

Which Findings Should Be Published?[†]

By ALEXANDER FRANKEL AND MAXIMILIAN KASY*

Given a scarcity of journal space, what is the optimal rule for whether an empirical finding should be published? Suppose publications inform the public about a policy-relevant state. Then journals should publish extreme results, meaning ones that move beliefs sufficiently. This optimal rule may take the form of a one- or two-sided test comparing a point estimate to the prior mean, with critical values determined by a cost-benefit analysis. Consideration of future studies may additionally justify the publication of precise null results. If one insists that standard inference remain valid, however, publication must not select on the study's findings. (JEL D61, D82, D83, L82)

Not all empirical findings get published. Journals may be more likely to publish findings that are statistically significant, as documented, for instance, by Franco, Malhotra, and Simonovits (2014); Brodeur et al. (2016); and Andrews and Kasy (2019). They may also be more likely to publish findings that are surprising or, conversely, ones that confirm some prior belief. Whatever its form, selective publication distorts statistical inference. If only estimates with large effect sizes were to be written up and published, say, then published studies would systematically overstate true effects. Such publication bias has been offered as one explanation for the perceived replication crisis in the social and life sciences.¹

In response to these concerns, there have been calls for reforms in the direction of nonselective publication. One proposal is to promote statistical practices that de-emphasize statistical significance, for instance by banning “stars” in regression tables. Another proposal is for journals to adopt Registered Reports in which preregistered analysis plans are reviewed and accepted prior to data collection (see Nosek and Lakens 2014 or Chambers et al. 2014). This has been implemented, for

*Frankel: University of Chicago Booth School of Business (email: afrankel@chicagobooth.edu); Kasy: Department of Economics, University of Oxford (email: maximilian.kasy@economics.ox.ac.uk). Leeat Yariv was coeditor for this article. We thank Isaiah Andrews, Victoria Baranov, Ben Brooks, Gary Chamberlain, Zoe Hitzig, Emir Kamenica, Matthew Notowidigdo, Jesse Shapiro, and Glen Weyl for helpful discussions and comments.

[†]Go to <https://doi.org/10.1257/mic.20190133> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹Worries about selective publication go back at least to Sterling (1959). Discussions of publication bias and other threats to the credibility and reproducibility of scientific output can be found in Ioannidis (2005, 2008) and in reviews including Simmons, Nelson, and Simonsohn (2011); Gelman and Loken (2014); and Christensen and Miguel (2016). Open Science Collaboration (2015) and Camerer et al. (2016) conduct large-scale replications of experimental studies in psychology and economics, giving insight into the extent to which published results are in fact reproducible.

instance, at the *Journal of Development Economics*. Registered Reports guarantee that publication will not select at all on findings—after a plan is accepted, the journal is committed to publishing the study, and the researcher has no flexibility over which results to write up.

This paper seeks the optimal rule for determining whether a study should be published given both its design (which determines a study's standard error) and its findings (the point estimate). Our analysis is from an instrumental perspective: the value of a study is that it informs the public about some policy-relevant state of the world before the public chooses a policy action. In this framework, we will show that nonselective publication is not, in fact, optimal. Some findings are more valuable to publish than others. Since selective publication distorts statistical inference, this implies a trade-off between policy relevance and statistical credibility.

In a world without constraints, the first-best rule would be for all results—or even better, all raw data—to be published. This paper solves for a second-best publication rule. In particular, we take as given that there is some cost of publication. One interpretation is that the cost is an opportunity cost of shifting a public's attention away from other studies. In that case, we can think of our paper as asking which findings should be published in top journals, i.e., ones where the results are more likely to be noticed.

The basics of our model are as follows. If a submitted study is published, the public observes its results and takes the optimal policy action given its updated belief. If a study is not published, the public never observes the results and does not necessarily know that a study was conducted; the public then takes a default action. This default action in the absence of publication is based on a default belief. We allow for either a Bayesian public whose default belief correctly accounts for publication bias or a naive public whose default belief always remains at its prior. The optimal publication rule is the one that maximizes the public's expected payoff from the eventual policy choice minus the publication cost.

The optimal publication rule will select on a study's findings. To understand why, observe that there is no instrumental value from publishing a study with a “null result” that doesn't move the policy away from the default action. Publishing a null result incurs a cost without a benefit: the same policy would have been chosen even if the study weren't published. The studies that are worth publishing are the ones that show a payoff gain from taking an action other than the default.

In order to generally characterize optimal publication rules, consider “supermodular” policy decisions in which the public's preferred policy action is monotonic in the state of the world. For example, the preferred investment in a public good increases in its expected return. In any supermodular environment, Theorem 1 establishes that it is more valuable to publish studies with more extreme results. These are the ones that move beliefs and actions further from the defaults.

For two canonical special cases, we give a more explicit description of optimal publication rules. In the first case, the public makes a continuous policy decision, such as the choice of a tax rate, and has quadratic losses in matching the policy to its ideal point. The optimal publication rule then takes the form of a “two-sided test”: the journal publishes estimates that are sufficiently far above or below the prior mean. In the second case, the public makes a binary policy choice, such as

whether to implement a job training program. Here, the optimal rule is a “one-sided test.” For instance, in the absence of publication, suppose that the public, by default, does not implement the program. The journal then publishes positive estimates of the program’s value, ones that convince the public to implement it, but not negative estimates. For these two-sided and one-sided tests, the critical values that determine publication come from a cost-benefit calculation. They do not correspond to a conventional significance level such as $p = 0.05$ against a null hypothesis of 0.

After characterizing optimal publication rules, we return to the distortions caused by selective publication. In Theorem 2, we show that under any rule in which the publication probability depends on the point estimate, common forms of frequentist inference will be invalid after conditioning on publication. Point estimates are no longer unbiased, for instance, and uncorrected likelihood-based estimation will be flawed. Moreover, when a study is not published, a naive public that maintains its prior will have a distorted belief relative to a Bayesian public that accounts for publication bias. If we desire that standard inference or naive updating be valid, we must impose a nonselective publication rule that does not depend at all on the point estimate (although journals may still publish studies with small standard errors over studies with large standard errors).

Putting these results together, we see that selectively publishing extreme results is better for policy relevance but leads to distorted inference. Therefore, a move away from the current (selective) publication regime toward nonselective publication in order to improve statistical credibility might have costs as well as benefits.

An abstraction in the model as described is that it considers a “static” environment with a single paper to be published and a single action to be taken. One may also be interested in the longer-term implications of publication rules, as in McElreath and Smaldino (2015) and Nissen et al. (2016).² To get some insight into these issues, we consider a dynamic extension to our model that appends a second period in which exogenous information arrives before another action is taken. The publication decision in the first period now affects the actions in both periods. Theorem 3 characterizes the optimal publication rule for this two-period model. Here, we find a new benefit of publishing null results that don’t change the current action. Publishing a null result today helps avoid future mistakes arising from the noise in the information that has yet to arrive.

In addition to the broad conclusions described above, the paper derives a number of comparative statics on optimal publication rules. Some are straightforward: for example, it is more valuable to publish a given point estimate when the standard error is smaller, and correspondingly, in the two-period model, it is more valuable to publish a precise null result than a noisy null result. We believe that some of these comparative statics, however, would have been difficult to intuit without a formal model. For instance, consider the canonical quadratic loss policy environment. Suppose that we fix the statistical significance of some result (the t -statistic of the point estimate relative to the prior mean) while varying the point estimate. Point estimates close to the prior

²McElreath and Smaldino (2015) and Nissen et al. (2016) provide dynamic models to study whether an academic publication process with publication bias will eventually converge to truthful estimates. Akerlof and Michailat (2017) perform a similar exercise for a more evolutionary form of the accumulation of academic knowledge.

move the posterior mean very little, and thus, naturally, publishing those results is not valuable. But point estimates too far from the prior also barely move the posterior. So at any specified statistical significance level, we only publish point estimates that are an intermediate distance from the prior mean. A similar nonmonotonicity occurs in the two-period model: publishing null results today is most valuable when the future information is expected to be neither too precise nor too noisy.

We wish to stress that the nature of our exercise is to solve for the socially optimal rule regarding whether to publish a study that has some given set of results. That is, our model does not consider the incentives of researchers or journals, and we are not attempting to characterize the equilibrium publication rule arising from a strategic interaction of these agents. As discussed in Glaeser (2006), researcher incentives play an important role in the publication process. Researchers make choices over the topics they study and their study designs and then may selectively submit or possibly even manipulate their findings.³ We do explore one way in which researcher study design choices may respond to journal publication rules in online Appendix C.3.

Our derivations of optimal publication rules rely on characterizing the value of information for specified decision problems. Most theoretical treatments of the value of information study the ex ante value of an experiment, i.e., the expected value prior to the realization; see classic treatments in Blackwell (1953); Lehmann (1988); or Persico (2000). These ex ante comparisons are relevant for a characterization of nonselective publication rules, as we examine in Proposition 3. However, we generally allow for publication to select on a study's findings. We are thus predominantly concerned with the ex post value of information given an experiment's realization, as studied in Frankel and Kamenica (2019).

The decision to reveal a signal, at a cost, based on its realization is also related to the analysis of the discretionary disclosure of product quality in Jovanovic (1982) or of accounting news in Verrecchia (1983) as well as in follow-up work. As in those papers, we find that information is disclosed only if it is sufficiently valuable ex post. Those papers, however, focus on a private value of disclosing information that may contain positive or negative news about one's type. We instead consider a social value of information in making better decisions.

The rest of the paper is structured as follows. Section I introduces our basic model of publication. Section II shows how to solve for the optimal publication rules and provides some characterizations of the solution. Section III addresses the distortions that arise from selective rather than nonselective publication. Section IV presents a two-period version of the model. Section V concludes and presents some extensions that we explore further in the online Appendix. We consider publication objectives that do not arise from policy-based welfare and instead derive from an inherent value of learning the state or publishing point estimates close to the truth, we show how a journal might adjust its publication rule if submitted study designs are endogenous to the publication rule, and we discuss the possibility that

³Furukawa (2018) looks at a model (without journals) in which researcher decisions to publish papers interact with a public policy choice, where in equilibrium, researchers choose to publish papers with extreme results. Müller-Iten (2016) looks at a competition between journals for prestige in which journals choose whether to publish a submitted study based on a signal of its quality while researchers choose which journal to submit to.

a study's findings may be informative about the quality of its design. Proofs are in the Appendix.

I. The Model of Publication

A. Setup

There is an uncertain state of the world whose value is relevant to some public policy decision. A study that reveals information about this state may or may not arrive, i.e., may or may not be submitted to a journal. If a study arrives, the journal decides whether to publish it. If it is published, the results of the study are observed by the public. Finally, the public chooses a policy.

Let $\theta \in \Theta \subseteq \mathbb{R}$ denote the state of the world, and suppose there is a common prior π_0 on θ shared by the public and the journal. The probability that a study arrives is $q \in (0, 1]$, independent of θ . A study is summarized by the random variables (X, S) , where S is drawn from some distribution F_S on $\mathcal{S} = \mathbb{R}_{++}$ and X is normally distributed on $\mathcal{X} = \mathbb{R}$ according to $X|\theta, S \sim \mathcal{N}(\theta, S^2)$. Call X the *point estimate* and S the *standard error*. In particular, S reflects the study's design, containing no direct information about the state θ . By contrast, X is the study's finding, which is informative about θ ; its information content depends on S .

If a study arrives, it will be evaluated by a journal that observes (X, S) and then decides whether to publish the study. The journal uses a publication rule $p: \mathcal{X} \times \mathcal{S} \rightarrow [0, 1]$, where $p(X, S)$ describes the probability that a study (X, S) is published. We say that the journal publishes a study when $p(X, S) = 1$ and does not publish a study when $p(X, S) = 0$. Let $D = 0$ denote the event that no study is published (because no study arrived or because one arrived but was not published) and $D = 1$ the event that a study arrives and is published.

After a study is published or not, the public's belief on θ updates to a posterior π_1 . When no study has been published ($D = 0$), π_1 is equal to some *default belief* π_1^0 . When a study has been published ($D = 1$), S and X are publicly observed, and π_1 is instead equal to the belief $\pi_1^{(X, S)}$. See Section IB below for the description of the belief updating process that determines π_1^0 and $\pi_1^{(X, S)}$.

Given updated beliefs π_1 , the public takes a policy action $a \in \mathcal{A} \subseteq \mathbb{R}$ to maximize its expectation of a utility function $U: \mathcal{A} \times \Theta \rightarrow \mathbb{R}$. Let $a^*(\pi_1) \in \operatorname{argmax}_a E_{\theta \sim \pi_1} [U(a, \theta)]$ indicate the chosen action when the public holds beliefs π_1 . We assume existence of this argmax for any relevant utility functions and posterior distributions, and we confirm existence for all of our examples. Let $a^0 = a^*(\pi_1^0)$ be the *default action*, i.e., the action taken under the default belief, whereas $a^*(\pi_1^{(X, S)})$ is the action taken if a study (X, S) is published.

Social welfare, corresponding to the shared objective of both the journal and public, is the action payoff net of a publication cost. Let $c > 0$ indicate the social cost of publication. The welfare $W(D, a, \theta)$ induced by publication D , chosen action a , and state of the world θ is

$$(1) \quad W(D, a, \theta) = U(a, \theta) - Dc.$$

We will search for the publication rule p that maximizes the ex ante expectation of welfare, which we call the *optimal* publication rule.

B. Belief Updating

If a study (X, S) is published, the public's posterior belief is $\pi_1 = \pi_1^{(X, S)}$. We assume that $\pi_1^{(X, S)}$ is derived according to Bayes' rule given the signal $X|\theta, S \sim \mathcal{N}(\theta, S^2)$. Denote the pdf of a standard normal distribution by φ . By Bayes' rule, recalling that S is independent of θ , the density of $\pi_1^{(X, S)}$ relative to the prior π_0 is given by

$$(2) \quad \frac{d\pi_1^{(X, S)}}{d\pi_0}(\theta) = \frac{\varphi((X - \theta)/S)}{\int \varphi((X - \theta')/S) d\pi_0(\theta')}.$$

Since the journal and the public share a common prior, we see that $\pi_1^{(X, S)}$ also represents the journal's Bayesian belief after it observes a submitted study (X, S) . As such, we often refer to $\pi_1^{(X, S)}$ as the *interim belief* that the journal holds when evaluating a paper for publication.

