Publication Design with Incentives in Mind*

Ravi Jagadeesan† Davide Viviano‡

March 20, 2025

Abstract

The publication process both determines which research receives the most attention, and influences the supply of research through its impact on the researcher's private incentives. We introduce a framework to study optimal publication decisions when researchers can choose (i) whether or how to conduct a study and (ii) whether or how to manipulate the research findings (e.g., via selective reporting or data manipulation). When manipulation is not possible, but research entails substantial private costs for the researchers, it may be optimal to incentivize cheaper research designs even if they are less accurate. When manipulation is possible, it is optimal to publish some manipulated results, as well as results that would have not received attention in the absence of manipulability. Even if it is possible to deter manipulation, such as by requiring researchers to conduct pre-registered experiments instead of quasi-experiments, it is suboptimal to do so when experiments entail high research costs.

*We thank Isaiah Andrews, Arun Chandrasekhar, Kevin Chen, Max Kasy, Toru Kitagawa, Paul Niehaus, Marco Ottaviani, Jesse Shapiro, Jann Spiess, Elie Tamer, Alex Tetenov, and Kaspar Wüthrich for helpful comments. We thank Karthik Seetharaman for exceptional research assistance.

†Department of Economics, Stanford University. Email address: ravi.jagadeesan@gmail.com.

‡Department of Economics, Harvard University. Email address: dviviano@fas.harvard.edu.

Publication decisions shape the process of scientific communication. By selecting what to publish, journals affect which findings receive the most attention and can inform the public about the state of the world. The design of publication rules has therefore motivated recent debates on how statistical significance should affect publication when the goal is to direct attention to the most informative results (e.g., Abadie, 2020; Frankel and Kasy, 2022).

However, the publication process also affects the supply of research by influencing researcher's incentives about how to conduct research. Researchers have many degrees of freedom about how to conduct their research, such as how and where to run an experiment (e.g., Allcott, 2015; Gechter and Meager, 2022); the size, cost, and effort associated with the study (e.g., Thompson and Barber, 2000; Grabowski et al., 2002); and which findings to report from a given study (Brodeur et al., 2020; Elliott et al., 2022). We refer broadly to their choices about each of these aspects of a study broadly as a *research design*.

Researcher's private incentives may influence how they choose their design. Yet, "while economists assiduously apply incentive theory to the outside world, we use research methods that rely on the assumption that social scientists are saintly automatons" (Glaeser, 2006). This raises the questions of when and how researchers' incentives should impact the design of publication processes, and, more broadly, the optimal allocation of attention to research.

This paper studies optimal publication decisions when researchers chooses the research design based on private costs and benefits. We frame this question as a mechanism design problem: a social planner (principal) optimizes a publication rule, taking into account the incentives of the researcher (agent). The social planner aims to use the publication process to allocate the attention of the audience to the most informative research findings.

More concretely, as in Frankel and Kasy (2022) (building in turn on Wald (1950)), we suppose that research results impact the decisions of an audience who has limits on how much attention they can devote to research. The social planner seeks to publish results that are most important for the audience, net of the cost of (or taking into account constraints on) publishing or attention. Due to attention constraints, not all results will be published, which leaves the planner with a non-trivial trade-off about which results and designs to publish. We then introduce a model in which publication decisions affect the supply of research in the first place (Section 1): given the publication rule, the researcher chooses the design that maximizes her value for publication (or other attention) net of research costs.

As a concrete example, consider a medical journal seeking to decide whether to publish results from a clinical study. The journal wants to convey accurate information about the drug's efficacy in the study (e.g., DeMaria, 2022; Ana, 2004) and direct the audience's attention to the most effective drugs. This objective is motivated by scientific practice. For example, the stated mission of the New England Journal of Medicine is

> "to publish the best research and information at the intersection of biomedical science and clinical practice and to present this information in understandable, clinically useful formats that inform health care practice and improve patient outcomes."

However, researchers may respond in how they conduct their studies through the size, length, cost of the studies, composition of the control groups (Thorlund et al., 2020), and—in some cases—in which specific findings to report (e.g., Riveros et al., 2013; Shinohara et al., 2015).

We draw a dichotomy between researchers' incentives about (i) whether or how to conduct a study and (ii) which findings to report (e.g., via data manipulation or selective reporting).

We first focus on (i) and abstract from data manipulation and selective reporting; we therefore suppose that the research design is observable and verifiable (Section 2). For example, a researcher could choose between (quasi)experiments with different mean-squared errors and costs and, by pre-committing to an analysis plan, truthfully report an unbiased estimate of a treatment effect. The planner can direct the attention of the audience to any executed study by publishing it; publication can depend on a study's design and results.

Returning to our example, after providing a new drug to a treatment group, scientists can evaluate its efficacy by comparing the treatment group to either an experimental placebo group or using a pre-specified synthetic control group obtained from historical medical records (Popat et al., 2022; Yin et al., 2022). Using a synthetic control group can have large impact on the supply of medical research and new drugs by decreasing research costs (Jahanshahi et al., 2021; Food and Drug Administration, 2023; Wong et al., 2014). However, using a synthetic control group may increase the estimate's mean-squared error due to lack of randomization (see Raices Cruz et al., 2022; Rhys Bernard et al., 2024), raising the question of whether a social planner should allow or incentivize their use.

We thus suppose that the publication process affects the supply of research through an individual rationality constraint: for a researcher to be willing to conduct a study, they must be compensated with a large enough *ex ante* publication probability. Without these individual rationality constraints, the first-best publication decision rule would ignore researchers' costs and publish only studies conducted with the lowest possible mean-squared error; due to attention constraints or publication costs, not all such studies would be published. When the feasible designs are (quasi)experiments that are inexpensive for the researcher to conduct, the first-best policy can be feasibly implemented by the planner; the individual rationality constraint does not bind. Thus, small research costs do not affect the planning problem.

However, when the research cost of accurate designs is above a tipping point, the first-best publication rule does not compensate the researcher enough to incentivize them to choose such a design. As a result, the planner faces a trade-off between providing attention

to results that are not valuable enough to deserve attention, and rewarding the researcher enough to make them willing to use a costly design in the first place.

This trade-off has implications for which designs the planner should publish. For example, suppose that clinical studies can be conducted based on either less expensive but noisy quasi-experiment, or a perfectly informative but costly experiment. Incentivizing researchers to conduct the costly experiment would require publishing results from such a study, even when doing so misdirects attention toward treatments with negligible effects on patients. In this case, the planner may prefer publishing quasi-experiments to avoid having to publish too many results from the experiment; researchers will then choose not to use costly experiments in equilibrium. This analysis suggests that due to the interaction between attention costs and supply effects, publishing medical studies with observational control groups is desirable when it is sufficiently costly for researchers to use experimental control groups.

We next turn in Section 3 to scenarios where the researcher, *after* observing the study's results, can report a biased statistic from such a study. The bias is chosen by the researcher and unobserved by the planner, but incurs a cost to the researcher that is increasing in the bias (e.g., capturing reputational costs). The audience is unaware of possible manipulations of published findings, taking results at face value. We think of this model as a stylized description of settings with potential data manipulation or selective reporting. For instance, in the absence of a precise pre-specification, we may be concerned that researchers select a synthetic control group from medical records based on their observed outcomes.

Each publication rule generates a different degree of manipulation. For example, suppose that the planner used the publication rule that would be optimal without manipulation—i.e., publishing if the reported statistic is above a cutoff. Then, researchers whose results would be close to the cutoff would manipulate their data or analyses to reach the cutoff.[1] Publishing substantially manipulated results would incur a substantial loss for the audience.

One approach that has been proposed is to completely deter manipulation. An extreme form of this approach would be to make publication dependent only on the design, and not on results.[2] However, this approach generally incurs substantial costs by directing the audience's attention to results that would not affect their decisions, and is not optimal.

By contrast, we show that the optimal publication rule under manipulability has three key features. First, it increases the cutoff for which findings always get published compared to settings without data manipulation. Second, just below this cutoff, it randomizes publication chances, making the researcher indifferent about whether to manipulating their results.

[1]This manipulation would lead to bunching at the cutoff, which is consistent with bunching that has been documented in settings with p-hacking (e.g., Elliott et al., 2022).

[2]In practice, this approach can be implemented by committing to publication based on pre-analysis plans, as the *Journal of Development Economics* does.

