Optimal Publication Rules for Evidence-Based Policy

Toru Kitagawa† Patrick Vu‡

October 23, 2023

Abstract

Empirical research can inform evidence-based policy choice but may be censored due to publication bias. How does this impact the decisions of policymakers who do not have, or are unwilling to use, prior beliefs about a policy's impact? For minimax regret policymakers, we characterize the optimal treatment rule with selective publication against statistically insignificant results. We then show that the optimal publication rule which minimizes maximum regret is non-selective. This contrasts with the optimal publication rule for Bayesian policymakers studied in the literature, where only 'extreme' results that sufficiently move the prior are published. Thus, in the minimax regret framework, the optimal publication regime for policy choice is consistent with valid statistical inference in scientific research.

Keywords: Statistical decision theory, minimax regret, meta-science.

†Department of Economics, Brown University. Email: toru_kitagawa@brown.edu
‡Department of Economics, Brown University. Email: patrick_vu@brown.edu

I. Introduction

Publication bias has been widely-documented across various fields and led to debates in the scientific community about reforming publication norms (Ioannidis, 2005; Franco et al., 2014; Nosek et al., 2015; Miguel and Christensen, 2018; Nosek et al., 2018; Andrews and Kasy, 2019). Proposals to combat publication bias are often aimed at mitigating selective publication of statistically significant findings. For example, launching journals dedicated to publishing null results (e.g. *PLOS One*); promoting preregistered analysis plans which are reviewed and published prior to data collection (e.g. the *Journal of Development Economics*); banning the use of stars to denote significance when presenting estimation results (e.g. the *American Economic Review*); and even abandoning statistical significance altogether (McShane et al., 2019).

However, non-selective publication may not necessarily be optimal from the perspective of a decision-maker who uses evidence from published studies to inform a policy decision. Frankel and Kasy (2022) develop a model of a Bayesian decision-maker who has a prior distribution over possible treatment outcomes and updates their beliefs using evidence from published studies before making a policy decision. When publication entails a cost (e.g. the opportunity cost of drawing attention away from other studies), the optimal rule is to publish only 'extreme' results that sufficiently move prior beliefs. This gives rise to a striking trade-off: selective publication enhances policy relevance while at the same time deteriorating statistical credibility.

While selective publication may be optimal for a Bayesian decision-maker, in many situations, policymakers may be unable or unwilling to base decisions on prior beliefs about treatment outcomes. For example, they may have insufficient information to form a reasonable prior, or if when decisions are made by a group, prior beliefs of different group members may conflict with one another. A common alternative to relying on prior beliefs is to introduce ambiguity on the treatment outcomes and pursue robust decisions.

In this paper, we consider a policymaker that aims to minimize maximum regret (Savage, 1951; Manski, 2004), where regret equals the difference between the highest possible expected welfare outcome given full knowledge of the true impact of all treatments and the expected welfare attained by the statistical decision rule. We first characterize the minimax regret decision rule of the policymaker in the presence of publication bias, and then derive the optimal publication rule that minimizes the value of minimax regret. In contrast to the Bayesian framework, we show that the optimal publication rule for minimax regret decision-makers is completely non-selective i.e. publication decisions do not depend on statistical significance. Importantly, non-selective publication implies valid statistical inference. Thus, in the minimax regret framework, there exists no trade-off between policy relevance and statistical credibility.

Following Manski (2004), Stoye (2009), and Tetenov (2012), our model considers a policymaker whose problem is to assign members of a population one of two treatments: a status quo treatment and an innovative treatment. A study about the relative effectiveness of the treatments is conducted. However, the study is only observed by the policymaker if it is published, which may depend on its statistical significance. We consider the case where t-ratios in a symmetric interval around zero are censored with probability $\beta_p \in [0, 1]$ e.g. statistical significance at the 5% level implies that t-ratios between -1.96 and 1.96 will be published with probability β_p. Additionally, publication may also entail a cost $c \geqslant 0$. We consider a policymaker who correctly accounts for publication bias when choosing their statistical treatment rule (and later consider a naive policymaker who does not account for it). If a study is published, the policymaker observes it and implements the innovative policy if its relative effect size is greater than a chosen threshold value T. Alternatively, if a study is not published, then the policymaker must act without evidence and implements the innovative treatment with probability δ_0. The policymaker chooses a statistical treatment rule – consisting of the threshold rule T and the default action δ_0 – that minimizes their maximum regret, that is, the expected welfare loss relative to optimal welfare attained with knowledge of the true treatment effect.

We show that the minimax regret decision rule implements the innovative treatment if and only if the published estimate of the relative efficacy of the innovative treatment is positive, and randomizes between treatments with equal probability when no study is published i.e. $(T^*, \delta_0^*) = (0, \frac{1}{2})$. The intuition for this result follows from the two key factors. First, the decision-maker's welfare equally weighs Type I errors (from mistakenly implementing an inferior treatment) and Type II errors (failing to implement the superior treatment). Second, the class of publication rules we consider censors insignificant empirical results symmetrically around zero. The first symmetry implies that the decision-maker will implement the innovative treatment when the published evidence supports the innovative treatment having a positive effect, and remain with the status quo treatment otherwise. Combined with the second symmetry, we can conclude that when the study is published, the sign of the estimate (i.e., $T^* = 0$) is sufficient for the decision-maker to infer the sign of the effect, and when no study is published, the decision-maker has no evidence regarding the sign of the relative treatment effect and will therefore randomize between treatments.

Given the minimax regret rule of the decision-maker, we optimize the value of minimax regret with respect to the publication rule. As the main result, we show that the resulting optimal publication rule is to *publish all results*. This accords well with common intuition: receiving evidence from a published study about the relative effectiveness of treatments allows the policymaker to do better than in the case where no study is published and they must randomize between treatments.

It is notable, however, that the opposite conclusion is reached when considering a Bayesian decision-maker, for whom the optimal publication rule censors relatively uninformative studies that do little to move prior beliefs of the decision-maker (Frankel and Kasy, 2022). Two differences account for this. First, publication costs enter the expected welfare function in the Bayesian framework of Frankel and Kasy (2022), while they do not appear in the expression for regret in our framework. This is because regret equals the probability of making an inferior treatment choice multiplied by the magnitude of the loss from doing so. Neither quantity is affected by publication costs. Put differently, publication costs are constant with respect to the decision rule and therefore have no impact on regret. The second difference is that Frankel and Kasy (2022) define null results in the Bayesian framework as those which do not move prior beliefs. By contrast, there is no notion of prior beliefs in the minimax regret framework. We instead use the common definition of null results as those which are statistically indistinguishable from zero. Accordingly, we consider a class of symmetric publication rules that yield no information about the sign of the true effect when studies are not published. Since published studies will always provide some evidence on the sign of the true effect, the optimal publication regime in terms of the regret criterion is to publish all the results.

Our results highlight that the optimal publication regime can change drastically depending on what optimality criterion the policymaker pursues for policy choice. Which optimality criterion is relevant in practice may depend on the factors such as behavioural axioms of the decision-makers, availability of the prior belief of the policy effect, and/or the form of publication bias relevant to the empirical literature of interest.

We consider three main extensions to the baseline model. Following Tetenov (2012), we first extend the model to incorporate decision criteria that asymmetrically weigh Type I error (from mistakenly implementing the inferior treatment) and Type II error (from mistaking rejecting a superior treatment). We provide numerical evidence consistent with the conjecture the optimal publication rule for minimax regret decision-makers with asymmetric regret criteria is also non-selective.

Second, we consider a naive policymaker who, unlike the sophisticated policymaker in the main analysis, does not account for publication bias when choosing their decision rule. Naive policymakers could in some cases be more realistic than sophisticated policymakers, because sophistication demands both knowledge of the publication rule and the ability to correctly adjust for it. In this model, the naive policymaker's expected welfare (and regret) is misspecified because they believe, erroneously, that there is no publication bias.[1] We evaluate their subsequent decision rule based on the worse case scenario under *correctly* specified regret. We show that minimax regret for the naive policymaker is weakly higher than for the sophisticated poli-

[1] This affects: (i) their beliefs about the distribution of the published estimates; and (ii) implies that they make no inferences about the size of the treatment effect when no study is published.

cymaker.[2] Thus, in general, the naive policymaker chooses a non-optimal decision rule because they fail account for publication bias.