If a study is not published, then the public's posterior belief is $\pi_1 = \pi_1^0$, the default belief. We consider the possibility of two distinct updating rules to determine π_1^0 . *Bayesian updating* is the sophisticated rule that correctly accounts for selection induced by the publication process. *Naive updating* is the unsophisticated rule that ignores the possibility of unpublished studies and fails to account for selection.

Bayesian Updating: When no study is published, the public understands that this event could have occurred because no study arrived (probability $1 - q$) or because a study arrived (probability q) and was unpublished (probability $1 - p(X, S)$, with $\theta \sim \pi_0$, $S \sim F_S$, and $X|\theta, S \sim \mathcal{N}(\theta, S^2)$). The public then updates beliefs on θ to π_1^0 according to Bayes' rule.⁴ Denote this Bayesian default belief under publication rule p by $\pi_1^{0,p}$; its density relative to the prior π_0 is given by

$$(3) \quad \frac{d\pi_1^{0,p}}{d\pi_0}(\Theta) = \frac{1 - qE[p(X, S)|\theta]}{1 - qE[p(X, S)]}.$$

Naive Updating: The public's default belief, π_1^0 , is equal to its prior: $\pi_1^0 = \pi_0$.

While Bayesian updating is "correct" in the fully specified model, we consider naive updating to be, in many cases, a realistic description of updating. One can interpret a naive public as having an incorrect model of the world: in the absence of seeing a publication, the public is unaware of the possibility that a study might have been submitted and rejected. Alternatively, naive beliefs arise as the limiting

⁴ If a publication rule publishes with probability one and the probability of study arrival is $q = 1$, then nonpublication is a zero probability event and beliefs are not pinned down by Bayes' rule. As a convention, we then let the Bayesian default belief be equal to the prior π_0 .

Bayesian beliefs when $q \rightarrow 0$, i.e., a rational public that did not expect a study to be submitted on this topic.

C. Leading Examples of Priors and Utility Functions

A typical state of the world θ estimated in an empirical economics study might be a demand or supply elasticity, the magnitude of a treatment effect, or the net benefit of implementing a program. Our leading example for a prior distribution will be the *normal prior*, in which case $\Theta = \mathbb{R}$ and π_0 is $\mathcal{N}(\mu_0, \sigma_0^2)$, with $\mu_0 \in \mathbb{R}$ and $\sigma_0 \in \mathbb{R}_{++}$. With a normal prior, the posterior belief after observing a study (X, S) is given by

$$(4) \quad \pi_1^{(X,S)} = \mathcal{N}\left(\frac{\sigma_0^2}{S^2 + \sigma_0^2}X + \frac{S^2}{S^2 + \sigma_0^2}\mu_0, \frac{S^2\sigma_0^2}{S^2 + \sigma_0^2}\right).$$

Two utility functions we consider are quadratic loss and binary action utility. The *quadratic loss* utility function has $\mathcal{A} = \mathbb{R}$ and $U(a, \theta) = -(a - \theta)^2$. This is a canonical utility function for a public that makes a continuous policy decision a , with the state θ representing the public's ideal point. Under quadratic loss utility, the maximizing action choice given belief π_1 is $a^*(\pi_1) = E_{\theta \sim \pi_1}[\theta]$.

The *binary action* utility function has $\mathcal{A} = \{0, 1\}$ and $U(a, \theta) = a \cdot \theta$. Here, there is a binary decision, such as implementing a program ($a = 1$) or not ($a = 0$). The state θ represents the net benefit of implementation. The chosen action is then $a^*(\pi_1) = \mathbf{1}(E_{\theta \sim \pi_1}[\theta] > 0)$, where $\mathbf{1}$ is the indicator function (taking $a = 0$ at indifference).

D. Interpretations of the Model

As mentioned in the introduction, not all research findings get published. The acceptance rate at top economics journals is now below 6 percent (Card and DellaVigna 2013). This paper can be viewed as solving for the optimal publication rule conditional on any fixed share of studies to be published. The cost c would then be a shadow cost on this publication constraint. Alternatively, there may be a genuine opportunity cost of the public's attention: if the public has a limited ability to process information, publishing one study pulls attention from others. To the extent that papers in high-ranking, selective journals receive disproportionate attention and influence, one can interpret our analysis as characterizing which papers should be published in these top journals.

In this model, the decision to publish is based on the objective of maximizing policy-based welfare. We discuss alternative social objectives in Section V and online Appendices C.1 and C.2. The model also assumes that information arrives only once, abstracting from the possibility that later studies would further change the public's beliefs and policies. Section IV extends the model to explore the decision to publish a study today when the journal expects more information to arrive in the future.

Throughout the paper, we make a simplifying assumption that a study is summarized by a normal signal with a point estimate X and standard error S . Note that one

might interpret the variable S to, in fact, be larger than a study's reported standard error, which only captures uncertainty due to sampling variation. That could occur if the study has limited external validity, meaning that the estimated parameter is only partially informative about the policy parameter of interest. Violations of the identifying assumptions required for internal validity could also add noise to the estimate. We discuss some considerations that may arise when S is not fully observed by the journal in online Appendix C.4.

II. Optimal Publication Rules

A. Characterizing the Optimum

Write out the ex ante expected welfare given some publication rule p and default action a^0 as $EW(p, a^0)$:

$$(5) \quad EW(p, a^0) = E \left[qp(X, S) \left(U(a^*(\pi_1^{(X, S)}), \theta) - c \right) + (1 - qp(X, S)) U(a^0, \theta) \right],$$

where the expectation is taken with respect to $\theta \sim \pi_0$, $S \sim F_S$, and $X|\theta, S \sim \mathcal{N}(\theta, S^2)$. Given a specified updating rule, the optimal publication rule p maximizes expected welfare $EW(p, a^0)$ for the appropriately determined default action a^0 . Under naive updating, p is chosen to maximize $EW(p, a^*(\pi_0))$. Under Bayesian updating, p is chosen to maximize $EW(p, a^*(\pi_1^{0,p}))$.

While we seek to maximize ex ante welfare, it is helpful to consider the journal's interim problem after a study has been submitted. The journal observes (X, S) and has interim belief π given by $\pi = \pi_1^{(X, S)}$. At this interim belief, the journal can evaluate the expected payoff from publication (leading to public belief $\pi_1 = \pi$ and action $a^*(\pi)$) and from nonpublication (leading to public belief $\pi_1 = \pi_1^0$ and action a^0). Denote by $\Delta(\pi, a^0)$ the gross interim benefit—not including publication costs—of publishing a study that induces interim belief π given default action a^0 :⁵

$$(6) \quad \Delta(\pi, a^0) = E_{\theta \sim \pi} [U(a^*(\pi), \theta) - U(a^0, \theta)].$$

Say that publication rule p is *interim optimal* given default action a^0 if it (almost surely) publishes a study when $\Delta(\pi_1^{(X, S)}, a^0) > c$ and does not publish when $\Delta(\pi_1^{(X, S)}, a^0) < c$.

Under naive updating, expected welfare is straightforwardly maximized by choosing the interim optimal publication rule given default action $a^0 = a^*(\pi_0)$. Under Bayesian updating, the default action a^0 depends on the publication rule. However, we find that the Bayesian optimal publication rule is also interim optimal against the default action that it induces.

⁵Frankel and Kamenica (2019) provide general characterizations of these gross interim benefit functions—the value of information, in an ex post sense—across decision problems.

LEMMA 1: *Under either naive or Bayesian updating, let p be an optimal publication rule, and let π_1^0 be the induced default belief. The publication rule p is interim optimal given default action $a^*(\pi_1^0)$.*

In other words, even under Bayesian updating, the journal's publication rule is a best response to the public's default action.⁶ To show this result, we first establish that for any fixed publication rule p , the Bayesian default action $a^*(\pi_1^{0,p})$ maximizes expected welfare $EW(p, a^0)$ over choice of a^0 . Under Bayesian updating, then, the journal and public can be thought of as engaging in a sequential game of common interest. The value of such a game is unchanged when the journal moves first (as we posit) or moves second: $\max_p \max_{a^0} EW(p, a^0) = \max_{a^0} \max_p EW(p, a^0)$.

Lemma 1 is a key result for characterizing optimal policies under Bayesian updating. First, the lemma implies that—as with naive updating—any characterization of interim optimal publication rules extends to optimal publication rules. Second, it simplifies the maximization program we use to solve for the optimal Bayesian policy: instead of maximizing over all publication rules, we can restrict attention to publication rules that are interim optimal given some default action.

One immediate corollary of Lemma 1 is as follows. Define a study (X, S) to be a *null result* if publishing the study does not change the optimal action from the default action, i.e., if $a^*(\pi_1^{(X,S)}) = a^0$.⁷

OBSERVATION 1 (Do Not Publish Null Results): *The gross interim benefit of publishing a null result is zero. Therefore, by Lemma 1, the optimal publication rule does not publish null results.*

Hence, it is only ever optimal to publish studies that move the public's beliefs—and its corresponding action—away from the default. Intuitively, we would expect that “extreme” findings—ones that move beliefs and actions further from the defaults—are more valuable to publish than “moderate” findings. The following condition on utility functions allows us to formalize this result.

Say that a utility function $U : \mathcal{A} \times \Theta \rightarrow \mathbb{R}$ is *supermodular* if for all $\underline{a} < \bar{a}$ and $\underline{\theta} < \bar{\theta}$, it holds that $U(\bar{a}, \bar{\theta}) + U(\underline{a}, \underline{\theta}) \geq U(\underline{a}, \bar{\theta}) + U(\bar{a}, \underline{\theta})$. Supermodular utilities guarantee that the public takes higher actions when it believes that the state is higher. Quadratic loss and binary action utilities are both supermodular.

⁶In a game with different timing in which the journal could not commit to a publication rule, one might define a publication rule p and default belief π_1^0 as constituting a Bayesian Nash equilibrium if they jointly satisfy (i) the publication rule p is interim optimal given default action $a^0 = a^*(\pi_1^0)$ and (ii) the default belief is Bayesian, i.e., $\pi_1^0 = \pi_1^{0,p}$. Our notion of optimality under Bayesian updating does not impose (i); nevertheless, Lemma 1 below clarifies that (i) will in fact be satisfied for an optimal publication rule. Hence, any optimal publication rule would induce a Bayesian Nash equilibrium, but there may exist Bayesian Nash equilibria that are not optimal.

⁷Whether a study is a “null result” in this sense depends on the default belief. Our definition differs from its common usage, which often refers to a point estimate that is not statistically significantly different from zero.

THEOREM 1: *Fix either updating rule. Let the utility function U be supermodular. Under an optimal publication rule, for every standard error $S = s$, there exists an interval $I \subseteq \mathbb{R}$ such that (X, s) is published if and only if $X \notin I$.⁸*

So with supermodular utilities, the journal publishes point estimates that are sufficiently high, sufficiently low, or both. Putting together Observation 1 and Theorem 1, we see that under these utility functions, the journal indeed publishes extreme findings—ones that lead to extreme beliefs and actions relative to the default.

The logic behind Theorem 1 is straightforward. Higher point estimates X yield higher interim beliefs on the state (in the sense of first-order stochastic dominance), because point estimates at a fixed standard error are ordered by the monotone likelihood ratio property. Under supermodularity, higher beliefs increase the benefit of taking higher actions. So for any low point estimate that yields an optimal action below the default, there would be a greater interim benefit of publishing an even lower point estimate. Likewise, for any high point estimate that yields an action above the default, there would be a greater interim benefit of publishing an even higher point estimate. Hence, it is interim optimal to publish point estimates that are sufficiently high or sufficiently low. Finally, Lemma 1 implies that this characterization of interim optimal rules extends to the optimal publication rule as well.

The next subsection derives explicit solutions for optimal publication rules for our leading examples of normal priors and quadratic loss or binary action utility functions.

B. Examples of Optimal Publication Rules

In this section, we present the optimal publication rules for our two leading example utility functions, quadratic loss and binary action, under the assumption of normal priors. Appendix A provides the derivations of these policies for the case of Bayesian updating; for naive updating, we can simply look for the interim optimal policy against the prior belief. It turns out that for these example utility functions and priors, Bayesian updating and naive updating yield identical optimal policies. Indeed, Appendix A presents more general conditions on priors under which Bayesian and naive updating yield identical publication rules for these two utility functions.

Quadratic Loss Utility with Normal Priors.—Under quadratic loss utility, welfare is $W(D, a, \theta) = -(a - \theta)^2 - Dc$ for $a \in \mathcal{A} = \mathbb{R}$. The public chooses an action equal to its posterior mean belief about the state. So when the default action is a^0 , the gross interim benefit of publishing a study (X, S) that induces a belief with mean $\mu_1^{(X,S)}$ evaluates to $(\mu_1^{(X,S)} - a^0)^2$. The interim optimal publication rule is therefore to publish if $|\mu_1^{(X,S)} - a^0| \geq \sqrt{c}$.

⁸It is possible that $I = \emptyset$, in which case all studies with $S = s$ are published, or that $I = \mathbb{R}$, in which no studies at $S = s$ are published.

Under normal priors (for which $\mu_1^{(X,S)}$ is given by (4)), Proposition 1, part (i) establishes that the above rule is optimal—for Bayesian as well as naive updating—with the default action $a^0 = \mu_0$. Parts (ii) and (iii) provide comparative statics.

PROPOSITION 1: *Suppose there is quadratic loss utility and a normal prior.*

- (i) *Under either Bayesian or naive updating, it is optimal to publish a study (X, S) if and only if $|X - \mu_0| \geq (1 + S^2/\sigma_0^2)\sqrt{c}$, i.e., $|X - \mu_0|/S \geq (1/S + S/\sigma_0^2)\sqrt{c}$.*
- (ii) *The publication cutoff $(1 + S^2/\sigma_0^2)\sqrt{c}$ in terms of the difference of the point estimate from the prior mean is independent of the study arrival probability q and the mean μ_0 . It is larger when the standard error S is larger, the prior variance σ_0^2 is smaller, or the cost of publication c is larger.*
- (iii) *The publication cutoff $(1/S + S/\sigma_0^2)\sqrt{c}$ in terms of the magnitude of the t -statistic is nonmonotonic and convex in the standard error S : it has a minimum at $S = \sigma_0$ and goes to infinity as $S \rightarrow 0$ or $S \rightarrow \infty$.*

The publication rule described in Proposition 1, part (i) corresponds to a “two-sided test” in which the journal publishes if the point estimate is sufficiently high or sufficiently low. Equivalently, we can restate the publication rule in terms of a two-sided test for the t -statistic $(X - \mu_0)/S$. See Figure 1.

