Working Paper 24-065

Business Experiments As Persuasion

Orie Shelef
Rebecca Karp
Robert Wuebker

Harvard Business School

Business Experiments As Persuasion

Orie Shelef
University of Utah

Rebecca Karp
Harvard Business School

Robert Wuebker
University of Utah

Working Paper 24-065

Funding for this research was provided in part by Harvard Business School.

Business Experiments As Persuasion

Orie Shelef[*] Rebecca Karp[†] Robert Wuebker[‡]

March 2024

ABSTRACT

Much of the prior work on experimentation rests upon the assumption that entrepreneurs and managers use—or should optimally adopt—a "scientific approach" to test possible decisions before making them. This paper offers an alternative view of experimental strategy, introducing the possibility that at least some business experiments privilege persuasion over generating unbiased information. In this view, actors may craft experiments designed to gain support for their ideas, even if doing so reduces the informativeness of the experiment. However, decision-makers are not naïve—they are aware that the results they are reviewing may be the product of a curated information environment. Using a formal model, this paper shows that under a wide range of conditions, actors prefer to enact a less than fully informative experiment designed to persuade—even when a fully-informative experiment is feasible at the same cost.

Keywords: Experimentation, Decision-Making, Entrepreneurship, Strategy, Bayesian Persuasion

JEL Codes: L26, M13, O31

Acknowledgements: The authors would like to thank: Jay Barney, Hart Posen, Chris Rider, Yue Maggie Zhou, and participants at University of Michigan's Ross School of Business Strategy Brown Bag and the David Eccles School of Business, University of Utah, Strategy and Entrepreneurship Brown Bag, and unnamed interviewees engaged in designing and interpreting experiments.

[*]Department of Entrepreneurship and Strategy, David Eccles School of Business, University of Utah
[†]Strategy Unit, Harvard Business School
[‡]Department of Entrepreneurship and Strategy, David Eccles School of Business, University of Utah

Introduction

There is growing consensus that experiments improve the way organizations learn about the performance of an idea, innovation, business practice, model, or routine (Gans, Stern, and Wu, 2019; Thomke, 2020; Zellweger and Zenger, 2023). This "scientific approach" to decision-making (Felin, Gambardella, Novelli, and Zenger, 2023; Murray and Tripsas, 2004) shows that experimental techniques, including (among many others) iterative hypothesis testing (Camuffo, Cordova, Gambardella, and Spina, 2020), A/B split tests (Koning, Hasan, and Chatterji, 2022) and piloting (Edmondson, Bohmer, and Pisano, 2001), can help many organizations improve their performance. As a result, there are increasing calls to make experimentation an integral part of organizational life—creating a "culture of experimentation...[where] anyone...can conduct or commission a test, all experiments are done ethically, data trumps opinion...and managers embrace a new model" (Thomke, 2020, p. 42).

However, a growing number of studies reveal that actors use business experiments for more than an aid to managerial decision making. For example, an entrepreneur may conduct an expensive, uninformative experiment for the purpose of enrolling stakeholders such as an early investor, potential first customer, or critical hire (Shelef, Wuebker, and Barney, 2023a). Beyond aiding a particular managerial decision, experiments are also used to identify previously-unconsidered critical decisions and market opportunities (Ehrig and Schmidt, 2022; Pillai, Goldfarb, and Kirsch, 2020). And, sometimes, actors explicitly design experiments intended to persuade a decision-maker to adopt (or reject) new technologies (Karp, 2023). Designing experiments to persuade might profoundly shift when and how actors experiment—that is, experimental strategy. Yet, little theorizing considers the role that these kinds of persuasive experiments might play in strategic decision making.

Drawing upon a formal model and qualitative insights, this paper explores how introducing the possibility of experiments designed to persuade, rather than merely inform, shapes an actor's experimental strategy. The results of the model show that, under a wide range of conditions, the optimal experiment is designed to persuade rather than maximize information—in other words, the optimal business experiment is a simulacrum of the scientific approach. Moreover, this result is not driven by the cost of information, any actor being fooled about how the information environment has been designed, or an actor manipulating experimental results. Further, the actor being persuaded by such an experiment is not harmed—and may benefit—from that experiment.

This paper begins by briefly discussing the literature on business experimentation and shows how this work fails to examine certain uses of business experiments that sit adjacent to the scientific approach to experimentation that dominates prior work. The paper then draws on the theoretical literature on Bayesian Persuasion (Kamenica and Gentzkow, 2011), developing a formal model to illustrate how and when experiments should be used to persuade. The paper concludes by discussing some of the implications of a view of experimental strategy in which sometimes business experimentation should scrupulously adhere to the scientific method, and other times where it is best that it only pretends to do so.

Related Literature

Business Experimentation

A growing literature on business experimentation defines an experiment as any act intended to assess the viability of a business idea (e.g. Camuffo et al., 2020; Luca and Bazerman, 2021; Shelef et al., 2023a; Thomke, 2020). This broad definition encompasses a wide variety of approaches to experimentation including formal A/B split testing, randomized control trials, iterative hypothesis testing, customer development, future probes, focus groups, crowdfunded product development, the release of "minimum viable products", et cetera. Within this literature, there is growing interest in the general method underpinning search, learning and decision-making—business experimentation's "scientific approach" (Camuffo et al., 2020; Felin et al., 2023; Thomke, 2020; Zellweger and Zenger, 2023).

Broadly, the scientific approach to business experimentation conceptualizes economic actors as scientists (Felin and Zenger, 2017) making strategic decisions by "engaging in causally inferential action by forming beliefs, testing these beliefs, and responding to the feedback received" (Zellweger and Zenger, 2023, p. 379). This work also observes that the entrepreneurial process "might be best viewed as an application of the scientific method to entrepreneurship...building a theory of the business and empirically testing the validity of hypotheses derived from the theory" (Blank and Eckhardt, 2023, p. 2).[1] This literature aims to assist entrepreneurs in running experiments to

[1] Like academic scientists, who advance knowledge by composing and testing theories about what is presently unknown, this growing literature argues that actors similarly craft hypotheses and then "testing their beliefs, from

find "product/market fit inside a market, and pivot when your hypotheses are incorrect" (Gruber and Tal, 2017), transforming "speculations...into formal hypotheses and tested using the scientific method" (Shepherd and Gruber, 2021, p. 977).

Likewise, a stream of research on business experimentation focuses on how actors strategically design experiments to maximize their value as a decision-making aid (e.g., Agrawal, Gans, and Stern, 2021; Camuffo, Gambardella, Maccheroni, Marinacci, and Pignataro, 2022; Chavda, Gans, and Stern, 2024; Chen, Elfenbein, Posen, and Wang, 2022b; Contigiani and Levinthal, 2019; Ehrig and Schmidt, 2022; Gans, 2023; Gans et al., 2019; Li, Lee, and Mahoney, 2023). For example, Gans et al. (2019) and Ehrig and Schmidt (2022) argue that not all information can be generated from experimentation, and as a result actors should design experiments that focus on information that is relevant for their particular decision. Agrawal et al. (2021) argues that validating ideas sometimes requires that actors tune their experiment such that it reveals information about one aspect of an idea, even at the expense of information about another aspect of that idea. Camuffo et al. (2022), Bigelow, Shelef, Wuebker, and Barney (2024) and Gans (2023) all show that less promising experiments may sometimes be preferred when such experiments are more relevant to the decision at hand.

While the findings from the literature on business experimentation underscore the strategic nature of experimental decisions (Leiblein, Reuer, and Zenger, 2018), they also imply many other opportunities for experimentation beyond its use as a decision-making aid—in particular, that some actors may prefer to design an experiment to achieve outcomes beyond revealing the most information at the least cost. Indeed, a growing number of studies suggest that experiments often do much more than serve as a workhorse for improving managerial decision-making. For example, Shelef et al. (2023a) shows that the presence of Heisenberg effects can cause an actor to prioritize an experiment's Heisenberg effects over its potential to generate information. Gans (2022) develops a two-actor model of experimentation where the actors have different priors, and shows that the experimenter sometimes wants to an experiment convince a resource provider of the experimenter's beliefs, rather than reveal the most information most useful to the experimenter. Karp (2023) finds that, when choosing between several potential valid experimental designs, actors sometimes select the experiment where the results—although not as informative—increase the likelihood that

which they progressively learn...and update accordingly" (Zellweger and Zenger, 2023, p. 380).

the experiment induces the decision that the experimenter desires. Finally, our interviews with practitioners reveal stories of sub-optimal experiments conducted for the benefit of (and, sometimes, the behest of) a stakeholder such as an early investor, potential first customer, or critical hire.[2]

With that said, there is also a voluminous theoretical and empirical literature which suggests that it is highly unlikely that managers, investors, and key stakeholders are unaware of the possibility of actors using experimentation to persuade, rather than inform. In a world where experiments can both inform and persuade, decision-makers are likely to scrutinize experimental designs with a critical eye. As a consequence, those charged with evaluating experimental results and making decisions based on them may engage in a form of "counter-strategy" where they attempt to extract useful conclusions from the results of experiments that others have carefully crafted.[3] In a loose sense, some business experimentation may be *kayfabe*—a ritualized performance of the scientific method collectively enacted by both the experiment designer and the decision-maker.[4]

Bayesian Persuasion

Taking a step back and thinking about the question of experimentation more broadly and abstractly—how should we think about information environments where one actor has the ability to influence the design of the information environment, and the other actor knows that the results from any experiment has been subjected to this design process? The general formulation of this kind of "kayfabe experimentation" might be: an informed actor has to design an information environment that discloses enough of the right amount of information to influence the behavior of

[2]For example, following the recommendations to "get out of the building" (Blank and Dorf, 2020; Sarasvathy, 2009) a biotech startup we interviewed met with a potential customer that promised significant orders if their needs were met. The team subsequently conducted a series of experiments to modify their drug formula, diverting resources away from broader experimentation. This choice led to a narrow, customer specific product with limited market appeal. In an interview with the founder of a drone company, he shared the company tested a SaaS business model and ultimately used it at the behest of a investor interested in attracting more investment. Though popular, the business model was inappropriate for their business. In another example, a Fintech startup we interviewed hastily incorporated product development experiments containing elements of the blockchain protocol into their financial services platform in an explicit attempt to secure investment from a seed-stage investor with capital specifically allocated to Web3 companies.

