

Three Essays in Economics

by

Chishio Furukawa

B.A. and Sc.B., Brown University (2013)

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Economics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2020

© Chishio Furukawa, MMXX. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Author

Department of Economics

August 15, 2020

Certified by

Abhijit V. Banerjee

Ford Foundation International Professor of Economics

Thesis Supervisor

Certified by

Stephen Morris

Peter A. Diamond Professor in Economics

Thesis Supervisor

Accepted by

Amy Finkelstein

John & Jennie S. MacDonald Professor of Economics

Chairman, Departmental Committee on Graduate Studies

Three Essays in Economics

by

Chishio Furukawa

Submitted to the Department of Economics
on August 15, 2020, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Economics

Abstract

This thesis consists of three essays on diverse topics but shared emphasis on statistical models with theory and empirics. The first and third essay examines the role of cognitive limitations in understanding biases in communication and learning. The second essay, joint with Masao Fukui, highlights the role of distributional assumptions of infection rates for epidemiological predictions, responding to the recent COVID-19 outbreak.

The first chapter considers the effects of aggregation frictions on scientific communication and shows that publication bias emerges even when researchers are unbiased and communicate their findings optimally for readers. Specifically, when readers are cognitively constrained, they may only consider the binary conclusions rather than the estimates of the papers. Under such aggregation frictions of readers, researchers are shown to omit noisy null results and inflate marginal results. This chapter presents evidence consistent with this prediction, and develops a new bias correction method, called stem-based correction method, that is robust under the prediction of this and other models of publication selection processes.

The second chapter examines the role of infection rate distributions for aggregate epidemiological dynamics in Susceptible-Infectious-Recovered (SIR) models. Specifically, we show that superspreading events (SSEs) of recent coronavirus outbreaks, including SARS, MERS, and COVID-19, follow a power law distribution with fat tails, or infinite variance. When embedding this distribution to stochastic SIR models, we find that idiosyncratic variations in SSEs generate important uncertainties in aggregate epidemiological dynamics. This result stands in contrast with the existing literature on stochastic SIR models that have assumed thin tailed distributions, and thus concluded that the idiosyncratic uncertainties are unimportant when the population is large.

The third chapter considers the impact of imperfect recall on experimentation decisions and resulting inferences. When a Bayesian experimenter has an imperfect recall over past actions and information, her decisions depend not only on a confidence level but also on the expectation the future self will hold for today's action. This expectation arises from the persistent prior belief, and leads to the biases to conform to

it. Meditation, to regulate one's attention with focus on the present, is shown to have an ameliorating effect: when the attention is focused, prior belief becomes essentially diffused so that the self-imposed expectation over behaviors becomes agnostic.

Thesis Supervisor: Abhijit V. Banerjee

Title: Ford Foundation International Professor of Economics

Thesis Supervisor: Stephen Morris

Title: Peter A. Diamond Professor in Economics

Acknowledgments

No words of appreciation can carry the gravity of all debts I have accumulated through the past seven years as a graduate student. Acknowledgements herein are limited to my mentors, colleagues, and other staff members in the economics community with whom I had sustained relationships. Yet I also received countless inspirations from spontaneous discussions; I am deeply grateful even though listing them is impossible.

For professional guidance, I am indebted to Abhijit Banerjee for his patience and insights throughout my time as a graduate student; Stephen Morris for his optimism and generosity; Anna Mikusheva for her kindness and care. While my interests are yet to materialize, I also benefited immensely from my discussions with Iván Werning and Jim Poterba for projects on public economics. For their instructive advice, I am also grateful to Bob Gibbons, Juuso Toikka, and Muhamet Yildiz. I have also gained important perspectives from my discussions with Esther Duflo and Frank Schilbach. While I was at the University of Tokyo between April 2018 and January 2020, I benefited from my discussions with Daiji Kawaguchi.

For a close personal companion, I was fortunate to have office mates, Alonso Bucarey, Román Andrés Zárate, and Alan Olivi; for many discussions on applied theory and beyond, I thank Harry Di Pei and Arda Gitmez; for many encouragements and personal support, I think Kosti Takala, Fei Song, and Cara Nikolaus; Mahvish Shaukat, Marco Tabellini, Ariel Zucker, Josh Dean, Ben Roth, Matt Lowe, and Hoai-Luu Nguyen; Dan Waldinger, Jon Petkun, Ray Kluender, Michael Stepner, Ben Marx, Maddie McKelway, and Otis Reid; Michael Wong, Sourav Saukar, and Sophie Sun; Yuhei Miyauchi, Masao Fukui, and Yoko Okuyama. From my college, I was fortunate to continue my friendships with Daniel Prinz, Yunan Ji, Neil Thakral, and Fatima Aqueel. In Tokyo, I had many inspiring discussions with Taiyo Fukai.

Yet maintaining my focus on research was impossible without administrative support: I thank Gary King, Emily Gallagher, Kim Scantlebury, and Keesler Welch. In Tokyo, I received numerous support from Tsukasa Atsuya. For editing support to improve expositions, I thank Varanya Chaubey and Philip McLellan.

For the generous financial support, I thank the Japan-IMF Scholarship Program and the Schultz Fund.

This Ph.D. thesis is but a progress report of my personal growth. The dispersed sets of essays contained here suggest there remains much work to be completed for various inquiries as a part of my coherent sets of interests. I wish to repay in many years ahead by building on the foundations my mentors and colleagues have laid in me.

JEL codes: D83, D91, C14.