The optimal publication rule in the main analysis assumes that the policymaker and the journal have the same preferences, namely, to minimize maximum regret. In the third extension, we consider the optimal publication rule under misaligned preferences. In particular, we consider the case where the policymaker chooses their decision rule to minimize maximum regret, but where the journal chooses the publication rule to maximize welfare (under some prior distribution for the policy's effect). The main result shows that the journal's optimal action takes the form a simple threshold rule: publish all results if the cost c is sufficiently low; otherwise, censor all null results. Thus, in the case where publication costs are low, the optimal publication rule under misaligned preferences is the same as with aligned preferences. However, when publication costs are high, it is possible that censoring all null results is optimal.

Related Literature. This article contributes to the literature on statistical decision theory (Manski, 2004; Stoye, 2009; Tetenov, 2012). It generalizes the canonical model in the minimax regret framework to incorporate publication bias against null results. We characterize the optimal decision rule that minimizes maximum regret and extend results to the case where Type I and Type II error are weighted asymmetrically. Our model coincides with the canonical model in the special case where there is no publication bias.

This article also contributes to the meta-science literature on publication bias and optimal publication rules (Ioannidis, 2005; Andrews and Kasy, 2019; Miguel and Christensen, 2018; Frankel and Kasy, 2022). It is most closely related to Frankel and Kasy (2022), who examine a similar problem in a Bayesian framework. In contrast to a Bayesian framework, where the optimal publication rule selects only 'extreme' results for publication, we show in a minimax regret framework that the optimal publication rule is completely non-selective.

II. Model

A. Setup

The policymaker's problem is to assign two treatments to a population with observationally identical members: the status quo treatment ($t = 0$) and the innovative treatment ($t = 1$). Following Manski (2004), suppose that each member j in population J has a treatment response function $y_j(\cdot) : \{0, 1\} \to Y$ mapping treatments into outcomes. The population is a probability space (J, Ω, P). The probability distribution $P[y(\cdot)]$ of the random vector $y(\cdot)$ describes

[2] Minimax regret for the naive policymaker is strictly higher when the Type I and Type II error are unequally weighted.

treatment response across the population. The population is "large" in the sense that J is uncountable and $P(j) = 0$. Next, let $\mathbb{E}[y(1)] - \mathbb{E}[y(0)] \equiv \theta \in \Theta \subseteq \mathbb{R}$ be the unknown average treatment effect of the innovative treatment relative to the status quo treatment, with the status quo treatment normalized to zero. When $\theta > 0$, the innovative treatment is preferred; otherwise, the status quo treatment is preferred.

Evidence about θ may be observed by the policymaker in the form of a published study. However, not all studies are necessarily published. Consider first a *latent study* (published or unpublished), which is characterized by (X, σ), where X is the *estimated treatment effect* and σ is the known *standard error*. We assume X is normally distributed on $\mathcal{X} = \mathbb{R}$ and normalize $\sigma = 1$, so that $X|\theta \sim N(\theta, 1)$. This assumption is motivated by the fact that study estimates are widely assumed to be approximately normal in practice. The normalization is for notational convenience, since σ is known and fixed. The journal observes the latent study $(X, 1)$ and decides the probability of publication according to their publication rule, $p : \mathcal{X} \to [0, 1]$. Let $D = 1$ denote the event when a study is published and $D = 0$ the event when it is not. We consider the class of publication rules where absolute t-ratios below a critical threshold t_α may be published with a lower probability than those above that threshold:

Assumption 1 (Publication Selection Function)**.** *Let* $p(X) = 1 - (1 - \beta_p) \cdot \mathbb{1}[|X| < t_\alpha]$ *with* $\beta_p \in [0, 1]$.

The form of publication bias in Assumption 1 implies that published estimated treatment effects follow a mixture of truncated normal densities, where the region below the critical threshold of the density is down-weighted and the region above it is up-weighted. Denote the cdf as

$$F(x|\theta, D = 1) \equiv \frac{\int_{-\infty}^{x} p(y)\phi(y - \theta)dy}{\int p(y)\phi(y - \theta)dy}, \tag{1}$$

where $\phi(x) = (2\pi)^{-1/2} \exp(2^{-1}x^2)$ is the probability density function of the standard normal distribution.

The policymaker's decision rule must cover two possible realizations of the publication process: the event when the study is published $(D = 1)$ and event when it is not $(D = 0)$. Let $Z = X \cdot D + \{\text{missing}\}(1 - D)$ and the *statistical treatment rule* be $\delta : Z \to [0, 1]$, with

$$\delta(X, D) = \begin{cases} \delta_1(X) & \text{if } D = 1 \\ \delta_0 & \text{if } D = 0 \end{cases} \tag{2}$$

That is, $\delta(X, D)$ maps study outcomes to treatment assignment proportions when the study is published, and assigns a default action $\delta_0 \in [0, 1]$ when it is not.

We first consider a sophisticated policymaker who knows the exact form of publication bias and correctly accounts for it when choosing their optimal decision rule. For example, a sophisticated policymaker could estimate $p(\cdot)$ from a sample of studies in the published literature (e.g. by using the Andrews and Kasy (2019) model). Their utility from treatment rule $\delta(X, D)$ with treatment effect θ and observed data X is given by

$$U\big(\delta, \theta\big) = \theta D \delta_1(X) + \theta(1 - D)\delta_0 - Dc \tag{3}$$

where $c \geqslant 0$ represents the cost of publishing an article. Following Frankel and Kasy (2022), we interpret this cost as the opportunity cost of directing the public's limited attention away from other studies. Welfare for a statistical decision rule $\delta(X, D)$ corresponds to a shared objective by the policymaker and the journal. Expected welfare is obtained by integrating over possible study outcomes:

$$\begin{aligned} W\big(\delta, \theta\big) = &\int U\big(\delta, \theta\big) f(x'|\theta) dx' \\ =&\theta \cdot \mathbb{P}[D = 1|\theta] \cdot \mathbb{E}[\delta_1(X)|\theta, D = 1] + \theta \cdot \big(1 - \mathbb{P}[D = 1|\theta]\big)\delta_0 - \mathbb{P}[D = 1|\theta] \cdot c \\ =&W_1(\delta_1, \theta) + W_0(\delta_0, \theta) - \mathbb{P}[D = 1|\theta] \cdot c \end{aligned} \tag{4}$$

where $W_1(\delta_1, \theta) = \theta \cdot \mathbb{P}[D = 1|\theta] \cdot \mathbb{E}[\delta_1(X)|\theta, D = 1]$ is the welfare for the case that the study is published, and $W_0(\delta_0, \theta) = \theta \cdot \big(1 - \mathbb{P}[D = 1|\theta]\big)\delta_0$ is the welfare for the case that the study is not published.

Finally, regret is given by the difference between the highest possible expected welfare conditional on θ and the expected welfare under the treatment rule. Let $W^*(\theta)$ be the welfare attained by the oracle rule $\delta_1 = \delta_0 = \mathbb{1}(\theta > 0)$. Then regret is given by

$$\begin{aligned} R\big(\delta, \theta\big) &= W^*(\theta) - W\big(\delta, \theta\big) \\ &= \begin{cases} -\theta\Big(\mathbb{P}[D = 1|\theta] \cdot \mathbb{E}[\delta_1(X)|\theta, D = 1] + (1 - \mathbb{P}[D = 1|\theta])\delta_0\Big) & \text{if } \theta \leqslant 0 \\ \theta\Big(\mathbb{P}[D = 1|\theta] \cdot \big(1 - \mathbb{E}[\delta_1(X)|\theta, D = 1]\big) + (1 - \mathbb{P}[D = 1|\theta])(1 - \delta_0)\Big) & \text{if } \theta > 0 \end{cases} \end{aligned} \tag{5}$$

That is, regret equals the magnitude of the loss $|\theta|$ multiplied by the expected probability of assigning the inferior treatment choice. The expected probability of assigning the wrong treatment is a weighted average of making the incorrect decision, where weights correspond to different realizations of the publication process. Two points are worth noting. First, the publication cost does not appear in the expression for regret because it is constant with respect to the policymaker's decision rule. Second, this expression for regret reflects a sophisticated policymaker who has complete knowledge of publication bias. In particular, the sophisticated policymaker correctly accounts for publication when considering the distribution of the estimated treatment effect X, and the probability that a study is or is not published. In a later section, we study a naive policymaker who does not account for publication bias.

