the Institute of Scientific Information, usually only available in the panel member's institutional library.

Does a heavy focus on bibliometrics affect the panel's decision when it comes to supporting risky research? The work by Wang, Veugelers and Stephan (2017) suggests that the answer could be yes: They find that novel research is systematically less likely to be published in high Impact Factor journals.[27] Moreover, as noted above, novel research takes longer to be a top hit than does non-novel research. Such a bibliometric bias against novel research can lead panels to select against individuals with a history of novel (risky) research, especially when the applicant is young and has a short history of citations. More generally, a focus on bibliometrics shifts the basis of decisions away from the substance of the proposal to an easily accessible metric.

How workload affects the selection of novel R&D projects has not been studied for funding agencies; it has, however, been studied in R&D departments of for-profit firms. The

[25] https://grants.nih.gov/grants/peerreview22713webv2.pdf. Accessed January 2, 2021. For Study Sections see https://public.csr.nih.gov/StudySections/StandingStudySections. Accessed January 2, 2021.
[26] Average for 2008-2013.
[27] The fact that these impact factors are calculated using relatively short citation windows, coupled with the supra mentioned finding of Wang et al. (2017) that it takes a longer time window for novel research to become highly cited may explain why journal editors, striving for good scores on their impact factor, may be biased against novel research.

authors of one study find that a high panel workload reduces the panel's preference for novel research (Criscuolo et al. 2017).

"Risk-biased panel members" hypothesis

Reviewers are often selected based on the excellent expertise of their research profile. Selection of top/senior experts specialized in the exact same niche area of the proposal are typically assumed to be the best choice for reviewers. But are top experts also the best reviewers for assessing risky research proposals? Which kind of reviewers are needed for an unbiased assessment of risky research? Does it require experience with risky research to be more willing to take risks in evaluating the proposals of others?

Unfortunately, we lack specific knowledge illuminating which kind of reviewer's expertise is best suited at assessing risky research, both in terms of willingness to take risks and in terms of capacity to accurately assess risk. Only a handful of studies have looked at reviewers' characteristics and related preferences for funding research proposals. Looking at the intellectual distance between the background of the reviewers and the research being proposed, Li (2017) finds that NIH proposals which were related to the research of the panelist were scored more favorably than those unrelated; they were also scored more accurately. DellaVigna and Pope (2018), however, studying the accuracy of reviewers when predicting the outcomes of experiments in behavioral social sciences, find that scholars with expertise in an area were not more accurate than scholars who were not experts. They also find that more senior scholars, and academics with more citations were no more accurate. Looking at assessment of novel proposals specifically, Boudreau and colleagues (2014), find that novel proposals were evaluated less favorably than non-novel ones. This was not explained by the degree of intellectual proximity of the evaluator. In conclusion, the limited evidence available suggests that reviewers who are expert in the niche area are not necessarily more accurate in assessing research proposals and may be more prone to fund research in their area.

Figure 3 Incentives and opportunities regarding risky research: A summary.