[3]For example, in one interview we conducted, one venture capitalist remarked that their firm never believes a word that entrepreneurs say or any numbers in the presentation they share.

[4]Kayfabe, a term originating in American professional wrestling refers to the portrayal of events, relationships, and actions as "real," even though they are staged or scripted. Today, fans, wrestlers, promoters, and the media are largely aware of the scripted nature of professional wrestling—but there is a general agreement by all involved to not "break kayfabe"—i.e., to collectively pretend that that what is occurring in and around the ring is real. More broadly, the concept of kayfabe can be applied beyond wrestling to any situation where there's a suspension of disbelief or a maintained illusion for the sake of entertainment or another purpose.

a self-interested decision-maker.

Kamenica and Gentzkow (2011) inaugurated a literature in theoretical economics—Bayesian Persuasion—that explores this question. In particular, this literature focuses on a information designer trying to influence the actions of a decision-maker by strategically designing how information is collected, analyzed, and characterized. In this framework both the information design and the information that design reveals are known to the decision-maker, and thus, their decisions reflect accurate Bayesian updating. Said another way, the information designer is attempting to influence—but not deceive—through the design of the information environment, and the decision-maker is a *smark*,[5] aware of the possibility that the information environment has been curated to shape their beliefs and subsequent actions.[6] Kamenica and Gentzkow (2011) demonstrate that across a very broad set of environments that a designer, can, in fact, shape beliefs and thus subsequent actions. Further, they highlight how optimally shaping actions requires designing an experiment that reveals enough information to be credible, while still trying to shape the decision-maker's actions.

The Bayesian Persuasion literature is large and active, spawning work in two main streams (for a review, see Kamenica, 2019). The first stream builds increasingly complicated economic theory relaxing assumptions and refining mathematical techniques used in calculating the optimal information design. Studies in this literature consider (among other topics) asymmetric information (Hedlund, 2017; Perez-Richet, 2014), data manipulation (Lipnowski, Ravid, and Shishkin, 2022) and multiple decision-makers (for a review, see Bergemann and Morris, 2019). A second stream extends the basic model to varied settings including securities design (Szydlowski, 2021), genetic testing (Schweizer and Szech, 2018), employee feedback (Habibi, 2020; Smolin, 2021), bank stress-testing (Goldstein and Leitner, 2018), drug approval (Kolotilin, 2015), platforms (Romanyuk and Smolin, 2019), and private equity (Monnet and Quintin, 2017).[7]

[5]In the context of professional wrestling, the terms "smark" and "mark" refer to two different types of fans, each with a distinct attitude and perspective towards the sport. In wrestling, a *mark* is a fan who believes that the performances and storylines are real or at least buys into the scripted nature of the events as if they were real. In contrast, a *smark*—a portmanteau of "smart mark"—a fan who is knowledgeable about the inner workings of professional wrestling, aware that the event and storylines are scripted and choreographed.

[6]While there is debate in the strategy and entrepreneurship fields about restricting attention to settings where beliefs and updating are possible (Bylund and Packard, 2022; Ehrig and Foss, 2022; Packard, Clark, and Klein, 2017; Zellweger and Zenger, 2022, 2023), formal modeling, including in the Bayesian persuasion literature, focuses on those.

[7]In addition to adding strategy and entrepreneurship to the applications of Bayesian Persuasion, this paper also responds to calls within that core literature to "synthesize existing results, rather than simply add" to the set of assumptions made in the original formulation and relaxed in subsequent work (Kamenica, 2019).

There has been limited application of this framework within strategy and entrepreneurship.[8] This paper draws on the theoretical Bayesian Persuasion literature to explore how and when experiments should be used to persuade. First, a stylized model generates intuition of how a founder might persuade a venture capitalist. Then we develop a general model considering an experiment designer and a decision-maker. After general results from the model, we consider several extensions particularly relevant to strategy and entrepreneurship.

Stylized Model

Consider the challenge of a founder attempting to convince a venture capitalist that their idea is promising enough to warrant investment.[9] All things equal, when the founder's idea is good, providing evidence about the idea to the investor will tend to help the entrepreneur's pitch. When the idea is bad, revealing evidence about the idea will not forward the founder's cause. Can the founder structure their experimentation to increase the probability of investment? Kamenica and Gentzkow (2011) propose that the answer to this question is "yes".

To illustrate why, suppose the venture capitalist must choose to either *not invest* or *invest* in an idea. Ideas are either *good* or *bad*. The venture capitalist profits from making the "correct" investment decision—we can suppose they earn $g - i$ for investing in good ideas, $-i$ for investing in bad ideas, and 0 for not investing. The founder values investment regardless of their idea, but especially so if their idea is good—that is, they earn $G + I$ if the venture capitalist invests in a good idea, I if the venture capitalist invests in a bad idea, and 0 if the venture capitalist does not invest. The founder and the venture capitalist both share a prior belief that the probability the idea is good is θ. We assume that $\theta g - i < 0$—that is, the venture capitalist will not invest absent more information.

The founder conducts a costless experiment whose full design and outcome is visible to the venture capitalist. We formalize an experiment as a distribution $\pi(\cdot|good)$ and $\pi(\cdot|bad)$ on some set of experimental results. The entrepreneur designs π and must truthfully report the experimental results to the venture capitalist. If there is no communication, no experimentation, or, equivalently,

[8] The only instance of Bayesian Persuasion in the strategy and entrepreneurship literature that we could find is a remark in Ehrig and Schmidt (2022, p. 1314), who offer a sentence suggesting that future work might integrate their formal model of theory-based learning with the Bayesian Persuasion literature.

[9] This stylized model is intentionally built on that of Kamenica and Gentzkow (2011).

if π is completely uninformative, the venture capitalist will not invest. If the founder designs a fully informative experiment—one that fully reveals the value of the idea—the venture capitalist invests with probability θ. However, in this paper, we show the founder can do better by designing a *less informative* experiment!

Because there are only two possible actions—invest or not invest—we can think about the possible experiments as sending two signals: "success" or "failure". Because there are also only two possible states of the world—good or bad—we can think about the possible experimental design as four probabilities: $\pi(failure|bad)$, $\pi(success|bad)$, $\pi(failure|good)$, and $\pi(success|good)$. The fully informative experiment sets these probabilities to the extreme: $\pi(failure|bad) = 1$, $\pi(success|bad) = 0$, $\pi(failure|good) = 0$, and $\pi(success|good) = 1$. The intuition driving our result is that the founder can do better than the fully informative experiment if they contaminate the "success" signal—designing the experiment such that not only good ideas send the success signal, but also some bad ideas. This is effective despite that the venture capitalist is aware of the contamination—they see the experiment design, understand that contamination is happening, and understand that the "success" experimental result is contaminated—but the presence of a small amount of bad ideas does not cause the venture capitalist to stop investing. This idea—that the founder can design an experiment that leads the venture capitalist to take an action the founder likes more often by designing an experiment that intentionally contaminates the signal that leads to the action the founder likes—is common to the more general and abstract model we present.

Formally, we can think of contamination as ρ and that the experiment design sends the "failure" signal, when the idea is bad with probability $1 - \rho$, and otherwise, sends the "success" signal. That is, good ideas always succeed, and so do ρ share of bad ideas: $\pi(failure|bad) = 1 - \rho$, $\pi(success|bad) = \rho$, $\pi(failure|good) = 0$, and $\pi(success|good) = 1$. Because the founder would like the venture capitalist to invest as often as possible, the founder would like to have ρ be as high as possible. However, the design still has to ensure that the venture capitalist will decide to invest following a "successful" experiment. Thus, to find the optimal amount of contamination, the founder considers the venture capitalist's payoff. The venture capitalist's payoff from investing after success is $\theta(g - i) + (1 - \theta)\rho(-i)$. Increasing ρ makes the payoff from investing after success lower but leads to a higher probability of investing. Solving this optimization, the founder sets ρ such that this payoff is equal to the payoff from not investing 0. That yields $\rho^* = \frac{\theta(g-i)}{(1-\theta)i}$. As a result, the venture capitalist invests in $\frac{q}{i}\theta$ ideas, even though the venture capitalist knows that only θ ideas are worthy of investing. And the venture capitalist makes these investments knowing full well that the experiment was designed by the founder to maximize the probability of investment. Particularly, in this example, the venture capitalist knows that the experiment was designed such that some bad ideas succeed. Moreover, it is clear that the founder would, in fact, be willing to pay more—specifically, by $(1-\theta)\rho I$—for this less-than-fully informative experiment, even if that fully informative experiment were available for free.

Model Setup

In this section we describe a broad model of experimentation grounded in the foundational work by Kamenica and Gentzkow (2011). This model is abstract such that it can capture a wide range of experimental settings including by entrepreneurs and inside of organizations. We first describe the characteristics of our two main actors, then their misalignment, the environment they are in, what an experiment is, and the steps in their interaction.

Experiment Designer. The *designer*, akin to the founder in our stylized model, an actor with purview over experimental strategy such as an entrepreneur, a product manager, or product evaluator designs the information environment. The designer possesses a thorough understanding of the relevant trends and distributions for the particular setting and the preferences of other actors. This knowledge is critical as it informs the design of an appropriate experimental strategy and the accurate interpretation of results. Crucially, the designer does not directly control decisions made as an outcome of an experiment, but through their experimental design choices may shape the posterior beliefs of those who do.