As a result, with nonzero probability, the planner publishes findings that are not truly valuable enough to deserve attention, and would not be published without the possibility of manipulation. Third, some manipulation *does* occur in equilibrium.

To gain some intuition for these features of the optimum, consider first the optimal publication rule without manipulation. The planner can eliminate the researcher's incentive to manipulate by publishing results below the cutoff with a positive probability. However, some results that do not affect the audience's decision do get published as a result, which is undesirable in the presence of attention costs or constraints. To counteract this effect, it then is optimal for the planner to increase the cutoff to guarantee publication.

With a more stringent cutoff, some results that would be published without manipulation are no longer published with probability one. The planner then encourages researchers with results in this range to engage in a limited degree of manipulation to increase their publication chances. The loss from publishing some slightly manipulated studies is second-order relative to the (first-order) gain from publishing results that do affect the audience's decisions.

To prove that the optimal publication rule takes the described form, we formulate the social planner's problem as a mechanism design problem with "false moral hazard" due to the researcher choosing a manipulation after learning the true results. The absence of direct transfers and the inability to reward the researcher with a publication probability above one make the mechanism design problem effectively one with limited transfers. As a result, standard methods as in Mirrlees (1971) and Myerson (1981) do not apply. We solve the mechanism design problem by identifying the precise pattern of binding incentive constraints.

As a final exercise in Section 4, we combine these two models to ask whether the planner should mandate researchers to send a costly signal that deters them from manipulating data. For example, the planner could mandate that researchers run costly experiments that adhere to pre-analysis plans, rather than also allowing for equivalently informative (but cheaper) observational studies. Without accounting for the effects of making research more costly on the supply of research, the planner would always require such signals to ensure that results are unbiased (Spiess, 2024; Kasy and Spiess, 2023). For example, the planner would not publish observational studies if equally informative experiments are feasible, no matter their costs.

However, to ensure that the researcher is actually willing to conduct the experiment, the planner needs to reward the researcher with enough attention through a high enough publication probability. With large attention costs (or binding publication constraints), this makes incentivizing experiments costly enough for the planner, that they prefer to publish observational studies, even though manipulation may occur. Therefore, taking into account the supply effects of the publication process, the planner may prefer to drive attention towards studies whose results may be manipulable, and hence possibly biased.

Related literature. This paper connects to a growing literature that studies economic models to analyze statistical protocols. In the context of scientific communication, Andrews and Kasy (2019), Abadie (2020), Andrews and Shapiro (2021), Kitagawa and Vu (2023), and (most closely related to our paper) Frankel and Kasy (2022) have analyzed how research findings are or should be reported to inform the public. Our analysis builds on this literature by introducing a model that incorporates researcher's incentives and studying how these incentives shape the optimal design of the optimal publication process.

We connect to a broad literature on statistical decision theory (e.g., Wald, 1950; Savage, 1951; Manski, 2004; Hirano and Porter, 2009; Tetenov, 2012; Kitagawa and Tetenov, 2018), focusing in particular on settings with private incentives of researchers. Other work in this line includes Chassang et al. (2012), Tetenov (2016), Banerjee et al. (2017), Henry and Ottaviani (2019), Banerjee et al. (2020), Di Tillio et al. (2021), Williams (2021), Bates et al. (2022, 2023), Libgober (2022), McCloskey and Michaillat (2022), Yoder (2022), Kasy and Spiess (2023), Spiess (2024) and Viviano et al. (2024). Different from these references, we both analyze settings where the researcher may choose the research design absent private information, or choose the design and manipulate reported findings with private information. This allows us to formally study ideas when/whether unsurprising results should be published, and whether manipulation should occur in equilibrium (Glaeser, 2006).

In particular, an important distinction from some of these models studying approval decisions, such as Tetenov (2016), Bates et al. (2022, 2023), and Viviano et al. (2024), is that these papers assume that researchers truthfully report the statistics sampled from their study (abstracting from questions about data manipulation or design choice studied in this paper). Spiess (2024) and Kasy and Spiess (2023) study models of scientific communication, without, however, focusing on optimal publication rules studied here. Different from these references, here we incorporate the researcher's costs of the design (and misreporting), which, we show, leads to qualitatively different solutions in the amount of misreporting. Di Tillio et al. (2021) study the different question of selective sampling, and Henry and Ottaviani (2019) study the question of decisions with continuous and sequential access to the data, different from the question of selective design choice studied here.

Finally, a large empirical and econometric literature has documented several aspects of the research process, including selective reporting, data manipulation, specification search, as well as site selection bias and observational studies' bias (e.g., Allcott, 2015; Gechter and Meager, 2022; Rosenzweig and Udry, 2020; Elliott et al., 2022; Brodeur et al., 2020; Miguel, 2021; Olken, 2015; Banerjee et al., 2020; Rhys Bernard et al., 2024). Our contribution here is to provide a formal theoretical model that studies how incentives interact with some of these choices, shedding light on qualitative aspects of optimal publication rules.

1 Setup

Consider three agents: a researcher, a (representative) audience, and a social planner. The audience and social planner are interested in learning a parameter $\theta_0 \in \mathbb{R}$, such as the average treatment effect of a given intervention. All agents share a common prior $\theta_0 \sim \mathcal{N}(0, \eta_0^2)$.

A researcher conducts a study to inform the audience about θ_0. A study is summarized by (X, Δ), where Δ denotes the design and

$$X(\Delta)|\theta_0 \sim \mathcal{N}\Big(\theta_0 + \beta_\Delta, S_\Delta^2\Big) \tag{1}$$

the results. Here, β_Δ and S_Δ^2 are the average bias, and variance, of design Δ, respectively.

If a study is conducted, it will be evaluated according to a publication rule $p(X, \Delta)$ with values in $[0, 1]$. Here, $p(X, \Delta)$ represents the probability of publishing the study, which is assumed to be a Borel measurable function of (X, Δ).

Conditional on publication, the audience forms posterior beliefs about θ_0 using Bayes' Rule assuming $\beta_\Delta = 0$. Conditional on non-publication, the audience's posterior mean equals its prior mean (zero), as would arise from Bayes' Rule with $p(\cdot)$ symmetric in X and β_Δ symmetric around 0. Based on these beliefs, the audience takes an action to minimize mean-squared distance from θ_0; thus, the audience's action $a_p^\star(X, \Delta)$ is the posterior mean

$$a_p^\star(X, \Delta) = \begin{cases} \frac{X\eta_0^2}{S_\Delta^2 + \eta_0^2} & \text{if the study is published} \\ 0 & \text{otherwise} \end{cases}.$$

Given results $X = X(\Delta)$ for a design Δ, and a parameter θ_0, the planner incurs a loss

$$\mathcal{L}_p(X, \Delta, \theta_0) = \mathbb{E}_p\Big[\big(\theta_0 - a_p^\star(X, \Delta)\big)^2\Big] - c_p p(X, \Delta), \tag{2}$$

where $\mathbb{E}_p$ denotes expectation with respect to any stochasticity in the publication decision rule. Thus, the planner minimizes the expected loss of the audience, net of a cost of publication or attention. The quantity c_p captures the publication or attention costs or constraints.

Given a design Δ, a publication rule p, and results X, the researcher's expected payoff is

$$v_p(X, \Delta) = p(X, \Delta) - C_\Delta \tag{3}$$

where $C_\Delta \leq 1$ is the researcher's cost of executing design Δ.[3] As is standard practice, whenever the researcher is indifferent between two designs, we implicitly assume she chooses the design that minimizes the planner's expected loss.

[3] We assume $C_\Delta \leq 1$ simply to rule out trivial cases in which design Δ is never chosen by the researcher.

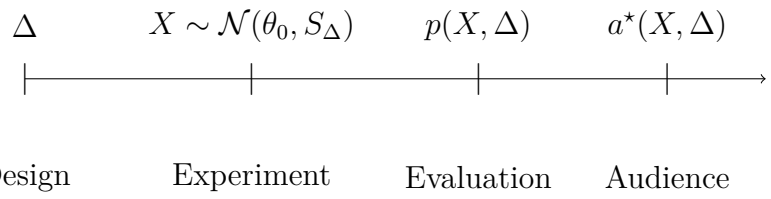


Figure 1: Illustration of the variables in the model under observable and verifiable designs. First, researchers pre-specify the population of interest. Second, they run an experiment and draw a statistic X. Finally, a social planner evaluates the experiment based on a decision rule $p(X, \Delta) \in [0, 1]$. Finally, the audience forms a posterior about the estimand of interest.