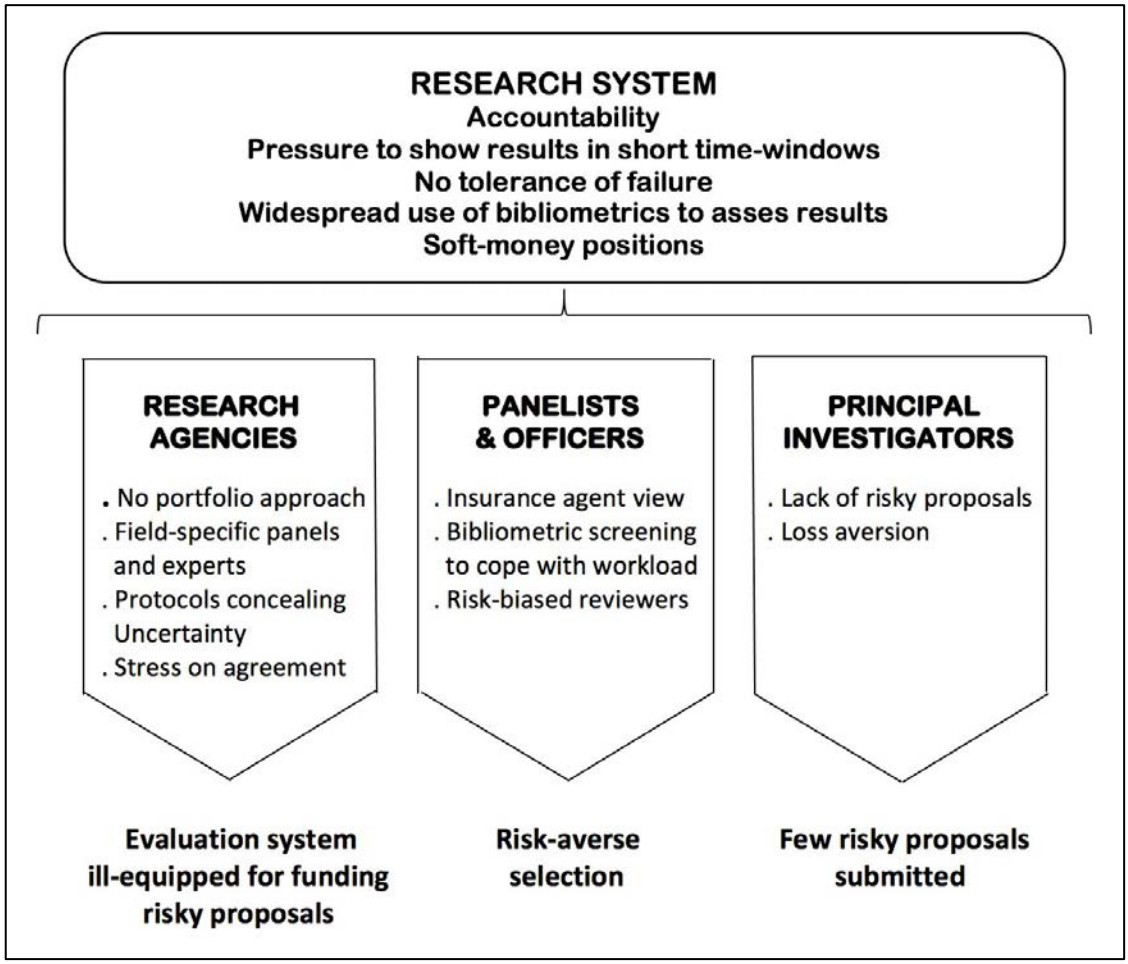

5. Conclusion and suggestions for encouraging risk-taking among science funders

Society needs scientific breakthroughs to tackle many of the challenges it faces. Because many of the paths to such breakthroughs are risky, its science system, and particularly its public science funding system, need to ensure that risk taking is encouraged or, at a minimum, that the system is not biased against risky research. The previous sections have made clear that we cannot take for granted that this is the case. The findings and hypotheses that we have explored also suggest possible ways for moving forward.

The Karikó -Weissman mRNA case already provides some initial thoughts regarding ways that risky research could be promoted. Their early joint work was partially supported

on a grant Weissman obtained that did not directly relate to designer mRNA. Such a "backburner" strategy is not uncommon in science. Scientists regularly rely on funding obtained to support another research objective, particularly in the very early exploratory stages of the other research, when it is still highly risky and without enough preliminary findings to apply for dedicated funding.

The mRNA example also shows the importance luck plays in pathbreaking research: if they had not met at the copy machine, Karikó and Weissman might never have formed the collaboration that led to mRNA vaccines. The probability that such lucky encounters occur, can be at least partly "engineered". A work environment enabling more open serendipitous encounters has the potential of leading to more risky research built on new unprecedented connections of knowledge pieces. The mRNA example also underlines the importance of taking an interdisciplinary approach: Karikó was trained as a biochemist, Weissman as an immunologist, a powerful combination to address the critical bottleneck for mRNA to be effective and safe in humans. An open environment enabling cross-disciplinary connections could thus already take away an important impediment for risky research. Moving the example from the mRNA novel scientific insights supported in part by NIH funding, into the development of mRNA-based vaccines for the market, supported in part by SBIR funding, DARPA, and eventually BARDA and Warp Speed, shows the importance of staged funding of new approaches.