Decision-Maker. The *decision-maker*, akin to the investor in the stylized model, or more broadly, an actor with decision rights, such as a key stakeholder or experimental committee member (Fabijan, Dmitriev, McFarland, Vermeer, Holmström Olsson, and Bosch, 2018; Kohavi and Thomke, 2017; Xifara, 2021), controls a decision that is relevant to their outcome and those of the experiment designer. The decision-maker observes and is fully informed about the design of the experiment, but has no role in its design or enactment. The decision-maker reviews results from the experiment designed by the experiment designer in light of the relevant trends and distributions for the particular setting, and then makes their decision. Their role lies not in designing experiments, but in interpreting the results.

The Decision. The decision that the decision-maker controls is one in which the decision-maker and experiment designer are not perfectly aligned. In other words, there are competing incentives and goals between the experiment designer and decision-maker that introduce a complex interplay of experimental strategy and actor motivations (Shelef, Wuebker, and Barney, 2023b). In particular, a few reasonable assumptions about the nature of the decision are made. First, there is some possible experimental result that would lead the decision-maker to make a different decision than what the decision maker would do absent that information—in other words, the decision-maker is responsive to information. Second, if a trivial amount of information is revealed, the decision-maker makes the same decision they would make absent any information.[10] Third, reflecting our assumption of competing goals between the decision-maker and experiment designer, the default decision the decision-maker makes is not the experiment designer's most preferred decision.[11] Fourth, recognizing that this misalignment persists even with full information, following a fully informative experiment, under some state, the experiment designer would prefer the decision-maker take an action that some other experimental result leads them to take. In sum, these assumptions capture that the decision-maker responds to information, the actors are not fully aligned, and, finally, that this misalignment (however mild) is not merely hypothetical.

Experiments. The experiment designer controls the information environment by deciding which of the possible underlying states of the world yield which signal and with what probability.[12] A fully informative experiment has a distinct signal for each possible state, while an uninformative experiment yields the same signal regardless of the state.[13] For example, one possible choice for an experimental design would be to decide a cutoff above which the "go" signal is sent, and below which the "no go" signal is sent. Experiments are truthful in the sense that if the designer designed a particular state to yield a particular signal, then that signal reliably occurs. However, those signals may not be trusted—just because the experiment designer labelled a particular signal "go" does not mean that the decision-maker will want to choose "go" in response. Instead, the decision-maker forms accurate posteriors as a result of the signal and then acts in their own best interest.

[10] Any categorical or binary decision meets this criteria such as business decisions to "go or not go," "invest or not invest," or "adopt or not adopt".

[11] This assumption ensures that the designer would like to shape the decision-maker's decision.

[12] We consider an experiment designer with more constrained experimental choices in an extension.

[13] This is effectively the same as not experimenting.

Game Play. Both the decision-maker and experiment designer begin with accurate, but not necessarily precise, prior beliefs. The experimental designer first designs the experiment. The experiment is then enacted. The decision-maker then observes the experiment and its results. Using their comprehensive understanding of the design of the experiment and the results, they incorporate this new information and update their beliefs.[14] The decision-maker then proceeds to make a decision that aligns with their revised perspective. This decision reflects a synthesis of the insights gained from the experiment and their original beliefs employed to make the decision the decision-maker most prefers.

As a result of the actors' competing goals and the experiment designer's control of experimental strategy, the designer may try to shape the decision of the decision-maker through the experiment. For example, the experiment designer may choose to sway the decision-maker towards a specific outcome by selecting an experimental design that aligns with their own interests. Their goal is to shape the decision-maker's perception in a manner that maximizes their own benefit, which might involve selectively generating information or framing data in a strategically favorable light. However, the experiment designer faces two important constraints on their behavior. First, the decision-maker is fully informed about the experimental strategy the designer enacts and the resulting information environment. Second, while the designer controls the experimental strategy, they do not control the experimental results. In other words, they decide what experiment to run, but cannot fake the data.[15] Or, in the language of contracting, the experimental results are verifiable. The question in this paper is if and how the experiment designer can shape the information environment to lead the decision-maker to make decisions that are more aligned with the experiment designer's goals.

Main Results

As implied by both Kamenica and Gentzkow (2011) and our stylized model above, it is always possible for an experiment designer to select an experiment that shapes the information environment in a way that yields a better outcome for the experiment designer. Further, we show there is always an experiment available to an experiment designer that would yield a better expected outcome for

[14]Given their common understanding of the environment, the design of the experiment, and the results from that experiment, both the decision-maker and experiment designer share the same view of the results—framed in terms of our model, both agree that objects resulting in a particular experimental result have a specific underlying distribution.
[15]We relax this assumption in an extension, allowing the experiment designer to also manipulate results.

the experiment designer than a fully-informative experiment would (or not experimenting at all). This result holds even though the experiment designer is not fooling the decision-maker about the nature of the experiment or the results it yields.

Proposition 1. *There always exists an experiment that induces the decision maker to make decisions that the experiment designer strictly prefers over a fully informative experiment and over an uninformative experiment.*

The stylized model above offers intuition about Proposition 1. With a fully informative experiment, the decision-maker only invests in projects they value. Proposition 1 shows that the experiment designer has an available experiment that, if enacted, would increase the likelihood that the decision-maker would invest. How could this be achieved? Instead of fully revealing all information, the experiment designer could design an experiment that yields "success" not only for the projects that the decision maker values, but also for some projects the decision-maker does not value. By providing less precise information, the decision-maker is then left with the choice of whether or not to invest in projects that yield "success". As long as there are not too many projects that the decision-maker does not value in the pool, the decision-maker invests in all projects that yield "success" from that experiment.

Thus, by designing the information environment—for which projects the experiment yields "success"—the experiment designer can increase the likelihood of investment. Notably, this is achieved despite the decision-maker having full knowledge that the experiment designer has intentionally selected a less informative experiment.

Corollary 1. *The experiment designer always chooses a less than fully informative experiment.*

Corollary 1 follows because the experiment designer always chooses an experiment of the sort that Proposition 1 shows always exists. There is never a world where the best experiment to design is the most informative experiment possible—it is always better to engage in kayfabe experimentation. Stepping back to the literature, while some information is valuable to the experiment designer, more information—even free information—is not always valuable. An experiment designer never wants to design a perfectly informative business experiment. It is not just that the "scientific approach" to experimentation is sometimes incomplete. The quest for the most informative experiment possible is never the optimal choice in situations where there are imperfectly aligned incentives between the designer and decision-maker.

Corollary 2. *The decision-maker is at least as well off as a result of the experiment designer's intentionally less-informative experiment than they would be absent experimentation.*

One would reasonably wonder whether the decision-maker is worse off since that the experiment designer has selected a less-than-fully informative experiment. At first blush, it would seem that they might indeed be worse off—after all, the experiment designer has shaped the information environment to influence the decision-maker's choice. Corollary 2 follows from the decision-maker's knowledge of the design of the experiment and subsequent choice. The decision-maker remains free to ignore the experimental results and make the decision they would have if there had been no experiment at all. Put differently, the information shaping the of the designed experiment induces the decision-maker to voluntarily make a decision; it does not fool or strong-arm the decision-maker.

Moreover, while Corollary 2 shows that the decision-maker is at least as well off, it can be the case that the designed experiment can be better for the decision-maker than no experiment. Indeed, in the stylized example the venture capitalist is neither harmed nor benefits from an optimally designed experiment. However, they may benefit. Consider, for example, if the designed experiment pools only a few projects the decision-maker does not value. In this situation, the decision-maker is making decisions that are nearly their ideal and much better than they would absent the experiment. Such an experiment could be the optimal experiment if the experiment designer did not value investment in some particularly poor quality ideas. Corollary 2 is consistent with a world in which decision-makers may not be excited about kayfabe experimentation, but are still willing to participate in the spectacle.

Corollary 3. *The decision-maker is strictly less well off as a result of the experiment designer's less-informative experiment than they would be from a fully informative experiment.*

A designed experiment is valuable for the experiment designer exactly because it induces the decision-maker to sometimes make a choice different than the decision-maker would if the experiment were fully informative. As a result, a designed experiment is not as good as a fully informative experiment would be from the perspective of the decision-maker. Despite this, the decision-maker is not fooled, does not ignore the experimental results, does not have any incentive to stop experimentation, and may still benefit from experimentation.

Implications for When to Experiment

The possibility of designing experiments increases the value of experimentation for the experiment designer—Proposition 1 says that designed experiments are more valuable to the experiment designer than fully informative ones. As a result, when considering when to experiment, the possibility of designing experiments makes it more likely that the experiment designer finds the value of experimentation exceeds its costs than if the experiment designer could only enact a fully informative experiment. Moreover, this holds even if the fully informative experiment is not more costly than the less informative, designed, experiment.

Proposition 2. *The possibility of designing business experiments that include persuasion elevates the value of experimentation for the experiment designer.*

This suggests that—for the experiment designer—the value of experimentation in the entrepreneurial process is higher than it otherwise would be, and as a result an experiment designer enacts more experiments.[16] Taken in view of the experimental strategy literature's question as to "when to experiment", allowing for the possibility of designing experiments to persuade means that the answer is to experiment more often than it would be in a world where a scientific approach to experimentation were the only possibility.[17]

Implications for How to Experiment

While Proposition 1 says that the experiment designer can always design an experiment that is better than a fully-informative one, there is almost always more than one experiment that an experiment designer would prefer over a fully informative experiment. The stylized model illustrates this—the designer can pool just a few projects in as "success" or can pool more projects in as "success". Any experiment that pools at least one project that the decision-maker does not value,

[16] We elide variation in the costs of different experiments. However, as long as experiments that are more informative are weakly more expensive, experimental cost does not impact our results.

[17] We elaborate on these and the decision-maker's welfare implications in the Discussion.

but not so many that the decision-maker ignores the experiment, is (for the experiment designer) more valuable than a fully-informative experiment.

