Industry and Innovation

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/ciai20

Randomisation as a tool for organisational decision-making: a debatable or debilitating proposition?

Gary Dushnitsky, J.P. Eggers, Chiara Franzoni & Florenta Teodoridis

To cite this article: Gary Dushnitsky, J.P. Eggers, Chiara Franzoni & Florenta Teodoridis (2023) Randomisation as a tool for organisational decision-making: a debatable or debilitating proposition?, Industry and Innovation, 30:10, 1275-1293, DOI: 10.1080/13662716.2023.2281983

To link to this article: https://doi.org/10.1080/13662716.2023.2281983

Published online: 04 Dec 2023.

INDUSTRY AND INNOVATION
2023, VOL. 30, NO. 10, 1275–1293
https://doi.org/10.1080/13662716.2023.2281983

Randomisation as a tool for organisational decision-making: a debatable or debilitating proposition?

Gary Dushnitsky[a], J.P. Eggers[b], Chiara Franzoni[c] and Florenta Teodoridis[d]

[a]Strategy and Entrepreneurship Area, London Business School, London, UK; [b]Stern School of Business, New York University, New York City, New York, US; [c]School of Management, Politecnico Di Miliano, Milan, Italy; [d]Marshall School of Business, University of Southern California, Los Angeles, US

ABSTRACT
The role of experts has been called into question recently. Scholarly works debate whether expert judgement is given excessive reliance on innovation, science, and entrepreneurial decision-making. Increasingly, there are arguments that managers, founders, and funders would be better off relying on randomisation to a much higher degree. This article sheds light on the integration of randomisation in decision-making, presenting the pros and cons of expert advice, on the one hand, and randomisation, on the other hand. The discussion goes beyond the Expert – Randomisation dichotomy and lays the foundation for thinking about decision-making in the modern era, and specifically the role of Artificial Intelligence. *This Version: 5 November 2023*

KEYWORDS
Experts; expertise; randomization; decision making; funding; science

JEL CLASSIFICATION
O3

1. Introduction

At least since the days of the Greeks, human societies have turned to the wise person – the sage – for advice, judgement, or prediction. Experts play an instrumental role in public and private decision-making. Yet, there is mounting evidence pointing at the limitations of experts especially in the face of irreducible uncertainty. Enter randomisation. There are recent calls – and reputable organisations that engage in the practice – for explicitly employing randomisation as a distinct stage of the decision-making processes. What is the logic underlying such an approach? What are the issues with expert decision-making that may stimulate the search for alternative approaches? And finally, where might we encounter such efforts and do they have merit?

This article draws on an animated debate at the 2023 DRUID conference to inform these questions. The debate tackled the following motion 'Let it be resolved that this conference believes that expert judgement is given excessive reliance in innovation, science and entrepreneurial decision-making, and that managers and funders would be better off relying on randomization to a much higher degree.' We draw on the insights put forward at the debate to inform our understanding of experts, randomisation and the role they play in modern decision-making.

CONTACT Gary Dushnitsky gdushnitsky@london.edu Strategy and Entrepreneurship Area, London Business School, Sussex Place NW1 4SA, London, UK

The contribution of this article is twofold. First, it sheds light on the integration of some elements of randomisation in decision-making in the sciences, entrepreneurship, and even among established organisations. We present a structured discussion of the pros and cons of expert advice, on the one hand, and randomisation, on the other hand. Second, we believe that the insights go well beyond the Expert – Randomisation dichotomy. Rather, the discussion lays the foundation for thinking about decision-making in the modern era, and specifically the role of Artificial Intelligence. Building on the aforementioned insights, the article derives arguments on the role of AI in decision-making.

2. Experts and the value of their judgement

2.1. Experts

Society has long deferred to experts. Experts are usually individuals with significant time and experience in various domains of business, science, or the arts. Experts have often been the source of advice and judgement and have long served the role of arbiters of taste, quality, and appropriateness (e.g. Caves 2000; Ginsburgh 2003; Zuckerman 1999). Expert judgements are influential in a number of ways, including acting as gatekeepers for vital resources to new projects in the arts, sciences, or business (Ginsburgh 2003; Reinstein and Snyder 2005), as well as predictors of commercial success and popular opinion (Eliashberg and Shugan 1997).

Research on the potential upside of expertise recognises the potential benefits of relying on experts. Experts bring a 'schematic' perspective to decision-making, allowing them to handle complex and nuanced decisions (Larrick and Feiler 2015). When experts have knowledge that is directly relevant to the task at hand, their forecasts are more accurate (Stewart, Roebber, and Bosart 1997). For example, work in decision sciences tested experts' accuracy in forecasting future global events (e.g. Who would win the next political elections? Will the price of BRENT crude oil rise over a given amount in the next six months?). Over more than three decades of studies, they came to the conclusion that experts' predictions are, on average, better than chance, especially in the near-term future (Mellers et al. 2015; Tetlock and Gardner 2015).

Extant work documented the impact of experts across various settings. In the context of entrepreneurship, we often think of angels and venture capitalists (VCs) as individuals who possess expert judgement about the prospect of entrepreneurial ventures (Gompers and Lerner 2004; Huang and Pearce 2015). Further, the backing of a reputable angel or VC investor serves as a signal to others and unlocks access to further resources (Baum and Silverman 2004; Hellmann and Puri 2002). It follows that these investors serve as gatekeepers to vital resources. The arts constitute another setting where experts are instrumental in securing access to resources. In the film industry, expert endorsement is known to shape the success of smaller artistic movies (Reinstein and Snyder 2005). In the arts, where profitability is not the signal or dominant objective, experts play a gatekeeper role, as their judgement funnels funding from government, corporate, and foundations (Woronkowicz et al. 2012). In policy advisory, the advice of experts is sought to make key decisions for the economy and citizens, such as to raise or lower the national discount rate, provide or deny market approval to a new drug or treatment, or tell if a site is idoneous for building a nuclear power plant. Expert scientists are also appointed in panel committees for selecting research projects that will receive large subsidies from taxpayers. The list could continue: from health policies to financial investments, the recourse to experts and pundits is ubiquitous in contemporary societies.

2.2. Boundary conditions to expert judgement

The cited findings on the accuracy of expert judgement come, however, with important caveats. The evidence suggests that while experts on average beat chance, they often do so by a small margin, and do just a little better than novices, if at all (Camerer and Johnson 1991). Moreover, experts' judgement is not superior to standard forecasting algorithms that assume the continuation of current trends (Tetlock 2017). The superforecasters, i.e. the small number of people that systematically exhibit superior forecasting abilities, are not necessarily experts, in the sense that their performance is not associated with longer or deeper expertise in the subject matter, nor is associated to personal demographic features, such as seniority (Mellers et al. 2015; Tetlock and Gardner 2015).