We conclude with suggestions of ways to encourage risk-taking (or at a minimum avoid a bias against risk) in science, combining insights from the Karikó-Weissman mRNA research with insights from the admittedly limited evidence and research on risk avoidance as it relates to science funding reviewed in the previous sections. As discussed supra, the promotion of risk needs to be addressed within the entire science system: It cannot be solved by an individual program or funding agency. It requires a holistic perspective on the science enterprise, activating not only funders and their reviewers, but also universities and research centers, journal editors and their reviewers and, last but not least, researchers themselves. Without ignoring the holistic perspective, we focus the discussion on suggestions for science funders if they would like to augment their support for risky research. Where appropriate, we suggest further research and experiments designed with the goal of advancing our knowledge of ways to promote risk taking in science.

Deemphasize bibliometrics

Funding agencies, in order to advance innovation, could insist on multiple ways to assess applicants' research record, avoiding an overly focused use of bibliometrics. They could refrain from asking or inferring that grant applicants provide short-term citation counts, and indicators based on short term windows, such as the Journal Impact Factor and top field-cited articles. They could instruct panels to abstain from discussing such indicators in reviews or at a very minimum instruct panels of the potential biases which using such indicators entails.

Diversity in panel composition

Funders could balance panel composition with a sufficiently large number of panel members holding diverse and independent perspectives. They could avoid a panel design which is narrowly discipline-focused, and thus runs the danger of underappreciating the out-of-field impact from risky research. This is a more complicated task than selecting panel members based on top expertise in the field as the main panel selection criterium, as is commonly done, and which runs the danger of overly relying on the assumed superior assessment of experts regarding possible outcomes. More importantly, much could be learned concerning the causal impact of panel composition on risky research selection by using random control design to run experiments on panel composition.

Allow for disagreement

Alternatives to the commonly used consensus or average procedures should be considered. Aggregation/deliberation rules could be adapted to the nature of the science that the grant aims at sponsoring. Because more risky research is more prone to extreme outcomes, it matters not only to have reviewers willing to take risk, but also an accurate assessment of these extreme outcomes and their probability of occurrence. This requires a large enough number of sufficiently uncorrelated risk-unbiased opinions. In addition, the evidence that risky research may lead to more polarized views, warns against aggregation methods that rely on consensus (e.g., behavioral aggregation), or that assumes distributions of opinions according to a bell-shape (e.g., arithmetic average). To learn more about alternative methods that do not imply consensus, experiments with such alternative procedures could be conducted, using random trials so that we can properly evaluate their impact on the selection of risky research.

Portfolio approach

The 'one by one' approach typically used in panels works against selecting risky proposals. At a minimum, panels need to think about correlation among the proposals they are funding. One sure sign of high correlation in terms of low risk is requiring that all successful proposals have convincing preliminary findings. More generally, a portfolio approach to address risk aversion could require panels to put in different baskets highly risky and moderately risky proposals and provide a way to choose proposals from each. In practice such a portfolio approach could be quite challenging to implement for research projects. First, portfolio theory requires that the research paths be sufficiently uncorrelated. This may not hold within panels that are focused on specific subdisciplines that share risk factors. Correlation between research paths, in and of itself, can be hard to determine, particularly when covering vastly different goals across different fields and with different research approaches. Second, there is the question of fairness: in building a portfolio approach some proposals may have to be eliminated in an effort to balance or de-risk the portfolio.