The form of a two-sided test is, of course, familiar from the null hypothesis significance testing paradigm. However, we wish to highlight two ways in which our policy is distinct from two-sided tests as they are traditionally applied. First, we compare the point estimate X to the prior mean, not to some other point, e.g., a null hypothesis of $\theta = 0$. Second, the cutoff for publication is not given by a conventional value, such as a t -statistic of 1.96 corresponding to a p -value of 0.05. The cutoff is instead determined by a cost-benefit analysis.

Proposition 1, part (ii) finds that a given point estimate of X moves beliefs more and thus makes publication more likely—in the sense of a smaller cutoff value for $|X - \mu_0|$ —when the standard error S is smaller or when the prior uncertainty σ_0 is larger. Likewise, publication is more likely when the cost of publication c is lower.

When deciding to publish at a given t -statistic rather than a point estimate, Proposition 1, part (iii) finds a nonmonotonic publication threshold as a function of S . (For other parameters, the comparative statics in terms of the t -statistic would be identical to those on the point estimate.) For a precise study with a low standard error or an imprecise one with a high standard error, the journal requires a high t -statistic to be willing to publish; for a study of intermediate precision, the journal publishes at a lower t -statistic. The journal is most willing to publish a study at a given t -statistic when the standard error S is equal to σ_0 , the standard deviation of the prior.

To gain intuition on this nonmonotonic comparative static, recall that here the journal publishes studies that move the interim mean sufficiently far from the prior mean. Fix a prior mean and standard deviation of $\mu_0 = 0$ and $\sigma_0 = 1$, and consider




FIGURE 1. OPTIMAL PUBLICATION REGION SHADED FOR QUADRATIC LOSS UTILITY, NORMAL PRIOR

different studies that might arrive with a given t -statistic $t = X/S$, say, $t = 4$. If a very precise study ($S \simeq 0$) arrives with $t = 4$, the point estimate must have been close to 0, so it moves the mean very little. As we scale up the point estimate and standard error while keeping $t = 4$, the posterior mean moves higher, peaking at a mean of 2 when $X = 4$ and $S = 1$. Increasing the point estimate and standard error further, the mean falls back toward 0, because the study becomes too noisy (relative to the prior) to move beliefs much.⁹ In other words, fixing the “statistical significance” as measured by the t -statistic, the change in mean first grows and then declines in the “practical significance” as measured by the magnitude of the point estimate.

Binary Action Utility with Normal Priors.—Under binary action utility, welfare is given by $W(D, a, \theta) = a\theta - Dc$ for $a \in \mathcal{A} = \{0, 1\}$. The public chooses action $a = 0$ if its posterior mean belief about the state is weakly less than 0 and action $a = 1$ if the posterior mean is positive. So for the default action $a^0 = 0$, the gross interim benefit of publishing a study (X, S) inducing belief $\pi_1^{(X,S)}$ with mean $\mu_1^{(X,S)}$ evaluates to $\max\{0, \mu_1^{(X,S)}\}$. For the default action $a^0 = 1$, the gross interim benefit is instead $\max\{0, -\mu_1^{(X,S)}\}$.

Under normal priors with prior mean normalized to $\mu_0 \leq 0$, Proposition 2 establishes that the optimal rule—for Bayesian as well as naive updating—yields $a^0 = 0$. The corresponding (interim) optimal publication rule is therefore to publish if $\mu_1^{(X,S)} \geq c$. This publication rule corresponds to a “one-sided test.” At any given standard error, a study is published only if the point estimate is sufficiently high.

⁹The general formula for the change in mean given a t -statistic $t = (X - \mu_0)/S$ and standard error S —with a corresponding point estimate of $X = \mu_0 + tS$ —is $t(\sigma_0^2 S / (\sigma_0^2 + S^2))$. To understand why this change in mean falls to zero when we fix t and take $S \rightarrow \infty$, recall that the interim mean is a weighted average of the prior mean μ_0 and the point estimate X . The point estimate scales linearly with S , but the weight on the point estimate is proportional to $1/S^2$.




FIGURE 2. OPTIMAL PUBLICATION REGION SHADED FOR BINARY ACTION UTILITY, NORMAL PRIOR

PROPOSITION 2: *Suppose there is binary action utility and a normal prior with $\mu_0 \leq 0$. Then under either Bayesian or naive updating, it is optimal to publish a study if and only if $X \geq (1 + S^2/\sigma_0^2)c - (S^2/\sigma_0^2)\mu_0$, i.e., $(X - \mu_0)/S \geq (1/S + S/\sigma_0^2)(c - \mu_0)$.*

Proposition 3 in online Appendix D.1 provides comparative statics for how the publication cutoff varies with parameters. Most of the comparative statics are analogous to those from the quadratic loss publication rule in Proposition 1. However, the two policies depend differently on the prior mean. Suppose we fix a point estimate $X > 0$ and we consider prior means $\mu_0 < 0$. With quadratic loss utility, increasing μ_0 toward 0 would make the journal less willing to publish: there will be a smaller difference $X - \mu_0$, and therefore, the posterior mean will be closer to the prior mean. With binary actions, increasing μ_0 toward 0 makes the journal more willing to publish: the posterior mean will be higher in absolute terms, indicating that the benefit of switching from $a = 0$ to $a = 1$ is higher.

III. Selective Publication and Statistical Inference

The key conclusion from the previous section was that welfare-maximizing publication rules should selectively publish extreme findings over moderate findings. This conclusion, of course, contrasts with calls for reform aimed at eliminating selection. Selective publication is understood to distort inference and harm replicability. See, for instance, Rosenthal (1979) and Ioannidis (2005) on how standard inference from published results can be inaccurate when publication is based on statistical significance.

To study these issues in our model, say that a publication rule p is *nonselective* if $p(x, s)$ is constant in the point estimate x for each standard error s ¹⁰ and otherwise

¹⁰We mean by this statement that $p(x, s)$ is constant in x almost surely over realizations of X , i.e., that $\Pr(p(X, s) = E[p(X, s)|S = s]|\theta, S = s) = 1$ for all θ . Nothing changes if $p(x, s)$ may vary with x on sets of X that can only occur with zero probability given $\theta, S = s$.

is *selective*. That is, nonselective publication rules do not condition publication on the study's findings. Nonselective publication rules may still condition on the standard error (i.e., on the study's design), which is independent of the state.

We can now see how selective publication may distort inference. Recall that the point estimate X is drawn from the distribution $X|\theta, S \sim \mathcal{N}(\theta, S^2)$ with density $f_{X|\theta, S}(x|\theta, s) = \varphi((x - \theta)/s)/s$. Conventional statistical inference on θ from X and S is based on this density. Under publication rule p , however, the density of X conditional on publication ($D = 1$) is instead

$$(7) \quad f_{X|\theta, S, D=1}(x|\theta, s) = \frac{p(x, s)}{E[p(X, S)|\theta, S = s]} \cdot f_{X|\theta, S}(x|\theta, s).$$

If publication is nonselective, the density conditional on publication $f_{X|\theta, S, D=1}$ matches $f_{X|\theta, S}$. With selective publication, these densities may differ. In that case, conventional inferences would be flawed. We present a set of novel results below showing that a nonselective publication rule is not just sufficient but also necessary for the validity of standard inference in a number of senses.

A. Distortions from Selection

Let Φ denote the cdf of a standard normal distribution.

THEOREM 2: *Suppose that there is an open set $\Theta_0 \subseteq \mathbb{R}$ contained in the support of the prior distribution of θ . Fix some standard error $s > 0$. Each of the following conditions holds if and only if the publication rule p is nonselective:*

- (i) (Unbiasedness) *For each $\theta \in \Theta_0$, $E[X|\theta, S = s, D = 1] = \theta$.*
- (ii) (Publication probability constant in state) *The publication probability conditional on standard error, $E[p(X, S)|\theta, S = s]$, is constant over $\theta \in \Theta_0$.*
- (iii) (Undistorted naive updating) *For all distributions F_S on S , the Bayesian default belief $\pi_1^{0,p}$ is equal to the naive default belief π_0 .*

Suppose additionally that $\Theta = \mathbb{R}$ and that the publication rule takes the form of $p(x, s) = \mathbf{1}(x \notin I(s))$ for some interval $I(s) \subsetneq \mathbb{R}$, where nonselective publication corresponds to $I(s) = \emptyset$. Fix some critical value $z > 0$. Then the publication rule p is nonselective if and only if

- (iv) (Size control of confidence intervals) *For each $\theta' \in \Theta$, $\Pr(\theta' \in [X - zs, X + zs]|\theta = \theta', S = s, D = 1) \geq \Phi(z) - \Phi(-z)$.*

As suggested above, when the publication rule is nonselective, the conclusions of parts (i)–(iv) of Theorem 2—the validity of conventional inferences—hold fairly straightforwardly. The novel results of the theorem are the converses, that each part

in turn implies nonselective publication.¹¹ We will first go over the interpretation of each part, and we then discuss the intuition behind their proofs.

Interpreting Theorem 2.—Part (i) of Theorem 2 establishes that after conditioning on publication, selective publication implies that the point estimate X is a biased estimator for the state θ . For instance, suppose that a study is only published when $|X| > 1.96 \cdot S$; this is the conventional cutoff for rejecting a null hypothesis of $\theta = 0$ at the 95 percent statistical significance level. In that case, the bias $E[X|\theta, S = s, D = 1] - \theta$ will be positive if the state θ is above 0 and will be negative if the state is below 0. Figure 3, panel A illustrates the bias conditional on publication when $S = 1$ and when a study is only published if $|X| > 1.96$.¹²

Part (ii) establishes that if publication is selective ($p(x, s)$ varies with x), then the publication probability $E[p(X, S)|\theta, S = s]$ conditional on any standard error $S = s$ and θ varies with the state θ . From equation (7), if the publication probability varies with θ , then so too does the ratio $f_{X|\theta, S, D=1}(x|\theta, s) / f_{X|\theta, S}(x|\theta, s)$. Selective publication therefore renders invalid uncorrected likelihood-based inference, such as maximum likelihood estimation or likelihood ratio tests. For the example in which a study is published if and only if $|X| > 1.96 \cdot S$, Figure 3, panel B fixes $S = 1$ and then shows how the publication probability depends on θ . Here, we see that the publication probability falls as the state θ gets closer to 0, the center of the nonpublication interval.

Part (iii) says that the Bayesian default belief is equal to the naive default belief (i.e., the prior) for every possible distribution of standard errors if and only if publication is nonselective. To interpret this result, think of a “partially sophisticated” public. This public is aware that studies may sometimes go unpublished and thus that naive updating may lead to distorted beliefs. But it does not know the study arrival rate q or the distribution of standard errors F_S , and it may not even have a well-specified prior over these objects. Therefore, it does not know how to correct these distortions. Under a nonselective publication rule, though, the partially sophisticated public can be confident in updating naively. For any q and any F_S , the Bayesian default belief is guaranteed to be equal to the naive one.¹³

Figure 3, panel C illustrates how the Bayesian and naive default beliefs differ under the selective publication rule that publishes if $|X| > 1.96 \cdot S$. When no publication is observed, a Bayesian understands that there may have been a study with point estimate $X \simeq 0$ that was submitted but went unpublished. Hence, the Bayesian default belief places a higher posterior probability on θ close to 0 and a lower probability on θ far from 0.¹⁴ The figure assumes a prior of $\theta \sim \mathcal{N}(0, 4)$, study arrival rate $q = 1$, and that the standard error is drawn as $S = 1$ with certainty.

¹¹ For parts (i), (ii), and (iv) of the theorem, if publication is nonselective, then the results hold for each s and/or z , and if the results hold for any s and/or z , then publication is nonselective.

¹² Figures similar to Figure 3, panels A and D can be found in Andrews and Kasy (2019).

¹³ The arrival rate q does not affect whether naive beliefs are distorted. The Bayesian default belief is equal to the prior for some $q \in (0, 1]$ if and only if it is equal to the prior for all $q \in (0, 1]$.

¹⁴ Abadie (2018) demonstrates how a failure to pass a standard statistical significance threshold can be extremely informative when studies are precise. In such cases, the Bayesian default belief diverges greatly from the naive one.