This tension offers insight into which experiment an experiment designer should select. For the experiment designer, the best experiment contaminates "success" with as many projects the decision-maker does not value as possible until the decision-maker is essentially indifferent between investing following success or not. The experiment designer's solution of this tension helps describe their optimal choice.

Proposition 3. *Following an optimally designed experiment, if the decision-maker takes the experiment designer's least preferred action, the decision-maker knows that the decision is uncontaminated. Further, following an optimally designed experiment, if the decision-maker takes the experiment designer's most preferred action, the decision-maker knows that decision is contaminated.*

While the specifics of the optimal experimental design depends on the underlying details of settings, actions, and payoffs, Proposition 3 provides very useful characterization. In the stylized model where there are two possible actions—invest, or not invest—interpretation is easy. In an optimally designed experiment, "failure" is uncontaminated—no projects that the decision-maker values yield failure—and "success" is contaminated—there are some projects that yield success even though the decision-maker does not value them. Put another way, the best experiments are designed such that everyone agrees that "failure" is failure, but that there is some ambiguity about "success". *Ex post* ambiguity is thus an *intentional feature of optimally designed experiments*, not evidence of subpar experimentation.

Extensions

We now turn to extensions particularly relevant to entrepreneurship and strategy. First, we take seriously that the experiment designer knows something the decision-maker does not. Second, we consider an experiment designer that also values learning. Third, we recognize that when experimentation takes place, the particular decision-maker an experiment designer may wish to persuade has not yet been identified. Fourth, we realize that there may be *ex ante* constraints on experimental design. Finally, we consider the possibility that experimental results may be manipulated.

Asymmetric Information

In many settings, and in particular entrepreneurship, information asymmetry is rampant. Indeed, the possibility of information asymmetry between entrepreneurs and financiers is at the core of the conceptual literature in financial economics, driving much of the interaction between these actors. So far, the paper has assumed the experiment designer is as equally uncertain about their idea as the decision-maker. What happens if we relax this assumption and allow for the possibility that the experiment designer possesses private information?

Hedlund (2017) and Perez-Richet (2014) extend the setting of Kamenica and Gentzkow (2011) to consider when there is information relevant to the decision that the designer has, but the decision-maker does not—for example, when a founder knows more about the prospects of their business than the investor does prior to experimentation. Adding asymmetric information to the model presents significant complication to the equilibrium between the experiment designer and the decision-maker. Namely, to interpret the experimental results the decision-maker also forms a belief about what the experiment designer knows based on the experiment the designer designed. For example, did the designer design a very exacting test, but one in which they know a quirk of their product will lead it to succeed in that specific setting, but not in a more general test? These models generally have many equilibria and game theory has attempted to refine them.

At one level, the results of the basic model extend as is. Propositions 1 and 2 are maintained—there is always an equilibrium where, if the designer designs an experiment other than the one that is the result of the symmetric information game, the decision-maker believes the private information is very bad news, and thus the designer and the decision-maker ignore the private information.

However, the results of Hedlund (2017) go further. Even using the equilibria refinements in the economic literature, the spirit of Proposition 1 and Proposition 2 remain along with a question of strategic disclosure. In one type of possible refinement selected equilibria, some private information of the experiment designer is revealed by their choice of experiment. The decision-maker, by observing which experiment the designer has selected, has now gathered some insight about the private information that the experiment designer possesses. The results of Hedlund (2017) show that, in the equilibria that survive the additional restrictions, if the private information is sufficiently promising, the designer chooses a fully-informative experiment—effectively revealing their private information and getting the decision they prefer through the credible signal of a fully-informative experiment. In other words, an experiment designer with very promising private information always tells the whole truth—after all, there is no need to engage in kayfabe experimentation if a fully-informative experiment will always lead to the decision you want.

In contrast, experiment designers with less positive private information do not reveal the details of their private information. Instead, they pool with each other and enact the same sort of experiment described in Propositions 1 and 2—implicitly revealing by the fact that they did not design a fully informative experiment that their private information is not so promising. Notably, if there is very little private information then the equilibrium is essentially the persuading equilibrium of Proposition 1 and 2. With more private information designers with very promising private information may choose to distinguish themselves from other experimenters and projects by disclosing their private information and implementing a fully-informative experiment. Said another way, the presence of asymmetric information may change how persuasive experiments are designed, but it does not lead *all* experiment designers to *always* design fully informative experiments.

Learning and Persuasion

The base model in this paper assumes that the only benefit the experiment designer can get from an experiment is influencing the decision-maker. What if we allow for the possibility that, in addition to influencing the decision, the experiment designer also has another aim—learning more about their business idea. For example, the experiment designer may want to learn which of two go to market strategies to implement, how to improve the project, or whether to develop the project further or abandon the project altogether. What happens if we allow the experiment designer to engage in both learning, and persuasion, in the same experiment?

Some of the more straightforward implications about joint learning from an experiment have received treatment in the original formulation of Bayesian Persuasion and extensions (Kamenica, 2019). For example, if the possibility of learning from the experiment is valuable because it enables some future decision, and if a decision-maker other than the experiment designer controls that decision, then that learning is already embedded in the general Bayesian Persuasion frame and extensions that consider multiple decision makers. As a consequence, in those cases the results of Propositions 1 and 2 are unchanged.

We can also consider a few more variations of joint learning. For example, suppose that the experiment designer also controls a different decision: not only is the investor is deciding whether to invest or not invest; but also the experiment designer is deciding whether to add a feature to the offering or not. Suppose these are truly separate decisions: information about one is not relevant to the other, but information about each can be revealed from the same experiment. At that extreme, the designer can construct an experiment that is fully-revealing on the information relevant to their decision, and still implement the insights of Proposition 1 and 2 on the other dimension.[18]

If the decisions are not so separable, then the experiment designer may face a trade-off. Designing an experiment that reveals more information relevant to their decision may leak information that they may not want to surface to the decision-maker. While this may drive the experiment designer to design a more revealing experiment, as long as the decisions are not perfectly informative about each other there is still control over how informative it is with respect to the decision-maker's decision—even if the designer designs an experiment that is fully-revealing with respect to their own decision. Thus, the insights of Proposition 1 and 2 still influence experimental design. Moreover, the designer may choose to reveal less information relevant to their own decision to better shape the decision-maker's decision. If the decision-maker's decision is relatively unimportant to the designer, the designer may design an experiment that prioritizes information relevant to their decision instead of prioritizing persuading the decision-maker. If the decision-maker's decision is relatively important to the designer, the designer may designer an experiment that prioritizes persuasion.

Formally, we can view joint learning as replacing the decision-maker with a composite decision-maker—one that reflects the original decision-maker's decisions and the learning dependent decisions of the experiment designer. As long as there are a finite possible learning dependent actions for the experiment designer and a condition like that of assumption 4 holds—that there exists some fully revealed state under which the experiment designer would prefer the composite decision-maker take an action that the composite decision-maker actually takes under some other state—then the main results hold. This extended condition basically states that there is, at least possibly, enough misalignment relative to the value learning to the experiment designer that the experiment designer sometimes prefers to sacrifice the quality of their own learning dependent decision to persuade.

[18] For example, the experiment designer subdivide each signal from the experiment they would design absent learning into signals that also reveal the information relevant to their learning objects. If these are separate decisions, the decision-maker's decisions are the same following any of the subdivided signals as the original.

Future Decision-Maker

The model in this paper assumes that the experiment designer is already aware of the desires of the relevant decision-maker and desires to tailor their experiment to that particular decision-maker. What if we relax that assumption, and allow for the possibility that an experiment designer is crafting an experiment absent a particular decision-maker? For example, a founder or a nascent team could be running experiments at a very early stage, with a currently-unrealized aspiration that the results could be used to persuade key stakeholders in the future; or, that same founder or nascent team running experiments that are designed to secure growth capital from outside investors, but without a clear idea *ex ante* about who that investor will be? Said another way, should an experiment designer consider how future decision-makers might evaluate the results of an experiment that they are running today?

These examples are all versions of a broad literature connecting Bayesian Persuasion with multiple receivers and the economics literature on communication in games. Bergemann and Morris (2019) surveys and connects this work, showing that there are a wide variety of ways to model multiple decision-makers. Despite these complications, the insights of Proposition 1 and 2 remain.

For example, suppose that the experiment designer does not know which particular decision-maker will be making the decision, but understands that there are a variety of possible decision-makers and their likelihood and all decision-makers observe the same results. Kamenica (2019) notes that this is mathematically equivalent to a decision-maker that makes probabilistic decisions equal to the probabilities of the different decision-makers. As such, the insights of Proposition 1 and 2 persist, and the experiment designer will design their experiment with the average expected decision-maker in mind. Computing the exact optimum may be difficult because the literature not only considers the possibility that the decision-makers are probabilistic, but that the decision-makers may engage in complicated games, such as competing to be part of an investment syndicate, that shape who the decision-maker is and what they decide.[19]

To understand the impact of considering future decision-makers on the design of an experiment, consider our stylized example of a founder designing an experiment, but add the possibility of two

[19]For example,Bhaskar, Cheng, Ko, and Swamy (2016) show that this calculation may be NP-hard, while Cheng, Cheung, Dughmi, Emamjomeh-Zadeh, Han, and Teng (2015) show a set of these challenges that can be solved in polynomial time. That literature has proceeded to consider complications and solutions of calculating the optimal information design.

different potential investors each with different investment thresholds. At the time the experiment is designed, the founder does not yet know which investor's attention they will be able to attract enough for that investor to review their experimental results. The founder faces a trade-off: they can design an experiment that, if it has "successful" results will induce both investors to invest. Or, they can design a less stringent experiment that, if successful, will only convince the investor with the lower threshold to invest, but is successful more often. Both possible designs still persuade and are not fully revealing, and the choice between them reflects the trade-off between investment more often from a single, less stringent, investor or the possibility of investment from multiple investors. That same tension remains if there is, in fact, a wide variety of investors each with their own thresholds.