Staging

An approach to de-risking, commonly used for funding entrepreneurial projects in the Venture Capital industry, where it is referred to as the 'spray and pray' strategy (Lerner and Nanda 2020), is to fund in stages, where increasingly larger amounts of funding are allocated, depending on whether interim milestones are being met. Funding in stages can be combined to include a portfolio approach[28].

Although a staging strategy is used by DARPA,[29] by the SBIR program and was recently introduced into the European Innovation Council program in the EU, it is rarely used for the funding of basic research. Can such a staging approach also be used by science funding agencies, allowing them to take more risk? Interim evaluation is especially useful when initial estimates are unreliable, but can be quickly updated and when investments can start at small scale and can be later scaled up (Vilkkumaa et al. 2015). It is thus especially suitable for research that can make substantial steps forwards in a relatively short period of time and does not require large fixed-costs to be started (Ewens, Nanda, and Rhodes-Kropf

[28] Veugelers and Zachmann (2020), for example, proposed a combination of a staging and portfolio approach to fund vaccines projects and calculated what such an approach would cost to society to obtain a desired number of vaccines at the end.
[29] https://fas.org/sgp/crs/natsec/R45088.pdf. Accessed August 30, 2020.

2018).[30] Some research fields meet these conditions, but the conditions are more the exception than the norm in the natural sciences, where the share of research that requires expensive equipment and substantial effort is large (Stephan 2010).

Loose-play, early-stage ideas

In the mRNA example, it was crucial that Weissman could use some of the funding he had already obtained to support early-stage risky joint research with Karikó. Other researchers often do the same. Dedicated loose-play programs have a number of pluses: it decreases the amount of effort that goes into preparation of fully-fledged proposals, it is reasonably easy to administer, and it has the potential of "de-risking" a research agenda before it goes up for formal evaluation at a granting agency. But a necessary condition to doing so requires that the PI have existing funding that can be redirected. If these are funds obtained from regular science funding programs, this raises the question of selection bias in terms of who obtains such funding. Researchers with a track record in novel research may be biased against. Also, early career researchers do not have access to such funds; others, such as Karikó, have tried to get funding but have not succeeded. And even for those who have such funding, the rules of engagement may not allow using funds for other than the research described in the proposal.

Going beyond the "backburner" option, could *dedicated* loose-play funding for early-stage risky explorations be operationalized? One approach is to have early-stage funding readily and quickly available to researchers at their home institution. The California Institute of Technology, by way of example, had such a program whereby faculty could submit a short proposal to the Vice Provost for Research and get a decision in a matter of days. Funds ranged from $25,000 to $250,000 a year for a period of two years. The idea was to give faculty the wherewithal to engage in early-stage risky research that, given apparent risk aversion of granting agencies, was deemed not yet ready for submission. If the initial findings looked promising, and produced enough preliminary data, the faculty could then submit a full grant proposal. Other institutions, such as ETH Zurich, provide generous base research funding to all their chairs which they can deploy for research at their own initiative.

[30] In the VC industry, this has largely favored IT and digital companies.

Dedicated programs for early-stage risky research and base research funding require resources, which institutions often do not have, or may not want to redirect away from other programs. One way to provide institutions with the financial means for such funding schemes would be for federal funders to shift a trench of resources to local institutions with the goal of encouraging risk taking. But loose-play funding can also have downsides. It could, for example, promote favoritism at the local level. It requires willingness and capacity to support risky projects at the local level, or at least no bias against risky projects, something which may not be present, as argued supra. It also involves a willingness to provide salary support to applicants, especially applicants in soft money positions. A primary reason Karikó was turned down for the $10,000 Merck starter grant, administered by the University of Pennsylvania, that she applied for in 1991, was a request for salary support. The rejection letter singled this out, citing one reviewer who said that "the most substantial weakness is the use of the entire award for faculty salary support."[31] Politically, shifting funds from federal agencies to universities for such programs also involves granting agencies ceding some control to local institutions.[32]