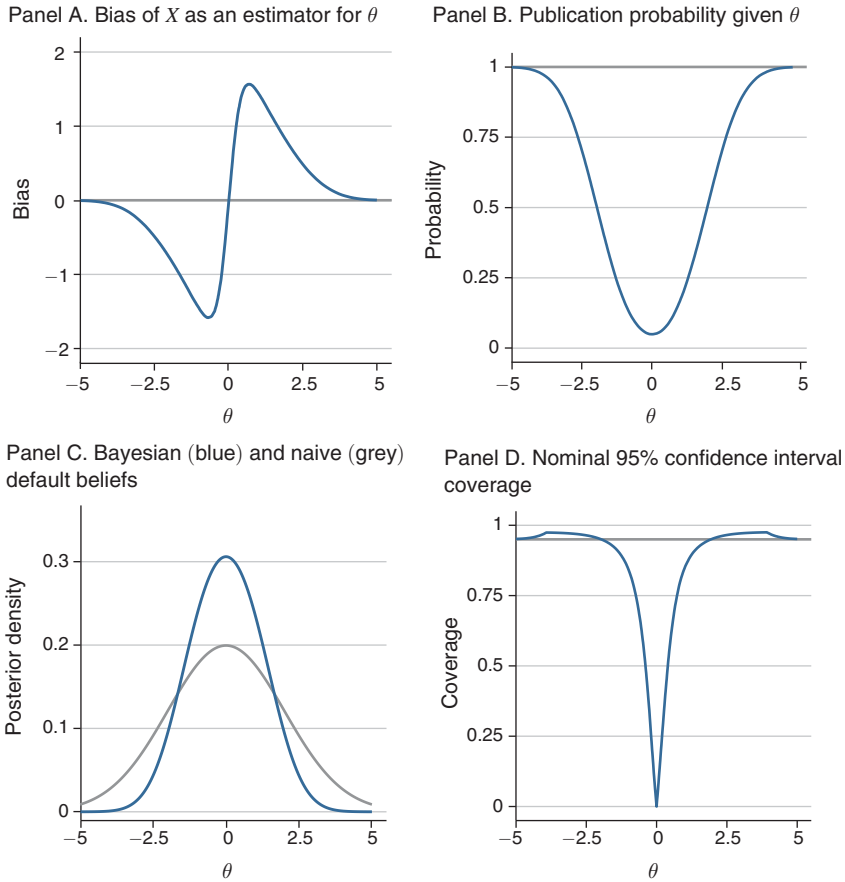


FIGURE 3. DISTORTIONS DUE TO SELECTIVELY PUBLISHING POINT ESTIMATES $|X| > 1.96$

Finally, part (iv) of Theorem 2 restricts attention to the form of publication rule found in Theorem 1, in which studies are published only if the point estimate falls outside of an interval. We find that selective publication of this form will fail to control the size of conventional frequentist confidence intervals. Given standard error $S = s$ and z -score z , there exists some realization of the state θ for which the probability that the interval $[X - zs, X + zs]$ contains θ is less than the nominal confidence level $\Phi(z) - \Phi(-z)$.¹⁵ Figure 3, panel D fixes $S = 1$ and shows the coverage probability of the nominal 95 percent confidence interval $[X - 1.96, X + 1.96]$ as a function of the state when the journal publishes only if $|X| > 1.96$. If the true state is $\theta = 0$, there is in fact zero probability conditional on publication that this

¹⁵In online Appendix D.2, we show that, fixing $z > 0$, there do exist selective publication rules outside of the form of Theorem 1 for which the coverage probability of the confidence interval $[X - zS, X + zS]$ does in fact always equal $\Phi(z) - \Phi(-z)$.

confidence interval contains the state. For some other values of θ , the probability that the confidence interval contains the state is higher than 95 percent.

Intuition for Proof of Theorem 2.—We now discuss the key intuitions behind the proof that parts (i)–(iv) of Theorem 2 each imply nonselective publication.

That part (ii) (a constant publication probability in the state, conditional on the standard error $S = s$) implies nonselective publication follows from the completeness of the normal location family of distributions; see, for instance, Theorem 6.22 in Lehmann and Casella (1998). Completeness is the continuous analog of a full rank condition, applied to the linear operator mapping g , a function of x , to $E[g(X)|\theta, S = s]$, a function of θ . Completeness implies that for any function $g(x)$ for which $E[g(X)|\theta, S = s]$ is constant in θ over an open set, it holds that $g(x)$ must be almost everywhere constant. Plugging in the function $g(x) = p(x, s)$, we see that if the conditional publication probability at state θ is constant in θ , then the publication probability $p(x, s)$ cannot vary with x .

To show that part (i) (unbiasedness) implies part (ii), let $S = 1$ without loss. We leverage “Tweedie’s formula” (Efron 2011), which holds because $\varphi'(x) = -x\varphi(x)$:

$$\int (x - \theta)\varphi(x - \theta)p(x, 1) dx = \partial_{\theta}E[p(X, S)|\theta, S = 1].$$

Unbiasedness implies that the left-hand side of this equality is 0, and thus, so is the right-hand side, yielding part (ii).

Next, to see that part (iii) (undistorted naive updating) implies part (ii), observe from equation (3) that the Bayesian default belief matches the prior—the naive default belief—only if the publication probability unconditional on the standard error, $E[p(X, S)|\theta]$, is constant in θ . If we desire that this hold for all possible distributions over the standard error S , then the publication probability conditional on any given $S = s$ must also be constant, which is exactly part (ii).

Finally, we show that part (iv) (size control of confidence intervals) implies part (ii). Suppose for the sake of contradiction that the publication rule is selective: at $S = s$, it only publishes point estimates X outside a nonempty interval $I(s)$. Then a straightforward calculation shows that there are states θ' in the interior of $I(s)$ for which the confidence interval contains θ' with probability less than the nominal level.

B. *Optimal Nonselective Publication*

The fact that selective publication distorts inference means that there is a trade-off between policy relevance and credibility. The policy relevance criterion pushes toward selectively publishing extreme results. But if we wish standard inference to remain valid, then we must restrict ourselves to nonselective publication rules. What is the *optimal nonselective* publication rule—the rule that maximizes utility subject to the constraint of being nonselective?

PROPOSITION 3: *There exists $\bar{s} \in \mathbb{R}_+ \cup \{\infty\}$ for which an optimal nonselective publication rule is to publish a study (X, S) if and only if $S < \bar{s}$.¹⁶ The optimal nonselective rule is the same under naive and Bayesian updating.*

When the journal is not allowed to screen on the point estimate, the only remaining option is to screen on the standard error. In that case, the journal should publish studies with smaller standard errors over those with larger standard errors. The result follows immediately from the fact that $X \sim \mathcal{N}(\theta, s^2)$ is a Blackwell more informative signal of the state θ and is thus more valuable ex ante, when the standard error s is smaller. Under a normal prior and quadratic loss utility, for instance, we can explicitly solve for the optimal nonselective publication rule. If $\sigma_0^2 \geq c$ (high prior uncertainty, low publication costs), then a study is published if $S \leq \bar{s}$, with $\bar{s} = \sigma_0 \sqrt{(\sigma_0^2/c) - 1}$; and if $\sigma_0^2 < c$ (low prior uncertainty, high costs), then no study is published.¹⁷

The publication region for the optimal nonselective rule neither nests nor is nested in that for the optimal selective rule. The nonselective rule publishes if $S \leq \sigma_0 \sqrt{(\sigma_0^2/c) - 1}$, and the selective rule from Proposition 1 for $S \leq \sigma_0 \sqrt{|X - \mu_0|/\sqrt{c} - 1}$. For X far from the prior mean, we see that the nonselective rule publishes less often.

IV. A Two-Period Model

The model introduced in Section I assumes that a single study is or is not published and then a policy is chosen. If an additional study were to arrive, though, the public might want to switch to a new policy. This new policy would depend on the results of the later study as well as those of the original one (if published). So if a journal expects additional studies to arrive in the future, it faces new considerations when deciding whether to publish a study today.

In order to explore such issues, this section introduces a *two-period model*. As before, there is an unknown policy-relevant state of the world θ , which we take to be persistent over time. The original model of publication and policy choice—which we now refer to as the one-period model—makes up the first period. The new second period captures, in reduced form, the impact of future studies: additional exogenous information arrives, and the public takes another action.

A. Setup of the Two-Period Model

At the start of the game, the common prior over θ is π_0 . In the first period, a study is submitted to a journal with probability q . If the study arrives, it has point estimate and standard error (X_1, S_1) with $S_1 \sim F_{S_1}$ and $X_1 | \theta, S_1 \sim \mathcal{N}(\theta, S_1^2)$. The study

¹⁶For $\bar{s} = 0$, no study would be published.

¹⁷If a nonselective publication rule is used and no publication is observed, then the default belief will be π_0 (under either updating rule) so the expected welfare will be $-\text{var}_{\theta \sim \pi_0}[\theta] = -\sigma_0^2$. Conditional on S but not X , the expected welfare of publishing can be solved for as $\sigma_0^2 \cdot (S^2/(S^2 + \sigma_0^2)) - c$. An optimal nonselective publication rule publishes a study if $\sigma_0^2 \cdot (S^2/(S^2 + \sigma_0^2)) - c > -\sigma_0^2$.

is published with probability $p(X_1, S_1)$. The public updates to belief π_1 as before, with $\pi_1 = \pi_1^{(X_1, S_1)}$ following publication ($D = 1$), or to default belief $\pi_1 = \pi_1^0$ following nonpublication ($D = 0$). The induced belief $\pi_1^{(X_1, S_1)}$ is given by Bayes' rule, and the default belief π_1^0 may be determined by either naive or Bayesian updating. Then the action a_1 is taken, with $a_1 = a^*(\pi_1) \in \operatorname{argmax}_a E_{\theta \sim \pi_1} [U(a, \theta)]$.

In the second period, an exogenous signal $X_2 \sim \mathcal{N}(\theta, s_2^2)$ —independent of D and of (X_1, S_1) given θ —is publicly observed. Beliefs update according to Bayes' rule from prior π_1 to posterior π_2 . Finally, the action a_2 is taken, with $a_2 = a^*(\pi_2) \in \operatorname{argmax}_a E_{\theta \sim \pi_2} [U(a, \theta)]$.

We assume that the standard error of the second-period signal, s_2 , is a parameter that is known by the journal at the start of the game. Our interpretation is that s_2 would be low (i.e., precise) when the journal expects that other high-quality studies on the topic in question will soon be performed. The parameter s_2 would be high (imprecise) when the journal expects future studies on the topic to be performed infrequently or be of low quality.

Let $\alpha \in [0, 1)$ describe the payoff weight on the first-period action relative to $1 - \alpha$ on the second-period action. Social welfare is the weighted sum of action payoffs minus a cost of publication $c > 0$ incurred if a study is published:

$$(8) \quad W(D, a_1, a_2, \theta) = \alpha U(a_1, \theta) - Dc + (1 - \alpha)U(a_2, \theta).$$

A *dynamically optimal* publication rule maximizes the ex ante expectation of welfare from (8). For a journal that has period 1 interim belief π_1^I after observing a study and faces default belief π_1^0 , the *gross interim benefit* of publishing the study, denoted $\Delta(\pi_1^I, \pi_1^0)$, is the journal's subjective belief about the increase in weighted action payoffs $\alpha U(a_1, \theta) + (1 - \alpha)U(a_2, \theta)$ if the study is published.¹⁸ A publication rule is *dynamically interim optimal* given default belief π_1^0 if it (almost surely) publishes a study when $\Delta(\pi_1^{(X, S)}, \pi_1^0) > c$ and does not publish when $\Delta(\pi_1^{(X, S)}, \pi_1^0) < c$.

Going forward in this section, we will restrict attention to quadratic loss utility. We explore binary action utility in online Appendix D.3.

B. General Properties of Dynamically Optimal Publication

We start our analysis by observing that the result of Lemma 1, that optimal publication rules are interim optimal, extends immediately to the two-period model.

¹⁸Unlike in the one-period model, the gross interim benefit here depends on the full default belief π_1^0 , not just the first-period default action $a^*(\pi_1^0)$. To expand out the formula for $\Delta(\pi_1^I, \pi_1^0)$, first denote by π_2^{I, X_2} and π_2^{0, X_2} the Bayesian updated beliefs after observing X_2 , starting from respective priors π_1^I and π_1^0 . We then have

$$\begin{aligned} \Delta(\pi_1^I, \pi_1^0) = E_{\theta \sim \pi_1^I, X_2 \sim \mathcal{N}(\theta, s_2^2)} & \left[\alpha \left(U(a^*(\pi_1^I), \theta) - U(a^*(\pi_1^0), \theta) \right) \right. \\ & \left. + (1 - \alpha) \left(U(a^*(\pi_2^{I, X_2}), \theta) - U(a^*(\pi_2^{0, X_2}), \theta) \right) \right]. \end{aligned}$$

LEMMA 1': Consider the two-period model. Under either naive or Bayesian updating, let p be a dynamically optimal publication rule, and let π_1^0 be the induced default belief. Then the publication rule p is dynamically interim optimal given default belief π_1^0 .

In other words, given the appropriate default belief, the journal will publish a submitted paper if the gross interim benefit is above the publication cost c .

Theorem 3, below, derives some key properties of gross interim benefit functions under quadratic loss utility. First, the theorem establishes that there is a positive benefit of publishing any study that changes the public's belief distribution from the default—even null results that don't change the mean belief.¹⁹ So unlike in the one-period model, null results may now sometimes be published. Second, the theorem looks at how the benefit of publishing null results depends on the informativeness of the future study, parameterized by standard error s_2 . It finds that the benefit of publishing null results goes to zero when the future study is either very precise or very imprecise. That is, it is more valuable to publish a null result when the precision of future information is neither too high nor too low.

To guarantee that the benefit of publication goes to zero as future information becomes imprecise, we introduce a mild sufficient condition on the prior distribution π_0 . Say that a belief π is *bounded by Pareto tails with finite variance* if there exist $K > 0$, $C > 0$, and $\gamma > 3$ such that for θ outside of the interval $[-K, K]$, π admits a density, and this density is bounded above by $C|\theta|^{-\gamma}$.²⁰

THEOREM 3: Consider the two-period model with quadratic loss utility. Given some prior π_0 with finite variance, let π_1^0 be the induced default belief either from naive updating or from Bayesian updating under some publication rule and some $q < 1$. Consider $\Delta(\pi_1^l, \pi_1^0)$, the gross interim benefit of publishing a study that induces period 1 interim belief π_1^l .

- (i) If $\pi_1^l \neq \pi_1^0$, then $\Delta(\pi_1^l, \pi_1^0)$ is strictly positive.
- (ii) Suppose further that π_1^0 and π_1^l have the same mean. Then:
 - (a) $\Delta(\pi_1^l, \pi_1^0)$ goes to 0 as s_2 goes to 0.
 - (b) Under the additional assumption that π_0 is bounded by Pareto tails with finite variance, $\Delta(\pi_1^l, \pi_1^0)$ goes to 0 as s_2 goes to infinity.

As mentioned above, the novel takeaway from part (i) is that in the two-period model, there is now a benefit from publishing a null result study, one for which the mean of the induced interim belief π_1^l is equal to the mean of the default belief π_1^0 . To prove this result, it suffices for us to show that whenever $\pi_1^l \neq \pi_1^0$, there exists

¹⁹In the context of the two-period model, we say a study is a “null result” if it leads to the default action in period 1.

²⁰As indicated by the terminology, the Pareto distribution with pdf decaying at a rate of $\theta^{-\gamma}$ has finite variance if and only if $\gamma > 3$ (corresponding to a standard Pareto shape parameter, usually denoted α , strictly greater than 2). Any distribution with compact support, normal tails, or exponentially decaying tails is bounded by Pareto tails with finite variance.

a positive measure of realizations of X_2 for which the posterior means of the second-period beliefs updated from these priors would differ. At these realizations of X_2 , the expected second-period utility $U(a_2, \theta)$ —evaluated by a journal with belief π_1^I —is higher for a public that updated from π_1^I than for a public that updated from π_1^0 .