Kamenica (2019) also consider an alternative conceptualization of multiple decision-makers where the designer can design the experiment such that each decision-maker views a subset of the experimental results designed for them. This conceptualization allows the experiment designer to design an experiment inline with Proposition 1 and 2 for each decision-maker.[20] Indeed, the computational literature on Bayesian Persuasion allows for more complicated conceptualizations—situations where there are multiple versions of experimental results viewed by different sets of decision-makers, and that the decisions of those decision-makers also interact. While the computation of the best experimental design in that setting is indeed *hard*, our intuitions persist.

Constrained Experimental Design

The model in this paper assumes that experimental design is quite flexible, and the experiment designer can control all aspects of the signals an experiment generates. Here we consider three simplifications of the set of possible experiments the experimenter can design that reflect the experimental strategy choices found in prior work. In each we show that the desire to induce the decision-maker to act can lead the optimal experimental design to differ from the experimental design that prior work would recommend.

For example, some prior work suggests that, sometimes, the central experimental strategy choice is which idea to experiment on (Agrawal et al., 2021; Bigelow et al., 2024; Camuffo et al., 2022).

[20]Babichenko and Barman (2016) consider the computational challenges when the decisions of these separate decision-makers interact.

This work focuses on selecting which idea to experiment on in order to maximize information about whether and which idea to develop. However, this work does not consider the possibility of a decision-maker that may be distinct from the experiment designer. Suppose that we add to these models a decision-maker considering whether to invest in the venture that is considering two ideas. Like the stylized example, the venture benefits from the investment on top of the value of the idea. The presence of this decision-maker, without changing the available experiments, shows that optimal experimental strategy can change. For example, in the model of Agrawal et al. (2021) they show that when there are two viable go to market strategies, it is often best to experiment on the less promising one. The logic driving that choice is that doing so leads to faster unambiguous failure, which is valuable. However, that experimental strategy choice reduces the probability of the experiment demonstrating promising, but not certain, experimental results. If the decision-maker is willing to invest if results are promising, even if they are not certain, the actor may instead prefer an experimental strategy that increases the probability of middling results and reduces the probability of total failure over experiments that "fail fast".

Another common experimental design choice in implementing experimental strategy is the definition of "success". Do our results persist if the experimental designer is constrained such that the only experimental design choice they make is defining the level of performance that is considered "success"? Indeed, they do. In many cases the optimal unconstrained persuasive experiment for the designer to design can be viewed as a threshold defining success.[21] As a result, in these cases, the Propositions above remain unchanged if the designer only designs a success threshold. Viewing experimentation in this way provides further insight on how persuasion shapes experimental design. For example, consider a committee deciding the definition of success for an experiment that will decide whether to adopt a vendor's product. If the committee would like to induce the organization to adopt the product, they can set a low bar for success, and induce the organization to adopt even some ineffective products. On the other hand, if they would like the dissuade the organization from adopting the product, they can set a high bar for success, and induce the organization to not adopt some effective products. Karp (2023) describes such experimental persuasion. By designing the bar, the committee can influence the decision without needing the organization to commit to follow the experimental results or being confused about what the design of the experiment was.

[21] For example, if projects can be ranked and the decision-maker has two possible actions.

A final common experimental design choice debated in implementing experimental strategy is the precision or the amount of noise in an experimental design. If the only choice a designer has is how much noise to introduce into an experiment, would they still want to design a less-than-fully informative experiment? They might. Consider our stylized model. While the optimal experimental design in that setting is to contaminate success or lower the bar, the experiment designer would prefer an experiment with some noise over a fully revealing one. Let $\epsilon \in [0, 1/2]$ be the noise parameter that the experiment designer can select, and the possible experiment designs thus be: $\pi(failure|bad) = 1-\epsilon$, $\pi(success|bad) = \epsilon$, $\pi(failure|good) = \epsilon$, and $\pi(success|good) = 1-\epsilon$. The venture capitalist's payoff from investing after success is $\theta(g-i)(1-\epsilon)+(1-\theta)\epsilon(-i)$. The founders payoff as a result is $\theta(G+I)(1-\epsilon)+(1-\theta)\epsilon(I)$. As long as $\theta(G+I) < (1-\theta)I$, the founder benefits from increasing ϵ. And, as above, the venture capitalist will continue to invest following success as long as there is not too much noise. While noise is less attractive of a way to persuade than contamination—adding noise leads an investor to pass on some good deals—when success is rare enough or when the benefit to the founder of the idea being good when invested in is small enough, it is still preferable to design noisy experiments over fully informative ones.

Manipulating Results

Generally, the model in this paper assumes that experimental results are verifiable. In other words, in our base model, lying is impossible. What happens if we relax this assumption and allow for the possibility for the experiment designer to not just design the information environment, but to manipulate experimental results? The Bayesian Persuasion literature views situations like this as a version of partial commitment, and generally model this challenge as considering what if, after designing and running the experiment, the experiment designer, with some probability, can return the results of their choice, and the decision-maker does understand this probability (Fréchette, Lizzeri, and Perego, 2022; Lipnowski and Ravid, 2020; Lipnowski et al., 2022; Min, 2021). The results are quite surprising.

In this stream of work, the first conclusion is introducing the possibility of manipulated experimental results unambiguously makes the experiment designer worse off. The intuition driving this result is that the ability to manipulate the results makes experiment results less credible, and thus the decision-maker is less inclined to believe them. Since, absent the ability to manipulate results, the designer could design the experiment to create the belief the designer most wanted, introducing the possibility of manipulated results harms the experiment designer's credibility more than the potential benefit of manipulated results. This is clear in the extreme—if the the results could always be manipulated, then the experiment would contain no information, and the decision-maker would always ignore it.

Next, this stream of work concludes that introducing the possibility of manipulated experimental results may sometimes make the decision-maker better off. The intuition driving this result is that the ability to manipulate the results removes some amount of credibility from the designer, and this lack of credibility can limit the ability of the designer to influence the decision-maker's beliefs following very specific experimental results. In some sense, the lack of credibility reduces the range of possible experiments that would be viewed as credible by the decision-maker. In particular, it may preclude experiments with very unlikely results from being believed. This constraint is sometimes good news for the decision-maker because of the actions that would be induced if those very unlikely results were believed.

One final conclusion from the stream of work focusing on manipulating results considers what happens to these relationships as the frequency of manipulated results increases. The first conclusion above said that more frequency manipulation always makes things worse for the designer. The second conclusion said that increases in frequency can, at least, at small amounts, be good for the decision-maker. Lipnowski et al. (2022) demonstrates that the benefit of experimentation does not simply increment down for each small increase in the frequency of manipulation, but at some point falls off a cliff. In our stylized example, even if there is a high likelihood of manipulated experimental results, an experiment may be believed—the decision-maker realizes that the results may be manipulated but acts as if the experiment is true—if the experiment in question is sufficiently informative to shape decision-making. In such cases, the decision-maker invests in all of the projects they would ideally invest in, and accepts that manipulated results imply that they are also investing in some number of projects that they would otherwise not have invested in. However, a small increase in the frequency of manipulation can sometimes flip that decision—leading the decision-maker to ignore experimental results, even though those results are still sometimes true and thus the results are informative, because the frequency of manipulation makes investment not worthwhile. In such cases, the decision-maker never invests. Said another way, a small change in the frequency of manipulation can transition a market from following the results of experimentation despite manipulation to completely eschewing experimentation.

Front and center in the above insights about the impact of manipulation is the assumption that the decision-maker is aware of the frequency of manipulation, but remains unaware as to whether the particular experimental results they are reviewing have been subject to such manipulation. In the above models, if a designer can manipulate the results without a corresponding change in the decision-maker's understanding of the frequency, they would, of course, do so. In practice, this could lead to an arms race of sorts, where designers endeavor to improve their capability to manipulate experimental results in ways that do not alert the decision-maker, and decision-makers' beliefs about the frequency of manipulation evolve in response. Each individual experimental designer in a market benefits from their manipulations and obfuscations, but is also harmed by concomitant increases in the beliefs of decision-makers about the frequency and efficacy of those manipulations and obfuscations.

Discussion

As I observed more than once at Facebook, and as I imagine is the case in all organizations from business to government, high-level decisions that affected thousands of people and billions in revenue would be made on gut feel, the residue of whatever historical politics were in play, and the ability to cater persuasive messages to people either busy, impatient, or uninterested (or all three). - Antonio Garcia Martinez, *Chaos Monkeys: Obscene Fortune and Random Failure in Silicon Valley*

Our engagement with managers, entrepreneurs, and investors revealed to us that actors routinely design sub-optimal experiments to (among other outcomes) influence decision-making, enroll key stakeholders, and secure financial resources. Moreover, actors on "both sides of the table"—those designing the experiments, and those reviewing the results—recognized the possibility of experiments designed to persuade. Building on these insights, through a formal model grounded in the economics literature on Bayesian Persuasion, this paper suggests that experiments designed to persuade may not be merely an occasional part of business experimentation, but rather much more common than prior work admits. Our model shows that experiment designers prefer to enact persuasive experiments rather than pursue the "scientific approach" because such experiments are—in almost all cases—better.

This paper primarily contributes to a large and active literature on business experimentation, joining an experimental strategy literature showing that experiments can serve important strategic purposes beyond providing information about an idea, product, or practice at some cost (Agrawal et al., 2021; Bigelow et al., 2024; Camuffo et al., 2022; Chen et al., 2022b; Contigiani and Levinthal, 2019; Ehrig and Schmidt, 2022; Gans, 2023; Li et al., 2023; Shelef et al., 2023a). While much of the prior work on business experimentation takes as axiomatic that the goal of experimentation is to improve decision-making by delivering information, this paper introduces the possibility that experiments not only can be, but also *should* be used to persuade. Our results do not imply that the recommendations of prior work on business experimentation is incorrect. Rather, prior work is bounded by an assumption that only a single actor is engaged in experimentation. Our work releases that boundary and shows that when more than a single actor is involved in a endeavor, such as when entrepreneurs engage with other stakeholders to mobile resources (or believes that they might eventually need resources), the optimal business experiment always privileges persuasion over maximizing relevant information.