Funding researchers rather than projects and for longer periods of time

Programs that fund researchers rather than projects for longer periods of time allow researchers to engage in more risky research. It gives the scope and time to researchers to redirect their research in case of failure. The example that readily comes to mind is the Howard Hughes Medical Institute (HHMI) that funds successful applicants for seven years, rather than for three to five years, as is common for most other funding organizations, and where selection is based more on the applicant and his longer-term research strategy, rather than a specific research project. HHMI, moreover, does not demand early results nor does it penalize researchers for early failure. Azoulay et al. (2011) compare the research output of HHMI investigators to a group of similarly accomplished NIH investigators using propensity scoring. They find that HHMI investigators use more novel keywords and produce more hits and more flops, compared to the NIH investigators. Although it is not clear whether these results depend upon the longer duration of grants and the practice of HHMI to not demand

[31] Letter to Karikó dated April 29, 1991.
[32] NIH already does this by awarding training grants to institutions to administer. In the early years of NIH, candidates for training awards were selected at NIH.

early results nor penalize researchers for early failure or for other variables, the results suggest that these practices encourage risk-taking.

Targeting science funding to risky breakthrough missions

In the Karikó -Weissman case, the vexing problem of immune response was a well-known scientific challenge, impeding the promising mRNA technology to be used as a modality for treating humans. As such, it could have been turned into a "mission" with dedicated funding. While the science funding system was either unable or unwilling to identify such a "mission" in the early research stages of mRNA research, a more mission-oriented approach was followed in later stages, when DARPA awarded Moderna up to $24.6 million in 2013 "to research and develop its messenger RNA therapeutics™" after having already awarded the company a "seedling" grant of $0.7 million to begin work on the project.[33] Subsequently, when the search for corona-virus vaccines became more pressing and BARDA and Warp Speed entered the picture, a more targeted approach for funding their development was used.

DARPA is frequently heralded for its successes in funding mission-oriented high-risk, high-reward research. Azoulay, Fuchs et al. (2019) identify the organizational flexibility and autonomy given to program directors as key success elements of the DARPA model. A key factor for DARPA's success is attracting program staff who are more like risk-taking, idea-driven entrepreneurs than like administrators. These program staff are given individual discretion to design project calls and select projects from across the distribution of reviewer scores, which is seen as antidote to the risk bias DARPA's reviewers may hold. The autonomy which program officers enjoy is combined with clear targets for which they are held accountable.

Can the DARPA model be replicated for avoiding the risk bias in funding basic research? Azoulay, Fuchs et al (2019) identify as "DARPAble" domain mission-motivated research on nascent technologies within an inefficient innovation system. These missions must be clearly identifiable, associated with quantifiable goals and have trackable progress metrics. Because a focus on basic research for improving understanding is not a clearly defined mission, the authors deem the DARPA model not appropriate for funding basic

[33] https://investors.modernatx.com/news-releases/news-release-details/darpa-awards-moderna-therapeutics-grant-25-million-develop. Accessed May 27, 2021.

research. While DARPA may not be a general model for funding basic research, it may nevertheless be inspirational for specific scientific challenges for which goals can be clearly defined.

Prizes

An alternative to a competitive-grant approach is to create prizes to encourage path-breaking research (e.g., Williams 2010, 2012). Although such prizes for risky breakthroughs can incentivize research, they shift the risk onto the shoulders of the researchers, given that prizes are only awarded conditional upon succeeding in the endeavor. It may thus only incentivize researchers who are risk lovers, who have high (perhaps overly high) estimates of their probability of success and have access to resources.

Conclusion

To conclude, science funding agencies should be encouraged to pave the way for promoting risk taking in scientific research, given that breakthrough research is often perceived as risky. The way forward is neither safe, nor clearly defined. It is risky in the sense that it is not clear which paths will work and which will not. Perhaps the most important contribution funding agencies can make would be to support research which builds knowledge on the *design of funding programs* and reviewing practices related to risky proposals that have the potential of delivering breakthroughs. This support could entail not only financing such research, but also granting access to data and championing experimental approaches to test alternative designs of research funding.