Restating the above logic, the benefit of publishing a null result study in period 1 is that it helps the public avoid mistakes when taking the period 2 action. Publishing a null result helps prevent the noisy signal X_2 from moving the public's mean belief in period 2 away from the truth. Theorem 3, part (ii)a observes that when the second-period signal is extremely precise ($s_2 \simeq 0$), there is actually no such benefit from publishing a null result. The signal X_2 will reveal the state very precisely, so to the extent that X_2 moves beliefs, it moves beliefs to the truth. Part (ii)b similarly points out that when the second-period signal is extremely imprecise ($s_2 \simeq \infty$), there is also no need to publish a null result: with high probability, observing X_2 will barely move beliefs. The period 2 studies that may cause mistakes by *moving the public's belief away from the truth* are those with an intermediate level of precision.

If the second-period signal were fully informative or completely uninformative, then the game would reduce to a one-period problem (with respective payoff weights α and 1) in which null results had no benefit. So parts (ii)a and (ii)b can essentially be reinterpreted as stating that the value of information is “continuous” as the second-period information approaches these limits via a normal signal.

To see how such continuity might fail in the absence of bounds on the tail behavior of the prior, consider the $s_2 \rightarrow \infty$ limit, and suppose that the prior and default belief are given by the improper uniform prior on the real line. Updating this improper prior with the second-period signal yields the posterior $\mathcal{N}(X_2, s_2^2)$. If a first-period result (X_1, S_1) is published, then the expected second-period action payoff is at worst $-S_1^2$, which is the expected payoff from taking $a_2 = X_1$. If the result is not published, though, then with posterior $\mathcal{N}(X_2, s_2^2)$, the public chooses $a_2 = X_2$ and gets expected payoff $-s_2^2$. So as $s_2 \rightarrow \infty$, the benefit of publishing the first-period result grows without bound. Similar behavior can arise for proper but heavy-tailed priors with well-defined means.

In order to prove continuity in the appropriate limit, we need to show that the second-period posterior mean after publication of a first-period null result converges in mean square to the second-period posterior mean after nonpublication (integrating over the realization of X_2). We are able to show this result for the $s_2 \rightarrow 0$ limit by assuming a finite variance for the default belief, in which case both posterior means converge to X_2 ; see Lemma 6 in Appendix BC. We are able to show this result for the $s_2 \rightarrow \infty$ limit by imposing the stronger assumption that the default belief is bounded by Pareto tails with finite variance, in which case both posterior means converge to the prior mean; see Lemma 7 in Appendix BC.

C. Example of Dynamically Optimal Publication

The following proposition gives an explicit formula for the gross interim benefit function of the two-period model under quadratic loss utility, normal priors, and naive updating. (Naive updating guarantees that when evaluating this interim

benefit, the default belief will be the prior π_0 .) It is dynamically optimal to publish a study if the gross interim benefit is above the publication cost c .

PROPOSITION 4: *In the two-period model with quadratic loss utility, normal priors, and naive updating, the gross interim benefit $\Delta(\pi_1^{(X_1, S_1)}, \pi_0)$ of publishing a study (X_1, S_1) is given by $\beta_0 + \beta_2 \cdot (X_1 - \mu_0)^2$, where*

$$(9) \quad \beta_0 = (1 - \alpha) \frac{\sigma_0^8 s_2^4}{(\sigma_0^2 + S_1^2) (\sigma_0^2 + s_2^2)^2 (\sigma_0^2 S_1^2 + \sigma_0^2 s_2^2 + S_1^2 s_2^2)} > 0,$$

$$(10) \quad \beta_2 = \frac{\sigma_0^4 (s_2^4 + 2\alpha \sigma_0^2 s_2^2 + \alpha \sigma_0^4)}{(\sigma_0^2 + S_1^2)^2 (\sigma_0^2 + s_2^2)^2} > 0.$$

The optimal publication region characterized in Proposition 4 is illustrated in Figure 4; the conclusion of Theorem 1, that the journal publishes point estimates outside of an interval, holds here as well. Proposition 4 finds that the interim benefit of publication can be broken up into two additively separable terms, $\beta_0 + \beta_2 \cdot (X_1 - \mu_0)^2$, where neither β_0 nor β_2 depends on the point estimate X_1 . We can interpret $\beta_2 \cdot (X_1 - \mu_0)^2$ as a benefit of publishing extreme findings, as is familiar from the one-period model. This benefit is larger when X_1 is further from the prior mean μ_0 . Then β_0 represents the new benefit of publishing null result studies with $X_1 = \mu_0$.

One insight from the formula for β_0 in (9) is that the benefit of publishing a null result decreases in its standard error S_1 : there is a bigger benefit of publishing *precise* null results. Indeed, as $S_1 \rightarrow \infty$, the benefit of publishing a null result goes to zero. Proposition 4 in online Appendix D.1 provides some additional comparative statics on the benefit of publishing a null result. It is more valuable to publish a null result when the prior uncertainty σ_0 is larger and the relative payoff weight on the first-period α is smaller. Moreover, in line with Theorem 3, part (ii), the benefit of publishing a null result increases, then decreases in the precision of the second-period signal, going to zero in the fully precise and imprecise limits.

V. Conclusion

Sections I and II of this paper presented and analyzed our benchmark model of publication. A submitted paper is to be published or not, and the social value of publication is derived from its impact on a public policy decision. There is thus an instrumental value in publishing some new result only insofar as it changes public policies. Broadly speaking, we argued for the publication of extreme results over moderate ones. It is more valuable to publish extreme results because they move public beliefs, and therefore public policies, further from the defaults.

As has been noted by many observers, there are reasons outside of this model to be concerned about selectively publishing only extreme results. Section III formalizes some of these concerns. Selective publication invalidates standard statistical inference and causes problems for a public that updates naively in the absence of publication.

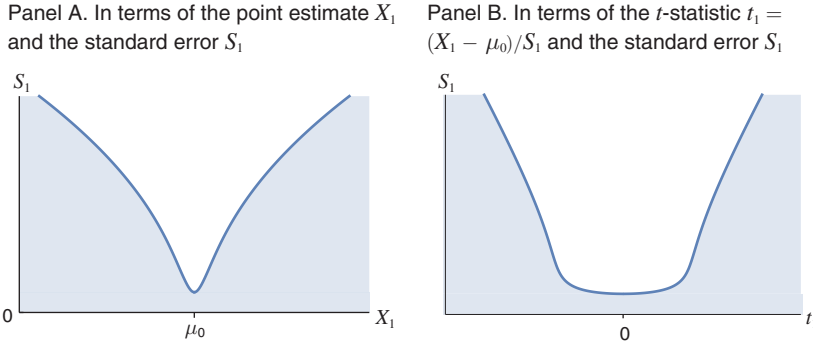


FIGURE 4. DYNAMICALLY OPTIMAL PUBLICATION REGION (SHADED) FOR QUADRATIC LOSS UTILITY, NORMAL PRIOR, NAIVE UPDATING

Notes: Parameter values $\sigma_0 = 1$, $\alpha = 1/2$, $s_2 = 1$, $c = 0.1$. Under different parameters such as a higher c , a study with a null result of $X_1 = \mu_0$ would not be published even for $S_1 \simeq 0$.

Putting these points together, we view the main contribution of this paper as highlighting an important trade-off between the statistical credibility and policy relevance of the publication process—a trade-off that has not been generally appreciated in some current debates focusing on replicability. Moreover, we believe that the simple model of publication we introduced can serve as a basis for further analysis. Section IV explored one such direction, looking at a two-period model. We now conclude the paper by describing a series of additional extensions, some of which are covered in greater detail in the online Appendix. Each of these illustrates how our results might change if we were to bring some additional consideration into the framework of our model.

Alternative Social Objectives: Consider publication rules that maximize social objectives other than policy-based welfare. Online Appendix C.1 presents a *learning* objective. When the social objective is to learn the true state of the world independently of any policy problem, we show that the form of the optimal publication rule may be essentially unchanged from our earlier analysis. The journal continues to publish extreme results. Next, online Appendix C.2 presents an *accuracy* objective. When the social objective is to publish accurate results that are as close as possible to the truth, the publication rule can reverse: the journal now prefers to publish unsurprising results.

Researcher Incentives and Endogenous Study Design: One assumption maintained throughout the paper was that the arrival of studies submitted to journals is exogenous. Online Appendix C.3 considers an extension in which researchers may alter their study designs in response to the publication rule. Specifically, the researcher chooses whether to perform a study on a given question and, if so, at what level of precision. The researcher receives a benefit if the study is published. Her cost of performing the study depends on its precision; e.g., she faces a higher cost to run an experiment with a larger sample size. Taking into account the researcher's incentives, we find that the journal optimally adjusts the publication rule in two ways: the journal

rejects imprecise studies regardless of their findings, and it becomes more willing to publish studies that are sufficiently precise. This modified publication rule induces the researcher to conduct studies that are more precise than she would otherwise choose. Nonetheless, extreme results are still published over moderate ones.

Imperfectly Observed Study Designs: In online Appendix C.4, we discuss the possibility that study designs may not be perfectly observed—a study may be a less reliable signal of the state than is indicated by its reported standard error. If that is the case, we will need to qualify our earlier claim about publishing extreme results: it would still be optimal to publish results that moved beliefs further, but those results might not be the ones with the most extreme point estimates. Extreme point estimates could be considered “implausible,” suggesting problems with the study rather than an extreme state.

Heterogeneous Policymakers: The audience for a study may consist of a number of heterogeneous policymakers, each with different beliefs about the state of the world or different preferences, as in Andrews and Shapiro (2021). Equivalently, there may be a single policymaker whose beliefs or preferences are uncertain to outsiders. In this case, we can get insight into the optimal publication rule by recalling the two-period model of Section IV: different policymaker types are like members of the public who take different actions in period 2 because they have observed different signal realizations X_2 . (The map between these models can be made exact when the heterogeneity is driven by private information.) As we have seen, there can now be a positive benefit of publishing a “null result” that doesn’t change the average belief, since it still moves the actions of some policymakers toward this mean. But there will be a larger benefit of more extreme findings.

APPENDIX A: SOLVING FOR BAYESIAN OPTIMAL PUBLICATION RULES

Under Bayesian updating, the optimal publication rule solves $\max_p EW(p, a^0)$ subject to $a^0 = a^*(\pi_1^{0,p})$; the default action a^0 changes with the rule p . One can simplify the problem by observing that for any fixed p , the induced Bayesian default action $a^0 = a^*(\pi_1^{0,p})$ maximizes expected welfare $EW(p, a^0)$ over choice of a^0 . This is because

$$\begin{aligned}
 (11) \quad \arg \max_{a^0} EW(p, a^0) &= \arg \max_{a^0} E[(1 - qp(X, S))U(a^0, \theta)] \\
 &= \arg \max_{a^0} E[U(a^0, \theta) | D = 0] \\
 &= \arg \max_{a^0} E_{\theta \sim \pi_1^{0,p}} [U(a^0, \theta)].
 \end{aligned}$$

The last equality holds by Bayesian updating, because the conditional distribution of θ given $D = 0$ is equal to $\pi_1^{0,p}$. Therefore, the Bayesian optimal publication rule p equivalently solves $\max_p \max_{a^0} EW(p, a^0)$, where $\max_p \max_{a^0} EW(p, a^0) = \max_{a^0} \max_p EW(p, a^0) = \max_{p, a^0} EW(p, a^0)$. Put differently, in a sequential game of common interest, the value is that of the “planner’s

solution” maximizing the objective over the joint choice of p and a^0 . Lemma 2 formally states this conclusion.

LEMMA 2: *Under Bayesian updating, let p be an optimal publication rule and let $a^0 = a^*(\pi_1^{0,p})$ be the induced default action. Then for any publication rule p' and any action a' , it holds that $EW(p, a^0) \geq EW(p', a')$.*

Proofs of all results stated in the Appendix are found in online Appendix E.

The following lemma provides a recipe for solving for Bayesian optimal publication rules, summarizing some implications of Lemmas 1 and 2. In the statement of this lemma, the optimal publication rule is characterized in terms of an interim optimal publication rule given a particular default action. While all interim optimal publication rules would, in fact, yield the same payoff, for concreteness, let $p^{I(a^0)}$ be the interim optimal publication rule given a^0 that deterministically publishes at indifference: $p^{I(a^0)} = 1$ if $\Delta(\pi_1^{(X,S)}, a^0) \geq c$, and $p^{I(a^0)} = 0$ otherwise.

LEMMA 3: *Suppose that the updating rule is Bayesian, in which case $\pi_1^0 = \pi_1^{0,p}$ and $a^0 = a^*(\pi_1^{0,p})$ under publication rule p .*

(i) *Let $\hat{a} \in \operatorname{argmax}_{a \in \mathcal{A}: a = a^*(\pi_1^{0,p^{I(\hat{a})})} EW(p^{I(a)}, a)$. Then $p^{I(\hat{a})}$ is an optimal publication rule.*

(ii) *Let $\hat{a} \in \operatorname{argmax}_{a \in \mathcal{A}} EW(p^{I(a)}, a)$. Then $p^{I(\hat{a})}$ is an optimal publication rule.*

Lemma 3 provides two alternative maximization programs that can be solved to find Bayesian optimal publication rules.

To understand part (i) of Lemma 3, recall that just as each action induces an interim optimal publication rule, so too does each publication rule induce a Bayesian default action. Lemma 1 establishes that the optimal publication rule is interim optimal with respect to its induced default action. In other words, the default action is a “fixed point” of the mapping from actions to publication rules and back to actions. When searching for an optimal publication rule, it is sufficient to maximize over interim optimal rules that are induced by some fixed point default action.