The scholarly evidence has many corollaries in the business world. Figure 1 illustrates the persistence of Fortune 500 companies between 1955 and 2015 (Liu 2020). We observe that the erosion rate – the rate by which Fortune 500 companies default – is increasing over the decades. Arguably, it underscores the limitation to expert CEOs judgement. Over the decades, companies face highly turbulent times due to ever intensifying uncertainty and conflicted views on target payoff. If the most successful companies – and the experts that lead them – end up failing at an increasing rate, it is clear that we must carefully think of the factors that underlie the limitation to expert judgement.

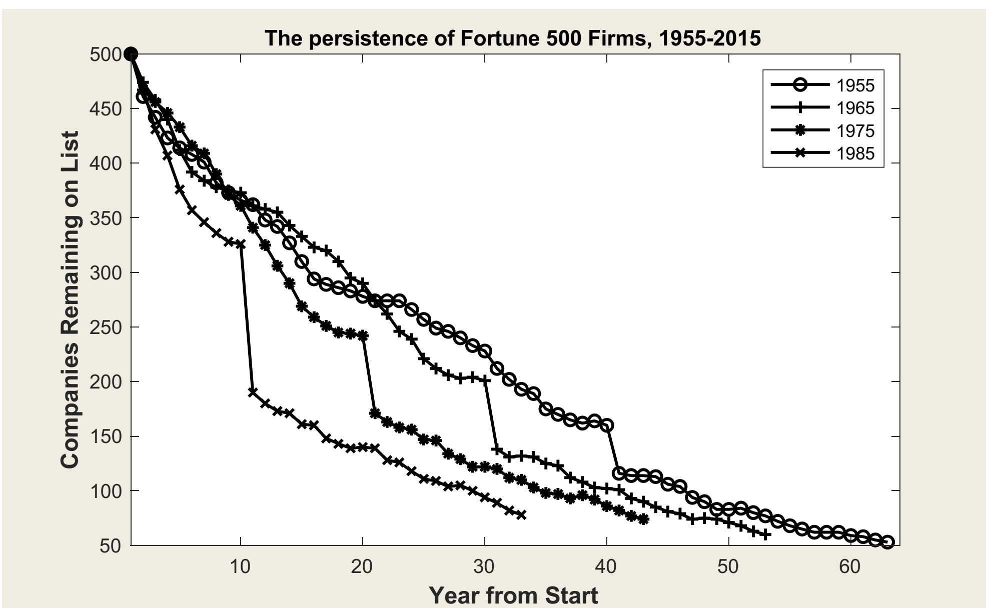

Figure 1. Liu (2020). Luck: a key idea for business and society. Oxford, UK: Routledge.

We therefore turn to consider the boundary conditions to superior expert judgement. Three sources of boundary conditions are identified: Knightian uncertainty regarding the value of inputs, expert biases in decision-making, and subjectivity regarding the prioritisation of multiple outputs.

First, consider the impact of Knightian uncertainty. In economics, Knightian uncertainty is a lack of knowledge about what can happen and/or a lack of information regarding the probability of event occurrence. It stands in contrast to the concept of risk, which refers to the presence of known possible outcomes and quantifiable probability of occurrence. In layperson terms, risk refers to 'known unknowns', whereas uncertainty is about 'unknown unknowns.' This view of uncertainty was first introduced by Frank Knight (Knight, 1921), and has played an important role in explaining entrepreneurial profits. When the uncertainty about future payoffs is due to 'unknown unknowns,' it may be almost impossible to know what types or magnitudes of inputs are necessary to drive payoff. Consequently, experts can hardly make a good job in forecasting which projects will maximise payoff. In this respect, Knightian uncertainty poses a boundary condition on the value of experts' judgement.

Second, judgement may be unwittingly shaped by extraneous factors and cognitive biases. For example, extant work suggests that early-stage investment decisions could be swayed by factors that have little to do with the actual payoff potential, such as sunshine at the time of making an investment (Dushnitsky and Sarkar 2022); or the gender (Balachandra et al. 2019; Bapna and Ganco 2021) or ethnicity of the founder (Pope and Sydnor 2011; Zhang et al., 2016). A similar pattern is observed for R&D projects where payoffs are also inherently uncertain; expert judgement is sensitive to the status of the project lead (Simcoe and Waguespack 2011) as well as the sequence of projects examined or overall level of workload (Criscuolo et al. 2021). The effect of extraneous factors is also documented in non-business settings. For example, judicial decisions were found to feature systematic patterns before and after judges go on a lunch break (Danziger and Avnaim-Pesso 2011). Experts are also on average more overconfident than non-experts (Mahajan 1992; Spence and Brucks 1997), suggesting that they have the tendency to overestimate their expertise, ability, and chances of success (Svenson 1981). This tendency, and the related implications for firm performance, has been largely investigated in managerial studies (see e.g. Burkhard et al. 2022 for a review). The conclusion is that overconfident CEOs are more likely to undertake risky projects (e.g. Galasso & Simcoe and Waguespack 2011; Engelen, Neumann and Schwens, 2015), which put companies under higher risk of both failure and success (Burkhard et al. 2022).

Third, consider the presence of multiple, unranked, outputs. This is increasingly the case, as societies and organisations contemplate multiple – competing or interrelated – objectives such as profitability, sustainability, and so on. Due to the multi-dimensional nature of alternatives, it is often the case that one alternative presents as superior to another in one dimension, but inferior to others in some other dimension. When this is the case, it is often difficult to rank-order all the alternatives without using criteria that are subjective and political and reflect the personal views and preferences of the experts. Nowhere is this more visible than in the selection of research proposals that compete for funding in large government agencies. Here, a set of experts, in this case scientists who serve as peer reviewers, evaluate the proposals along a number of criteria identified by the agency, for example: the ability of the proponent, the technical quality of the research proposed and the importance of the practical applications. The experts first score each criterion on an ordinal scale (e.g. 1 to 10), then give a final score that subsumes all. This process of converting multiple criteria scores into a single overall score is called *commensuration* and inevitably poses the experts in front of difficult dilemmas (Lee 2015). For example, is it more important to fund research on the neurobiology of Alzheimer's disease or of Cocaine abuse? Shall we prioritise the reduction of pollution from the burning of fossil-fuels or enhance the performance of solar panels? Is it better to fund a project that is feasible in technical terms, but not very ambitious in its findings or shall we rather fund research that could be disruptive, but may fail for a technical problem? The answers to these questions are arguably subjective and politically charged. As such, they may generate disagreement and criticism. Furthermore, experts' views may be different from the views of the general public. For example, recent studies suggest that experts seem to exhibit a strong preference for technical merit over practical impact (Eblen et al. 2016; Franzoni, Anders, and Stephan 2023). On the other hand, lay citizens are likely to feel the opposite (Franzoni, Sauermann, and Di Marco 2023). For some contexts and for some decision-makers, choices that are largely dependent on subjective or political considerations may be seen as unnecessarily arbitrary, inappropriate, or, at worst, illegitimate, unfair, or discriminatory. It follows that, even in the absence of Knightian uncertainty, expert judgement can face a notable boundary condition.