2 Publication rules under verifiable designs

This section studies optimal publication rules when the planner can observe the research design, and condition publication on it. We define a design Δ *unbiased on average* if $\beta_\Delta = 0$, and focus on such designs for simplicity in this section (and defer designs that are not unbiased on average to the following section). We suppose that the researcher chooses between two designs that have different mean-squared errors and research costs. The researcher selects the design before observing the results $X(\Delta)$, and the selected design Δ is observable to the planner (and verifiable as part of the publication process). Figure 1 illustrates.

Our analysis proceeds in three steps. We first characterize the optimal publication rule subject to the constraint of incentivizing the researcher to implement a particular design Δ. We then characterize which designs are worth incentivizing relative to an outside option. Last, we characterize the fully optimal publication rule with multiple possible designs.

2.1 Optimal publication rule that implements a particular design

As a first step, we characterize the constrained optimal publication when the planner must make implementing a particular design Δ individually rational for the researcher, i.e., when the planner must guarantee that the researcher always conducts the study.

Definition 1. A *constrained optimal publication rule* for a design Δ is a publication rule $p^\star_\Delta$ that minimizes $\mathbb{E}\left[\mathcal{L}_p(X(\Delta), \Delta, \theta_0)\right]$ subject to $\mathbb{E}[v_p(X(\Delta), \Delta)] \geq 0$.

Intuitively, a constrained publication rule is a publication rule that always guarantees that research is conducted. Our first result shows that the constrained optimal publication rule then takes a threshold form, where the threshold t^*_Δ for publication depends on the prior and the design Δ.

Proposition 1 (Individual rational publication rule)**.** *If Δ is unbiased on average, then a constrained optimal publication rule for Δ is the threshold rule $p^\star_\Delta(X) = 1\{|X| \geq t^\star_\Delta\}$, where*

$$t^\star_\Delta = \min\left\{\frac{S^2_\Delta + \eta_0^2}{\eta_0^2}\sqrt{c_p}, \left|\Phi^{-1}(C_\Delta/2)\right|\sqrt{S^2_\Delta + \eta_0^2}\right\}.$$

To understand the intuition behind Proposition 1, first suppose that the research cost is $C_\Delta = 0$, so there is no individual rationality constraint for the researcher. Then, as in Frankel and Kasy (2022), as the planner's publication cost c_p is nonzero, the planner will publish results that move the audience's optimal action enough to justify theattention (publication) costs c_p:[4] i.e., the planner will publish results $|X| \geq \gamma^*_\Delta$, where[5]

$$\gamma^*_\Delta = \frac{S^2_\Delta + \eta_0^2}{\eta_0^2}\sqrt{c_p}.$$

When this cutoff rule guarantees an *ex ante* publication probability of at least C_Δ, the individual rationality constraint does not bind. In this case, we say the design is cheap.

Definition 2. Let Δ be a design that is unbiased on average. Design Δ is *cheap* if $C_\Delta \leq \mathbb{P}(|X(\Delta)| \geq \gamma^*_\Delta)$ and *expensive* otherwise.

Note that higher publication costs c_p, and lower prior variances η_0^2, both raise the threshold γ^*_Δ and hence make designs more likely to be expensive. Intuitively, whether a design is cheap depends on how informative it is (relative to publication costs): more informative designs are more likely to be classified as cheap designs (holding fixed research costs).

For expensive designs, the cutoff rule from Frankel and Kasy (2022) does not provide a large enough *ex ante* publication probability to entice the researcher to conduct research in the first place. Hence, the planner needs to commit to publishing more results in order to satisfy the researcher's individual rationality constraint. It is best for the planner to publish results that move the audience's action the most, even if these results do not move the audience's action enough to justify the attention cost c_p. Hence, the planner sets a cutoff that ensures an *ex ante* publication chance of C_Δ—i.e., a cutoff of

$$\left|\Phi^{-1}(C_\Delta/2)\right|\sqrt{S^2_\Delta + \eta_0^2}.$$

[4]The dependence of the publication cutoff on publication or attention costs is unlike the approval frameworks of Tetenov (2016) and Viviano et al. (2024).

[5]Our framework directly extends to the case in which $\theta_0 \sim \mathcal{N}(\mu, \eta_0^2)$ for a prior mean μ. Here, μ is the audience's action in the absence of publication. In this case, the optimal constrained publication rule takes the form $|X - \mu| \geq t^\star_\Delta$ for the same threshold $t^\star_\Delta$ as in Proposition 1. Therefore, as in Frankel and Kasy (2022), one should interpret surprising results as ones that move the audience away from its default action.

This second cutoff is below γ^*_Δ for (and only for) expensive designs, and is the optimal cutoff for such designs. Thus, if research costs are large enough that the implemented design is expensive, the researcher's incentives play a central role in determining the optimal publication rule, unlike in Frankel and Kasy (2022): the publication process should become less stringent as research becomes more costly.

The following corollary summarizes and formalizes the preceding discussion.

Corollary 1. *(a) If Δ is a cheap design, then the cutoff $t^\star_\Delta$ for a constrained optimal publication rule is $t^\star_\Delta = \gamma^*_\Delta$.*

(b) If Δ is an expensive design, then $t^\star_\Delta = |\Phi^{-1}(C_\Delta/2)| \sqrt{S^2_\Delta + \eta^2_0}$ is the optimal cutoff for a constrained publication rule. This cutoff is strictly decreasing in the cost C_Δ.

In words, if the planner wants to allow for expensive experiments to be conducted, she must adjust the relevant threshold accordingly to make such designs individually rational.

2.2 Which designs are ever worth incentivizing

So far, our discussion provided intuition on optimal publication rules, under the constraint that researcher would conduct the study. We next study when a design is ever worth incentivizing, relative to the outside option of no study (i.e., relative to not making it individually rational for the researcher to conduct research with design Δ).

Let $\mathcal{L}^*_\Delta = \mathbb{E}\left[\mathcal{L}_{p^\star_\Delta}(X(\Delta), \Delta, \theta_0)\right]$ denote the optimal expected loss for the planner once implementing design Δ. (Here $p^\star_\Delta$ is a constrained optimal publication rule for Δ.) The expected loss if no research is published is the prior variance η^2_0. Comparing these two quantities determines whether a design is worth incentivizing in the first place.

Definition 3. A design Δ is *worthwhile* if $\mathcal{L}^*_\Delta \leq \eta^2_0$.

We next characterize which designs are worthwhile. If a design is cheap, then the planner can selectively publish only results that move the audience's beliefs enough to justify the publication or attention cost. Thus, the *ex post* loss from incentivizing the design (under the constrained optimal publication rule) is always lower than η^2_0, and similarly its *ex ante* loss.

Proposition 2 (When are cheap designs worthwhile?)**.** *Every cheap design Δ is worthwhile.*

Whereas *cheap* design are always worthwhile (regardless of their cost), for expensive designs, the situation is more delicate. Incentivizing the researcher to implement a design requires committing to publish results that the planner would *ex post* prefer not to publish. When attention costs c_p are large enough, the costs of publishing these marginal results

outweigh the benefits of publishing surprising results. How large c_p needs to be for this to occur depends on the design's cost and variance.

To formalize this intuition, it will be convenient to express our results in terms of the difference between the posterior and prior variances conditional on publication of the results of a design Δ, which we denote by

$$\text{PostVarRed}(\Delta) := \eta_0^2 - \frac{S_\Delta^2 \eta_0^2}{S_\Delta^2 + \eta_0^2} = \frac{\eta_0^4}{S_\Delta^2 + \eta_0^2}.$$

Here, $\text{PostVarRed}(\Delta)$ measures how informative a given design is in terms of posterior variance reduction.

Lemma 1. *(a) For each unbiased-on-average design Δ with $C_\Delta > 0$, there exists a threshold $c_p^\star(\Delta, \eta_0) > 0$ such that Δ is worthwhile if and only if $c_p \leq c_p^\star(\Delta, \eta_0)$.*

(b) This threshold is strictly decreasing in C_Δ and S_Δ, strictly increasing in η_0, and satisfies

$$\frac{\text{PostVarRed}(\Delta)}{C_\Delta} \geq c_p^\star(\Delta, \eta_0) \geq \frac{\text{PostVarRed}(\Delta)}{C_\Delta} - \eta_0^2 O\big((1 - C_\Delta)^3\big).$$

Intuitively, expensive designs are *less* likely to be worthwhile when their research costs C_Δ are higher because in this case more surprising results must be published for incentive reasons. Similarly, either (or both) the larger the variance S_Δ^2 or the smaller the prior variance η_0^2 the more likely the design is not worthwhile.