The expression for minimax regret can be further simplified by restricting the class of decision rules for $\delta_1(X)$ to threshold rules. As in Tetenov (2012) for the Gaussian signal case, this restriction is innocuous since in terms of welfare $W_1(\delta_1, \theta)$ for published case, the class of threshold rules is essentially complete, i.e., for any admissible rule $\delta_1(X)$ in terms of $W_1(\delta_1, \theta)$, its welfare level can be replicated by a threshold rule.

Lemma 1 (Threshold Rules are Essentially Complete). *Under Assumption 1, the class of threshold decision rules $\delta_1^T(X) = \mathbb{1}[X > T]$ is essentially complete in terms of the welfare of $W_1(\delta_1, \theta)$.*

With Lemma 1, any decision rule δ is fully characterized by a tuple (δ_1^T, δ_0). The first element corresponds to the threshold rule $\mathbb{1}[X > T]$ and is applicable when a study is published. The second element is a default action δ_0 and is applicable when a study is not published. With this simplification, we can rewrite regret for decision rule δ in equation (10) as

$$R\big((\delta_1^T, \delta_0), \theta\big) = \begin{cases} -\theta\Big(\mathbb{P}[D = 1|\theta] \cdot [1 - F(T|\theta, D = 1)] + (1 - \mathbb{P}[D = 1|\theta])\delta_0\Big) & \text{if } \theta \leqslant 0 \\ \theta\Big(\mathbb{P}[D = 1|\theta] \cdot F(T|\theta, D = 1) + (1 - \mathbb{P}[D = 1|\theta]) \cdot (1 - \delta_0)\Big) & \text{if } \theta > 0 \end{cases} \tag{6}$$

Finally, the optimal decision rule (T^*, δ_0^*) selects the rule which minimizes maximum regret:

$$(T^*, \delta_0^*) = \arg \min_{(T,\delta_0)\in\mathbb{R}\times[0,1]} \max_{\theta\in\mathbb{R}} R\big((\delta_1^T, \delta_0), \theta\big) \tag{7}$$

III. Optimal Publication Rule For Minimax Regret

In this section, we first characterize the optimal minimax regret decision rule for the sophisticated policymaker. Given this decision rule, we then show analytically that the optimal publication rule that minimizes the value of minimax regret is non-selective. Finally, we provide numerical evidence that this result generalizes to the case where the policy-maker's concerns over Type I and Type II error are asymmetric. Proofs are in Appendix A.

A. Optimal Minimax Regret Decision Rule

In the presence of publication bias, decision-makers must choose optimal actions for when studies are published and when they are not. The following lemma characterizes the optimal minimax regret decision rule:

Lemma 2 (Minimax Regret Decision Rule Under Publication Bias)**.** *Under Assumption 1 for the sophisticated policymaker, $(T^*, \delta_0^*) = \left(0, \frac{1}{2}\right)$ for any $\beta_p \in [0, 1]$.*

When a study is published, the optimal minimax decision rule implements the innovative treatment if the published estimate is positive; and when no study is published, the policymaker randomly choose between treatments with equal probability. With symmetric concern of Type I and Type II error, the policymaker will implement the innovative treatment when there is evidence that it is superior to the status quo treatment. When no study is published there exists no such evidence and hence the policymaker randomizes between treatments. The only information available to the policymaker when no study is published is that the difference in the efficacy of treatments is likely to be small, since publication bias censors small effect sizes. However, because publication bias is symmetric around zero, no information is gained about which treatment might be superior. When the study is published, a positive estimate is more likely to come from a positive true treatment effect, while a negative estimate is more likely to come from a negative true treatment effect. Hence, the policymakers' threshold rule implements the innovative policy if and only if the signal is positive.

It is noteworthy that the optimal minimax regret threshold decision rule in the presence of publication bias is identical to the case where there is no publication bias (Tetenov, 2012). This is a consequence of the symmetry of the problem when Type I error and Type II error are equally weighted by the policymaker.

B. Optimal Non-Selective Publication Rule

Given the minimax decision rule $(T^*, \delta_0^*) = \left(0, \frac{1}{2}\right)$, what publication rule minimizes the value of minimax regret? The following result provides the answer:

Proposition 1 (Non-Selective Optimal Publication Rule)**.** *Under Assumption 1, the value of minimax regret is minimized for the sophisticated policymaker when the publication rule is non-selective, that is, when $\beta_p = 1$.*

The optimal publication rule for a minimax regret policymaker publishes all results. Thus, under the optimal publication regime, the policymaker's problem collapses into the standard model with no publication bias in Tetenov (2012) where signals are normally distributed. The intuition behind this result is straightforward: publishing a study always provides useful information about the relative effectiveness of the treatments, which allows the policymaker to do better than in the case where no study is published and they randomize between choices.

This conclusion differs starkly from the optimal publication rule in a Bayesian framework. Frankel and Kasy (2022) show that the optimal publication rule in a Bayesian model only publishes extreme results, that is, results that move prior beliefs sufficiently. In this framework, null results are defined as those which do not change the policy-maker's prior. By contrast, the minimax regret framework does not rely on a prior distribution about treatment efficacy and thus this notion of 'null results' is not well-defined. Instead, we model publication bias based on the common definition of results being statistically indistinguishable from zero (Assumption 1).

A second key difference across these frameworks is the role of publication costs. In the Bayesian setting, publishing relatively uninformative results that do little to move the policymaker's prior belief yields no benefits, while at the same time incurring a cost; it is thus not optimal to publish such results. By contrast, in the minimax regret framework, the cost parameter c does not appear in the expression for regret, as can be seen in equation (6). This is because regret equals the size of the loss from making an inferior treatment choice, $|\theta|$, multiplied by the probability of this occurring. Since the expected cost of publication is the same irrespective of the decision rule, the expression for regret does not include it. Thus, publication costs have no impact on the minimax decision rule and therefore the optimal publication rule.

C. Type I Error Loss Aversion

Up until now, we have made the implicit assumption of equal weight for Type I error (of mistakenly implementing an inferior policy) and Type II error (of failing to implement the superior policy). However, in practice, policymakers may exhibit loss aversion and weigh the regret from Type I error higher relative to Type II error. In fact, Tetenov (2012) finds that classical hypothesis testing at the 5% level is consistent with a policymaker who weighs the regret from Type I error around 100 times more than the regret arising from Type II error.