2.3. Expert judgement as a 'decision ritual'

The discussion highlights the value of expert judgement, and the limitations therefore. At the same time, we recognise that modern societies and organisations are in search of a routine or process that can aid with the intense uncertainty and conflicted values they face. Hard decisions require lots of information and a deep understanding of complex subject matters and intricate situations. It is therefore natural to ask the support of experts, that is individuals who, having spent a long time to train and practice in a given field, are presumed to have well-grounded views. These are the roots of trust in expertise.

And so, although experts face notable boundary conditions, the ritual of deferring to their judgement persists (Eyal 2019). The legitimacy that expertise confers is powerful and appealing, but – some argue – also over-rated. Seen from within, many decisions are made despite big margins of uncertainty (Morgan 2014; Yaniv and Foster 1995), for the simple reason that they are made in the context of insufficient information, where uncertainty is irreducible (Klibanoff, Marinacci, and Mukerji 2005; Knight 1921). This is reflected by frequent and large disagreement among experts, a circumstance that is well documented in the studies of expert decisions (Mumpower and Stewart 1996). When uncertainty is large, personal opinions, preferences, as well as conscious and unconscious biases may end up playing a large role in final decisions, ultimately undermining the legitimacy of experts. This has led to warnings against excessive trust in experts. Back in 1975, Burton Malkiel inaugurated a more disenchanted look at expert ability, claiming in his influential book 'A Random Walk Down Wall-Street', that 'A blindfolded monkey throwing darts at a newspaper's financial pages could select a portfolio that would do just as well as one carefully selected by experts.' Similar criticisms emerged in other realms besides finance and pushed to considering alternatives (Tetlock 2005). Looking at the world around us, we observe many situations where decision processes are augmented with other routines. One approach that has gained increasing attention in the latest years is to embed phases of randomisation as part of the decision process, that is take some decisions on the basis of chance. Note that randomisation does not come to substitute expert judgement. Rather, it comes to complement it. The discussion begs the following question: Can expert judgement be augmented by randomisation processes? We explore this question in the following sections.

3. Randomisation as an approach to augment decision processes

Before discussing the pros and cons of randomisation, we review different settings where notable organisations have embraced randomisations as a way to augment expert decision processes.

3.1. Randomisation augments selection of research project in the sciences

Government, philanthropic organisations, and universities must provide resources to scientific projects. The traditional approach entails a lengthy process of evaluation and selection entirely led by experts. However, the traditional process is costly and its outcomes are lacking. Comparisons of expert scores and subsequent research performance indicate that expert opinions are at best moderately predictive (Li and Agha 2015) or not predictive at all (Danthi et al. 2014). Experts appeared to be good at separating wheat and chaff, that is identifying very good and conversely mediocre proposals, but not so good in judging the intermediate range (Fang, Bowen, and Casadevall 2016). As a result, proposals were difficult to arrange around the funding line.

This has led several funding agencies, including the Swiss National Science Foundation, the British Academy, and the Novo Nordisk Foundation to augment expert selection with a random draw in the final phases of selection (Fang and Casadevall 2016). At the British Academy, experts do not rank-order all proposals. They only judge whether research proposals are fundable or not. If the budget is not sufficient to support all fundable proposals, these are put in a raffle and given an equal chance to be drawn and funded. At The Swiss National Science Foundation, experts rank-order all proposals and use randomisation only to choose among the proposals that are in intermediate positions, near the funding line. In this way, outstanding proposals are sure to go funded, while those that are good but not outstanding have a fair chance. In both these cases, random draws are used in combination, not in replacement, of expert decisions, mimicking the idea that both skills and chance would both ultimately play a role in research outcomes.

3.2. Randomisation augments investment decisions in early-stage entrepreneurial ventures

Entrepreneurship is another setting that is characterised by a high level of Knightian uncertainty. Entrepreneurs, and those who back them, face substantial uncertainty about different facets of the new venture; is there product-market fit? Would the technology scale? Would regulations change? Of course, one approach is to avoid an entrepreneurial pursuit or funding those who do so. An alternative approach is to adopt a strategy that embraces the reality of early-stage startups.

Enter the 'spray and pray' strategy. This investment strategy is guided by the belief that attempting to expertly assess the prospects of an early-stage venture is futile. Rather, this VC strategy augments expert judgement with randomisation. The venture capitalists do not attempt to fully evaluate the nascent startups and simply invest at random in a large pool of startups that match certain minimum criteria. The traction of those who persist – as well as any information gleaned from those who failed – serves to stabilise the value of the portfolio. The investors process this information and make a judgement call as to which startups should receive a subsequent round of funding.

Among those who adopt this strategy are 500 Startups, a reputable early-stage investor. As the name suggests, this investor makes the point of funding a large pool (about 500) startups at random. Scholarly work reveals that this approach has come to dominate early-stage investments (Lerner and Nanda, 2020). Moreover, this is particularly the case in sectors that exhibit an intense level of Knightian uncertainty (Ewens, Nanda and Rhodes-Kropf, 2018).

3.3. Randomisation augments strategic decision-making

The introduction of randomisation can enhance strategic decisions of established firms as well. That is the insight at the core of the Mendelian view of strategic decision-making (Levinthal 2017). In a nutshell, the Mendelian view asserts the executives should *explicitly* augment decision processes with randomisation. This contrasts with a traditional binary view that sees executives as either undertaking actions solely guided by rational choice or being subjected to random variation and market-based selection, typical of 'Darwinian' processes.

Mendelian executives do not fit in either extreme. Rather, they harness randomisation to drive superior outcomes and payoffs. An important part of this view is the understanding that variation is associated not only with 'happenstance randomness' in nature. Rather, it can be deliberately injected by executives with 'artificial selection' (Levinthal 2017). This allows executives to acknowledge the limitations imposed by uncertainty on the efficacy of their judgement. As a result, they can turn to curate sets of random trials of concurrent strategies, learn from these and shape the firm's trajectory accordingly. Therefore, the role of executives is reconceptualised away from experts who know everything and are driven by rational judgement and more as individuals who engage in 'artificial selection' through the systematic culling and amplification of strategic initiatives (Levinthal 2017). In fact, one can draw parallels between the Mendelian view and the VC adaptation of randomisation through the 'spray & pray' investment strategy.

Anecdotal evidence illustrates the value of this approach. Consider one of the most influential business executives of our generation; Jeff Bezos, the founder of Amazon. Figure 2 captures two major strategic initiatives led by Bezos back when he was at the helm. The figure on the right depicts the cover of Time Magazine ahead of Amazon's launch of its cloud service AWS. The figure on the left showcases Bezos at the launch event of the Fire Phone, Amazon's proprietary cell phone. Upon inspection of the figure, one may reflect back and ponder what was the most natural strategy for Amazon to adopt, keeping in mind that to that time Amazon was predominantly in the business of

Figure 2. Jeff Bezos as a Mendelian Executive.

selling products to end users. Seen in this light, the decision to launch the AWS service seems to be a departure from its product selling strategy, especially as it calls for completely different competencies. Indeed, the title of the Time Magazine front page mirrors this reservation. Arguably, Bezos' actions portray him as a Mendelian executive that goes beyond the predictable trajectory of strategic expansion, and rather engages in explicit curation of strategic experimentation. It follows that executives who strive to be as successful as Jeff Bezos should also consider introducing variation and randomisation as part of their strategic decision process.