Part (ii) of Lemma 3 does not require solving for fixed points and instead maximizes over the full action space. Moreover, while the payoff of the publication rule $p^{I(a)}$ induced by action a would generally be given by $EW(p^{I(a)}, a^*(\pi_1^{0,p^{I(a)}}))$ —requiring one to solve for the Bayesian default action induced by $p^{I(a)}$ —we need only evaluate the simpler expression $EW(p^{I(a)}, a)$ since the maximizing action will be a fixed point.²¹

²¹ More precisely, there exists some action a such that $EW(p^{I(a)}, a)$ achieves the welfare of the optimal policy, $\max_{p,a^0} EW(p, a^0)$ (see Lemma 2). And if $EW(p^{I(a)}, a) = \max_{p,a^0} EW(p, a^0) p^{I(a)}$, then $p^{I(a)}$ is an optimal policy that does in fact induce welfare $EW(p^{I(a)}, a)$. In particular, by (11), for any action a , $EW(p^{I(a)}, a) \leq EW(p^{I(a)}, a^*(\pi_1^{0,p^{I(a)}}))$, with the right-hand side of the inequality being the ex ante welfare of publication rule $p^{I(a)}$. So if $EW(p^{I(a)}, a) = \max_{p,a^0} EW(p, a^0)$, then $EW(p^{I(a)}, a^*(\pi_1^{0,p^{I(a)}})) = \max_{p,a^0} EW(p, a^0)$ as well.

We now apply the two parts of Lemma 3 to solve for Bayesian optimal publication rules for quadratic loss and binary action utility under certain distributional conditions. Propositions 1 and 2, assuming normal priors, follow as corollaries.

For quadratic loss utility, we impose the distributional condition that the interim mean is single peaked and symmetric about the prior mean.

PROPOSITION 5: *Suppose that there is quadratic loss utility and that, conditional on a study arriving, the distribution of the interim mean $\mu_1^{(X,S)}$ is single peaked and symmetric about the prior mean μ_0 .²² Then the optimal publication rule under Bayesian updating is the same as under naive updating: publish if and only if $|\mu_1^{(X,S)} - \mu_0| \geq \sqrt{c}$.*

To prove Proposition 5, we show that under single peakedness and symmetry, the prior mean is the *only* fixed point default action under Bayesian updating. So by Lemma 3, part (i), it must be the default action for the optimal policy.

For binary action utility, normalize the prior mean of θ to be less than zero, meaning that the naive default action will be $a^0 = 0$. We then impose the condition that the ex ante distribution of interim expectations on the state is sufficiently “left leaning” relative to $\theta = 0$. (An analogous result would hold for $a^0 = 1$ and a sufficiently “right-leaning” distribution if the prior mean were above zero.)

PROPOSITION 6: *Let $\mu_0 \leq 0$. Suppose that there is binary action utility and that, conditional on a study arriving, the distribution of the interim mean satisfies $\Pr(\mu_1^{(X,S)} \leq -k) \geq \Pr(\mu_1^{(X,S)} \geq k)$ for all $k > 0$. Then the optimal publication rule under Bayesian updating is the same as under naive updating: publish if and only if $\mu_1^{(X,S)} \geq c$.*

To prove Proposition 6, we apply Lemma 3, part (ii). There are two possible default actions, and we confirm that the interim optimal publication rule for default action $a^0 = 0$ gives a higher payoff than that for default action $a^0 = 1$.

Note that the distributional assumption of Proposition 6 is strictly weaker than that of Proposition 5: given a prior mean $\mu_0 \leq 0$, any symmetric distribution of the interim mean $\mu_1^{(X,S)}$ is guaranteed to satisfy the condition of Proposition 6.

²²To be precise, by single peaked and symmetric, we mean that (i) the distribution of the random variable $\mu_1^{(X,S)}$ has a pdf that is symmetric about μ_0 and (ii) for any $\mu' < \mu'' \leq \mu_0$, it holds that if the pdf evaluated at μ' is strictly positive, then the pdf evaluated at μ'' is strictly larger than at μ' . (Symmetry implies the same result for $\mu_0 \leq \mu'' < \mu'$.)

APPENDIX B: PROOFS

A. Proofs for Section II (Optimal Publication Rules)

PROOF OF LEMMA 1:

For naive updating, the result follows from arguments in the text. For Bayesian updating, the result follows from Lemma 2 in Appendix A. In particular, if p is the Bayesian optimal publication rule and a^0 is the induced Bayesian default action, then Lemma 2 establishes that $EW(p, a^0) \geq \max_{p'} EW(p', a^0)$. Hence, $EW(p, a^0) = \max_{p'} EW(p', a^0)$, establishing that p is interim optimal given a^0 . ■

PROOF OF THEOREM 1:

We begin by stating a Lemma. The notation \geq_{FOSD} indicates an ordering of distributions according to first-order stochastic dominance.

LEMMA 4: *Let U be supermodular. Let beliefs π' , π'' , and π''' satisfy $\pi''' \geq_{FOSD} \pi'' \geq_{FOSD} \pi'$. Then for any default action a^0 , it holds that $\Delta(\cdot, a^0)$ is quasiconvex in the sense that $\Delta(\pi'', a^0) \leq \max\{\Delta(\pi', a^0), \Delta(\pi''', a^0)\}$.*

Fix standard error $S = s$, and consider ordered point estimates $x''' > x'' > x'$. To prove the theorem, it suffices to show that if study (x'', s) is published, then at least one of (x''', s) or (x', s) is published as well. By Lemma 1, it is, in turn, sufficient to show that the gross interim benefit of publishing the middle study (x'', s) cannot be strictly higher than that from both the lower study (x', s) and the higher study (x''', s) .

To see why this is the case, recall that at any fixed standard error $S = s$, higher point estimates are more likely than lower point estimates at higher states in the sense of the monotone ratio likelihood property (MLRP).²³ MLRP implies that for any fixed prior, the corresponding posteriors are ranked by first-order stochastic dominance according to their point estimates: $\pi_1^{(x''', s)} \geq_{FOSD} \pi_1^{(x'', s)} \geq_{FOSD} \pi_1^{(x', s)}$. Hence, Lemma 4 implies the result. ■

PROOF OF PROPOSITION 1:

- (i) This prior and signal structure satisfy the hypotheses of Proposition 5 in Appendix A: the distribution of the interim mean is normally distributed and therefore symmetric and single peaked. So the optimal policy, for either updating rule, is to publish if $|\mu_1^{(X, S)} - \mu_0| \geq \sqrt{c}$. By the normal updating formula (4), $|\mu_1^{(X, S)} - \mu_0| = (\sigma_0^2 / (S^2 + \sigma_0^2)) |X - \mu_0| = (1 + S^2 / \sigma_0^2)^{-1} |X - \mu_0|$.
- (ii) All comparative statics are immediate from the formula.
- (iii) The only comparative static that is not immediate is that for the t -statistic cutoff, $(1/S + S/\sigma_0^2) \sqrt{c}$, with respect to S . Taking straightforward limits

²³That is, for point estimates $x'' > x'$, the ratio $f_{X|\theta, S}(x''|\theta, s) / f_{X|\theta, S}(x'|\theta, s)$ is increasing in θ .

confirms that the cutoff goes to infinity as $S \rightarrow 0$ and $S \rightarrow \infty$. The derivative of the cutoff with respect to S is $(-1/S^2 + 1/\sigma_0^2)\sqrt{c}$, and the second derivative is $2\sqrt{c}/S^3$. Since the second derivative is positive, the cutoff is convex over $S \in \mathbb{R}_{++}$ and is minimized at the point where the first derivative is 0, which is $S = \sigma_0^2$. ■

PROOF OF PROPOSITION 2:

This prior and signal structure satisfy the hypotheses of Proposition 6 in Appendix A: the distribution of the interim mean is normally distributed and therefore symmetric. So the optimal policy, for either updating rule, is to publish if $\mu_1^{(X,S)} \geq c$. From the normal updating formula (4), that corresponds to $(\sigma_0^2/(S^2 + \sigma_0^2))X + (S^2/(S^2 + \sigma_0^2))\mu_0 \geq c$, i.e., $X \geq (1 + S^2/\sigma_0^2)c - (S^2/\sigma_0^2)\mu_0$. ■

B. Proofs for Section III (Selective Publication and Statistical Inference)

We begin this section with a lemma establishing that nonselective publication implies the key properties of parts (i)–(iv) of Theorem 2. For parts (i) and (iv), we prove something stronger than what is in Theorem 2. Part (i) establishes the unbiasedness of more general estimators than the estimator X for the state θ . Part (iv) establishes size control for arbitrary confidence sets.

LEMMA 5: *Suppose that the publication rule is nonselective and that $\Pr(D = 1) > 0$. Then $f_{X|\theta,S,D=1}(x|\theta,s) = f_{X|\theta,S}(x|\theta,s)$, and thus, the following properties hold.*

- (i) (Frequentist unbiasedness) *If the estimator $\hat{g}: \mathcal{X} \times \mathcal{S} \rightarrow \mathbb{R}$ for the estimand $g: \Theta \times \mathcal{S} \rightarrow \mathbb{R}$ satisfies $E[\hat{g}(X,S)|\theta,S = s] = g(\theta,s)$ for all θ and s , then $E[\hat{g}(X,S)|\theta,S = s, D = 1] = g(\theta,s)$ for all θ and s .*
- (ii) (Publication probability constant in state) *The publication probability $E[p(X,S)|\theta,S = s]$ is constant in θ for all s .*
- (iii) (Bayesian validity of naive updating) *The Bayesian default belief $\pi_1^{0,p}$ is equal to the naive default belief, i.e., the prior π_0 .*
- (iv) (Frequentist size control) *Fix a level $\alpha \in (0,1)$ and consider a confidence set C mapping from $\mathcal{X} \times \mathcal{S}$ to subsets of Θ . If $\Pr(\theta \in C(X,S)|\theta,S = s) \geq 1 - \alpha$ for all θ and s , then $\Pr(\theta \in C(X,S)|\theta,S = s, D = 1) \geq 1 - \alpha$ for all θ and s .*

PROOF OF THEOREM 2:

Nonselectivity implies the other parts by Lemma 5. Specifically, for part (i) of the Theorem, apply part (i) of Lemma 5 with $\hat{g}(X,S) = X$ and $g(\theta,S) = \theta$, and for part (iv) of the Theorem, apply part (iv) of Lemma 5 with $C = [X - zs, X + zs]$ and $1 - \alpha = \Phi(z) - \Phi(-z)$. We next show the reverse implications.

Part (ii) \Rightarrow Nonselective Publication: Fixing $S = s$, recall that X is a complete statistic for θ in the normal location model when Θ_0 contains an open set in \mathbb{R} ; see, for instance, Theorem 6.22 in Lehmann and Casella (1998). Completeness means that for any measurable function $g: \mathcal{X} \rightarrow \mathbb{R}$, if $E[g(X)|\theta, S = s] = 0$ for all $\theta \in \Theta_0$, then $\Pr(g(X) = 0|\theta, S = s) = 1$ for all $\theta \in \Theta_0$. Apply this definition to $g(x) = p(x, s) - E[p(X, s)|S = s]$. Assuming part (ii), that the publication probability is constant over $\theta \in \Theta_0$, it holds that the expectation of $g(X)$ is 0 for all $\theta \in \Theta_0$ and hence that $p(X, s) = E[p(X, s)|S = s]$ with probability 1 given θ and $S = s$, establishing nonselective publication.

Part (i) \Rightarrow Part (ii): To simplify notation, consider without loss of generality the case $s = 1$. Then the unbiasedness condition $E[X|\theta, S = 1, D = 1] = \theta$ can be written as

$$\frac{\int x\varphi(x - \theta)p(x, 1) dx}{\int \varphi(x - \theta)p(x, 1) dx} = \theta.$$

Equivalently, using the fact that $\varphi'(x) = -x \cdot \varphi(x)$,

$$\begin{aligned} 0 &= \int (x - \theta)\varphi(x - \theta)p(x, 1) dx = -\int \varphi'(x - \theta)p(x, 1) dx \\ &= \partial_\theta \left[\int \varphi(x - \theta)p(x, 1) dx \right] = \partial_\theta E[p(X, S)|\theta, S = 1]. \end{aligned}$$

If the final expression is equal to 0, then $E[p(X, S)|\theta, S = 1]$ is constant over θ in any open set contained in the support. The same argument applies for all other values of S .

Part (iii) \Rightarrow Part (ii): Restating (3), the relative density of the Bayesian default belief to the prior is given by

$$\frac{d\pi_1^{0,p}}{d\pi_0}(\theta) = \frac{1 - q \cdot E[p(X, S)|\theta]}{1 - q \cdot E[p(X, S)]}.$$

The Bayesian default belief is equal to the prior when, under the prior $\theta \sim \pi_0$, this relative density is almost surely constant in θ (in which case the ratio is identically equal to 1). In other words, it holds when $E[p(X, S)|\theta]$ is almost surely constant in θ . Moreover, note that $E[p(X, S)|\theta]$ must be continuous in θ since the signal density function $f_{X|\theta, S}(x|\theta, s)$ is a smooth function of θ for all x, s . Hence, if the Bayesian default belief is equal to the prior, then $E[p(X, S)|\theta]$ must be constant in θ over the support of the prior.

Now, highlighting the dependence of this publication probability on the distribution F_S ,

$$E[p(X, S)|\theta] = \int_{s \in \mathcal{S}} E[p(X, S)|\theta, S = s] dF_S(s).$$

We see that the LHS of this equation is constant over θ in the support of the prior for all distributions F_S if and only if, for all s , $E[p(X, S)|\theta, S = s]$ is constant over θ

in the support. (If there exists s' such that $E[p(X,S)|\theta, S = s']$ varies in θ , then the distribution F_S placing all probability mass on s' will have $E[p(X,S)|\theta]$ vary in θ .) So if the Bayesian default belief is equal to the prior for all F_S , then the publication probability is constant over θ in Θ_0 for all s .

Part (iv) \Rightarrow Nonselective Publication: Without loss of generality, fix $s = 1$. We show that if $I(1) \neq \emptyset$, then there exists θ' for which $\Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, D = 1) < \Phi(z) - \Phi(-z)$.