To incorporate this asymmetry in the concern over different types of error, we follow Tetenov (2012) and introduce a Type I error loss aversion parameter $K > 0$. With this, the policymakers utility is given by

$$U\big(\delta, \theta, c\big) = \begin{cases} K\Big(\theta D\delta_1(X) + \theta(1-D)\delta_0 - Dc\Big) & \text{if } \theta \leqslant 0 \\ \theta D\delta_1(X) + \theta(1-D)\delta_0 - Dc & \text{if } \theta > 0 \end{cases} \tag{8}$$

Expected welfare is given by

$$W\big(\delta,\theta,c\big) = \begin{cases} K\Big(\theta \cdot \mathbb{P}[D=1|\theta] \cdot \mathbb{E}[\delta_1(X)|\theta, D=1] + \theta \cdot \big(1-\mathbb{P}[D=1|\theta]\big)\delta_0 - \mathbb{P}[D=1|\theta] \cdot c\Big) & \text{if } \theta \leqslant 0 \\ \theta \cdot \mathbb{P}[D=1|\theta] \cdot \mathbb{E}[\delta_1(X)|\theta, D=1] + \theta \cdot \big(1-\mathbb{P}[D=1|\theta]\big)\delta_0 - \mathbb{P}[D=1|\theta] \cdot c & \text{if } \theta > 0 \end{cases} \tag{9}$$

and regret is equal to

$$R\big(\delta,\theta\big) = W\big(\mathbb{1}(\theta>0),\theta\big) - W\big(\delta,\theta\big)$$

$$= \begin{cases} -K\theta\Bigg(\mathbb{P}[D=1|\theta] \cdot \mathbb{E}[\delta_1(X)|\theta, D=1] + (1-\mathbb{P}[D=1|\theta])\delta_0\Bigg) & \text{if } \theta \leqslant 0 \\ \theta\Bigg(\mathbb{P}[D=1|\theta] \cdot \big(1-\mathbb{E}[\delta_1(X)|\theta, D=1] + (1-\mathbb{P}[D=1|\theta])(1-\delta_0)\big)\Bigg) & \text{if } \theta > 0 \end{cases} \tag{10}$$

What is the optimal publication rule for different levels of loss aversion for Type I error K? Figure 1 plots minimax regret as a function of β_p for different values of K, in addition to the optimal minimax decision rule in each case. These figures are computed numerically. As a benchmark, the first column shows the regime where $K = 1$. First, see that minimax regret is decreasing β_p, in line with Proposition 1. Second, see that the optimal minimax regret decision rule is $(T^*, \delta_0^*) = (0, \frac{1}{2})$ for all $\beta_p \in [0, 1]$, in line with Lemma 2.

Now consider the case where $K = 3$ i.e. the policymaker weighs the Type I error of implementing the inferior treatment three times larger than the Type II error of failing to implement the superior treatment. As in the case where $K = 1$, minimax regret is also decreasing in β_p. See that the threshold rule is increasing in β_p. Similarly, the default probability of implementing the innovative policy in the event that a study Just not published, δ_0^*, is decreasing in β_p (and weakly less than $\frac{1}{2}$). That is, as β_p gets larger, the policymakers decision rule becomes more conservative with respect to assigning the innovative treatment. The intuition behind this is that as β_p increases, the possibility of noisier small effect being published increases, which increases the risk of committing Type I error.

Finally, consider the case where $K = 102.4$, which is the value that rationalizes hypothesis testing at the 5% significance level (Tetenov, 2012). Again, the level of minimax regret decreases as a the relative probability of publishing null results increases. Given the very high level of Type I loss aversion, the no-data rule is essentially zero for any value of β_p. Again, the threshold rule is increasing in β_p, and at a faster rate than as for the case where $K = 3$.

These are two particular cases. Numerical results for other values of K show similar patterns, namely, that the value of minimax regret is a decreasing function of β_p. Based on this, we conjecture that the optimal publication rule minimizing maximum regret being non-selective generalizes to any $K \geqslant 1$, although we do not have at present an analytical proof.

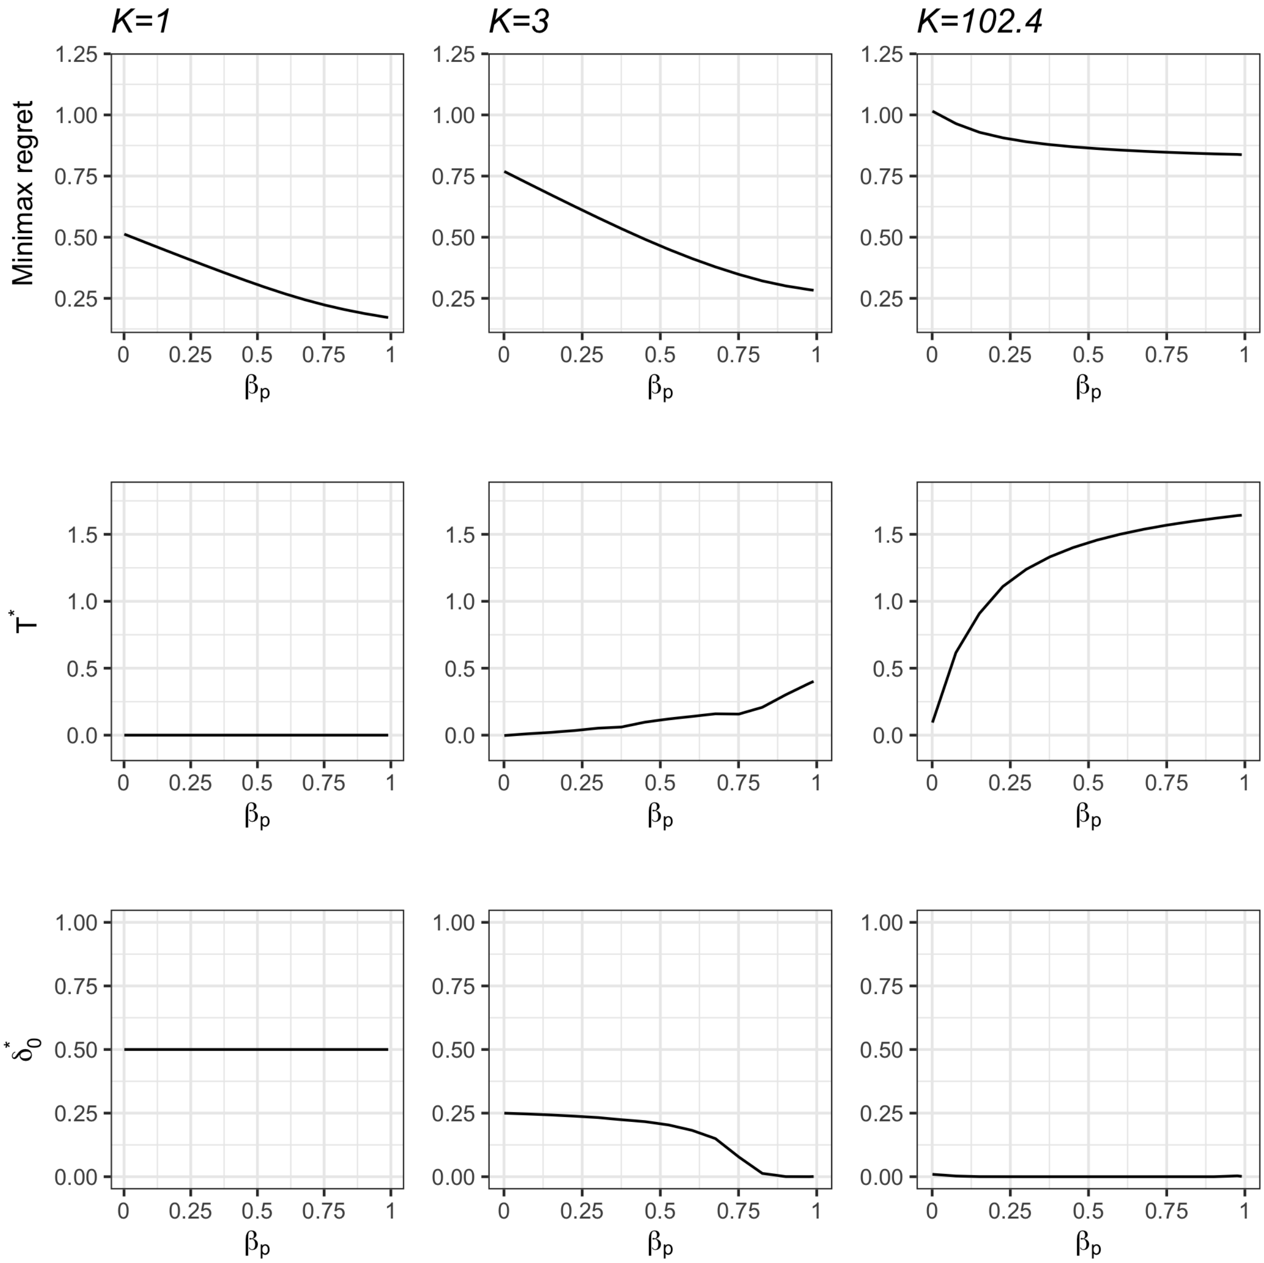

Figure 1. Minimax Regret and Optimal Decision Rule for Different Value of K

IV. Naive Policymakers

The sophisticated policymaker knows the exact form of publication bias and can accurately account for it. This is a strong assumption. As an alternative, we may consider a policymaker who naively chooses their decision rule without accounting for selective publication. This is perhaps more realistic, in the sense that most published research reports standard errors and assumes (approximately) normally distributed treatment effects for inference. 'Naiveity' impacts both realizations of the publication process. When a study is published, the policymaker erroneously believes it is normally distributed; and in the event that a study is not published, the naive policymaker fails to account for censoring in the publication process when choosing their default action. As in the sophisticated policymaker's problem, a decision rule δ is equivalent to the tuple (T, δ_0). The naive policymaker's *misspecified welfare* is equal to

$$\widetilde{W}\big((T,\delta_0),\theta\big) = \begin{cases} \theta[1-\Phi\big(T-\theta\big)] & \text{if } D=1 \\ \theta\cdot\delta_0 & \text{if } D=0 \end{cases} \tag{11}$$