While the exact characterization of worthwhile designs is not analytically simple for intermediate research costs, Lemma 1 provides us with an intuitive bound on the threshold $c_p^\star(\Delta, \eta_0)$ that is exact for large research costs C_Δ. To understand the Lemma 1 in the case of large research costs, consider a comparison when $C_\Delta \approx 1$. In this case, the researcher must be guaranteed to always publish for the study to be incentive compatible. The corresponding loss of the planner equals $\text{PostVarRed}(\Delta) - c_p$. Therefore, the design is worthwhile only if the expected ex-post variance reduction exceeds the attention cost.

But how does the decision of incentivizing a given design vary as we vary the design's cost C_Δ? Studying this question allows us to shed lights on the relevant trade-offs between variance reduction and cost of the experiment. We formalize the main relevant trade-offs in the following proposition.

Proposition 3 (When are expensive designs worthwhile?)**.** *Let Δ be unbiased on average.*

(a) If $\text{PostVarRed}(\Delta) \geq c_p$, then Δ is worthwhile.

(b) If $\text{PostVarRed}(\Delta) < C_\Delta c_p$, then Δ is not worthwhile.

Part (a) of the proposition formalizes the intuition that *no matter* how large the cost of the design C_Δ, any design whose posterior variance reduction exceeds the attention cost is always worthwhile. For example, suppose researcher are interested in studying an unexplored yet relevant research question, so that the prior variance $\eta_0 \to \infty$ and $c_p < \infty$. (For example, researchers may want to study the effect of large economic stimula on macro-economic fiscal multipliers via a large scale experiment as in Egger et al. (2022).) Then no matter how costly the design is, the design will always be worthwhile, provided that the prior variance is sufficiently large.

Part (b) of the proposition states that an experiment is not worthwhile if the posterior variance reduction *relative to the cost of the experiment* is smaller than the publication cost (which recall is normalized to be $C_\Delta \in (0, 1]$). That is, as C_Δ increases, the posterior variance reduction should increase *proportionally* to the cost for the experiment to be worthwhile.

The proposition provides us with explicit recommendation on when designs are not worthwhile as a function of the design's costs. For instance, suppose that research costs scale inversely with the design's variance (and therefore linearly in the sample size). Formally, we suppose that $C_\Delta = \frac{c_v}{S^2_\Delta} + c_f$, where c_v denotes a variable cost of sample size, and c_f denotes a fixed costs of a design. Then whenever the variance exceeds a threshold that depends on *both* the variable cost of the design and the attention cost such designs are not worthwhile.

Corollary 2. *Consider a design Δ that is unbiased on average with $C_\Delta = \frac{c_v}{S^2_\Delta} + c_f$. If $\eta_0^2 > c_p c_f$ and $S^2_\Delta > \frac{c_p c_v}{\eta_0^2 - c_p c_f}$, then design Δ is worthwhile.*

Corollary 2 illustrates in an example with variable and fixed costs how our results can be applied to obtain concrete expressions. In this example, the experiment is not worthwhile if the variance of the design is sufficient large (more precisely, larger than an explicit function increasing in the variable and fixed cost). Therefore, despite a smaller variance increases the cost of the experiment, its contribution to the posterior variance reduction makes a costly experiment with sufficiently small variance worthwhile.

2.3 Choosing which design to incentivize

We next turn to the structure of the optimal publication rule when there is more than one possible design. For simplicity, we focus on a setting with two, unbiased designs. Without loss of generality, we suppose that the one with a lower mean squared error has a higher research cost, so the planner faces a non-trivial problem about which design to incentivize.

Assumption 1. Researchers can choose between two designs E, O that are each unbiased on average. The designs have different mean squared errors $S^2_E < S^2_O$ and costs $C_E > C_O$.

We refer to $\Delta = E$ as an experiment and $\Delta = O$ as an observational study, using the choice between an experiment or observational study as a leading case. To understand this case, let us return to to our example of the choice of control in a clinical trial from the introduction. Here, $\Delta = E$ represents a medical experiment with a treatment and placebo control group, and $\Delta = O$ represents a medical study with a pre-specified synthetic control group. Design O has smaller cost, since under that design, researchers have to recruit participants for the treatment group but not for a control group. For simplicity, we suppose that the bias of the quasi-experimental control is unknown with mean zero; as we discuss below, we can then capture the bias as part of the mean-squared error.

Example 1 (Experiment versus observational study with irreducible non-sampling error)**.** Let $X(E) = \theta_0 + \varepsilon_E, X(O) = \theta_0 + b_O + \varepsilon_O$, where $\varepsilon_E \sim \mathcal{N}(0, S_E^2)$ denote the estimation noise from an experiment, $\varepsilon_O \sim \mathcal{N}(0, \sigma_O^2)$ denotes the idiosyncratic noise from an observational study and $b_O|\varepsilon_O \sim \mathcal{N}(0, \sigma_B^2)$ denotes the observational study's random effect which capture unobserved bias drawn from a fixed (Gaussian) distribution. For example, Rhys Bernard et al. (2024) empirically investigates the distribution of b_O through a meta-analysis, where σ_B^2 captures the impact of the random bias in statistics' distribution.[6] We then have that $X(O) \sim \mathcal{N}(\theta_0, S_O^2)$, where the mean-squared error $S_O^2 = \sigma_O^2 + \sigma_B^2$ includes both sampling uncertainty σ_O^2 and irreducible error σ_B^2 arising the variance of the bias. □

We can also intepret the design O more literally as a low-cost experiment.

Example 2 (Low-cost and costly experiment)**.** Suppose that E corresponds to an experiment with costly implementation C_E, where O is an experiment with possibly different variance, but smaller cost of implementation. In this case $\mathbb{E}[X(E)|\theta_0] = \mathbb{E}[X(O)|\theta_0] = \theta_0$ while the two experiments may have different costs (and variances). □

We next use Proposition 1 to study the optimal choice between the experiment and the observational study. Because Δ is observable by the planner (and verifiable as part of the publication process), the planner can incentivize their preferred design by setting

$$p^\star(X, \Delta) = 1\{\Delta = \Delta^{\text{planner}}\} p^\star_\Delta(X) \qquad \text{where} \qquad \Delta^{\text{planner}} \in \operatorname*{arg\,min}_{\Delta \in \{E,O\}} \mathcal{L}^*_\Delta. \tag{4}$$

It is immediate that $p^\star(X, \Delta)$ minimizes the planner's expected loss. We therefore study the optimal design choice by comparing the minimized loss of the social planner when under implementing the experiment versus under implementing the observational study. More generally, we can use similar logic to compare the effectiveness of any two designs.

[6] Whenever b_O has non zero expectation we can think of X recentered by its expectation, that can be learned through meta-analyses (see, e.g., Rhys Bernard et al. (2024)).

Definition 4. Design Δ is *planner-preferred* to design Δ' if $\mathcal{L}^*_\Delta < \mathcal{L}^*_{\Delta'}$.

It is immediate that if a design is planner-preferred to a worthwhile design Δ', then Δ is worthwhile. In particular, Proposition 2 implies that if a design Δ is planner-preferred to a cheap, unbiased design Δ', then Δ is worthwhile.

If the experiment is cheap, then its higher research cost is irrelevant to the planner. Therefore, the experiment is planner-preferred to the observational study.

Proposition 4. *Under Assumption 1, if E is cheap, then E is planner-preferred to O.*

This result suggests that any experiment that is cheap is always planner-preferred no matter of their cost, since here, by construction we assume that the mean-squared error S^2_E of the experiment is smaller than the one of the observational study S^2_O.

When the experiment is expensive, the situation is more delicate. Implementing the experiment requires committing to publish more results, which may be costly for a planner. When the publication or attention costs c_p are large enough, the costs of publishing more results outweighs the benefits of a more precise design.