Our results reveal that there is always an experiment available to the experiment designer that, if selected, would be better for the experiment designer than a fully-informative experiment would. Thus, an experiment designer will always choose such an experiment—one that contains enough information to be believed, but designed to persuade—and, the actor responsible for evaluating the results of that experiment will be at least as well off as a result of the experiment designer's choice than they would be absent experimentation. From the perspective of the decision-maker, even an experiment that the decision-maker knows has been designed to persuade contains valuable information. Thus, the decision-maker, while ideally desiring a fully-informative experiment, does not "break kayfabe". She accepts the simulacrum of a scientific experiment crafted by the experiment designer and incorporates the information generated by that experiment in her decision-making.

While the base model in this paper makes a number of assumptions, we show in a series of extensions that our main results also hold when the experiment designer possesses private information; seeks to both learn and inform; conducts experiments in advance of knowing who the decision-maker will be; has limited flexibility in the design of an experiment; or manipulates the results of the experiment. These extensions demonstrate that the key insights of the paper persist, and also offer additional insight. For example, one practical insight is that in many common cases an experiment designer only needs to control the definition of "success" in order to optimally persuade.

Another practical insight revealed by this paper is that in situations in which the decision-maker is aware of the possibility that results are subject to manipulation, the ability to manipulate results does not actually help an actor—if that actor is able to design the experiment. Indeed, critics of experimentation focus attention on the potential for manipulated results, fabricated experiments, or selective disclosure. Moreover, critics worry that experimenters trying to influence decision makers want to manipulate results. This paper shows that actors can get all of the benefit that they would get from these and other kinds of experimental trickery—and more—simply by controlling how the experiment is designed and then credibly and accurately communicating the results.

Our overall approach contributes to aspirations to integrate stakeholders into experimentation in strategy and entrepreneurship (e.g., Chen, Elfenbein, Posen, and Wang, 2022a; Gans, 2022; Shelef et al., 2023b) by offering a model that has broad reach and applicability. While some work on experimentation in an entrepreneurial context takes as central that entrepreneurs have beliefs that differ from others (e.g., Gans, 2023) or are characterized by behavioral biases (e.g., Chen et al., 2022a; Shelef et al., 2023b), a strength of our approach is that, as part of a literature on experimental strategy, it is not limited to experimentation by individual entrepreneurs. Our approach is not restrictive, suitable to analyze experimentation by individual actors, group of actors operating within an organization, and experimentation by organizations as long as there is some misalignment between the experiment designer and a decision-maker. Further, our modeling approach in this paper is ecumenical—it does not, for example, require additional intellectual commitments such as overconfidence, risk aversion, or distinct perspectives that, even if accurate are more difficult to generalize beyond the individual (Åstebro, Herz, Nanda, and Weber, 2014).

One broad conclusion from this paper is that attempts to embed a scientific approach to experimentation in a social context has limits. As Merton (1979) would suggest, science is rarely value free. Much of the prior work in the business experimentation literature is implicitly about the pursuit of truth—"doing science". Thus, the literature's "solution" to the "problem" of poor business decision-making is to improve experimental designs, train the experimenter, or govern the experimental process to overcome organizational and cognitive frictions that encourage sub-optimal decision-making.

This paper recognizes that embedding experimentation in a social context introduces the possibility of experiments that mimic the scientific method but privilege persuasion. Given this possibility, the model in this paper demonstrates that this possibility is, actually, an eventuality. If a persuasive experiment exists (and, as is shown in this paper, one *always* exists) an experiment designer optimally selects such an experiment—even when a fully-informative experiment is also available. And the results of this experiment, despite having been designed to persuade, contains valuable information for the decision-maker. Said another way, business experiments are optimally designed to privilege persuasion—even if there were such a thing as a "scientific" experiment.

Further, the quixotic quest for fully-informative experiments may be a red herring in the entrepreneurial process. Prior work has long emphasized that due to cognitive bias, entrepreneurs may (among other things) design less informative experiments. As a result, some scholars exhort entrepreneurs to craft experiments carefully, studiously avoiding running experiments that tell them what they want to hear.[22] Our results show that, in many cases, such "poorly designed" experiments may be quite valuable precisely *because* they are designed to persuade—in other words, biased experimentation may, at least sometimes, not be a problem to solve. An optimistic founder may design an experiment that results in "success" more often than warranted. However, a founder engaged in optimal persuasion also designs an experiment that results in "success" more often than warranted. Likewise, a product manager who is pessimistic about incorporating a new technology into an existing offering (for example, adding self-driving features to passenger cars) may design an experiment that tells them what they want to hear (for example, that the technology does not have a perfect safety record). However, a founder engaged in optimal persuasion who prefers the technology not be adopted also designs an experiment with a too high bar.

Welfare Implications

Although the results from this paper suggest that while persuasive experimentation may indeed be the prevailing form of business experimentation, this is not a necessarily a problem to resolve. As a prominent venture investor we interviewed observed, "the fact that founding teams do this [design experiments to persuade] is a feature, not a bug". We show that even when the experiment designer uses experimentation strategically to shape the information environment in ways that

[22]"It doesn't matter how beautiful your theory is, it doesn't matter how smart you are. If it doesn't agree with experiment, it's wrong...The first principle is not to fool yourself—and you are the easiest person to fool." - Richard Feynman

favor a particular decision and the decision-maker is aware of the fact that they are being "sold something", the experiment designer is strictly better off than even a fully informative experiment, and the decision-maker is at least as well off as absent experimentation. As a result, if experiment designers engage in persuasive experimentation, the decision to experiment is still jointly optimal and improves aggregate welfare.

It may be helpful to understand the aggregate implications by viewing persuasive experiments as a mechanism to transfer welfare from decision-makers to experiment designers through experimentation. This creates an incentive for actors to engage in experimentation more than they otherwise would. Thus, the possibility of persuasive experimentation implies that some ideas will be experimented on that would not have been suitable for experimentation if only fully informative experiments was available. In other words, the possibility of persuasive experiments generates more, but less informative, experiments.

As a consequence, the net effect of persuasive experiments is, in aggregate, unclear. Returning to the stylized model, the possibility of persuasive experimentation means that more ideas get evaluated, but also more ideas that do not meet the investor's bar get funded. The impact of the extensive margin increase on the decision-maker may be positive, if, for example, they get some rents from those extra ideas experimented on. The net effect, even on the decision-maker is unclear. The impact of persuasive experiments on society at large is also unclear. The impact of the trade-off on society depends on how the investor's preferred decision relates to the socially optimal choice. If their marginal project is socially valuable, inducing more investment may be socially optimal. Likewise, the impact of the trade-off on society also depends on how the experiment designer's choice to experiment relates to the socially optimal choice. If their marginal decision to experiment generates socially valuable information, then experimenting more broadly may create more social value. In other words, neither persuasive experimentation nor fully-informative experimenting may maximize social outcomes, as under each, the actors make choices that maximize their own well-being.

Finally, persuasive experimentation also has implications what happens to different quality ideas in aggregate. Recall the intuition that persuasive experimental designs often are designed such that some bad ideas cannot be distinguished from good ideas. In other words, fully informative experimentation may reveal more low-quality ideas as bad than persuasive experiments would. Thus, we might expect actors induced to engage in fully informative experimentation would fail faster than actors engaged in persuasive experimentation (Camuffo et al., 2020; Leatherbee and Katila, 2020). However, this faster failure may not be accompanied by increased success conditional on not failing, and actors engaged in fully informative experimentation are less likely to get the investment they desire than actors engaged in persuasive experiments. Our results may help explain some of the empirical findings in the literature on entrepreneurial experimentation that persistently finds treatments that lead to faster failure but less consistent results on other outcomes. Introducing the possibility of persuasive experimentation suggests we should not be quick to conclude implications for either experiment designers or social welfare from faster failure rates. On the other hand, because of the tendency to not distinguish between good ideas, persuasive experimentation may be less likely to reveal exceptional ideas as truly exceptional.

Implications Beyond Experimentation

While our analysis has focused on the implications for experimental strategy this paper has several other broader implications. Indeed, prior work's definition of an experiment—any act intended to assess the viability of a business idea—and the possibility of misalignment between those who collect information and those who make some relevant decision in numerous settings, suggests the possibility that persuasion can be an important goal that leads to intentional information design that is less than fully informative in a number of ways.

First, the possibility of imperfect alignment between those that intentionally collect information inside of organizations and others in those organizations that make decisions based on that information has implications that span several literatures. For example, those designing the information environment for evidence-based management (e.g., Barends and Rousseau, 2018; Pfeffer and Sutton, 2006a,b; Rousseau, 2006; Rousseau and ten Have, 2022) may design less informative environments to shape decisions. Similarly, those using data-driven decision making (e.g., Brynjolfsson and McElheran, 2016; McElheran and Brynjolfsson, 2016) may systematically not surface all the information in the available data. Finally, prior work studying organizations engaged in purposeful learning (e.g., Alcacer and Oxley, 2014; Bingham, Heimeriks, Schijven, and Gates, 2015; Cohen and Levinthal, 1990; Howard, Steensma, Lyles, and Dhanaraj, 2016; Lane, Teplitskiy, Gray, Ranu, Menietti, Guinan, and Lakhani, 2022; Posen, Leiblein, and Chen, 2018) often view the organization as a whole. One promising avenue for future work is to focus attention on the dynamics and frictions that organizations present, and how individuals within organizations may be shaping the information they generate in order to persuade their organization.