4. Advantages of augmenting expert decision-making with randomisation

We advance four insights that shed light on the value of randomisation in enhancing organisational benefits or 'buy in.'

4.1. Acknowledging uncertainty, allowing exploration

A good deal of decisions involves irreducible uncertainty or involves large margins of errors. When information is lacking or where chance plays a large role, any decision lays on unsolid ground and cannot be taken with sufficient confidence. While no expert can resolve irreducible uncertainty, the use of experts may certainly lead to the erroneously belief of certainty, with the result that decisions are taken with excessive confidence and experts take unnecessary responsibilities. In the present word of distributed knowledge, fast information and post-truth, decisions are under heavy scrutiny, and can quickly lead to criticism and unpopularity. In this respect, randomisation allows acknowledging the role of chance and enhances transparency of the decision-making processes and the burden than comes with it.

Acknowledging the role of chance has another major benefit: It gives a 'licence to fail.' Introducing an explicit randomisation element to the process implies a recognition that there is no fault or stigma associated with failure. As a consequence, it may also increase the willingness to engage in exploration.

The corporate entrepreneurship literature echoes this point. Insights from the work of Burgelman and Van de Ven suggest that employees can be encouraged to explore by curating an organisational environment where one can experiment and fail. This Mendelian approach is nicely illustrated by William McKnight, a former 3 M CEO, who was known for his willingness to support employees who experience a well-intentioned failure. A similar approach is responsible, in part, to the rise of corporate venturing efforts where incumbent firms collaborate and invest in innovative startups (Dushnitsky, 2012). There is an explicit understanding that being close to the certain parts of the entrepreneurial ecosystem is akin to 'curating experimentations.' And while many of the startups may fail, there are higher-order learning to be gained (Maula, Keil and Zahra, 2013).

4.2. Minimising biases with procedural justice

As discussed above, when decision-making occurs among alternatives that differ along multiple dimensions, which are impossible to rank-order without using subjective judgements, the resulting decisions may look unnecessarily arbitrary. For example, alternative investments may differ in the amount of capital that they require, in their environmental implications, time horizon, exit cost. Research proposals may differ in their level of feasibility and transformative potential, and so on. In modern society, characterised by distributed knowledge, fast information, and post-truth, such decisions are under heavy scrutiny, and can quickly lead to criticism and unpopularity for those who take it. In these contexts, randomisation provides a tool of procedural justice that serves to dial-back conflicts and minimise criticism. A decision-maker, such as a CEO, or appointed expert can use its competences to identify a short-list of fair alternatives, then propose to randomise the choice among those alternatives as a means of *procedural justice*. The benefits are twofold. First, the integration of a randomisation phase will give a sense of fairness and trust in the process. The key is that randomisation is deployed only for the subset of those who have clearly passed a high threshold. Concerns of unconscious bias or 'foul play' often arise when one qualified person is selected over another. This is particularly likely to happen when the provision of a full and transparent explanation is impossible due to privacy or competitive reasons. In such cases, selection through randomisation can alleviate the concerns. Enhanced inclusion constitutes the second benefit, as we discuss in Section 4.4.

4.3. Economising on the opportunity cost of experts

Our economic and innovation systems are taxed by the excess reliance on experts. The workload of experts is amazing. By way of example, consider the cost of scientists selecting research proposals at the main funding institutions. In a rare attempt to quantify the workload, the National Science Foundation is estimated to have spent approximately 360 person-years full-time in 2015 alone to perform its reviews. Each proposal required about 3.9 hours of work to the expert, not considering travels and meetings (NSF, 2015, p.104). For an order of magnitude, consider that the National Institute of Health handles another 80,000 applications annually and the European Research Council about 8,000. Even those who see no fault in expert evaluation would agree that dispensing time and effort towards evaluation and selection imposes constraints on the experts. One might argue that a productive use of an expert's time is to be leading their own innovation effort. In fact, the reason for which we believe those individuals are best placed to evaluate innovative projects – their unique expertise – is a reason to avoid taxing their time.

It follows that any part of an evaluation process that can be meaningfully delegated to randomisation can be viewed as a net benefit. That is because it frees up the most capable researchers and inventors – the experts – to engage in work that propels the world forward. Note that this benefit is independent from any concerns about bias or 'foul play.' Even in a world where experts are benevolent, there is real value in not over taxing their time with administrative roles.

4.4. Stimulate participation and enhance inclusion

A large body of work suggests that selection and evaluation processes are plagued by homophily, bias, Matthew effect, and other frictions. As a result, those who believe themselves to be most susceptible to adverse bias may opt out from the process to begin with. They may 'lean out.' Randomisation and the procedural justice associated with it can regain the confidence of those individuals. It may therefore enhance the inclusion of communities that may otherwise forego participation altogether. By way of example, the British Academy reports that, after they introduced partial randomisation in their selection, they received many more applications from smaller and more peripheral institutions, including some many first-ever applicants.

5. Limitations of randomisation

The possibility to include randomisation in idea selection can be appealing as a way to reduce biases in cases where it is difficult to *ex ante* discern the underlying quality of an idea. However, the potential benefits of randomisation in science funding, venture funding, and corporate R&D allocations may be difficult to achieve due to the challenges of randomisation. Specifically, randomisation may affect behaviour around developing new ideas, deciding where to pursue support for an idea, and developing the idea towards completion. In addition, when considering the entire portfolio of ideas that are being supported, and not just each idea in isolation, the benefits of randomisation may weaken.

The question then is understanding the tradeoffs in how far to cull out bad projects versus letting randomisation make allocation decisions. Consider a process to approve five projects from 100 submissions (grant proposals, R&D project ideas, venture investments). On the one hand, if you cull to seven submissions before randomly choosing five, randomisation likely fails to remove any biases. On the other hand, randomly choosing five out of 100 has problems evident from taking a broader perspective. These problems fall into three buckets – adverse selection (especially in science funding), portfolio strategy (across all populations), and the role of advising the further development of the idea (mostly for firms & VCs). Our core message is that randomisation cannot be a panacea, and that the optimal threshold for how much to embrace randomisation may be far lower than initially expected.

5.1. Adverse selection

Let us assume two scientists, each considering proposing an idea for grant funding. One believes they have a high-quality idea, while the other knows this is not their best work, though it is still solid and we have two potential funding agencies, one that uses randomisation and one that tries to select the best idea. If these scientists know (even noisily) the quality of their own ideas, and know that review process A will use a threshold and then randomise, while review process B will seek to pick the best ideas, the scientist with the stronger idea will logically screen themselves out and select review process B. The result is a market for lemons (Akerlof, 1978), where the only ideas submitted to randomised processes will be low-quality ones.