Theorem 1 (Main comparison between experiment and observational study)**.** *Under Assumption 1, suppose that O is expensive.*

(a) If $\text{PostVarRed}(E) - \text{PostVarRed}(O) \geq \left(1 - \frac{C_O}{C_E}\right)c_p$, *then E is planner-preferred to O.*

(b) If $\text{PostVarRed}(E) - \text{PostVarRed}(O) \leq \left(C_E - \frac{1+C_O}{2}\right)c_p$, *then O is planner-preferred to E.*

If instead O is cheap, then (a) and (b) hold with $P(|X(O)| \geq \gamma^*_O)$ in lieu of C_O.

Intuitively the comparison between an experiment and observational study must depend on the posterior variance reduction of each study (which itself depends on their "quality" captured through their mean-squared error) and the *costs* of each study and of attention.

As the cost of attention c_p increases, the planner's preference shifts from a more accurate design to a *less* accurate design with a smaller cost. This is in contrast with the intuition that, for higher attention cost, we should published studies with lower mean squared error. This is because, for costly studies, the planner's must internalize not only the effect of the mean-squared error on the audience's loss function, but also the research cost associated with the study. Whenever she can publish fewer results (c_p is larger), more costly experiments impose more stringent constraints on the publication rules, making those undesiderable for the planner. This is suggestive that medical studies with synthetic control groups may be preferred over placebo groups when the cost of the placebo group is sufficiently large.

Table 1 summarizes the implication of our findings as the cost of the experiment C_E and the cost of attention c_p vary. Similarly, Figure 2 shows how the planner's preference

Cost of experiment	Cost of publication	MSE comparison	Chosen design
C_E small	c_p any	$S_E^2 - S_O^2 < 0$	Experiment
C_E any	c_p any	$S_E^2 - S_O^2 > 0$	Obs studies
C_E any	c_p small	$S_E^2 - S_O^2 < 0$	Experiment
C_E large	c_p any	$S_E^2 - S_O^2 \ll 0$	Experiment
C_E large	c_p large	$S_E^2 - S_O^2 \not\ll 0$	Obs studies

Table 1: Practical implications of Theorem 1 when O is a cheap design. The table summarizes the intuition that for cheap experiments, we should prefer experiments over observational studies if their mean squared error is smaller than the one of an observational. For expensive experiments, comparisons also depend on the cost of the experiment.

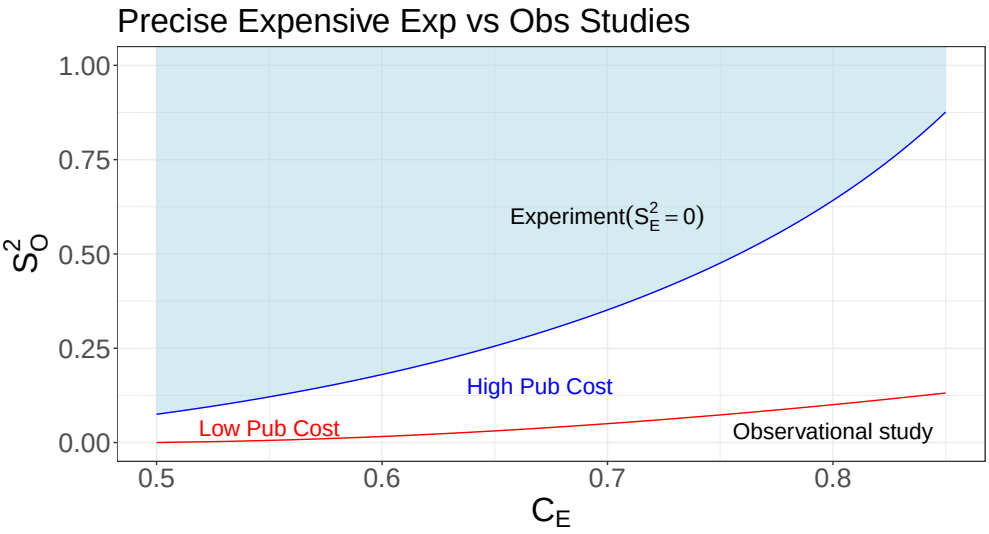


Figure 2: The figure reports on the x-axis the cost of a non-cheap experiment with no variance ($S_E^2 = 0$), and on the y-axis the rescaled squared error S_O^2/η_0^2 of a cheap observational study using the exact expression. The red line denotes the frontier of values for cost of publication $c_p/\eta_0^2 = 0.5$, and the blue lines for the higher cost of publication $c_p/\eta_0^2 = 1$. The region above the blue line denotes the set of values under which an experiment is preferred for a large cost of publication, and the area above the red line denotes the low cost of publication. The figure shows that an observational study may be preferred over an experiment, even if its mean squared error is larger (but not too large), whenever the experiment is sufficiently costly. As the cost of publication increases, the planner prefers more the observational study.

shifts towards noisy observational studies even when the experiment has *no* variance and the observational study is a cheap study.

These results also have implications for the choice of optimal accuracy in an experiment. Suppose that the sample size of the experiment can be chosen freely, and affects the research costs linearly—i.e., $C_\Delta = \frac{c_v}{S_\Delta^2} + c_f$, where c_v denotes a variable cost of sample size, and c_f

denotes a fixed costs of a design. Then, the optimal design is expensive, but involves a sample size that is bounded in terms of the costs.

Corollary 3. *Consider a family of designs Δ that are unbiased on average, where $C_\Delta = \frac{c_v}{S^2_\Delta} + c_f$. Each design Δ in the family that minimizes $\mathcal{L}^*_\Delta$ is expensive and satisfies $S^2_\Delta > \frac{c_v}{1-c_f}$.*

Intuitively, for cheap designs, it is always optimal to increase the sample size to obtain more precise results (Proposition 4). By contrast, once the sample size is large enough (i.e., variance small enough) that the design is expensive, increasing the sample size while preserving individual rationality inefficiently directs attention to unsurprising results. As a result, the optimal sample size is bounded above.

3 Publication rules under non-verifiable designs

In this section, we turn to optimal publication rules when researchers choose the research design Δ using information of the statistics drawn in the experiment. Here, the design Δ (and its corresponding bias β_Δ) are unknown to the social planner but not to the researcher. Specifically, we consider the following model.

Assumption 2. Consider a class of designs $\Delta \in \mathbb{R}$, with $\beta_\Delta = \Delta$ known to the researcher but not observable to the social planner, and $X(\Delta) = \theta_0 + \beta_\Delta + \varepsilon, \varepsilon|\theta_0 \sim \mathcal{N}(0, S^2_E)$. The researcher observes $\theta_0 + \varepsilon$ and chooses Δ (and so β_Δ) to maximizes her realized payoff $v_p(X(\Delta), \Delta)$ in Equation (3), after observing $X(\Delta)$. Let $C_\Delta = c_d|\beta_\Delta| + C_0$, where $c_d < \infty$ and $C_0 < 1$. The social planner chooses $p(X)$ as a function of X only (and not Δ).

Figure 3 illustrates the model: the researcher chooses deterministically the bias of the reported statistic. They, however, pay a cost C_Δ increasing in the bias. We think of the researcher's action as a stylized description of data manipulation or selective reporting. In particular, researchers, after looking at the data, can change their specification by, e.g., changing the covariates in a regression, winsorizing the data in particular ways, or making other design choices functions of the statistics. In our stylized description, this is approximated by defining X as the sum of $\theta_0 + \varepsilon$ plus a bias arising from manipulation. The component $c_d|\beta_\Delta|$ of the cost C_Δ captures reputational or computational costs associated with the manipulation, assumed to be increasing and linear in the magnitude $|\beta_\Delta|$ of the bias. The researcher observes $\theta_0 + \varepsilon$, and hence, she maximizes her *realized* utility conditional on the observed statistics when choosing Δ.

As we discuss in Section 1, the audience (but not the planner) is unaware of the possibility that published findings may have some data manipulation.

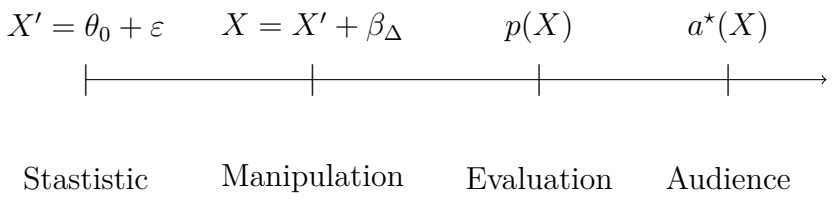


Figure 3: Illustration of the variables in the model. First, researchers observe the vector of statistics. They then manipulate the design by introducing a bias into the statistics and maximize their private utility. The social planner does not observe the bias, and evaluate the study based on a publication function $p(X)$ that only depends on the statistics X.