6. References

Althaus, Catherine E. 2005. “A Disciplinary Perspective on the Epistemological Status of Risk.” *Risk Analysis* 25(3):567–88.

Aven, Terje. 2011. “On Some Recent Definitions and Analysis Frameworks for Risk, Vulnerability, and Resilience.” *Risk Analysis* 31(4):515–22.

Azoulay, Pierre, Erica Fuchs, Anna P. Goldstein, and Michael Kearney. 2019. “Funding Breakthrough Research: Promises and Challenges of the ‘ARPA Model.’” *Innovation Policy and the Economy* 19(1):69–96.

Azoulay, Pierre, J. Graff Zivin, and Gustavo Manso. 2011. “Incentives and Creativity: Evidence from the Howard Hughes Medical Investigator Program.” *The Rand Journal of Economics* 42(3):527–54.

Azoulay, Pierre, Joshua Graff Zivin, and Gustavo Manso. 2012. *NIH Peer Review: Challenges and Avenues for Reform*. 18116. Cambridge, MA.

Banal-Estañol, Albert, Inés Macho-Stadler, and David Pérez-Castrillo. 2019. “Evaluation in Research Funding Agencies : Are Structurally Diverse Teams Biased Against ?” *Research Policy* 48:1823–40.

Banerjee, A. V. 1992. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics* 107(3):797–817.

Bolger, Fergus and Gene Rowe. 2015. “The Aggregation of Expert Judgment: Do Good Things Come to Those Who Weight?” *Risk Analysis* 35(1):5–11.

Boudreau, Kevin, Eva Guinan, Karim R. Lakhani, and Christoph Riedl. 2014. “Looking Across and Looking Beyond the Knowledge Frontier: Intellectual Distance and Resource Allocation in Science.” *Management Science* 62(10):2765–2783.

Bromham, Lindell, Russell Dinnage, and Xia Hua. 2016. “Interdisciplinary Research Has Consistently Lower Funding Success.” *Nature* 534(7609):684–87.

Clemen, Robert T. and Robert L. Winkler. 1999. “Combining Probability Distributions From Experts in Risk Analysis.” *Risk Analysis* 19(2):187–203.

Cohen, Wesley M., Henry Sauermann, and Paula Stephan. 2020. “Not in the Job Description: The Commercial Activities of Academic Scientists and Engineers.” *Management Science* 66(9):4108–4117.

Cooke, Roger M. 1991. “Experts in Uncertainty. Opinions and Subjective Probabilities in Science.” New York Oxford: Oxford University Press.

Criscuolo, Paola, Linus Dahlander, Thorsten Grohsjean, and Ammon Salter. 2017. “Evaluating Novelty: The Role of Panels in the Selection of R&D Projects.” *Academy of Management Journal* 60(2):433–60.

DellaVigna, Stefano and Devin Pope. 2018. “Predicting Experimental Results: Who Knows

What?" *Journal of Political Economy* 126(6):2410–56.

Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf. 2018. "Cost of Experimentation and the Evolution of Venture Capital." *Journal of Financial Economics* 128(3):422–42.

Fang, Ferric C. and Arturo Casadevall. 2016. "Research Funding: The Case for a Modified Lottery." *MBio* 7(2).

Foster, Jacob G., Andrey Rzhetsky, and James A. Evans. 2015. "Tradition and Innovation in Scientists' Research Strategies." *American Sociological Review* 80(5):875–908.

Franzoni, Chiara and Paula Stephan. 2021. *Uncertainty and Risk-Taking in Science: Meaning, Measurement and Management*.

Gross, Kevin and Carl T. Bergstrom. 2019. "Contest Models Highlight Inherent Inefficiencies of Scientific Funding Competitions." *PLoS Biology* 17(1).

Hansson, Sven Ove. 2018. "Risk." in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. Stanford University.

Harris, Richard. 2013. "Scientists Win Nobel For Work On How Cells Communicate." Retrieved (https://www.npr.org/templates/story/story.php?storyId=230192033).