This gives rise to two misspecified regret expressions. First, in the event that a study is published

$$\widetilde{R}_1\big(T,\theta\big) = \begin{cases} -\theta[1-\Phi\big(T-\theta\big)] & \text{if } \theta \leqslant 0 \\ \theta\Phi\big(T-\theta\big) & \text{if } \theta > 0 \end{cases} \tag{12}$$

and second, in the event that no study is published,

$$\widetilde{R}_0\big(\delta_0,\theta\big) = \begin{cases} -\theta\delta_0 & \text{if } \theta \leqslant 0 \\ \theta(1-\delta_0) & \text{if } \theta > 0 \end{cases} \tag{13}$$

Misspecified regret in equation (12) when a study is published is equivalent to the expression for regret in the model in Tetenov (2012) with normally distributed signals. This expression is misspecified because the policymaker does not account for the fact that selective publication distorts the distribution of estimated treatment effects. Misspecified regret when no study arrives, in equation (13), is simply a function of the default action δ_0 and the true effect θ. It is misspecified in that it ignores that possibility that a study was not published because of selective publication. For the minimax problem to be well-defined, we need to impose bounds on θ. For the naive policymaker, we impose the following assumption:

Assumption 2 (Symmetric Bounds on Average Treatment Effect)**.** *Let the support of* Θ *be* $[-B, B]$ *for some* $B > \theta^* > 0$*, where* $\theta^* = \arg\max_{\theta>0}\big\{\theta\cdot\Phi(0-\theta)\big\}$.

The technical condition that the bound is larger than $\theta^* = \arg\max_{\theta>0}\big\{\theta\cdot\Phi(0-\theta)\big\}$ ensures that the minimax problem when a study is published is not constrained by the bound.[3] The naive policymaker has, in effect, two decision problems, one for each realization of the publication process.

$$T^* = \arg\min_{T\in\mathbb{R}} \max_{\theta\in[-B,B]} \widetilde{R}_1(T,\theta) \tag{14}$$

[3] Tetenov (2012) shows that the maximum θ^* is attained on a closed interval $[0, H]$ for some $H > 0$.

$$\delta_0^* = \arg\min_{\delta_0 \in [0,1]} \max_{\theta \in [-B,B]} \widetilde{R}_0\big(\delta_0, \theta\big) \tag{15}$$

While (T, δ_0) are chosen by the naive policymaker under misspecified beliefs about the DGP, regret of any decision rule is assessed against the 'true' worst-case scenario which accounts for publication bias. That is, regret for any decision rule (T, δ_0) is identical to regret in the sophisticated policymaker's problem in equation (6).

To compare the 'cost' of naivity with respect to publication bias, we make the following calculation for some fixed K and assuming that $t_\alpha = 1.96$:

$$100 \cdot \left(\frac{\mathrm{MMR}^*_{Naive}(K)}{\mathrm{MMR}^*_{Soph}(K)} - 1 \right) \tag{16}$$

where $\mathrm{MMR}^*_{Naive}(K)$ is the value of minimax regret for the naive policymaker and $\mathrm{MMR}^*_{Soph}(K)$ is the value of minimax regret for the sophisticated policymaker.

Figure 2 illustrates the cost of naivity when $K = 3$. Results show that the cost of naivity if weakly positive. This is to be expected, since the naive planner chooses their decision rule under misspecified beliefs. Interestingly, the results show that the costs of naivety are highest when publication bias is moderate. When there is no publication bias, the cost of naivety is zero because the naive policymaker belief that there is no publication bias is correct in this special case. More surprisingly, the cost of naivety is also zero when there is extreme publication bias, such that no insignificant results are published. This is because the optimal threshold rule when the study is published is set identified and the solution for the naive policymaker and the sophisticated policymaker both fall within this set. In particular, any threshold rule above which the innovative treatment is implemented in the range of (-1.96σ, 1.96σ) will be effectively identical, because no insignificant studies within this range are ever published.

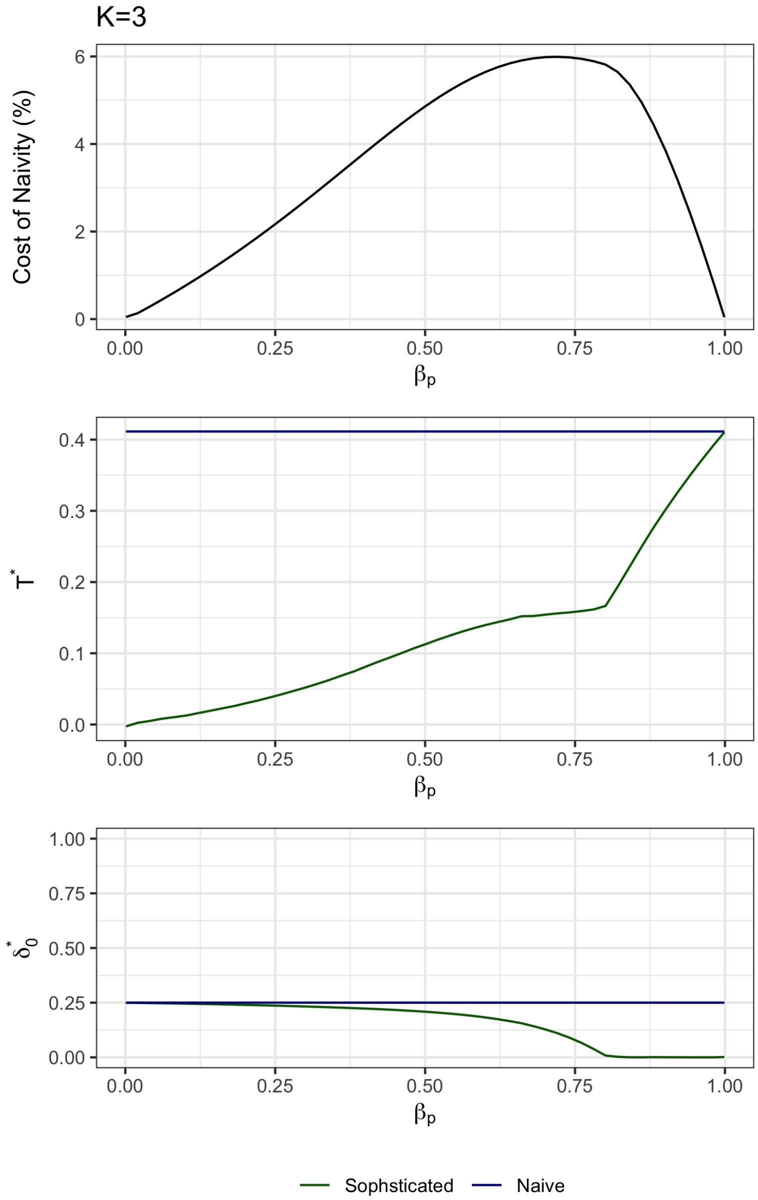

Figure 2. Cost of Naivity ($K = 3$)

V. Misaligned Preferences

In the main analysis, the policymaker chooses a decision rule to minimize maximum regret, and we consider the optimal publication rule of a journal editor who chooses $\beta_p \in [0, 1]$ with the same preferences. In this extension, we consider what happens when the policymaker and the journal editor do not have aligned preferences. In particular, we continue to assume that the policymaker optimizes using minimax regret, but instead consider a journal editor who chooses $\beta_p \in [0, 1]$ to maximize welfare under a Bayesian prior. Since the policymakers' decision rule could in theory depend on the journal editor's choice for β_p, we can view the equilibrium outcome as resulting from a game between the editor and the policymaker. However, since Lemma 2 shows that the minimax decision rule is the same for any value of $\beta_p \in [0,1]$, there are no strategic considerations at play. Throughout, we assume that Type I and Type II error are equally weighted ($K = 1$).