Finally, note that we assume that the variance of the residual noise ε is independent of Δ and equal to S_E^2. Our results will be valid for any S_E^2 including $S_E^2 \approx 0$ as in large observational studies. We interpret this assumption as stating that standard errors are verifiable as part of the publication process; we therefore focus on manipulation that introduces unobserved bias in the reported results.

3.1 Optimal publication rule under manipulation

The planner cannot observe or verify β_Δ, knows S_E^2, and minimizes expected loss over (θ_0, ε) taking into account the researcher's (endogeneous) incentives to manipulate their results. That is, we define an optimal publication rule as

$$p^\star \in \arg\max_{p \in \mathcal{P}} \mathbb{E}_{X,\theta_0}\left[\mathcal{L}_p(X(\Delta_p^\star), \Delta_p^\star, \theta_0)\right], \quad \Delta_p^\star = \arg\max_{\Delta} v_p(X(\Delta), \Delta) \tag{5}$$

where $\mathcal{P}$ is the set of all Borel measurable functions $p(X, d)$ constant in d (i.e., do not depend on the design), which we write without loss as $p(X)$ (note that p may implicitly depend on S_E^2). Here, X satisfies Assumption 2 and $\Delta_p^\star$ denotes the optimal researcher's response. The publication rule only depends on X but not on Δ, as this is chosen privately by the researcher. Before introducing our next theorem we introduce the following definition.

Definition 5. A *linearly smoothed cutoff rule* with cutoff X^* and slope m is defined by

$$p_{X^*,m}(X) = \begin{cases} 0 & \text{if } |X| \leq X^* - \frac{1}{m} \\ 1 - m(X^* - |X|) & \text{if } X^* - \frac{1}{m} < |X| < X^* \\ 1 & \text{if } |X| \geq X^* \end{cases}.$$

A linearly smoothed cutoff rule considered here is a deterministic publication rule above and below thresholds $(X^* - \frac{1}{m}, X^*)$ for given m and it randomizes the publication chances between these two values, with publication probability increasing in the value of the reported statistic $|X|$. To gain further intuition, consider the threshold $\gamma_E^* = \frac{S_E^2+\eta_0^2}{\eta_0^2}\sqrt{c_p}$. Then $X^* = \gamma_E^*$ and $m = \infty$ corresponds to a publication rule in Frankel and Kasy (2022), i.e., a publication rule for cheap experiments. When $m < \infty$ and $X^* > \gamma_E^*$, the linearly smoothed cutoff rule publishes with probability one more surprising results than rules in the absence of manipulation, while randomizing around the threshold. In the following theorem we characterize the optimal publication probability in contexts with manipulation.

Theorem 2 (Optimal publication probability)**.** *Under Assumption 2:*

(a) There exists a cutoff $X^* \in \left(\gamma_E^*, \gamma_E^* + \frac{1-c_e}{c_d}\right)$ *such that the linearly smoothed cutoff rule* p_{X^*,c_d} *is optimal.*

(b) For each optimal publication rule p*, there exists* $X^* \in \left(\gamma_E^*, \gamma_E^* + \frac{1-C_0}{c_d}\right)$ *such that* $p^\star(X) = p_{X^*,c_d}(X)$ *(resp.* $p^\star(X) \leq c_e$*) for almost all* $X \geq 0$ *with* $p_{X^*,c_d}(X) > C_0$ *(resp.* $p_{X^*,c_d}(X) \leq C_0$*).*

Theorem 2 characterizes the optimal publication rule under manipulation. The rule is a linearly smoothed cutoff rule that: (i) does not publish results below a certain cutoff $X^* - \frac{1}{c_d}$; (ii) randomizes whenever X is above the $X^* - \frac{1}{c_d}$ but below X^*; and (iii) publishes with probability 1 for results $X \geq X^*$.

To gain further intuition, it is instructive to compare this with the optimal publication rule for a *cheap* experiment in Proposition 1. For simplicity, suppose $C_0 = 0$ so that we abstract from fixed research costs. Consider first a scenario in which the social planner ignores manipulation. Then we would observe *bunching* around the cutoff to publish γ_E^*, as researchers with $\theta_0 + \varepsilon \in (\gamma_E^* - \frac{1}{c_d}, \gamma_E^*)$ would introduce a bias to publish. Researchers with $\theta_0 + \varepsilon < \gamma_E^* - \frac{1}{c_d}$ would find it nonprofitable to introduce any bias (as the cost would not compensate the benefits) and therefore would not publish.

Next, suppose that the planner introduces randomization in the publication rule whenever $X \in (\gamma_E^* - \frac{1}{c_d}, \gamma_E^*)$ as described in Theorem 2. It follows that the randomization device in Theorem 2 makes the researcher indifferent between manipulating and not manipulating the data, at the cost of publishing some unsurprising results. However, this choice is sub-optimal as the planner may publish too many unsurprising results.

The last step is to increase the threshold X^* to compensate the loss from publishing unsurprising results. In summary, optimal publication rules with manipulation (a) increase the cutoff for publication; (b) allow for randomization at the margin below the cutoff.

Action	Testable observation	Published findings	Manipulation
Select optimal cutoff rule without manipulation	Large bunching	Only surprising findings are published	Large
Add randomization below cutoff	No bunching	Many unsurprising findings are published	None
Increase the cutoff + randomize below	Some bunching	Some unsurprising findings are published	Some

Table 2: Sequence of actions to decrease the loss function. We start from a simple cutoff rule for cheap experiment wrongly assuming no data manipulation. Under this rule, we observe bunching of X around the cutoff. We then randomize below the cutoff to make researcher indifferent between manipulation and not. Finally, we increase the cutoff to publish less unsurprising results.

Table 2 illustrates main features associated with each action. The first action is equivalent to ignoring manipulation, as in settings considered in Frankel and Kasy (2022). Manipulation in this scenario has testable implications since it introduces large bunching around the cutoff, see, e.g., the discussion in Elliott et al. (2022) among others. Through sufficient randomization below the cutoff, we can guarantee no bunching (and no manipulation). The last step is then to increase the cutoff. We should then observe just some bunching above the old cutoff and below the new one (see Proposition 5). This corresponds to the planner's preferred publication rule.

3.2 Some implications

In this section, we explore some of the implications of Theorem 2. The first implication we study is whether, in the *planner's preferred equilibrium*, some researcher would manipulate their findings.

Proposition 5 (Manipulation in equilibrium)**.** *Under the model in Assumption 2, under the publication rule in Theorem 2 and the planner's preferred equilibrium for* $|\theta_0 + \varepsilon| \in (\gamma_E^*, X^*)$*, we have* $|\beta_{\Delta_{p^\star}^\star}| > 0$*.*

Proposition 5 states that some manipulation should occur just below the threshold. Intuitively, if researchers could not manipulate their findings, then at the optimum, we would have published findings above the threshold γ_E^*. However, because manipulation can occur, the social planner increases the threshold from γ_E^* to X^*. It follows that manipulation is beneficial in the interval between the old threshold γ_E^* and the "new" threshold X^*. It guarantees that findings that should have been published in the absence of manipulation do get published, while the cost of manipulation is second order. This result suggests that in equilibrium, we should not force manipulation to be exactly zero. A practical implication of

Proposition 5 is that we observe a discontinuous distribution of X below X^* and above γ_E^* *at the optimum.*

As a second exercise, we study the role of fixed costs in the presence of manipulation.

Proposition 6 (Implementation costs)**.** *Under the model in Assumption 2, for $C_0 \geq 1 - c_d\gamma_E^*$, the set of optimal cutoffs X^* in Theorem 2 is strictly decreasing in C_0 in the strong set order.*

Proposition 6 shows that for sufficiently costly studies, the social planner *lowers* the cutoff as the fixed cost increases. This result is suggestive that less surprising results may be published more when the study is sufficiently costly. It differs from what we found in Theorem 1 in the absence of manipulation: in the absence of manipulation, the planner may force the researcher to run cheap over expensive studies. However, with manipulation, fixed costs imply a lower chance of manipulation, motivating increasing the chance of publication for such studies.