These concerns are especially valid under several conditions. For example, when quality gaps are larger, such that the best ideas will be the ones that are most likely to select away from randomisation. The idea generator is overconfident in their idea quality, which we know is a common issue for creators, inventors, entrepreneurs who tend to fall in love with their ideas. Finally, they may arise when there is a higher degree of randomisation. For example, if getting into the final 100 only gives you a 5% chance of success, this is very discouraging.

There are two additional implications that apply even if *all* decision-making processes use randomisation, such that scientists cannot select out of randomised processes. First, randomisation diminishes incentives to fully develop one good idea, and instead incents proposing multiple half-baked ideas. This could be beneficial in cases where quality is highly uncertain *ex ante*, but on the margin, it is hard to believe this improves outcomes. Second, randomisation rewards those subfields that churn out 'good enough' ideas at the fastest rate. Consider how such behaviour would play out in the journal publication process, and which subfields would benefit the most if we randomised journal acceptances.

5.2. Implementation & development

Particularly for firms and venture investing, this is not just 'invest and forget', but the beginning of a long-term relationship to develop and implement the funded ideas (e.g. mentorship). In such cases, randomisation increases the chances of a poor fit between the advising/implementing body and the idea.

Advising is a key role for venture capital investing. VCs often end up with board seats, and those VC board members typically offer advice to the venture — Bacon-Gerasymenko and Eggers (2019) show that investors typically advise on eight issues per company, ranging from strategy to hiring to tech stack. Their possessing relevant expertise is particularly important and helpful to offering advice.[1] We also note that

[1]Despite the belief that VCs often feel comfortable offering opinions on nearly anything.

research on co-ethnic matches between VCs and ventures improves performance (Hegde & Tumlinson, 2014) based on the ability to communicate and trust each other, which often builds on shared knowledge and experience.

Meanwhile, firms do not just invest in innovative ideas, they both develop and (typically) implement them. If firms are not careful, they may randomly fund ideas for which they do not possess the proper complementary assets. Gaining access to complementary assets is an important bottleneck that often discourages market entry decisions (Chang, Eggers, and Keum 2022). Enlisting appropriate advisors after randomisation can be difficult if the advisors do not currently exist within the organisation – this leads to substantial additional costs down the road to build out the team as a reaction to the portfolio. The firm likely builds a capability that may not be useful after that one project.

5.3. Portfolio strategy

Organisations often need a portfolio strategy and should not just let their portfolio of investments evolve based on the whims of what is submitted. Portfolio strategies require taking into account the interdependencies between projects within a strategic portfolio. Some projects will be complements, substitutes, or orthogonal – funding decisions across those should be made with the portfolio in mind. Firms and investors often invest in portfolios to reduce risk, but randomisation may complicate portfolio strategies, as random choices would have to be made only within groups of projects with given types of interdependencies. If the projects are chosen randomly, without considering their interdependencies, there is the potential to increase risk.

For example, firms may want to invest in three substitute pathways, to ensure that ONE of them is successful, such as pursuing a Covid vaccine. Other firms may want to invest in complementary pathways, such as when a VC invests in multiple startups around a given problem area, knowing that the ventures may help one another over time. Science funding organisations often make portfolio investments to ensure 'coverage' or spread of resources in line with strategy. Decision makers must balance the tradeoffs of different types of projects and make intentional choices – randomisation runs counter to this intentionality.

One potential solution would be to create funding 'buckets' of projects with similar correlations and interdependencies across the other buckets, then randomise within each bucket. However, projects are not always plentiful. Their number may be so small that, in order to compare only apples-to-apples within each bucket, the buckets may get so small as to make randomisation impossible or almost irrelevant. For example, imagine 'randomly' choosing five of seven projects in a given bucket.

5.4. Use cases for randomisation

An obvious assumption may be that an intermediate solution – randomising five of 20 or 40 or 60 – might be a viable approach. The potential viability depends on underlying assumptions about the shape of benefits and costs from randomising. One can easily build models such that the benefits of randomisation to remove biases disappear with *any* randomisation, based on assumptions such as:

- Strong biases, where even initial selection will weed out projects that do not conform to those biases.
- Strong inventor beliefs about idea quality, such that ANY randomisation is extremely demotivating.
- Different project types, where the negatives of randomising across the portfolio may be greater.

This suggests boundary conditions for when randomisation makes sense, namely when (a) the biases being removed through randomisation are minor, (b) inventors have no choice but to submit their ideas to a randomised funding process, and (c) the submissions are all exceptionally similar to one another. It follows that those considering adopting randomisation would be well served to understand the potential downsides.

6. Alternative solution – improve expert decision-making

This suggests an important problem for those seeking to improve efficacy and reduce biases in decision-making – as discussed above, there are three important limitations to expert decision-making, but randomisation as a solution creates an additional set of issues. An alternative that is potentially complementary to randomisation is to use science and technology to improve expert decision-making.

We focus on two approaches. First, expert decision-making could be enhanced through interventions that aim to improve team-level decision-making through access to a broader set of knowledge. Second, expert decision-making could be improved with access to technology that can increase the accuracy of prediction. More specifically, the recent and significant advancements in machine learning (ML), which spearheaded artificial intelligence developments (AI), are, fundamentally, a prediction technology.

6.1. Improving expert team decision-making

Psychologists and decision scientists have worked extensively at ways to improve the accuracy of expert judgements. They have developed practices and protocols that allow to minimise biases, such as anchoring, or overconfidence, avoid groupthink and negativity biases, minimise misunderstanding created by ambiguous language and extrapolate richer and more accurate opinions. Many tools are ready off the shelf, but they are used only rarely in high-level policy advisory. A broader adoption would be one easy way to improve expert judgement. For example, Franzoni et al. (2022) provide some guidelines on how expert decision-making could be improved in scientific research funding.

The authors also discuss the importance of understanding the scientific production process and the limitations of various metrics used to evaluate success in science. For example, citations as a measure of evaluating prior success in research and hence as a predictor of future success are subject to limitations because, ultimately, citations are a social construct. Citation behaviour varies across domains of science, with some groups tending to cite larger volumes of prior work than others. Citations to certain academic papers that outline statistical methodologies or define theoretical concepts is expected in some domains, whereas in others such fundamentals are considered widely known. Moreover, the number of citations to an academic paper is correlated with the number of authors on the paper (Vakili, Teodoridis and Bikard 2022; Vakili, Teodoridis and Bikard, 2022. The correlation is, in part, explained by social network diffusion patterns rather than by an accurate evaluation of research quality. Similarly, science of science scholars (e.g. Stephan, 2012; Wang and Barbabasi, 2021) draw attention to the disproportionate emphasis placed on star scientists. Certainly, past success carries some predictive power for future success; however, several scholars have warned about the dangers of the Matthew Effect (Merton, 1957; 1973; Azoulay, Stuart and Wang, 2014) where reputable scientists are given disproportionate credit for their ideas or contributions to projects when compared to scientists who have not yet had a chance to make their mark.