More formally, for the policymaker's decision rule δ and publication cost c, the Bayesian journal editor's problem is given by

$$\max_{\beta_p \in [0,1]} \int W\big(\delta, \theta, c\big)\pi(\theta)d\theta \tag{17}$$

where welfare is given by equation 4 and $\pi(\cdot)$ denotes the prior belief distribution of the journal editor. We assume that the prior satisfies the following regularity conditions:

Assumption 3 (Support of Journal Editor's Prior)**.** *Let the prior distribution $\pi(\cdot)$ have support on an open subset of the real line.*

Recall the policymaker's optimal minimax rule from Lemma 2 and that it is identical for under both sophisticated and naive beliefs when $K = 1$. The following Proposition gives the optimal publication rule of the Bayesian journal editor:

Proposition 2 (Optimal Bayesian Publication Rule)**.** *Suppose the policymaker implements the optimal minimax regret decision rule $(T^*, \delta_0^*) = (0, \frac{1}{2})$. Under Assumptions 1 and 3, the Bayesian journal editor's optimal publication rule for any $c \geqslant 0$ is given by*

$$\beta_p^* = \begin{cases} 1 & \text{if } c \leqslant T \\ 0 & \text{if } c > T \end{cases} \tag{18}$$

where

$$T = \frac{\frac{1}{2}\int \theta\Big(\big[\Phi(t_\alpha - \theta) - \Phi(-\theta)\big] - \big[\Phi(-\theta) - \Phi(-t_\alpha - \theta)\big]\Big)\pi(\theta)d\theta}{\int \big[\Phi(t_\alpha - \theta) - \Phi(-t_\alpha - \theta)\big]\pi(\theta)d\theta} > 0$$

The journal's optimal action takes the form a simple threshold rule: publish all results if publication costs are sufficiently low; otherwise, censor all null results. Thus, when publication costs are low, the optimal publication rule under misaligned preferences is the same as with aligned preferences, namely, it is non-selective. However, when publication costs are high, it will be optimal to censor all null results.

For the Bayesian policymaker in the Frankel and Kasy (2022) model, the optimal publication rule recommends censoring results which do not sufficiently move the prior. In other words, the journal does not publish ‘unsurprising’ findings close to its prior beliefs on a given research question (which is assumed to be shared by the public). Our result in Proposition 2 differs because we consider the class of publication rules which censor statistically insignificant results. The rationale behind this is that the censoring of null results is the most common form of publication bias highlighted in the literature.

VI. Conclusion

This paper studies treatment choice in the presence of publication bias in the case where policymakers are unwilling or unable to rely on prior beliefs about relative treatment efficacy. We show that the optimal publication rule which minimizes maximum regret is non-selective. This holds whether or not policymakers account for publication bias in choosing their treatment rule i.e. whether they are sophisticated or naive in their beliefs about the DGP. This contrasts with the Bayesian policymaker studies in the literature, where the optimal publication rule for policy choice censors results close to the decision-maker’s prior. Thus, the optimal publication regime – and hence the statistical credibility of published research – can vary drastically depending on the optimality criterion pursued by the policymaker and journals. In the minimax framework, the publication regime which is optimal for treatment choice also delivers valid statistical inference.

References

Andrews, Isaiah and Maximilian Kasy, "Identification of and Correction for Publication Bias," *American Economic Review*, 2019, *109* (8), 2766–2794.

Ferguson, Thomas S., *Mathematical Statistics: A Decision Theoretic Approach*, New York, Academic Press, 1967.

Franco, Annie, Neil Malhotra, and Gabor Simonovits, "Publication bias in the social sciences: Unlocking the file drawer," *Science*, 2014, *345* (6203), 1502–1505.

Frankel, Alexander and Maximilian Kasy, "Which Findings Should Be Published?," *American Economic Journal: Microeconomics*, 2022, *14* (1), 1–38.

Ioannidis, John P.A., "Why Most Published Research Findings Are False," *PLoS Med*, 2005, *2* (8).

Karlin, Samuel and Herman Rubin, "The Theory of Decision Procedures for Distributions with Monotone Likelihood Ratio," *Annals of Mathematical Statistics*, 1956, *27*, 272–299.

Lehmann, Erlich L. and Joseph P. Romano, *Testing Statistical Hypotheses*, Springer, 2005.

Manski, Charles F., "Statistical Treatment Rules for Heterogeneous Populations," *Econometrica*, 2004, *72* (4), 1221–1246.

McShane, Blakeley B., David Gal, Andrew Gelman, Christian Robert, and Jennifer L. Tackett, "Abandon Statistical Significance," *The American Statistician*, 2019, *73* (1), 235–245.

Miguel, Edward and Garret Christensen, "Transparency, Reproducibility, and the Credibility of Economics Research," *Journal of Economic Literature*, 2018, *56* (3), 920–980.

Milgrom, Paul and Ilya Segal, "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 2002, *70* (2), 583–601.

Nosek, Brian A., Charles R. Ebersole, Alexander C. DeHaven et al., "The Preregistration Revolution," *Proceedings of the National Academy of Sciences*, 2018, *115* (11), 2600–2606.

_ , George Alter, George C. Banks et al., "Promoting an Open Research Culture," *Science*, 2015, *348* (6242), 1422–1425.

Savage, Leonard J., "The Theory of Statistical Decision," *Journal of the American Statistical Association,*, 1951, *46* (253), 55–67.

Stoye, Jörg, "Minimax Regret Treatment Choice With Finite Samples," *Journal of Econometrics*, 2009, *151* (1), 70–81.

Tetenov, Aleksey, "Statistical treatment choice based on asymmetric minimax regret criteria," *Journal of Econometrics*, 2012, *166*, 157–165.

A. Proofs

Proof of Lemma 1: We focus on the welfare function $W_1(\delta_1, X)$ for the published case. Karlin and Rubin (1956) shows that if the distribution of sufficient statistics for θ satisfies the monotone likelihood ratio property, the class of threshold decision rules is essentially complete for a class of loss functions including the current one. Under Assumption 1, $F(X|\theta, D=1)$ is an exponentially family distribution with pdf

$$C(\theta)h(x)\exp(x\theta), \tag{19}$$

where $C(\theta) = \frac{\exp(-\theta^2/2)}{\sqrt{2\pi}\int p(t)\phi(t-\theta)dt}$ and $h(x) = p(x)\exp(-x^2/2)$, and X being a sufficient statistics for θ. Since the exponential family distribution satisfies the monotone likelihood ratio property (see, e.g., Section 3.4 in Lehmann and Romano (2005)), the current lemma follows. □

Proof of Lemma 2: The proof follows two main steps. First, we solve the minimax problem for the sophisticated policymaker. In the second step, we show that the naive policymaker, who optimizes under misspecified beliefs about the DGP, nonetheless arrives at the optimal solution.

Sophisticated policymaker.—First, we show that the optimal decision rule for the sophisticated policymaker is $(T^*, \delta_0^*) = \left(0, \frac{1}{2}\right)$. To do this, we use the following theorem (for reference, see Theorem 1 in section 2.11 (pg 90) in Ferguson (1967)):

Lemma A.1. *Suppose δ minimizes Bayes risk under π:*

$$\delta \in arg \min_{\delta' \in \mathcal{D}} \int_\theta R(\delta', \theta)d\pi(\theta)$$

and

$$R(\delta, \theta) \leqslant \int_\theta R(\delta, \theta)d\pi(\theta)$$

for all $\theta \in \Theta$. Then δ is a minimax rule and π is least favourable.