Next, we states that some findings below the optimal cutoff γ_E^* *in the absence* of data manipulation, do get published in the planner's preferred equilibrium under manipulation.

Proposition 7 (Some unsurprising results are published)**.** *Under the model in Assumption 2, under the publication rule in Theorem 2 and the planner's preferred equilibrium, for some values of $|\theta_0 + \varepsilon| < \gamma_E^*$, we have $p^\star(X) > 0$ but $\beta_{\Delta_{p^*}^*} = 0$.*

Proposition 7 shows that some unsurprising results that would have not been published without manipulation do get published under the optimal publication rule with manipulation. This feature is not due to manipulation of these results, but rather to deter manipulation. It is in stark contrast with the case without manipulation, as in Frankel and Kasy (2022).

4 Implications for alternative publication mechanisms

In this section we study the implications of our results for alternative publication mechanisms, such incentivizing more observational studies with possible data-manipulation or costly experiments with no data manipulation (e.g., when enforcing pre-analysis plans) or introducing a registry of non-surprising results.

4.1 Planner's choice between possibly manipulated design or costly designs with no manipulation

We first study the implications of Theorem 2 for pre-analysis plans.

Definition 6 (Pre-analysis vs possible manipulation)**.** Consider the following two possible families of designs, all of which have variance S^2.

(A) Experiment with pre-analysis plan: Experiment E is unbiased on average and worthwhile. Researchers cannot manipulate their findings ($c_d = \infty$), truthfully report $X = \theta_0 + \varepsilon$, and pay a resaerch cost C_E. The social planner chooses a publication rule $p^\star_E$ as in Definition 1 and incurs a loss $\mathcal{L}^*_E$

(B) Possible manipulation (observational studies): Assumption 2 holds, so there is a family of designs $\Delta \in \mathbb{R}$. The manipulation cost is $c_d < \infty$, and researchers can manipulate their findings after observing $\theta_0 + \varepsilon$. The fixed research cost is $C_0 = 0$. The social planner chooses a publication rule $p^\star$ as in Equation (5) and incurs corresponding expected loss $\mathcal{L}^*_M = \mathbb{E}_{X,\theta_0}\left[\mathcal{L}_{p^\star}(X(\Delta^\star_{p^\star}), \Delta^\star_{p^\star}, \theta_0)\right]$ under the planner's preferred equilibrium.

Scenario (A) states that the researcher cannot manipulate her findings but pays a fixed cost c_e, interpreted as the cost that guarantees no bias in the study. This setting may correspond to the cost of conducting an experiment (including time for grant approval, research assistants, etc.).[7] Scenario (B) allows for manipulation of the findings, with the researcher not required to write a pre-analysis plan. In this case, we assume for simplicity no fixed costs $c_e = 0$, but possible (reputational) costs associated with manipulation. We think of (B) as scenarios where an experiment or observational study is already available to the researcher (and therefore, its cost is sunk).

Proposition 8 (Value of pre-analysis plan)**.** *Under (A) and (B) in Definition 6:*

(i) If E is a cheap design, then $\mathcal{L}^*_E < \mathcal{L}^*_M$.

(ii) If E is an expensive design, then

$$\frac{1}{\eta_0^2}\Big(\mathcal{L}^*_E - \mathcal{L}^*_M\Big) \geq \frac{c_p}{\eta_0^2}(c_e - \lambda^*_E) - \frac{1}{\eta_0^2 c_d^2} + O\Big((1 - C_E)^3 + \frac{c_p^2}{\eta_0^4}\Big)$$

where $\lambda^*_E = 1 - \frac{1}{\sqrt{2\pi}3}\frac{c_p^{1/2}}{\eta_0}\sqrt{\frac{S_O^2}{\eta_0^2} + 1}$.

In Proposition 8, we show that an experiment (with a pre-analysis plan) is preferred over an observational study with possible manipulation (with no pre-analysis plan) if the cost of the observational study is sufficiently low.

When the cost of the experiment is high, the trade-off depends on (i) its cost c_e and (ii) the cost of data manipulation c_d. Clearly, if $c_d = 0$, then an experiment with a pre-analysis

[7] For simplicity, we consider here experiments that are worthwhile.

plan plan may be preferred. Intuitively, with a low cost of manipulation, the planner must publish with positive probability a larger set of unsurprising results, at a possibly large publication cost.

Suppose instead $c_d < \infty$. Then, an observational study with possible manipulation is preferred for sufficiently high researcher's perceived cost. This is despite the cost of the experiment c_e being private and paid by the researcher and not by the planner.

Different from (and complementary to) models where either no experimentation or no reputational costs occur (e.g., Kasy and Spiess, 2023; Spiess, 2024), these results illustrate *trade-offs* in the choice of unbiased vs. biased studies. A larger experiment cost may decrease the supply of research, making the planner prefer possible manipulation.

4.2 Registry of unsurprising results

A core theme is that if attention is costly, publication decisions must take into account the *private* research cost. It is natural to ask whether we can think of publication mechanisms that reward studies without paying such a cost of attention.

For instance, consider a scenario where the researcher conducts a (possibly) expensive study with cost C_E in the model in Section 2 – i.e., abstracting from manipulation for simplicity. Because attention is costly, we only want to publish sufficiently surprising results. Suppose that we implement the first-best policy, ignoring the researcher's costs. This takes the form $p_E^\star = 1\{|X| \geq \gamma_E^*\}$ of a threshold crossing rule as described in Section 2.1.

In our model in Section 2, if the experiment is costly, however, the researcher may choose the outside option. To guarantee implementability of $p_E^\star$ suppose now we can reward the researcher whose results are not published in the journal by publishing the results in a repository of unsurprising results. This repository may have no cost c_p and guarantee a reward R to the researcher. The researcher's objective function upon conducting the experiment reads as

$$\mathbb{E}[v_{p_E^\star}(X(E), E)] = \mathbb{P}(|X(E)| \geq \gamma_E^*) + R\mathbb{P}(|X(E)| \leq \gamma_E^*) - C_E.$$

When choosing $R \geq \frac{C_E+1}{\mathbb{P}(|X(E)|\geq\gamma_E^*)} - 1$, the researcher is always indifferent between conducting or not the research study. Rewarding results through a "repository of unsurprising results", with a reward increasing in the cost of the experiment guarantees that the first-best policy is implementable at no attention cost for the audience.

5 Conclusion

This paper studies how researcher's incentives shape the optimal design of the scientific process. Ignoring the researcher's incentives, it is optimal to publish the most surprising results (Frankel and Kasy, 2022). When researcher's incentives matter, we show that optimal publication rules depend on private costs of research and incentives for research manipulation.

As a first exercise, we show that, in the absence of manipulation, the planner prefers (quasi)experiments with lower accuracy (larger mean-squared error) over sufficiently costly experiments. In medical studies, a pre-specified synthetic control group obtained from medical records (Jahanshahi et al., 2021; Food and Drug Administration, 2023) may be preferred over a sufficiently expensive placebo control group, despite lack of randomization.

With manipulation, we show that it is optimal to (i) publish some unsurprising results and (ii) knowlingly allow for manipulation (biased studies) at the margin. Observationally, the optimal policy would reduce the bunching of the findings (e.g., t-tests) around the publication cutoff that would be optimal if the planner naively assumed researchers do not manipulate. However, the optimal policy does not completely remove bunching. Even when the planner can mandate a signal to enforce no manipulation, this may not be her preferred policy when the signal entails a large perceived burden on the researcher.

Our model disentangles the contribution of design choice and data manipulation to optimal publication decisions. Future research should study more complex communication strategies. For example, in contexts with pre-analysis plans, the planner may allow multiple signals to decrease the burden on the authors or allow for the publication of non-prespecified findings. Similarly, researchers may report multiple findings in their study. As we point to in our results in Section 4, studying more complex action spaces can shed light on alternative mechanisms relevant to scientific communication.

Future research should also study the implications of some of these conclusions on empirical research. Our framework relies on parameters that capture costs and benefits for the researcher. In the absence of manipulation, such parameters can be learned using cost data for medical trials (e.g., Tetenov, 2016) and field experiments (e.g., Viviano et al., 2024). With manipulation, and possible reputational costs associated with it, such parameters may be learned using information from meta-analyses through the distribution of the submitted findings (different form, but in the spirit of Andrews and Kasy, 2019). In this case, experimentation or different norms across different journals can help to identify such parameters.