The decision-making skills of expert evaluators could be improved if they were to be made aware of all the various limitations and intricacies of metrics used to evaluate past success in science. While this might sound trivial, humans crave the simplicity of metrics and often take past practices for granted either because updating their knowledge is costly or because better metrics are not readily available (Wu, 2023).

The burden of knowledge hypothesis (Jones, 2009) suggests another avenue towards improving the decision-making of expert evaluator teams. It states that knowledge accumulation leads to experts being knowledgeable in narrower and narrower domains, a fact which increases returns to collaboration. Thus, the decision-making of experts could be improved by focusing on carefully selecting experts with a diverse breadth of skills and knowledge, rather than aiming to maximise domain expertise. However, collaboration that needs to bridge across individuals with diverse expertise is costlier than that of more homogeneous teams. Scholars have found that generalists, individuals with broader knowledge across various domains, albeit at the cost of shallower depth, play a central role in lowering collaboration costs by coordinating between specialists (Teodoridis 2018). It follows that the decision-making of expert teams could be improved not only by recognising the benefits of expertise diversification in teams of evaluators but also by ensuring a mix between specialist evaluators that bring essential domain knowledge and generalist evaluators who ensure coordination between specialists and hence help maximise the benefits of diversification as a mechanism to increase the accuracy of predicting a project's likelihood of success.

6.2. ML/AI as complement to expert decision-making

One of the earlier identified limitations of expert decision-making is the observation that judgement can be influenced by extraneous factors and cognitive biases. These limitations could impede prediction in cases where patterns of the past are predictive of patterns of the future ('known known'), but also when evaluating risky proposals ('known unknowns') and under conditions of uncertainty ('unknown unknowns').

ML algorithms are fundamentally prediction algorithms (Agrawal, Gans and Goldfarb, 2018). They have the ability to generate more accurate predictions than humans, especially in conditions of known-knowns. In such situations, ML algorithms can analyse large volumes of historical data to identify patterns and trends that human decision-makers might overlook. For example, in the context of early-stage investments, ML models can analyse past investment data and outcomes to identify the key factors that actually influence success, helping investors make more objective data-driven decisions. ML algorithms can also consider a broader range of relevant variables that might be difficult for human experts to handle. This reduces the influence of irrelevant factors and cognitive biases that might influence expert decision-making, such as those driven by weather conditions, gender, or ethnicity. Beyond the potential for ML to complement expert judgement by reducing biases, the reverse is also true – experts are more likely to be aware of the nature of biases that may exist in the data, and thus improve outcomes by combining expertise and technology (Choudhury, Evan, and Agarwal 2020).

The benefits of identifying patterns in large datasets extends to risky proposals. ML algorithms can observe patterns in the data that might be difficult to identify by experts. By analysing vast amounts of data, ML can also uncover hidden patterns and relationships previously unknown to human experts. For example, given the increasing importance of interdisciplinary research and the difficulty of evaluating such proposals, ML could aid in identifying interdisciplinary connections. This could be done by analysing a vast corpus of scientific literature across interdisciplinary domains in order to reveal trends, recent breakthroughs, and gaps in the existing knowledge base.

Moreover, unlike expert prediction, ML prediction comes with a confidence interval. Thus, even if ML prediction with known unknowns might have limitations, those limitations are numerically captured in the prediction confidence interval. Such quantification would help reduce the impact of cognitive biases for expert decision-making when evaluating risky proposals.

ML algorithms can also be trained to identify outliers or anomalies in the data. This capability is crucial in identifying unexpected events or 'unknown unknowns' that might affect the outcome of a project or decision. Such outliers might be missed by human experts. ML models can also simulate various scenarios under conditions of uncertainty. By generating a range of possible outcomes based on historical data or assumptions, decision-makers would be able to explore a broader spectrum of possibilities and make more informed judgements.

Moreover, ML models can assess the expertise and overconfidence of experts by analysing their past decisions and performance. By providing feedback on their decision accuracy and highlighting areas where they may be overconfident, experts can become more aware of their biases and make more informed decisions. In addition, the algorithms could be set-up to provide real-time feedback. This continuous learning and feedback loop can help decision-makers adjust their judgements over time and become more objective.

Most importantly, ML algorithms can learn over time. This allows them to adapt to changing conditions and incorporate new information as it becomes available. This adaptability is crucial when dealing with predictions under conditions of risk and uncertainty.

7. Conclusions

The discussion reveals that experts are not without fault, yet randomisation is not a panacea. To fully understand the role – and limitations – of randomisation, one should look beyond the technical application and rather consider the organisational and social contexts within which they are deployed. Case in point is the discussion of adverse selection, and the very valid concern that the most able entrepreneurs or researchers will opt out of evaluation processes that heavily rely on randomisation. At the same time, the value of integrating elements of randomisation may be best realised when viewed in such organisational settings. The goal of this article is to summarise some of the organisational benefits and challenges associated with the deployment of randomisation as a distinct stage of organisational decision-making processes. Our discussion focused on three distinct organisational settings; the funding of grants for scientists (university/science setting), venture funding for new firms (small entrepreneurial firm setting), and within-firm investments in innovative ideas (established firm setting).

We believe that this article contributes a fresh lens to the discourse on the role of AI in modern organisations. The rise of AI can significantly augment decision processes and expert judgement, as clearly articulated in the previous section. It can also ease the administrative demands discussed immediately above. At the same time, we are cognisant that AI prediction is not without, sometime significant, limitations. For example, AI is still struggling with some lack of transparency about how the prediction is calculated. On the bright side, the growing field of Explainable AI (XAI) is engaged in significant meaningful work towards alleviating this 'black box' concern (e.g. Arrieta et al. 2020 for a review). Future work can build on these conversations to chart the applicability of experts, randomisation and AI across different settings, and explore the ideal mix for organisations and society as a whole.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

Agrawal, A. J. Gans, and A. Goldfarb. 2018. "Prediction Machines: The Simple Economics of Artificial Intelligence." *Harvard Business Review Press*.

Agrawal, A., G. Joshua, and G. Avi. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.

Akerlof, G. A.1978. "The market for "lemons": Quality uncertainty and the market mechanism." In *Uncertainty in economics*, 235–251. Academic Press.

Arrieta, A. B., N. Diaz-Rodrigues, J. Del Ser, A. Bennetot, S. Tabik, A. Barbado, S. Garcia, et al. 2020. "Explainable Artificial Intelligence (XAI): Concepts, Taxonomies." *Opportunities and Challenges Towards Responsible AI, Information Fusion* 58:82–115. https://doi.org/10.1016/j.inffus.2019.12.012.

Bacon-Gerasymenko, V., and J. P. Eggers. 2019. "The Dynamics of Advice Giving by Venture Capital Firms: Antecedents of Managerial Cognitive Effort." *Journal of Management* 45 (4): 1660–1688. https://doi.org/10.1177/0149206317721626.