Indeed, scholars of organizations often conceptualize them as engaging in purposeful learning—i.e., intentional learning and exploration—in order to generate information that improves decision-making and, ultimately, performance. Our results imply that these same learning efforts may also have a second purpose: to persuade other stakeholders such as lenders, investors, or policy-makers to make decisions in a manner that benefits the organization. This paper explicitly considers the implications of this second purpose for experimentation and learning. However, there are important differences between an organization's stakeholders. Another promising avenue for future work would be to explore how different stakeholders influence the design and interpretation of experiments.

Finally, the results of this paper begin to build a bridge from the literature on information and business experimentation in entrepreneurship and strategy to the work on narrative and storytelling within entrepreneurship and the broader management literature. For example, Garud, Schildt, and Lant (2014) finds that constructing a promising narrative sets the stage for future disappointment when the truth is eventually revealed. Likewise, our model shows that ideas with "successful" experimental results may not be as promising as naïve consumers of experimental results may believe. Similarly, both narrative construction and persuasive experimentation benefit from understanding the perspective of the decision-maker (Falchetti, Cattani, and Ferriani, 2022; Manning and Bejarano, 2017). It is a small step to see persuasive experiments as a possible input to a narrative.

References

Agrawal, Ajay, Joshua S. Gans, Scott Stern. 2021. Enabling Entrepreneurial Choice. *Management Science*, **67**(9) 5510–5524.

Alcacer, Juan, Joanne Oxley. 2014. Learning by supplying. *Strategic Management Journal*, **35**(2) 204–223.

Babichenko, Yakov, Siddharth Barman. 2016. Computational Aspects of Private Bayesian Persuasion. *arXiv*.

Barends, Eric, Denise M. Rousseau. 2018. *Evidence-Based Management: How to Use Evidence to Make Better Organizational Decisions*.

Bergemann, Dirk, Stephen Morris. 2019. Information Design: A Unified Perspective. *Journal of Economic Literature*, **57**(1) 44–95.

Bhaskar, Umang, Yu Cheng, Young Kun Ko, Chaitanya Swamy. 2016. Hardness Results for Signaling in Bayesian Zero-Sum and Network Routing Games. *Proceedings of the 2016 ACM Conference on Economics and Computation*. EC '16, 479–496.

Bigelow, Lyda, Orie Shelef, Robert Wuebker, Jay B Barney. 2024. Experimental Strategy With Entangled Ideas. *Working Paper*.

Bingham, Christopher B., Koen H. Heimeriks, Mario Schijven, Stephen Gates. 2015. Concurrent learning: How firms develop multiple dynamic capabilities in parallel. *Strategic Management Journal*, **36**(12) 1802–1825.

Blank, Steve, Bob Dorf. 2020. *The startup owner's manual: The step-by-step guide for building a great company*.

Blank, Steve, Jonathan T Eckhardt. 2023. The lean startup as an actionable theory of entrepreneurship. *Journal of Management* 01492063231168095.

Brynjolfsson, Erik, Kristina McElheran. 2016. The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, **106**(5) 133–139.

Bylund, Per L., Mark D. Packard. 2022. Subjective value in entrepreneurship. *Small Business Economics*, **58**(3) 1243–1260.

Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, Chiara Spina. 2020. A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial. *Management Science*, **66**(2) 564–586.

Camuffo, Arnaldo, Alfonso Gambardella, Fabio Maccheroni, Massimo Marinacci, Andrea Pignataro. 2022. Microfoundations of low-frequency high-impact decisions. *SSRN*.

Chavda, Ankur, Joshua S. Gans, Scott Stern. 2024. Theory-Driven Entrepreneurial Search. *SSRN*.

Chen, John S., Daniel W. Elfenbein, Hart E. Posen, Ming Zhu Wang. 2022a. The Problems and Promise of Entrepreneurial Partnerships: Decision-Making, Overconfidence, and Learning in Founding Teams. *Academy of Management Review*, **47**(3) 489–520.

Chen, John S, Daniel Walter Elfenbein, Hart E. Posen, Ming zhu Wang. 2022b. Programs of Experimentation and Pivoting for (Overconfident) Entrepreneurs. *Academy of Management Review*.

Cheng, Yu, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, Shang-Hua Teng. 2015. Mixture Selection, Mechanism Design, and Signaling. *Foundations of Computer Science*. 1426–1445.

Cohen, Wesley M., Daniel A. Levinthal. 1990. Absorptive Capacity: A New Perspective on Learning and Innovation. *Administrative Science Quarterly*, **35**(1) 128–152.

Contigiani, Andrea, Daniel A Levinthal. 2019. Situating the construct of lean start-up: Adjacent conversations and possible future directions. *Industrial and Corporate Change*, **28**(3) 551–564.

Edmondson, Amy C., Richard M. Bohmer, Gary P. Pisano. 2001. Disrupted Routines: Team Learning and New Technology Implementation in Hospitals. *Administrative Science Quarterly*,

46(4) 685–716.
Ehrig, Timo, Nicolai J. Foss. 2022. Why we need normative theories of entrepreneurial learning that go beyond Bayesianism. *Journal of Business Venturing Insights*, **18** e00335.
Ehrig, Timo, Jens Schmidt. 2022. Theory-based learning and experimentation: How strategists can systematically generate knowledge at the edge between the known and the unknown. *Strategic Management Journal*, **43**(7) 1287–1318.
Fabijan, Aleksander, Pavel Dmitriev, Colin McFarland, Lukas Vermeer, Helena Holmström Olsson, Jan Bosch. 2018. Experimentation growth: Evolving trustworthy A/B testing capabilities in online software companies. *Journal of Software: Evolution and Process*, **30**(12) e2113.
Falchetti, Denise, Gino Cattani, Simone Ferriani. 2022. Start with "Why," but only if you have to: The strategic framing of novel ideas across different audiences. *Strategic Management Journal*, **43**(1) 130–159.
Felin, Teppo, Alfonso Gambardella, Elena Novelli, Todd Zenger. 2023. A Scientific Method for Startups. *SSRN*.
Felin, Teppo, Todd R. Zenger. 2017. The Theory-Based View: Economic Actors as Theorists. *Strategy Science*, **2**(4) 258–271.
Fréchette, Guillaume R., Alessandro Lizzeri, Jacopo Perego. 2022. Rules and Commitment in Communication: An Experimental Analysis. *Econometrica*, **90**(5) 2283–2318.
Gans, Joshua S. 2022. A Theory of Visionary Disruption. *SSRN*.
Gans, Joshua S. 2023. Experimental Choice and Disruptive Technologies. *Management Science*.
Gans, Joshua S., Scott Stern, Jane Wu. 2019. Foundations of entrepreneurial strategy. *Strategic Management Journal*, **40**(5) 736–756.
Garud, Raghu, Henri A. Schildt, Theresa K. Lant. 2014. Entrepreneurial Storytelling, Future Expectations, and the Paradox of Legitimacy. *Organization Science*, **25**(5) 1479–1492.
Goldstein, Itay, Yaron Leitner. 2018. Stress tests and information disclosure. *Journal of Economic Theory*, **177** 34–69.
Gruber, Marc, Sharon Tal. 2017. *Where to play: 3 steps for discovering your most valuable market opportunities*.
Habibi, Amir. 2020. Motivation and information design. *Journal of Economic Behavior & Organization*, **169** 1–18.
Hedlund, Jonas. 2017. Bayesian persuasion by a privately informed sender. *Journal of Economic Theory*, **167** 229–268.
Howard, Michael, H. Kevin Steensma, Marjorie Lyles, Charles Dhanaraj. 2016. Learning to collaborate through collaboration: How allying with expert firms influences collaborative innovation within novice firms. *Strategic Management Journal*, **37**(10) 2092–2103.
Kamenica, Emir. 2019. Bayesian Persuasion and Information Design. *Annual Review of Economics*, **11**(1) 249–272.
Kamenica, Emir, Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review*, **101**(6) 2590–2615.
Karp, Rebecca. 2023. Working within the Digital Implementation Line: How Organizations Manipulate the Process of Piloting to Shape the Use and Design of AI. *Working Paper*.
Kohavi, Ron, Stefan Thomke. 2017. The Surprising Power of Online Experiments. *Harvard Business Review* 74–82.
Kolotilin, Anton. 2015. Experimental design to persuade. *Games and Economic Behavior*, **90** 215–226.
Koning, Rembrand, Sharique Hasan, Aaron Chatterji. 2022. Experimentation and start-up performance: Evidence from A/B testing. *Management Science*, **68**(9) 6434–6453.
Lane, Jacqueline N., Misha Teplitskiy, Gary Gray, Hardeep Ranu, Michael Menietti, Eva C. Guinan,

Karim R. Lakhani. 2022. Conservatism Gets Funded? A Field Experiment on the Role of Negative Information in Novel Project Evaluation. *Management Science*, **68**(6) 4478–4495.

Leatherbee, Michael, Riitta Katila. 2020. The lean startup method: Early-stage teams and hypothesis-based probing of business ideas. *Strategic Entrepreneurship Journal*, **14**(4) 570–593.

Leiblein, Michael J, Jeffrey J Reuer, Todd Zenger. 2018. What makes a decision strategic? *Strategy Science*, **3**(4) 558–573.

Li, Yong, Ma Lee, Joe Mahoney. 2023. Entrepreneurial flexibility in business model experimentation: A theoretical foundation. *UNLV Working Paper*.

Lipnowski, Elliot, Doron Ravid. 2020. Cheap talk with transparent motives. *Econometrica*, **88**(4) 1631–1660.

Lipnowski, Elliot, Doron Ravid, Denis Shishkin. 2022. Persuasion via Weak Institutions. *Journal of Political Economy*, **130**(10) 2705–2730.

Luca, Michael, Max H. Bazerman. 2021. *The power of experiments: Decision making in a data-driven world*.