Using this lemma, we first propose a guess for δ and π and then show that this guess satisfies the sufficient conditions in Theorem A1 which imply that δ is the minimax regret decision rule.

Our guess is that the minimax regret decision rule is $(T^*, \delta_0^*) = (0, \frac{1}{2})$. Regret under this proposed rule for any θ is equal to:

$$R\big((0,0.5),\theta\big) = \begin{cases} -\theta\Big(\mathbb{P}[D=1|\theta]\cdot[1-F(0|\theta,D=1)] + (1-\mathbb{P}[D=1|\theta])\frac{1}{2}\Big) & \text{if } \theta \leqslant 0 \\ \theta\Big(\mathbb{P}[D=1|\theta]\cdot\big(F(0|\theta,D=1) + (1-\mathbb{P}[D=1|\theta])\frac{1}{2}\big)\Big) & \text{if } \theta > 0 \end{cases} \tag{20}$$

Next, guess that Nature's least favorable prior is equal to

$$\pi = \begin{cases} \theta_+^* & \text{with probability } \frac{1}{2} \\ -\theta_+^* & \text{with probability } \frac{1}{2} \end{cases} \tag{21}$$

where $\theta_+^* = \arg\max_{\theta>0} R\big((0,0.5),\theta\big)$. We know that $\theta_+^* \in (0,\infty)$ because $R\big((0,0.5),0\big) = 0$; $R\big((0,0.5),\theta\big) \to 0$ as $\theta \to \infty$; and $R\big((0,0.5),\theta\big) > 0$ for any $\theta > 0$. The first and third claims can be seen directly from equation (20). To see why the second claim is true see that

$$\lim_{\theta\to\infty}\Big\{\theta\cdot\mathbb{P}[D=1|\theta]\cdot F(0|\theta,D=1)\Big\} + \frac{1}{2}\lim_{\theta\to\infty}\Big\{\theta\cdot(1-\mathbb{P}[D=1|\theta])\big)\Big\} \tag{22}$$

The first term equals zero because

$$\lim_{\theta\to\infty}\Big\{\theta\cdot\mathbb{P}[D=1|\theta]\cdot F(0|\theta,D=1)\Big\} < \lim_{\theta\to\infty}\Big\{\theta\cdot\Phi(0-\theta)\Big\} = \lim_{\theta\to\infty}\Big\{\theta^2\cdot\phi(0-\theta)\Big\} = 0 \tag{23}$$

where the first inequality follows because $F(.|\theta,D=1)$ is a truncated normal cdf and $\theta > 0$; the second last equality follows from L'Hôpital's rule; and the final equality follows from the fact that $\theta^2\cdot\phi(0-\theta)$ has finite moments. The second term also equals zero since we have

$$\theta\cdot(1-\mathbb{P}[D=1|\theta])\big) = (1-\beta_p)(2\pi)^{-1}\int_{-t_\alpha}^{t_\alpha}\theta\exp\left(-\frac{1}{2}(t-\theta)^2\right)dt \tag{24}$$

and $\lim_{\theta\to\infty}\theta\exp\left(-\frac{1}{2}(t-\theta)^2\right) = 0$ at every $t \in [-t_\alpha, t_\alpha]$, and apply the dominated convergence theorem.

Next, we will show that $(T^*,\delta_0^*) = (0,\frac{1}{2})$ minimizes Bayes risk with respect to π. For any decision rule (T,δ_0), Bayes risk equals

$$\int_\theta R\big((T,\delta_0),\theta\big)d\pi(\theta) = \frac{1}{2}\cdot\theta_+^*\Big(\mathbb{P}[D=1|\theta_+^*]\cdot F(T|\theta_+^*,D=1) + (1-\mathbb{P}[D=1|\theta_+^*])(1-\delta_0)\Big)$$

$$+\frac{1}{2}\cdot\theta_+^*\Big(\mathbb{P}[D=1|\theta_+^*]\cdot[1-F(T|\theta_+^*,D=1)]+(1-\mathbb{P}[D=1|\theta_+^*])\delta_0\Big)$$

$$=\frac{1}{2}\cdot\theta_+^*\Big(1-\mathbb{P}[D=1|\theta_+^*]\Big)+\frac{1}{2}\cdot\theta_+^*\mathbb{P}[D=1|\theta_+^*]\Big(F(T|\theta_+^*,D=1)+F(-T|\theta_+^*,D=1)\Big) \quad (25)$$

Note that any δ_0 is optimal, so we can choose $\delta_0^*=\frac{1}{2}$. We will show that $T^*=0$ minimizes Bayes risk by showing that $F(T|\theta_+^*,D=1)+F(-T|\theta_+^*,D=1)$ is minimized when $T=0$. To do this, we will show that any other choice of T leads to higher regret. Since the Bayes risk under π (25) is symmetric in T, without loss of generality, we assume $T>0$. Consider first the case where $T>t_\alpha>0$. We have

$$\begin{aligned}F(-T|\theta_+^*,D=1)+F(T|\theta_+^*,D=1)=\frac{1}{C}\Big(&\Phi(-T-\theta_+^*)+\\ &+\Phi(-T-\theta_+^*)+\big[\Phi(-t_\alpha-\theta_+^*)-\Phi(-T-\theta_+^*)\big]\\ +\beta_p\big[\Phi(0-\theta_+^*)-\Phi(-t_\alpha-\theta_+^*)\big]+\beta_p\big[\Phi(t_\alpha-\theta_+^*)-\Phi(0-\theta_+^*)\big]&+\big[\Phi(T-\theta_+^*)-\Phi(t_\alpha-\theta_+^*)\big]\Big)\\ >\frac{2}{C}\Big(\Phi(-t_\alpha-\theta_+^*)+\beta_p\big[\Phi(0-\theta_+^*)-\Phi(-t_\alpha-\theta_+^*)\big]\Big)&=2\cdot F(0|\theta_+^*,D=1) \quad (26)\end{aligned}$$

where $C=\int p(z)\phi(z-\theta_+^*)dz$ is the normalization constant of the truncated normal distribution. The case where $t_\alpha>T>0$ follows a similar argument. Thus, $(T^*,\delta_0^*)=(0,\frac{1}{2})$ minimizes Bayes risk with respect to π.

Finally, see that for any $\theta\in\mathbb{R}$, we have that

$$R\big((0,0.5),\theta\big)\leqslant R\big((0,0.5),\theta_+^*\big)=\frac{1}{2}R\big((0,0.5),\theta_+^*\big)+\frac{1}{2}R\big((0,0.5),-\theta_+^*\big)=\int_\theta R(\delta,\theta)d\pi(\theta) \quad (27)$$

The first inequality follows from the construction of θ_+^*. The next equality uses the symmetry of the regret function with respect to θ around zero. From Theorem A1, it then follows that the minimax regret decision rule for the sophisticated policymaker is $(T^*,\delta_0^*)=(0,\frac{1}{2})$ and the least favourable prior is π in equation (21).

Naive policymaker.—Next, we show that the naive policymaker arrives at the same decision rule, despite ignoring selective publication. The naive policymaker's optimal decision rule consists of two problems, when a study is published and when it is not. When a study is published, the policymaker (erroneously) believes the signal is normally distributed. This is equivalent to the problem in Tetenov (2012), who proves that the optimal solution in the symmetric case is $T^*=0$.

Next, consider the case where no study is published. Misspecified regret is equal to

$$\widetilde{R}_0(\delta_0, \theta) = \begin{cases} -\theta\delta_0 & \text{if } \theta \leqslant 0 \\ \theta(1-\delta_0) & \text{if } \theta > 0 \end{cases} \tag{28}$$

and thus misspecified worse-case regret given bounds in Assumption 2 is given by $\max_{\theta \in [-B,B]} \widetilde{R}_0(\delta_0, \theta) = \max\{B\theta_0, B(1-\delta_0)\}$. The minimax regret decision rule equalizes the arguments in the max operator, giving $\delta_0^* = \frac{1}{2}$.