Balachandra, L., T. Briggs, K. Eddleston, and C. Brush. 2019. "Don't Pitch Like a Girl!: How Gender Stereotypes Influence Investor Decisions." *Entrepreneurship Theory and Practice* 43 (1): 116–137. https://doi.org/10.1177/1042258717728028.

Bapna, S., and M. Ganco. 2021. "Gender Gaps in Equity Crowdfunding: Evidence from a Randomized Field Experiment." *Management Science* 67 (5): 2679–2710. https://doi.org/10.1287/mnsc.2020.3644.

Baum, J. A. C., and B. S. Silverman. 2004. "Picking Winners or Building Them? Alliance, Intellectual, and Human Capital as Selection Criteria in Venture Financing and Performance of Biotechnology." *Journal of Business Venturing* 19 (3): 1–38. https://doi.org/10.1016/S0883-9026(03)00038-7.

Burkhard, B., C. Siren, M. van Essen, D. Grichnik, and D. A. Shepherd. 2022. "Nothing Ventured, Nothing Gained: A Meta-Analysis of CEO Overconfidence, Strategic Risk Taking, and Performance." *Journal of Management* 49 (8): 1–38. https://doi.org/10.1177/01492063221110203.

Camerer, C. F. and E. J. Johnson. 1991. "The process-performance paradox in expert judgment: How can experts know so much and predict so badly?." In *Toward a general theory of expertise: Prospects and limits*, edited by K. A. Ericsson and J. Smith, 195–217. Cambridge, UK: Cambridge University Press.

Caves, R. E. 2000. *Creative Industries: Contracts Between Art and Commerce*. Cambridge, Mass: Harvard University Press.

Chang, S., J. P. Eggers, and D. Keum. 2022. "Bottleneck Resources, Market Relatedness, and the Dynamics of Organizational Growth." *Organization Science* 33 (3): 1049–1067. https://doi.org/10.1287/orsc.2021.1479.

Choudhury, P., S. Evan, and R. Agarwal. 2020. "Machine Learning and Human Capital Complementarities: Experimental Evidence on Bias Mitigation." *Strategic Management Journal* 41 (8): 1381–1411.

Criscuolo, P., L. Dahlander, T. Grohsjean, and A. Salter. 2017. "Evaluating Novelty: The Role of Panels in the Selection of R&D Projects." *Academy of Management Journal* 60 (2): 433–460. https://doi.org/10.5465/amj.2014.0861.

Criscuolo, P. L. Dahlander, T. Grohsjean, and A. Salter. 2021. "The Sequence Effect in Panel Decisions: Evidence from the Evaluation of Research and Development Projects." *Organization Science 2021* 32 (4): 987–1008.

Danthi, N., C. O. Wu, P. Shi, and M. Lauer. 2014. "Percentile Ranking and Citation Impact of a Large Cohort of National Heart, Lung, and Blood Institute-Funded Cardiovascular R01 Grants." *Circulation Research* 114 (4): 600–606. https://doi.org/10.1161/CIRCRESAHA.

Danziger, S. J. L., and L. Avnaim-Pesso. 2011. "Extraneous Factors in Judicial Decisions." *Proceedings of National Academy of Sciences* 108 (17): 6889–6892. https://doi.org/10.1073/pnas.1018033108.

Dushnitsky, G. 2012. "Corporate Venture Capital in the Twenty-First Century: an Integral Part of Firms' Innovation Toolkit." In *The Oxford Handbook of Venture Capital, Oxford Handbooks*, edited by D. Cumming. Oxford Academic. Accessed 15 Nov 2023. https://doi.org/10.1093/oxfordhb/9780195391596.013.0006.

Dushnitsky, G., and S. Sarkar. 2022. "Here Comes the Sun: The Impact of Incidental Contextual Factors on Entrepreneurial Resource Acquisition." *Academy of Management Journal* 65 (1): 66–92. https://doi.org/10.5465/amj.2019.0128.

Eblen, M. K., R. M. Wagner, D. RoyChowdhury, K. C. Patel, and K. Pearson. 2016. "How Criterion Scores Predict the Overall Impact Score and Funding Outcomes for National Institutes of Health Peer-Reviewed Applications." *PLoS ONE* 11 (6): e0155060. https://doi.org/10.1371/journal.pone.0155060.

Eliashberg, J., and S. M. Shugan. 1997. "Film critics: Influencers or predictors?" *Journal of Marketing* 61 (2): 68–78. https://doi.org/10.1177/002224299706100205.

Engelen, A., C. Neumann, and C. Schwens 2015. ""Of Course I Can": The Effect of CEO Overconfidence on Entrepreneurially Oriented Firms." *Entrepreneurship Theory and Practice* 39 (5): 1137–1160. https://doi.org/10.1111/etap.12099 .

Ewens, M., R. Nanda, and M. Rhodes-Kropf. 2018. "Cost of Experimentation and the Evolution of Venture Capital." *Journal of Financial Economics* 128 (3): 422–442. https://doi.org/10.1016/j.jfineco.2018.03.001.

Eyal, G. 2019. *The Crisis of Expertise*. John Wiley & Sons.

Fang, F. C., A. Bowen, and A. Casadevall 2016. "NIH Peer Review Percentile Scores are Poorly Predictive of Grant Productivity" *Elife 2016* 5: e13323. https://doi.org/10.7554/eLife.13323.

Fang, F. C., A. Bowen, and A. Casadevall. 2016. "NIH Peer Review Percentile Scores are Poorly Predictive of Grant Productivity."*Elife 2016* 5: e13323. https://doi.org/10.7554/eLife.13323.

Franzoni, C., B. Anders, and P. Stephan. 2023. “Implicit Criteria Weighting and Evaluators’ Background in Grant Peer Review.” Unpublished woeking paper.

Franzoni, C., H. Sauermann, and D. Di Marco. 2023, August 15. “When Citizens Judge Science: Social Impact, Representativeness, and Personal Interest.” Available at SSRN. 2023(1): https://doi.org/10.5465/AMPROC.2023.13860abstract.

Franzoni, C., P. Stephan, and R. Veugelers. 2022. “Funding Risky Research.” *Entrepreneurship and Innovation Policy and the Economy* 1:103–133. https://doi.org/10.1086/719252.

Ginsburgh, V. 2003. “Awards, Success and Aesthetic Quality in the Arts.” *Journal of Economic Perspectives* 17 (2): 99–111. https://doi.org/10.1257/089533003765888458.

Gompers, P. A., and J. Lerner. 2004. *The Venture Capital Cycle*. 2nd ed. Cambridge: MIT Press.

Hegde, D. and J. Tumlinson. 2014. “Does Social Proximity Enhance Business Partnerships? Theory and Evidence from Ethnicity's Role in US Venture Capital.” *Management Science* 60 (9): 2355–2380.