Manning, Stephan, Thomas A Bejarano. 2017. Convincing the crowd: Entrepreneurial storytelling in crowdfunding campaigns. *Strategic Organization*, **15**(2) 194–219.

McElheran, Kristina, Erik Brynjolfsson. 2016. The Rise of Data-Driven Decision Making Is Real but Uneven. *Harvard Business Review*.

Merton, Robert K. 1979. *The Sociology of Science: Theoretical and Empirical Investigations*.

Min, Daehong. 2021. Bayesian persuasion under partial commitment. *Economic Theory*, **72**(3) 743–764.

Monnet, Cyril, Erwan Quintin. 2017. Rational Opacity. *The Review of Financial Studies*, **30**(12) 4317–4348.

Murray, Fiona, Mary Tripsas. 2004. The exploratory processes of entrepreneurial firms: The role of purposeful experimentation. *Advances In Strategic Management*. 45–75.

Packard, Mark D., Brent B. Clark, Peter G. Klein. 2017. Uncertainty Types and Transitions in the Entrepreneurial Process. *Organization Science*, **28**(5) 840–856.

Perez-Richet, Eduardo. 2014. Interim Bayesian Persuasion: First Steps. *American Economic Review*, **104**(5) 469–474.

Pfeffer, Jeffrey, Robert I. Sutton. 2006a. Evidence-Based Management. *Harvard Business Review*.

Pfeffer, Jeffrey, Robert I. Sutton. 2006b. Management Half-Truths and Nonsense: How to Practice Evidence-based Management. *California Management Review*, **48**(3) 77–100.

Pillai, Sandeep D., Brent Goldfarb, David A. Kirsch. 2020. The origins of firm strategy: Learning by economic experimentation and strategic pivots in the early automobile industry. *Strategic Management Journal*, **41**(3) 369–399.

Posen, Hart E, Michael J Leiblein, John S Chen. 2018. Toward a behavioral theory of real options: Noisy signals, bias, and learning. *Strategic Management Journal*, **39**(4) 1112–1138.

Romanyuk, Gleb, Alex Smolin. 2019. Cream Skimming and Information Design in Matching Markets. *American Economic Journal: Microeconomics*, **11**(2) 250–276.

Rousseau, Denise M. 2006. Is there Such a thing as "Evidence-Based Management"? *Academy of Management Review*, **31**(2) 256–269.

Rousseau, Denise M., Steven ten Have. 2022. Evidence-based change management. *Organizational Dynamics*, **51**(3) 100899.

Sarasvathy, Saras D. 2009. *Effectuation: Elements of entrepreneurial expertise*.

Schweizer, Nikolaus, Nora Szech. 2018. Optimal Revelation of Life-Changing Information. *Management Science*, **64**(11) 5250–5262.

Shelef, Orie, Robert Wuebker, Jay B Barney. 2023a. Heisenberg Effects In Experiments On Business Ideas. *Academy of Management Review*.

Shelef, Orie, Robert Wuebker, Jay B. Barney. 2023b. The Impact of Entrepreneurial Cognitive Biases and Stakeholder Competing Goals on Experimental Strategy Choices.

Shepherd, Dean A, Marc Gruber. 2021. The lean startup framework: Closing the academic–practitioner divide. *Entrepreneurship Theory and Practice*, **45**(5) 967–998.

Smolin, Alex. 2021. Dynamic Evaluation Design. *American Economic Journal: Microeconomics*, **13**(4) 300–331.

Szydlowski, Martin. 2021. Optimal Financing and Disclosure. *Management Science*, **67**(1) 436–454.

Thomke, Stefan H. 2020. *Experimentation works: The surprising power of business experiments*.

Xifara, Tatiana. 2021. Designing Experimentation Guardrails.

Zellweger, Thomas, Todd Zenger. 2022. Entrepreneurs as scientists, Bayesian inference, and belief revision. *Journal of Business Venturing Insights*, **18** e00350.

Zellweger, Thomas, Todd Zenger. 2023. Entrepreneurs as scientists: A pragmatist approach to producing value out of uncertainty. *Academy of Management Review*, **48**(3) 379–408.

Åstebro, Thomas, Holger Herz, Ramana Nanda, Roberto A Weber. 2014. Seeking the roots of entrepreneurship: Insights from behavioral economics. *Journal of Economic Perspectives*, **28**(3) 49–70.

Appendix

This Appendix provides proofs of the claims made in the text. We first formally state the assumptions made in the Model Setup.

Assumption 1. *There is some possible experimental result that would lead the decision-maker to make a different decision than what the decision maker would do absent that information.*

Assumption 1 requires that the decision-maker's choice is one that depends on the state of the world, and it is possible that an experiment reveals a state of the world that leads the decision-maker to make a different choice than they would absent experimentation. In other words, the decision-maker is responsive to at least some information from at least some experiments. This assumption also implies that that the decision-maker has some decision relevant uncertainty about the underlying state and more than one action that the decision-maker might make.

Assumption 2. *If a trivial amount of information is revealed, the decision-maker makes the same decision they would make absent any information.*

Assumption 2 requires that, absent information, the decision-maker is not indifferent between possible decisions. Any categorical decision (including binary choices) meets this criteria generically such as business decisions to "go or not go," "invest or not invest," or "adopt or not adopt".[23]

Assumption 3. *The default decision the decision-maker makes is not the experiment designer's most preferred decision. Formally:*

Let action a_1 be the action the decision-maker takes absent more information. There exists state s and action a_2 such that the experiment designer strictly prefers a_2 over a_1 absent more information.

Note that assumptions 1 and 2 and 3 together satisfy the assumptions in Proposition 2 of Kamenica and Gentzkow (2011) which demonstrates that the experiment designer benefits from experimentation.

Assumption 4. *Under some fully revealed state, the experiment designer would prefer the decision-maker take an action that some other fully revealed state would lead the decision-maker to take. Formally:*

Let actions a_1 and a_2 be the actions the decision-maker takes for states $\boldsymbol{S_1}$ and $\boldsymbol{S_2}$. There exists $s' \in \boldsymbol{S_2}$ such that the experiment designer strictly prefers a_1 over a_2 in state s' and that the decision-maker strictly prefers action a_1 over the alternatives for states $\boldsymbol{S_1}$.

Together, assumptions 3 and 4 reflect competing goals between the decision-maker and the experiment designer. Assumption 3 reflects that the competing goals between the decision-maker and experiment designer are strong enough that, absent experimentation, the experiment designer would prefer a different decision in the status quo. Assumption 4 reflects competing goals between the decision-maker and the experiment designer if an experiment reveals the underlying state. Assumption 4 requires that the misalignment under full information is not "hypothetical"—the experiment designer must prefer the decision-maker to make a choice that the decision-maker strictly prefers to under some *other* state. While these conditions may appear complicated when stated carefully they reflect quite common misalignment.[24]

[23]We use generic in the mathematical sense—infinitesimal variation in parameters will cause this to hold, or if the parameters are drawn from distributions without mass points, the probability that this holds is 1.

[24]Requiring strict preferences in assumption 4 is akin to Assumption 1 of Kamenica and Gentzkow (2011).

Finally, to simplify our analysis and following Kamenica and Gentzkow (2011) and Kamenica (2019) we assume that there are finite possible actions and a finite set of underlying states. Also, like Kamenica and Gentzkow (2011) do in their Assumption 1, we rule out ties in utility functions that would be be present under a perturbation of preferences.

We now consider Proposition 1

Proof. Assumptions 1 2 and 3 are sufficient, following Kamenica and Gentzkow (2011) Proposition 2, that there exists an informative experiment that is better for the experiment-designer than no experiment. However, under Kamenica and Gentzkow (2011) this may be a fully informative experiment. Using Assumption 4 we now show that the fully informative experiment is not the experiment designer's best choice either. That is, we prove there exists an experiment that is better for the experiment-designer than a fully informative experiment.

Consider an experiment that returns signal $A1$ for states $\boldsymbol{S_1}$ and with probability δ for state $s' \in \boldsymbol{S_2}$, and otherwise returns the full information signal.

Following assumption 4, there exists $\delta > 0$ small enough, that the decision-maker takes action $a1$ following signal $A1$, and the same actions as under full information otherwise. By assumption 4 the experiment designer prefers this experiment to the full information one. Thus, the full information experiment is not the experiment-designer-optimal. Taken together, the experiment-designer-optimal experiment is *partially* informative. It is neither uninformative nor fully informative.

□

Corollary 1 follows as the experiment designer is able to select it. Corollary 2 follows because the decision-maker could take its default action following the experiment. Corollary 3 follows by construction of the proof of proposition 1.

We now consider Proposition 2.

Proof. The proof of Proposition 1 demonstrates that the experiment designer strictly prefers an experiment other than a fully-informative one. Thus, designing experiments that are neither uninformative nor fully informative creates more value for the designer than they would create faced with only the option of a fully informative experiment. □

We now consider Proposition 3.

We first define most and least preferred actions. We say an action is the experiment designer's most preferred action if, for every state, the experiment designer strictly prefers that action over any other action. Likewise, an action is the experiment designer's least preferred action if for every state the experiment designer strictly prefers all other actions. With those definitions, we turn to the proof:

Proof. The claim in the proposition that when the least preferred action is taken the decision is uncontaminated follows from Kamenica and Gentzkow (2011) Proposition 4.

Contamination of the signals that lead to the most preferred action has intuition similar to the logic of Proposition 1. If the signal that leads to the most preferred action was not contaminated, then it could be contaminated some by shifting some probability mass to it from another induced belief and improve outcomes for the experiment designer. Formally, Kamenica and Gentzkow (2011), Proposition 5 demonstrates that when taking the experiment designer's most-preferred action following an optimally designed experiment, the decision-maker must be indifferent between that and some other decision. That, combined with the no-ties assumption, must mean that the signal is contaminated. □