First, consider the case of a bounded interval $I(1)$. Then there exist θ' (the midpoint of the interval) and $y > 0$ (the radius) such that $I(1) = [\theta' - y, \theta' + y]$. If $y > z$, $\Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, D = 1) = 0$, and the result follows. If $y \leq z$, applying the law of iterated expectations and letting $R = 1$ denote the event of study submission,

$$\begin{aligned} \Phi(z) - \Phi(-z) &= \Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, R = 1) \\ &= \Pr(D = 0 | \theta = \theta', S = 1, R = 1) \\ &\quad \times \Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, R = 1, D = 0) \\ &\quad + \Pr(D = 1 | \theta = \theta', S = 1, R = 1) \\ &\quad \times \Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, R = 1, D = 1). \end{aligned}$$

Conditional on a study submitted but not published, it holds that $X \in [\theta' - y, \theta' + y]$ and, therefore, since $y \leq z$, that $\theta' \in [X - z, X + z]$:

$$\Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, R = 1, D = 0) = 1.$$

Therefore, $\Phi(z) - \Phi(-z)$ is equal to a weighted average of 1 and $\Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, D = 1, R = 1)$ —with positive weights on both—and hence, $\Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, D = 1, R = 1) < \Phi(z) - \Phi(-z)$, yielding the desired result.

Consider, finally, the case of unbounded $I(1)$. If $I(1) = (-\infty, y]$ for some y , then for $\theta' < y - z$, $\Pr(\theta' \in [X - z, X + z] | \theta = \theta', S = 1, D = 1) = 0 < \Phi(z) - \Phi(-z)$. A symmetric argument holds for $I(1) = [y, \infty)$ and $\theta' > y + z$, concluding our proof. ■

PROOF OF PROPOSITION 3:

As stated in the text, the result follows from the fact that the signal $X|S = s$ distributed according to $\mathcal{N}(\theta, s^2)$ is a Blackwell more informative signal of θ when s is smaller. Blackwell more informative signals have higher expected value

to a decision-maker regardless of utility function U or prior π_0 . The updating rule is irrelevant because with nonselective publication, Bayesian updating and naive updating are identical. ■

C. Proofs for Section IV (A Two-Period Model)

PROOF OF LEMMA 1':

The argument closely follows that for Lemma 1, with updated notation. First, given the welfare function (8), define $EW(p, \pi_1^0)$ for the two-period model analogously to the earlier definition of $EW(p, a^0)$ from (5) for the one-period model. For this definition, let π_2^{0, X_2} and $\pi_2^{(X_1, S_1), X_2}$ be the Bayesian updated beliefs following observation of X_2 from the period 2 priors of either π_1^0 (if no period 1 study was published) or $\pi_1^{(X_1, S_1)}$ (if (X_1, S_1) was published at period 1):

$$\begin{aligned} EW(p, \pi_1^0) = E \bigg[& qp(X_1, S_1) \left(\alpha U \left(a^* \left(\pi_1^{(X_1, S_1)} \right), \theta \right) \right. \\ & \left. - c + (1 - \alpha) U \left(a^* \left(\pi_2^{(X_1, S_1), X_2} \right), \theta \right) \right) \\ & + (1 - qp(X_1, S_1)) \left(\alpha U \left(a^* \left(\pi_1^0 \right), \theta \right) \right. \\ & \left. \left. + (1 - \alpha) U \left(a^* \left(\pi_2^{(0, X_2)} \right), \theta \right) \right) \right]. \end{aligned}$$

A publication rule p is dynamically interim optimal given default belief π_1^0 if it solves $\max_{p'} EW(p', \pi_1^0)$.

For the case of naive updating, the publication rule doesn't affect the default belief. So as in the one-period model, it is immediate that a dynamically optimal publication rule is dynamically interim optimal.

For the case of Bayesian updating, we can essentially follow the one-period argument of Appendix A. By the same argument as in that section (see equation (11)), for any fixed publication rule p , the Bayesian default belief $\pi_1^{0, p}$ maximizes $EW(p, \pi_1^0)$ over choice of default beliefs π_1^0 . Hence, the analog of Lemma 2 extends to the two-period model: at the Bayesian dynamically optimal rule p and the corresponding Bayesian default beliefs $\pi_1^{0, p}$, we have $EW(p, \pi_1^{0, p}) = \max_{p'} \max_{\pi_1^0} EW(p', \pi_1^0) = \max_{p', \pi_1^0} EW(p', \pi_1^0)$. In particular, $EW(p, \pi_1^{0, p}) = \max_{p'} EW(p', \pi_1^{0, p})$, and so p is interim optimal. ■

PROOF OF THEOREM 3:

The proof of part (i) holds for any distributions $\pi_1^I \neq \pi_1^0$. For part (ii), the proofs rely on the fact that both distributions arise from the same prior π_0 (implying, for instance, that they share a common support) and that $q < 1$ if updating is Bayesian.

- (i) Write mean beliefs at the first period when a study (X_1, S_1) is published or not by $\mu_1^{(X_1, S_1)}$ and μ_1^0 and in the second period conditional on

X_2 by $\mu_2^{(X_1, S_1), X_2}$ and μ_2^{0, X_2} . The gross interim benefit of publishing a study (X_1, S_1) can be expressed as follows as the first-period action benefit plus the expected second-period action benefit:

$$\alpha \left(\mu_1^{(X_1, S_1)} - \mu_1^0 \right)^2 + (1 - \alpha) E_{\theta \sim \pi_1^{(X_1, S_1)}, X_2 \sim \mathcal{N}(\theta, s_2^2)} \left[\left(\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2} \right)^2 \right].$$

The first term, the first-period action benefit, is nonnegative (and strictly positive when the means $\mu_1^{(X_1, S_1)}$ and μ_1^0 differ). So it suffices to show that when $\pi_1^{(X_1, S_1)} \neq \pi_1^0$, the second term, the expected second-period action benefit, is strictly positive. In turn, it suffices to show that when $\pi_1^{(X_1, S_1)} \neq \pi_1^0$, there exists X_2 for which $\mu_2^{(X_1, S_1), X_2} \neq \mu_2^{0, X_2}$. The second-period action benefit is nonnegative and is continuous in X_2 , and X_2 has full support given any first-period interim belief $\pi_1^{(X_1, S_1)}$. So if the second-period action benefit is strictly positive at some X_2 , then it is strictly positive in expectation.

The claim thus follows if we can show that if $\mu_2^{(X_1, S_1), X_2} = \mu_2^{0, X_2}$ holds for all X_2 , then $\pi_1^{(X_1, S_1)} = \pi_1^0$.

Without loss of generality, normalize $s_2 = 1$ so that $X_2 \sim \mathcal{N}(\theta, 1)$. Define

$$m(x; \pi) = E_{\theta \sim \pi} [\theta | X_2 = x]$$

as the posterior mean of θ under π when $X_2 = x$. We seek to show that if $m(x; \pi) = m(x; \tilde{\pi})$ for almost all $x \in \mathbb{R}$, then $\pi = \tilde{\pi}$.

Taking φ to be the pdf of the standard normal, define $\pi * \varphi$ to be the marginal density of X_2 given $\theta \sim \pi$, which always exists:

$$(\pi * \varphi)(x) = \int_{\mathbb{R}} \varphi(x - \theta) d\pi(\theta).$$

It then holds that

$$\begin{aligned} (12) \quad \frac{\partial \log((\pi * \varphi)(x))}{\partial x} &= \frac{1}{(\pi * \varphi)(x)} \frac{\partial (\pi * \varphi)(x)}{\partial x} = \frac{\int_{\mathbb{R}} \varphi'(x - \theta) d\pi(\theta)}{\int_{\mathbb{R}} \varphi(x - \theta) d\pi(\theta)} \\ &= \frac{\int_{\mathbb{R}} (\theta - x) \varphi(x - \theta) d\pi(\theta)}{\int_{\mathbb{R}} \varphi(x - \theta) d\pi(\theta)} \\ &= E_{\theta \sim \pi} [\theta | X_2 = x] - x = m(x; \pi) - x, \end{aligned}$$

where the equality on the second line follows from the identity $\varphi'(x) = -x\varphi(x)$. (This equation is also known as ‘‘Tweedie’s formula,’’ cf. Efron 2011.) Integrating the left- and right-hand sides yields $(\pi * \varphi)(x) = C \cdot \exp\left(\int_0^x (m(x; \pi) - x) dx\right)$ for a constant of integration C pinned down by the fact that $\pi * \varphi$ integrates to 1. The same

formula holds for $\tilde{\pi} * \varphi$, replacing π by $\tilde{\pi}$ on the right-hand side. We can therefore conclude that if $m(x; \pi) = m(x; \tilde{\pi})$ for all x , then $(\pi * \varphi)(x) = (\tilde{\pi} * \varphi)(x)$ for all x as well.

So suppose that $m(x, \pi) = m(x, \tilde{\pi})$ for all $x \in \mathbb{R}$ and hence that $(\pi * \varphi)(x) = (\tilde{\pi} * \varphi)(x)$. For any distribution π of θ , denote its characteristic function (Fourier transform) by $\psi_\pi(t) = E_{\theta \sim \pi}[\exp(it\theta)]$. The fact that $\pi * \varphi = \tilde{\pi} * \varphi$ implies

$$\psi_\pi(t) \cdot \exp(-t^2/2) = \psi_{\tilde{\pi}}(t) \cdot \exp(-t^2/2)$$

for all t , where $\exp(-t^2/2)$ is the characteristic function of the standard normal distribution. This holds because the Fourier transform maps convolutions of random variables into products of their characteristic functions. It immediately follows that $\psi_\pi(\cdot) = \psi_{\tilde{\pi}}(\cdot)$, since $\exp(-t^2/2)$ is different from 0 for all t , so that the characteristic function of π is equal to the characteristic function of $\tilde{\pi}$. Equality of their characteristic functions implies equality of π and $\tilde{\pi}$, by Lemma 2.15 in Van der Vaart (2000).

- (ii) Let μ_1 denote the shared mean of π_1^0 and of π_1^I . Throughout this proof, it will be convenient to highlight the dependence of the distribution of the signal X_2 on the standard error parameter s_2 , and so we will write the signal as $X_2^{(s_2)}$. In particular, $X_2^{(s_2)} | \theta \sim \mathcal{N}(\theta, s_2^2)$. Furthermore, let

$$m(x; \pi, s_2) = E_{\theta \sim \pi}[\theta | X_2^{(s_2)} = x]$$

be the public's period 2 expectation of θ given period 1 belief π followed by period 2 observation $X_2^{(s_2)} = x$. As a final notational point, in this proof and the proofs of the corresponding lemmas, any integral is to be interpreted as a definite integral over the domain \mathbb{R} unless otherwise specified.

Since the two beliefs π_1^0 and of π_1^I yield the same period 1 action of $a_1 = \mu_1$, the interim gross benefit of publishing the study is the expected benefit in the second period, which can be written as

$$(13) \quad (1 - \alpha) E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^I, s_2) - m(X_2^{(s_2)}; \pi_1^0, s_2) \right)^2 \right].$$

We seek to show that under the appropriate conditions, the expression (13) goes to 0 as $s_2 \rightarrow 0$ (for part (ii)a) and as $s_2 \rightarrow \infty$ (for part (ii)b). We will apply Lemma 6 below to show part (ii)a of the Theorem and Lemma 7 to show part (ii)b.

LEMMA 6: *If distribution π has a finite mean and variance, then*

$$\lim_{s_2 \rightarrow 0} E_{\theta \sim \pi} \left[\left(m(X_2^{(s_2)}; \pi, s_2) - X_2^{(s_2)} \right)^2 \right] = 0.$$

LEMMA 7: *If distribution π has mean μ_1 and is bounded by Pareto tails with finite variance, then*

$$(14) \quad \lim_{s_2 \rightarrow \infty} E_{\theta \sim \pi} \left[\left(m(X_2^{(s_2)}; \pi, s_2) - \mu_1 \right)^2 \right] = 0.$$

The proofs of both parts will also reference the following technical result.

LEMMA 8: *Given any π_1^0 and π_1^I as derived under the hypotheses of Theorem 3, there exists $C' > 0$ such that for all $s_2 > 0$ and all functions $y : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, it holds that $E_{\theta \sim \pi_1^I} [y(X_2^{(s_2)})] \leq C' E_{\theta \sim \pi_1^0} [y(X_2^{(s_2)})]$.*

We now complete the proofs of Theorem 3 parts (ii)a and (ii)b.

Part (ii)a: First, observe that the distributions π_1^0 and π_1^I both have a finite variance. To see that this holds for π_1^I , recall that $\pi_1^I = \pi_1^{(x_1, s_1)}$ is a posterior distribution updated after observing a normal signal $(X_1, S_1) = (x_1, s_1)$. The posterior distribution (from any prior) after observing a normal signal has a finite variance. To see that this holds for π_1^0 , recall that π_1^0 arises as a default belief from the prior π_0 with a finite variance. In the case of naive updating, $\pi_1^0 = \pi_0$, so the result is immediate. In the case of Bayesian updating, observe from (3) that $(d\pi_1^0/d\pi_0)(\theta) \leq 1/(1 - q)$ for all θ and therefore $\pi_0 \geq (1 - q)\pi_1^0$, so if π_1^0 had an infinite variance, then so too would π_0 .

Plugging $\pi = \pi_1^I$ into Lemma 6, we have that

$$\lim_{s_2 \rightarrow 0} E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^I, s_2) - X_2^{(s_2)} \right)^2 \right] = 0.$$

Applying Lemma 8, we also have that there exists a constant $C' > 0$ such that

$$\begin{aligned} 0 &\leq \lim_{s_2 \rightarrow 0} E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^0, s_2) - X_2^{(s_2)} \right)^2 \right] \\ &\leq \lim_{s_2 \rightarrow 0} C' E_{\theta \sim \pi_1^0} \left[\left(m(X_2^{(s_2)}; \pi_1^0, s_2) - X_2^{(s_2)} \right)^2 \right]. \end{aligned}$$

Plugging $\pi = \pi_1^0$ into Lemma 6, we have that the right-hand expression is equal to 0. Hence,

$$\lim_{s_2 \rightarrow 0} E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^0, s_2) - X_2^{(s_2)} \right)^2 \right] = 0.$$

In other words, both $m(X_2^{(s_2)}; \pi_1^I, s_2)$ and $m(X_2^{(s_2)}; \pi_1^0, s_2)$ converge to $X_2^{(s_2)}$ in mean square as $s_2 \rightarrow 0$ under $\theta \sim \pi_1^I$. Therefore, they converge to each other in mean square, establishing the desired conclusion that the expression (13) goes to 0 as $s_2 \rightarrow 0$ as long as the three variables $m(X_2^{(s_2)}; \pi_1^I, s_2)$, $m(X_2^{(s_2)}; \pi_1^0, s_2)$, and $X_2^{(s_2)}$ are all square integrable under $\theta \sim \pi_1^I$.