Hora, Stephen C., Benjamin R. Fransen, Natasha Hawkins, and Irving Susel. 2013. "Median Aggregation of Distribution Functions." *Decision Analysis* 10(4):279–291.

Ismail, Sharif, Alice Farrands, and Steven Wooding. 2008. *Evaluating Grant Peer Review in the Health Sciences: A Review of the Literature*. Santa Monica.

Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2):263–92.

Kaplan, David, Nicola Lacetera, and Celia Kaplan. 2008. "Sample Size and Precision in NIH Peer Review." *PLoS ONE* 3(7).

Knight, Frank. 1921. *Risk, Uncertainty and Profit.* (1964 repr. New York, NY: Sentry Press.

Lamont, Michele. 2009. *How Professors Think. Inside the Curious World of Academic Judgement*. Cambridge, MA: Harvard University Press.

Lanoë, Marianne. 2019. “The Evaluation of Competitive Research Funding: An Application to French Programs.” L’Universite de Bordeaux.

Larrick, Richard P. and Jack B. Soll. 2006. “Intuitions About Combining Opinions : Misappreciation of the Averaging Principle.” *Management Science* 52(1):111–27.

Laudel, Grit. 2017. “How Do National Career Systems Promote or Hinder the Emergence of New Research Lines?” *Minerva* 55(3):341–69.

Lerner, Josh and Ramana Nanda. 2020. “Venture Capital’s Role in Financing Innovation: What We Know and How Much We Still Need to Learn.” *Journal of Economic Perspectives* 34(3):237–61.

Li, Danielle. 2017. “Expertise versus Bias in Evaluation: Evidence from the NIH.” *American Economic Journal: Applied Economics* 9(2):60–92.

Linton, Jonathan D. 2016. “Improving the Peer Review Process: Capturing More Information and Enabling High-Risk/High-Return Research.” *Research Policy* 45(9):1936–38.

List, Christian. 2012. “The Theory of Judgment Aggregation: An Introductory Review.” *Synthese* 187:179–207.

Martini, Carlo and Jan Sprenger. 2018. “Opinion Aggregation and Individual Expertise.” Pp. 180–201 in *Scientific Collaboration and Collective Knowledge*, edited by T. Boyer-Kassem, C. Mayo-Wilson, and M. Weisberg. Oxford University Press.

Mazzucato, Mariana. 2015. *The Entrepreneurial State. Debunking Public vs. Private Sector Myths*. New York: Anthem Press.

Mengel, Friederike. 2019. *Gender Biases in Opinion Aggregation*.

Merton, Robert K. 1957. “Priorities in Scientific Discovery: A Chapter in the Sociology of Science.” *American Sociological Review* 22(6):635–59.

Morgan, M. Granger. 2014. “Use (and Abuse) of Expert Elicitation in Support of Decision Making for Public Policy.” *Proceedings of the National Academy of Sciences* 111(20):7176–84.

Morgan, Millett Grager and Max Henrion. 1990. *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis 6th Edition*. Cambridge University

Press.

Nelson, Richard R. 1959. “The Simple Economics of Basic Scientific Research.” *Journal of Political Economy* 67(3):297–306.

Nicholson, Joshua M. and John P. A. Ioannidis. 2012. “Research Grants: Conform and Be Funded.” *Nature* 492(7427):34–36.

Ottaviani, Marco and Peter Norman Sorensen. 2015. *Combining Reviews in Expert Panels*.

Packalen, Mikko and Jay Bhattacharya. 2018. *Does the NIH Fund Edge Science?* 24860. Cambridge, MA.

Pier, Elizabeth L., Markus Brauer, Amarette Filut, Anna Kaatz, Joshua Raclaw, Mitchell J. Nathan, Cecilia E. Ford, and Molly Carnes. 2018. “Low Agreement among Reviewers Evaluating the Same NIH Grant Applications.” *Proceedings of the National Academy of Sciences of the United States of America* 115(12):2952–57.