Proof of Proposition 1: We have shown that for any $\beta_p \in [0,1]$, both the sophisticated and naive policymakers' optimal minimax decision rule is $(T^*, \delta_0^*) = (0, \frac{1}{2})$. It remains to show that $\beta_p = 1$ is the optimal publication rule, in the sense that it minimizes minimax regret.

Denote the value of minimax regret as a function of parameter β_p:

$$\begin{aligned} V(\beta_p) &\equiv \max_{\theta > 0} \left\{ \theta \Big(\mathbb{P}[D=1|\theta] F(0|\theta, D=1) + \big(1 - \mathbb{P}[D=1|\theta]\big) \frac{1}{2} \Big) \right\} & (29) \\ &= \max_{\theta > 0} \left\{ \theta \int_{-\infty}^{0} p(y)\phi(y-\theta)dy + \frac{\theta}{2} \int_{-\infty}^{\infty} [1-p(y)]\phi(y-\theta)dy \right\}, \\ &= \max_{\theta > 0} f(\theta, \beta), & (30) \end{aligned}$$

where $f(\theta, \beta) = \theta \int_{-\infty}^{0} p(y)\phi(y-\theta)dy + \frac{\theta}{2} \int_{-\infty}^{\infty} [1-p(y)]\phi(y-\theta)dy$ and its dependence on β_p is only through $p(\cdot)$.

Note that the value function inside the maximum operator is continuously differentiable in β_p with an integrable envelope over the domain of $\beta_p \in [0,1]$. Hence, by the generalized envelope theorem (Theorem 2 in Milgrom and Segal (2002)), $V(\beta_p)$ is absolutely continuous and admits the following integral representation:

$$V(\beta_p) = V(0) + \int_0^{\beta_p} f_{\beta_p}(\theta^*(\beta_p'), \beta_p') d\beta_p', \tag{31}$$

where $f_{\beta_p}(\cdot,\cdot) = \frac{\partial}{\partial \beta} f(\theta, \beta)$ and $\theta^*(\beta_p)$ is a maximizer of $f(\theta, \beta_p)$ in θ given β_p. Note that for $\theta > 0$, we can show

$$f_{\beta_p}(\theta, \beta_p) = \frac{\theta}{2} \left[\int_{-t_\alpha}^{0} \phi(y-\theta)dy - \int_0^{t_\alpha} \phi(y-\theta)dy \right] < 0. \tag{32}$$

To see this inequality holds, consider two cases. First, suppose that $\theta \geqslant t_\alpha$. Then we immediately get the desired result because $\phi(z-\theta)$ is strictly increasing over $(-t_\alpha, t_\alpha)$.

Next consider the case where $\theta \in (0, t_\alpha)$. Then $\int_0^\theta \phi(y-\theta)dy > \int_{-\theta}^0 \phi(y-\theta)dy$ since $\phi(y-\theta)$ is strictly increasing in y for any $y < \theta$. And we also have that $\int_\theta^{t_\alpha} \phi(y-\theta)dy = \int_{2\theta-t_\alpha}^\theta \phi(y-\theta)dy > \int_{-t_\alpha}^{-\theta} \phi(y-\theta)dy$, where the first equality uses symmetry of the normal distribution about θ and the second equality again uses the fact that $\phi(y-\theta)$ is strictly increasing in y for any $y < \theta$. Taking these two inequalities together leads to the inequality of (32).

Combining (31) and (32), we conclude that $V(\beta_p)$ is a monotonically decreasing function, and $\beta_p = 1$ minimizes the value of minimax regret.

Proof of Proposition 2: Fix the optimal minimax rule for the policymaker: $\delta^* = (T^*, \delta_0^*) = (0, \frac{1}{2})$. Then

$$\int W\big(\delta^*, \theta, c\big)\pi(\theta)d\theta = \int \theta \cdot \mathbb{P}[D=1|\theta,\beta_p]\big[1 - F(0|D=1,\theta,\beta_p)\big]\pi(\theta)d\theta$$

$$+\frac{1}{2}\int \theta \cdot \big(1 - \mathbb{P}[D=1|\theta,\beta_p]\big)\pi(\theta)\theta - c\int \mathbb{P}[D=1|\theta,\beta_p]\pi(\theta)d\theta$$

Now see that
$$\frac{\partial}{\partial \beta_p}\bigg(\mathbb{P}[D=1|\theta,\beta_p]\bigg) = \Phi(t_\alpha - \theta) - \Phi(-t_\alpha - \theta)$$

$$\frac{\partial}{\partial \beta_p}\bigg(F(0|D=1,\theta,\beta_p)\cdot \mathbb{P}[D=1|\theta,\beta_p]\bigg) = \Phi(-\theta) - \Phi(-t_\alpha - \theta)$$

which implies that

$$\frac{\partial}{\partial \beta_p}\bigg[\int W\big(\delta^*, \theta, c\big)\pi(\theta)d\theta\bigg] = \frac{1}{2}\int \theta\Big(\big[\Phi(t_\alpha - \theta) - \Phi(-\theta)\big] - \big[\Phi(-\theta) - \Phi(-t_\alpha - \theta)\big]\Big)\pi(\theta)d\theta$$

$$-c\int \big[\Phi(t_\alpha - \theta) - \Phi(-t_\alpha - \theta)\big]\pi(\theta)d\theta$$

It is clear that the integral in the second term multiplied by c is positive. If the integral in the first term is strictly positive, then the desired result clearly follows. That is, for sufficiently low c, the derivative will be positive and the optimal rule will be $\beta_p^* = 1$. Conversely, for sufficiently high c, the derivative will be negative and the optimal publication rule will be $\beta_p^* = 0$.

In the remainder of the proof, we show the integral is indeed positive. For clarity, define the integrand $g(\theta) \equiv \theta\big([\Phi(t_\alpha - \theta) - \Phi(-\theta)] - [\Phi(-\theta) - \Phi(-t_\alpha - \theta)]\big)$. First, see that $g(0) = 0$. However, Assumption 3 implies that there exists some $\theta \neq 0$ on the support of $\pi(\cdot)$. Thus, to show that the integral is positive, it suffices to show that $g(\theta) > 0$ for all $\theta \neq 0$.

To show this, first note that $g(\cdot)$ is symmetric about zero i.e. $g(\theta) = g(-\theta)$. We can therefore restrict our attention to $\theta > 0$. Consider two cases. First, suppose $t_\alpha - \theta \leqslant 0$. Then

$g(\theta) > 0$ if and only if $[\Phi(t_\alpha - \theta) - \Phi(-\theta)] - [\Phi(-\theta) - \Phi(-t_\alpha - \theta)] > 0$, which clearly holds because the normal density is increasing over $(-\infty, 0)$.

Next, suppose that $t_\alpha - \theta > 0 \iff t_\alpha > \theta > 0$. Then breakup up the integral and using the symmetry of the normal density, we have

$$g(\theta) = [\Phi(t_\alpha - \theta) - \Phi(-\theta)] - [\Phi(-\theta) - \Phi(-t_\alpha - \theta)]$$
$$= \Big([\Phi(t_\alpha - \theta) - \Phi(0)] + [\Phi(0) - \Phi(-\theta)]\Big) - \Big([\Phi(-\theta) - \Phi(-2\cdot\theta)] + [\Phi(-2\cdot\theta) - \Phi(-t_\alpha - \theta)]\Big)$$
$$= \Big([\Phi(0) - \Phi(-\theta)] - [\Phi(-\theta) - \Phi(-2\cdot\theta)]\Big) + \Big([\Phi(t_\alpha - \theta) - \Phi(0)] - [\Phi(t_\alpha + \theta) - \Phi(2\cdot\theta)]\Big) > 0$$

where the inequality follows because both differences in the parentheses are strictly positive. The first difference is positive because the normal density if increasing over $(\infty, 0)$. The second difference is positive because the normal density if decreasing over $(0, \infty)$. $\square$