Hellmann, T., and M. Puri. 2002. “Venture Capital and the Professionalization of Start-Up Firms: Empirical Evidence.” *The Journal of Finance* 57 (1): 169–197. https://doi.org/10.1111/1540-6261.00419.

Huang, L., and J. L. Pearce. 2015. “Managing the Unknowable: The Effectiveness of Early-Stage Investor Gut Feel in Entrepreneurial Investment Decisions.” *Administrative Science Quarterly* 60 (4): 634–670. https://doi.org/10.1177/0001839215597270.

Ilan, Y. and D. P. Foster. 1995. “Graininess of Judgment Under Uncertainty: An Accuracy-Informativeness Trade-Off.” *Journal of Experimental Psychology: General* 124 (4): 424.

Jane, W. 2023. “Innovation for Dummies: Exploring the Role of Metrics in Automotive Safety”, working paper.

Jones, B.2009. “The Burden of Knowledge and the Death of the Renaissance Man: Is Inno- Vation Getting Harder?.” *The Review of Economic Studies* 76 (1): 253–281.

Klibanoff, P. M. Marinacci, and S. Mukerji. 2005. “A Smooth Model of Decision Making Under Ambiguity.” *Econometrica* 73 (6): 1849–1892.

Knight, F. H.1921. *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin Company.

Larrick, R. P., and D. C. Feiler. 2015. “Expertise in decision making.” *The Wiley Blackwell Handbook of Judgment and Decision Making* 2:696–721. https://doi.org/10.1002/9781118468333.

Lee, C. J. 2015. “Commensuration Bias in Peer Review.” *Philosophy of Science* 82: 1272–1283.

Lerner, J. and R. 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): 237–261.

Levinthal, D. A. 2017. “Mendel in the C-Suite: Design and the Evolution of Strategies.” *Strategy Science* 2 (4): 282–287. https://doi.org/10.1287/stsc.2017.0047.

Li, D. and L. Agha. 2015. “Big Names or Big Ideas: Do Peer-Review Panels Select the Best Science Proposals?.” *Science* 348 (6233): 434–438.

Liu, C. 2020. *Luck: A Key Idea for Business and Society*. Oxford, UK: Routledge.

Mahajan, J. 1992. “The Overconfidence Effect in Marketing Management Predictions.” *Journal of Marketing Research* 29 (3): 329–342. https://doi.org/10.1177/002224379202900304.

Maula, M. V. T. Keil, and S. A. Zahra. 2013. “Top Management’s Attention to Discontinuous Technological Change: Corporate Venture Capital as an Alert Mechanism.” *Organization Science* 24 (3): 926–947.

Mellers, B., E. Stone, P. Atanasov, N. Rohrbaugh, S. E. Metz, L. Ungar, M. M. Bishop, M. Horowitz, E. Merkle, and P. Tetlock. 2015. “The Psychology of Intelligence Analysis: Drivers of Prediction Accuracy in World Politics.” *Journal of Experimental Psychology: Applied* 21 (1): 1–14. https://doi.org/10.1037/xap0000040.

Merton, R. 1973. *The Sociology of Science: Theoretical and Empirical Investigations*. University of Chicago Press.

Merton, R. K. 1957. “Priorities in Scientific Discovery: A Chapter in the Sociology of Science.” *American Sociological Review* 22 (6): 635–659.

Morgan, M. G. 2014. “Use (And Abuse) of Expert Elicitation in Support of Decision Making for Public Policy.” Proceedings of the National Academy of Sciences 111(20):7176–7184.

Mumpower Jeryl, L. and R. S. Thomas. 1996. "Expert Judgement and Expert Disagreement." *Thinking & Reasoning* 2 (2/3): 191–211.

Philip, E. T. 2017. *Expert Political Judgment: How Good is It? How Can We Know?*. Princeton University Press.

Pope, D. G., and J. R. Sydnor. 2011. "Implementing Anti-Discrimination Policies in Statistical Profiling Models." *American Economic Journal: Economic Policy* 3 (3): 206–231. https://doi.org/10.1257/pol.3.3.206.

Reinstein, D. A., and C. M. Snyder. 2005. "The Influence of Expert Reviews on Consumer Demand for Experience Goods: A Case Study of Movie Critics." *Journal of Industrial Economics* 53 (1): 27–51. https://doi.org/10.1111/j.0022-1821.2005.00244.x.

Simcoe, T. S., and D. M. Waguespack. 2011. "Status, Quality, and Attention: What's in a (Missing) Name?" *Management Science* 57 (2): 274–290. https://doi.org/10.1287/mnsc.1100.1270.

Spence, M. T., and M. Brucks. 1997. "The Moderating Effects of Problem Characteristics on Experts' and Novices' Judgments." *Journal of Marketing Research* 34 (2): 233–247. https://doi.org/10.1177/002224379703400204.

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

Stewart, T. R., P. J. Roebber, and L. F. Bosart. 1997. "The Importance of the Task in Analyzing Expert Judgment." *Organizational Behavior and Human Decision Processes* 69 (3): 205–219.

Svenson, O. 1981. "Are we all less risky and more skillful than our fellow drivers?" *Acta Psychologica* 47 (2): 143–148. https://doi.org/10.1016/0001-6918(81)90005-6.

Teodoridis, F. 2018. "Understanding Team Knowledge Production: The Interrelated Roles of Technology and Expertise." *Management Science* 64 (8): 3625–3648. https://doi.org/10.1287/mnsc.2017.2789.

Tetlock, P. E. .2005. *Expert Political Judgment: How Good is It? How Can We Know?*. Princeton University Press.

Tetlock, P. E., and D. Gardner. 2015. *Superforecasting: The Art and Science of Prediction*. Crown Publishers/Random House.

Vakili, K., F. Teodoridis, and M. Bikard. 2022. "Detrimental Collaborations in Creative Work: Evidence from Economics." *Organization Science* 33 (5): 1741–1755. https://doi.org/10.1287/orsc.2021.1501.

Vakili, K., F. Teodoridis, and M. Bikard. 2022. "Detrimental Collaborations in Creative Work: Evidence from Economics." *Organization Science* 33 (5): 1701–2083. https://doi.org/10.1287/orsc.2021.1501.

Wang, D. and A. L. Barbabasi. 2021. *"The Science of Science"*. Cambridge University Press.

Woronkowicz, J., D. C. Joynes, P. Frumkin, A. Kolendo, B. Seaman, R. Gertner, and N. Bradburn. 2012. "Set in Stone: Building America's New Generation of Arts Facilities, 1994-2008."

Zhang, J., P. K. Wong, and Y. P. Ho. 2016. "Ethnic Enclave and Entrepreneurial Financing: Asian Venture Capitalists in Silicon Valley." *Strat. Entrepreneurship J* 10: 318–335. https://doi.org/10.1002/sej.1230.

Zuckerman, E. W. 1999. "The Categorical Imperative: Securities Analysts and the Illegitimacy Discount." *American Journal of Sociology* 104 (5): 1398–1438. https://doi.org/10.1086/210178.