Roumbanis, Lambros. 2019. “Peer Review or Lottery? A Critical Analysis of Two Different Forms of Decision-Making Mechanisms for Allocation of Research Grants.” *Science, Technology, & Human Values* 44(6):994–1019.

Rzhetsky, Andrey, Jacob G. Foster, Ian T. Foster, and James A. Evans. 2015. “Choosing Experiments to Accelerate Collective Discovery.” *Proceedings of the National Academy of Sciences* 112(47):14569–74.

Sahin, Ugur, Katalin Karikó, and Özlem Türeci. 2014. “MRNA-Based Therapeutics-Developing a New Class of Drugs.” *Nature Reviews Drug Discovery* 13(10):759–80.

Snell, Richard R. 2015. “Menage a Quoi? Optimal Number of Peer Reviewers.” *PLoS ONE* 10(4).

Stephan, Paula. 2007. “Early Careers for Biomedical Scientists: Doubling (and Troubling) Outcomes Presentation Harvard University.” in *Presentation Harvard University*. Boston, MA.

Stephan, Paula. 2012. *How Economics Shapes Science*. Cambridge, MA: Harvard University Press.

Stephan, Paula. 2019. “Practices and Attitudes Regarding Risky Research.” Stanford

University: Metascience Symposium.

Stephan, Paula E. 2010. “The Economics of Science.” Pp. 217–73 in *Handbook of Economics of Innovation*, edited by B. H. . Hall and N. Rosenberg. Amsterdam, NL: North Holland.

Stephan, Paula and Sharon Levin. 1992. *Striking the Mother Lode in Science: The Importance of Age, Place, and Time*. Oxford, UK: Oxford University Press.

Stephan, Paula, Reinhilde Veugelers, and Jian Wang. 2017. “Blinkered by Bibliometrics.” *Nature* 544(7651):411–12.

Tversky, Amos and Daniel Kahneman. 1991. “Loss Aversion in Riskless Choice: A Reference-Dependent Model.” *Quarterly Journal of Economics* 106(4):1039–61.

Uzzi, Brian, Satyam Mukherjee, Michael Stringer, and Ben Jones. 2013. “Atypical Combinations and Scientific Impact.” *Science* 342(6157):468–72.

Veugelers, Reinhilde, Paula Stephan, and Jian Wang. 2021. “Excess Risk-Aversion at ERC.” *Research Policy* forthcoming.

Veugelers, Reinhilde and Georg. Zachmann. 2020. “Racing against COVID-19: A Vaccines Strategy for Europe.” *Policy Contribution* (April).

Vilkkumaa, Eeva, Ahti Salo, Juuso Liesiö, and Afzal Siddiqui. 2015. “Fostering Breakthrough Technologies - How Do Optimal Funding Decisions Depend on Evaluation Accuracy?” *Technological Forecasting and Social Change* 96:173–90.

Viner, Neil, Philip Powell, and Rod Green. 2004. “Institutionalized Biases in the Award of Research Grants: A Preliminary Analysis Revisiting the Principle of Accumulative Advantage.” *Research Policy* 33(3):443–54.

Wagner, Caroline and Jeffrey M. Alexander. 2013. “Evaluating Transformative Research Programmes: A Case Study of the NSF Small Grants for Exploratory Research Programme.” *Research Evaluation* 22(3):187–97.

Wang, Jian, Reinhilde Veugelers, and Paula Stephan. 2017. “Bias against Novelty in Science: A Cautionary Tale for Users of Bibliometric Indicators.” *Research Policy* 46(8):1416–36.

Williams, Heidi. 2010. *Incentives, Prizes, and Innovation*.

Williams, Heidi. 2012. “Low-Performing Schools Attract and Keep Academically Talented Teachers ? Evidence.” *Journal of Policy Analysis and Management* 31(3):752–76.

Wolff, JA, RW Malone, P. Williams, W. Chong, G. Acsadi, A. Jani, and PL. Felgner. 1990. “Direct Gene Transfer into Mouse Muscle in Vivo.” *Science* 23(247):1465–68.