References

Abadie, A. (2020). Statistical nonsignificance in empirical economics. *American Economic Review: Insights 2*(2), 193–208.

Allcott, H. (2015). Site selection bias in program evaluation. *Quarterly Journal of Economics 130*(3), 1117–1165.

Ana, J. (2004). The role of a general medical journal. *BMJ 328*(7439), 591.

Andrews, I. and M. Kasy (2019). Identification of and correction for publication bias. *American Economic Review 109*(8), 2766–2794.

Andrews, I. and J. M. Shapiro (2021). A model of scientific communication. *Econometrica 89*(5), 2117–2142.

Banerjee, A., S. Chassang, and E. Snowberg (2017). Chapter 4 – Decision theoretic approaches to experiment design and external validity. In A. V. Banerjee and E. Duflo (Eds.), *Handbook of Field Experiments*, Volume 1 of *Handbook of Economic Field Experiments*, pp. 141–174. North-Holland.

Banerjee, A., E. Duflo, A. Finkelstein, L. F. Katz, B. A. Olken, and A. Sautmann (2020). In praise of moderation: Suggestions for the scope and use of pre-analysis plans for RCTs in economics. Working paper, National Bureau of Economic Research.

Banerjee, A. V., S. Chassang, S. Montero, and E. Snowberg (2020). A theory of experimenters: Robustness, randomization, and balance. *American Economic Review 110*(4), 1206–1230.

Bates, S., M. I. Jordan, M. Sklar, and J. A. Soloff (2022). Principal-agent hypothesis testing. arXiv:2205.06812.

Bates, S., M. I. Jordan, M. Sklar, and J. A. Soloff (2023). Incentive-theoretic Bayesian inference for collaborative science. arXiv:2307.03748.

Brodeur, A., N. Cook, and A. Heyes (2020). Methods matter: P-hacking and publication bias in causal analysis in economics. *American Economic Review 110*(11), 3634–3660.

Chassang, S., G. Padro I Miquel, and E. Snowberg (2012). Selective trials: A principal-agent approach to randomized controlled experiments. *American Economic Review 102*(4), 1279–1309.

DeMaria, A. N. (2022). The role of a medical journal. *Structural Heart 6*(4), 100079.

Di Tillio, A., M. Ottaviani, and P. N. Sørensen (2021). Strategic sample selection. *Econometrica 89*(2), 911–953.

Egger, D., J. Haushofer, E. Miguel, P. Niehaus, and M. Walker (2022). General equilibrium effects of cash transfers: Experimental evidence from Kenya. *Econometrica 90*(6), 2603–2643.

Elliott, G., N. Kudrin, and K. Wüthrich (2022). Detecting p-hacking. *Econometrica 90*(2), 887–906.

Food and Drug Administration (2023). Considerations for the design and conduct of externally controlled trials for drug and biological products. Available at https://www.fda.gov/regulatory-information/search-fda-guidance-documents/considerations-design-and-conduct-externally-controlled-trials-drug-and-biological-products.

Frankel, A. and M. Kasy (2022). Which findings should be published? *American Economic Journal: Microeconomics 14*(1), 1–38.

Gechter, M. and R. Meager (2022). Combining experimental and observational studies in meta-analysis: A debiasing approach. Working paper, Pennylvania State University and London School of Economics.

Glaeser, E. L. (2006). Researcher incentives and empirical methods. Working paper, National Bureau of Economic Research.

Grabowski, H., J. Vernon, and J. A. DiMasi (2002). Returns on research and development for 1990s new drug introductions. *Pharmacoeconomics 20*, 11–29.

Henry, E. and M. Ottaviani (2019). Research and the approval process: the organization of persuasion. *American Economic Review 109*(3), 911–55.

Hirano, K. and J. R. Porter (2009). Asymptotics for statistical treatment rules. *Econometrica 77*(5), 1683–1701.

Jahanshahi, M., K. Gregg, G. Davis, A. Ndu, V. Miller, J. Vockley, C. Ollivier, T. Franolic, and S. Sakai (2021). The use of external controls in FDA regulatory decision making. *Therapeutic Innovation & Regulatory Science 55*(5), 1019–1035.

Kasy, M. and J. Spiess (2023). Optimal pre-analysis plans: Statistical decisions subject to implementability. Working paper, University of Oxford and Stanford University.

Kitagawa, T. and A. Tetenov (2018). Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica 86*(2), 591–616.

Kitagawa, T. and P. Vu (2023). Optimal publication rules for evidence-based policy. Working paper, Brown University.

Libgober, J. (2022). False positives and transparency. *American Economic Journal: Microeconomics 14*(2), 478–505.

Manski, C. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica 72*(4), 1221–1246.

McCloskey, A. and P. Michaillat (2022). Incentive-compatible critical values. Working paper, National Bureau of Economic Research.

Miguel, E. (2021). Evidence on research transparency in economics. *Journal of Economic Perspectives 35*(3), 193–214.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies 38*(2), 175–208.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–73.

Olken, B. A. (2015). Promises and perils of pre-analysis plans. *Journal of Economic Perspectives 29*(3), 61–80.

Popat, S., S. V. Liu, N. Scheuer, G. G. Hsu, A. Lockhart, S. V. Ramagopalan, F. Griesinger, and V. Subbiah (2022). Addressing challenges with real-world synthetic control arms to demonstrate the comparative effectiveness of pralsetinib in non-small cell lung cancer. *Nature Communications 13*(1), 3500.

Raices Cruz, I., M. C. Troffaes, J. Lindström, and U. Sahlin (2022). A robust Bayesian bias-adjusted random effects model for consideration of uncertainty about bias terms in evidence synthesis. *Statistics in Medicine 41*(17), 3365–3379.

Rhys Bernard, D., G. Bryan, S. Chabé-Ferret, J. De Quidt, J. Fliegner, and R. Rathelot (2024). How much should we trust observational estimates? accumulating evidence using randomized controlled trials with imperfect compliance. Working Paper, Toulouse School of Economics.

Riveros, C., A. Dechartres, E. Perrodeau, R. Haneef, I. Boutron, and P. Ravaud (2013). Timing and completeness of trial results posted at clinicaltrials. gov and published in journals. *PLoS medicine 10*(12), e1001566.

Rosenzweig, M. R. and C. Udry (2020). External validity in a stochastic world: Evidence from low-income countries. *Review of Economic Studies 87*(1), 343–381.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical association 46*(253), 55–67.

Shinohara, K., A. Tajika, H. Imai, N. Takeshima, Y. Hayasaka, and T. A. Furukawa (2015). Protocol registration and selective outcome reporting in recent psychiatry trials: New antidepressants and cognitive behavioural therapies. *Acta Psychiatrica Scandinavica 132*(6), 489–498.

Spiess, J. (2024). Optimal estimation when researcher and social preferences are misaligned. Working paper, Stanford University.

Tetenov, A. (2012). Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics 166*(1), 157–165.

Tetenov, A. (2016). An economic theory of statistical testing. Working Paper, University of Bristol.

Thompson, S. G. and J. A. Barber (2000). How should cost data in pragmatic randomised trials be analysed? *BMJ 320*(7243), 1197–1200.

Thorlund, K., L. Dron, J. J. Park, and E. J. Mills (2020). Synthetic and external controls in clinical trials–a primer for researchers. *Clinical Epidemiology*, 457–467.

Viviano, D., K. Wuthrich, and P. Niehaus (2024). A model of multiple hypothesis testing. arXiv:2104.13367.

Wald, A. (1950). *Statistical Decision Functions*. Wiley.

Williams, C. (2021). Preregistration and incentives. Working paper, Durham University.

Wong, H.-H., A. Jessup, A. Sertkaya, A. Birkenbach, A. Berlind, and J. Eyraud (2014). Examination of clinical trial costs and barriers for drug development final. *Office of the Assistant Secretary for Planning and Evaluation, US Department of Health & Human Services*, 1–92.

Yin, X., P. S. Mishra-Kalyan, R. Sridhara, M. D. Stewart, E. A. Stuart, and R. C. Davi (2022). Exploring the potential of external control arms created from patient level data: a case study in non-small cell lung cancer. *Journal of Biopharmaceutical Statistics 32*(1), 204–218.

Yoder, N. (2022). Designing incentives for heterogeneous researchers. *Journal of Political Economy 130*(8), 2018–2054.