The three variables are indeed square integrable, as they each have finite means and variance. To see that, observe that the posterior mean $m(X_2^{(s_2)}; \pi_1^I, s_2)$ has mean equal to μ_1 and, by the Law of Total Variance, variance less than $\text{var}_{\theta \sim \pi_1^I}$: the variance of the posterior mean given some signal is bounded above by the variance of the prior. The other posterior mean variable $m(X_2^{(s_2)}; \pi_1^0, s_2)$ has a finite mean and variance under the distribution $\theta \sim \pi_1^0$ by the same arguments and, therefore, finite mean and variance under the distribution $\theta \sim \pi_1^I$ by Lemma 8.²⁴ Finally, the mean of $X_2^{(s_2)}$ is μ_1 , and the variance is $\text{var}_{\theta \sim \pi_1^I}(\theta) + s_2^2$.

Part (ii)b: First, observe that the distributions π_1^0 and π_1^I are both bounded by Pareto tails with finite variance since they arise from the prior π_0 that is bounded by Pareto tails with finite variance. To see that this holds for π_1^I , recall that $\pi_1^I = \pi_1^{(x_1, s_1)}$ is a posterior distribution updated after observing a normal signal $(X_1, S_1) = (x_1, s_1)$. It holds that $d\pi_1^I(\theta)/d\pi_0(\theta)$ is equal to a constant times $\varphi((x_1 - \theta)/s_1)$ and hence the tails of π_1^I decay at a rate at least as fast as those of π_0 . To see that this holds for π_1^0 in the case of naive updating, $\pi_1^0 = \pi_0$, and so the result is immediate. To see that this holds for π_1^0 in the case of Bayesian updating, observe from (3) that $(d\pi_1^0/d\pi_0)(\theta) \leq 1/(1 - q)$ for all θ , and therefore $\pi_0 \geq (1 - q)\pi_1^0$, so if π_1^0 were not bounded by Pareto tails with finite variance, then neither would π_0 .

Plugging $\pi = \pi_1^I$ into Lemma 7, we have that

$$\lim_{s_2 \rightarrow \infty} E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^I, s_2) - \mu_1 \right)^2 \right] = 0.$$

Applying Lemma 8, we also have that there exists a constant $C' > 0$ such that

$$\begin{aligned} 0 &\leq \lim_{s_2 \rightarrow \infty} E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^0, s_2) - \mu_1 \right)^2 \right] \\ &\leq \lim_{s_2 \rightarrow \infty} C' E_{\theta \sim \pi_1^0} \left[\left(m(X_2^{(s_2)}; \pi_1^0, s_2) - \mu_1 \right)^2 \right]. \end{aligned}$$

Plugging $\pi = \pi_1^0$ into Lemma 7, we have that the right-hand expression is equal to 0. Hence,

$$\lim_{s_2 \rightarrow \infty} E_{\theta \sim \pi_1^I} \left[\left(m(X_2^{(s_2)}; \pi_1^0, s_2) - \mu_1 \right)^2 \right] = 0.$$

²⁴To see that $m(X_2^{(s_2)}; \pi_1^0, s_2)$ has a finite mean under $\theta \sim \pi_1^I$, recall $E_{\theta \sim \pi_1^0} [m(X_2^{(s_2)}; \pi_1^0, s_2)]$ is finite if and only if $E_{\theta \sim \pi_1^I} [m(X_2^{(s_2)}; \pi_1^0, s_2)]$ is finite, and the latter being finite implies by Lemma 8 that $E_{\theta \sim \pi_1^I} [m(X_2^{(s_2)}; \pi_1^0, s_2)]$ and hence $E_{\theta \sim \pi_1^I} [m(X_2^{(s_2)}; \pi_1^0, s_2)]$ are finite. Call $\tilde{\mu}$ the mean of $m(X_2^{(s_2)}; \pi_1^0, s_2)$ under $\theta \sim \pi_1^I$; the fact that $m(X_2^{(s_2)}; \pi_1^0, s_2)$ has a finite variance under $\tilde{\theta} \sim \pi_1^0$ means that $E_{\theta \sim \pi_1^0} [m(X_2^{(s_2)}; \pi_1^0, s_2) - \tilde{\mu}]^2$ is finite, and thus by Lemma 8, $E_{\theta \sim \pi_1^I} [m(X_2^{(s_2)}; \pi_1^0, s_2) - \tilde{\mu}]^2$ is also finite.

In other words, both $m(X_2^{(s_2)}; \pi_1^I, s_2)$ and $m(X_2^{(s_2)}; \pi_1^0, s_2)$ converge to μ_1 in mean square as $s_2 \rightarrow 0$ under $\theta \sim \pi_1^I$. Therefore, they converge to each other in mean square, establishing the desired conclusion that the expression (13) goes to 0 as $s_2 \rightarrow 0$ as long as they are both square integrable under $\theta \sim \pi_1^I$; that was established in the proof of the previous part. ■

PROOF OF PROPOSITION 4:

Suppose a study (X_1, S_1) arrives at period 1. Let μ_1^0 indicate the posterior mean at period 1 in the absence of publication and $\mu_1^{(X_1, S_1)}$ the posterior mean at period 1 if the study is published. Let μ_2^{0, X_2} indicate the posterior mean at period 2 if the study had not been published, and then the second-period signal is observed to be X_2 , and $\mu_2^{(X_1, S_1), X_2}$ the posterior mean at period 2 if the study had been published. We can calculate these posterior means as follows:

$$\begin{aligned}\mu_1^0 &= \mu_0, \\ \mu_1^{(X_1, S_1)} &= \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{S_1^2}} \left(\frac{\mu_0}{\sigma_0^2} + \frac{X_1}{S_1^2} \right), \\ \mu_2^{0, X_2} &= \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{s_2^2}} \left(\frac{\mu_0}{\sigma_0^2} + \frac{X_2}{s_2^2} \right), \\ \mu_2^{(X_1, S_1), X_2} &= \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{S_1^2} + \frac{1}{s_2^2}} \left(\frac{\mu_0}{\sigma_0^2} + \frac{X_1}{S_1^2} + \frac{X_2}{s_2^2} \right).\end{aligned}$$

Consider the interim stage, at which (X_1, S_1) has been observed by the journal and not yet published and hence at which the journal has interim belief $\pi_1^1(X_1, S_1)$. From this interim perspective, publication has a cost of c . It then delivers a benefit toward the first-period action payoff and a benefit toward the second-period action payoff.

The benefit of publication toward the first-period payoff is $\alpha(\mu_1^{(X_1, S_1)} - \mu_1^0)^2$, which simplifies to

$$(15) \quad \alpha(\mu_1^{(X_1, S_1)} - \mu_1^0)^2 = \alpha \frac{\sigma_0^4}{(\sigma_0^2 + S_1^2)^2} (X_1 - \mu_0)^2.$$

The period 2 action is $\mu_2^{(X_1, S_1), X_2}$ if the study is published and is μ_2^{0, X_2} otherwise. Hence, conditional on X_2 , the benefit of first-period publication toward the second-period payoff is $(1 - \alpha)(\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2})^2$. At the interim stage, then, the expected second-period payoff is the expectation of that value over the random

variable X_2 given beliefs $\theta \sim \pi_1^{(X_1, S_1)}$ and $X_2 \sim \mathcal{N}(\theta, s_2^2)$. Writing out this expectation and simplifying,

$$(16) \quad E \left[(1 - \alpha) \left(\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2} \right)^2 \middle| X_1, S_1 \right] \\ = (1 - \alpha) \left(E \left[\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2} \middle| X_1, S_1 \right]^2 + \text{var} \left[\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2} \middle| X_1, S_1 \right] \right).$$

Next, observe that given X_1 and S_1 , the conditional distribution of X_2 is

$$X_2 \middle| (X_1, S_1) \sim \mathcal{N} \left(\mu_1^{(X_1, S_1)}, \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{S_1^2}} + s_2^2 \right).$$

Plugging this conditional distribution into the various terms of (16) and simplifying,

$$(17) \quad (1 - \alpha) E \left[\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2} \middle| X_1, S_1 \right]^2 \\ = (1 - \alpha) \left(\frac{\sigma_0^2 s_2^2}{(\sigma_0^2 + S_1^2)(\sigma_0^2 + s_2^2)} \right)^2 (X_1 - \mu_0)^2,$$

$$(18) \quad (1 - \alpha) \text{var} \left[\mu_2^{(X_1, S_1), X_2} - \mu_2^{0, X_2} \middle| X_1, S_1 \right] \\ = (1 - \alpha) \frac{\sigma_0^8 s_2^4}{(\sigma_0^2 + S_1^2)(\sigma_0^2 + s_2^2)^2 (\sigma_0^2 S_1^2 + \sigma_0^2 s_2^2 + S_1^2 s_2^2)}.$$

The gross interim payoff of publication is the sum of the right-hand sides of (15), (17), and (18). To get the form stated in the proposition, we add up the coefficients on $(X_1 - \mu_0)^2$ in (15) and (17):

$$\alpha \frac{\sigma_0^4}{(\sigma_0^2 + s_1^2)^2} + (1 - \alpha) \left(\frac{s_2^2 \sigma_0^2}{(\sigma_0^2 + s_1^2)(\sigma_0^2 + s_2^2)} \right)^2 = \frac{\sigma_0^4 (s_2^4 + 2\alpha \sigma_0^2 s_2^2 + \alpha \sigma_0^4)}{(\sigma_0^2 + S_1^2)^2 (\sigma_0^2 + s_2^2)^2}. \blacksquare$$

REFERENCES

- Abadie, Alberto.** 2020. "Statistical Nonsignificance in Empirical Economics." *American Economic Review: Insights* 2 (2) 193–208.
- Akerlof, George, and Pascal Michailat.** 2017. "Beetles: Biased Promotions and Persistence of False Belief." NBER Working Paper 23523.
- Andrews, Isaiah, and Maximilian Kasy.** 2019. "Identification of and Correction for Publication Bias." *American Economic Review* 109 (8): 2766–94.
- Andrews, Isaiah, and Jesse M. Shapiro.** 2021. "A model of scientific communication." *Econometrica* 89 (5) 2117–42.
- Blackwell, David.** 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24 (2): 265–72.
- Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg.** 2016. "Star Wars: The Empirics Strike Back." *American Economic Journal: Applied Economics* 8 (1): 1–32.

- Camerer, Colin F., Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, et al. 2016. "Evaluating Replicability of Laboratory Experiments in Economics." *Science* 351 (6280): 1433–36.
- Card, David, and Stefano Della Vigna. 2013. "Nine Facts about Top Journals in Economics." *Journal of Economic Literature* 51 (1): 144–61.
- Chambers, Christopher, Eva Ferredoes, Suresh Muthukumaraswamy, and Peter Etchells. 2014. "Instead of 'Playing the Game' It Is Time to Change the Rules: Registered Reports at AIMS Neuroscience and Beyond." *AIMS Neuroscience* 1 (1): 4–17.
- Christensen, Garret S., and Edward Miguel. 2018. "Transparency, Reproducibility, and the Credibility of Economics Research." *Journal of Economic Literature* 56 (3) 920–80.
- Efron, Bradley. 2011. "Tweedie's Formula and Selection Bias." *Journal of the American Statistical Association* 106 (496): 1602–14.
- Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. "Publication Bias in the Social Sciences: Unlocking the File Drawer." *Science* 345 (6203): 1502–05.
- Frankel, Alexander, and Emir Kamenica. 2019. "Quantifying Information and Uncertainty." *American Economic Review* 109 (10): 3650–80.
- Furukawa, Chishio. 2018. "Unbiased Publication Bias: Theory and Evidence." Unpublished.
- Gelman, Andrew, and Eric Loken. 2014. "The Statistical Crisis in Science." *American Scientist* 102 (6): 460.
- Glaeser, Edward L. 2006. "Researcher Incentives and Empirical Methods." Harvard Institute of Economic Research (HIER) Discussion Paper 2122.
- Ioannidis, John P.A. 2005. "Why Most Published Research Findings Are False." *PLoS Medicine* 2 (8): e124.
- Ioannidis, John P.A. 2008. "Why Most Discovered True Associations Are Inflated." *Epidemiology* 19 (5): 640–48.
- Jovanovic, Boyan. 1982. "Truthful Disclosure of Information." *Bell Journal of Economics* 13 (1): 36–44.
- Lehmann, E.L. 1988. "Comparing Location Experiments." *Annals of Statistics* 16 (2): 521–33.
- Lehmann, E.L., and George Casella. 1998. *Theory of Point Estimation*. Vol. 31, *Springer Texts in Statistics*. New York: Springer.
- McElreath, Richard, and Paul E. Smaldino. 2015. "Replication, Communication, and the Population Dynamics of Scientific Discovery." *PLoS ONE* 10 (8): e0136088.
- Müller-Iten, Michèle. 2016. "Gatekeeping under Asymmetric Information." https://economics.nd.edu/assets/186597/paper_gatekeepers_main.pdf.
- Nissen, Silas Boye, Tali Magidson, Kevin Gross, and Carl T. Bergstrom. 2016. "Publication Bias and the Canonization of False Facts." *eLife* 5: e21451.
- Nosek, Brian A., and Daniël Lakens. 2014. "Registered Reports: A Method to Increase the Credibility of Published Results." *Social Psychology* 45 (3): 137–41.
- Open Science Collaboration. 2015. "Estimating the Reproducibility of Psychological Science." *Science* 349 (6251): aac4716.
- Persico, Nicola. 2000. "Information Acquisition in Auctions." *Econometrica* 68 (1): 135–48.
- Rosenthal, Robert. 1979. "The File Drawer Problem and Tolerance for Null Results." *Psychological Bulletin* 86 (3): 638–41.
- Simmons, Joseph P., Leif D. Nelson, and Uri Simonsohn. 2011. "False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant." *Psychological Science* 22 (11): 1359–66.
- Sterling, Theodore D. 1959. "Publication Decisions and Their Possible Effects on Inferences Drawn from Tests of Significance—or Vice Versa." *Journal of the American Statistical Association* 54: 30–34.
- Van der Vaart, Aad W. 2000. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Verrecchia, Robert E. 1983. "Discretionary Disclosure." *Journal of Accounting and Economics* 5: 179–94.