Contents

1 Unbiased Publication Bias:	
From Communication Model to New Correction Method	13
1.1 Introduction	13
1.2 A Communication Model of Publication Bias	19
1.2.1 Set-up	19
1.2.2 Analyses	20
1.2.3 Main Result 1. Omission of Insignificant Results	22
1.2.4 Main Result 2. Inflation of Marginally Insignificant Results	28
1.2.5 Additional Result. Amplification of Small Bias of a Researcher	32
1.2.6 Discussions of Key Assumptions	35
1.3 An Empirical Test of the Communication Model	37
1.3.1 Various Models of Publication Bias	37
1.3.2 A Test to Distinguish the Various Models	40
1.3.3 Application	41
1.4 A New “Stem-based” Bias Correction Method	43
1.4.1 Main Argument	44
1.4.2 Estimation	46
1.4.3 Assessment	50
1.4.4 Final Remarks	53
1.5 Conclusion	54
2 Power Laws in Superspreading Events: Evidence from Coronavirus Outbreaks and Implications for SIR Models	103

2.1	Introduction	103
2.2	Evidence	109
2.2.1	Statistical model	109
2.2.2	Data	110
2.2.3	Estimation	111
2.2.4	Results	113
2.3	Theory	117
2.3.1	Stochastic SIR model with fat-tailed distribution	117
2.3.2	Effects of fat-tailed distribution on uncertainty	119
2.3.3	Effects of fat-tailed distribution on average	122
2.3.4	Lockdown policy targeted at SSEs	124
2.4	Estimation methods	125
2.4.1	Limitations of sample means	126
2.4.2	Using power law exponents to improve inference	127
2.5	Conclusion: implications for COVID-19 pandemic	130
2.5.1	Proof that I^{\max} is concave in \mathcal{R}_0 if and only if $\mathcal{R}_0 > \frac{1}{S_0} \exp(0.5)$	149
3	A Model of Reflection and Meditation in Experimentation under Imperfect Recall	153
3.1	Introduction	153
3.2	Set-up	159
3.2.1	Environment	159
3.2.2	Equilibrium	161
3.3	Main Analyses	164
3.3.1	Behaviors and Beliefs	164
3.3.2	Role of Reflection and Meditation	167
3.3.3	Coordination Problems between Two Selves	169
3.4	Conclusion	172

List of Figures

1-1	Equilibrium publication decision	15
1-2	Equilibrium thresholds and policy decisions for $N = 2, c = 0$	25
1-3	$\bar{\beta}(\sigma)$ and $\underline{\beta}(\sigma)$ thresholds	31
1-4	Selection models described in funnel plots	39
1-5	KS-type test illustrated in funnel plots	43
1-6	Meta-analysis estimates in funnel plots	45
1-7	A visualization of stem-based method	47
2-1	Log cluster size vs log rank for COVID-19 worldwide	105
2-2	Log size vs log rank for COVID-19	114
2-3	Comparison of power law and negative binomial distributions	116
2-4	Ten sample paths from simulation	120
2-5	Histogram from 1000 simulation	121
2-6	An example of sample mean estimates	127
2-7	An example of “plug-in” estimates	129
2.5.1	Log size vs log rank for Superspreading Events in SARS 2003	146
3.1.1	A rapid increase in evidence and practice of meditation	155
3.2.2	Time line	160
3.3.3	Experimentation with bias towards prior belief	166

List of Tables

1.1	Simulation of bias correction methods across various models	53
2.1	Estimates of power law exponent ($\hat{\alpha}$) and their fit with data	115
2.2	Probabilities of extreme SSEs under each distribution	117
2.3	Parameter values	119
2.4	Summary statistics for epidemiological outcomes	132
2.5	Coverage probability of 95% confidence interval	133
2.5.1	Estimates of power law exponent: robustness	147
2.5.2	Examples of superspreading events	148
2.5.3	Peak infection under targeted lockdown policy	151

Chapter 1

Unbiased Publication Bias: From Communication Model to New Correction Method

1.1 Introduction

In many settings, professionals often rely on scientific publication when making important decisions. For example, suppose doctors are deciding whether to prescribe a new medication. Suppose they find 10 studies, and out of them, 7 report positive results whereas 3 report negative¹ results. Since most of the evidence is positive, they decide to use the new drug. Later they read a *New York Times* article that writes “researchers and pharmaceutical companies never published about one third of drug trials that they conducted to win the government approval, misleading the doctors and consumers regarding the drug’s true effectiveness.” The article² reports that 94 percent of positive results were published whereas only 14 percent of negative results were published. Moreover, while there were 14 studies with negative results according to conventional statistical thresholds, 11 of them emphasized positive aspects of their conclusions. These forms of publication bias, both omission and inflation, are widely documented in economics and other areas of science.

Concern about publication bias plays a prominent role in discussions of scientific reporting and meta-analyses today (Christensen and Miguel 2018, Andrews and Kasy 2019). The basis of these discussions is a belief that, if the aim of research is to as-

¹By “negative”, this paper refers to any non-positive results, including both results that are not statistically significant and statistically significant negative.

²This is based on an actual article titled “Anti-depressant unpublished” by Cray 2008.

sist policymakers with their decisions, then all results must be published accurately. However, since researchers are known to omit or inflate many negative results, they must have biased objectives that compromise readers' welfare (or, social welfare in the case of policymaker). When experiments depend financially on funding from industry, researchers may wish to report positive results that support them; when careers depend on publications, researchers may wish to report significant results that surprise journal editors and referees. While the bias may come from a variety of incentive structures, these concerns have led to a number of initiatives to reduce publication bias, ranging from developing registries of experimental protocols to introducing journal policies that de-emphasize statistical significance.³ In addition, most commonly used bias correction methods are based on either of these interpretations and assume that any sufficiently positive estimates or any statistically significant ones will have a higher likelihood of publication than all other estimates.

This paper proposes an alternative approach to publication bias and meta-analyses by re-examining the belief that underlies the above concerns: to maximize readers' welfare, should researchers publish all their results just as they are? When policymakers can process all the information contained in the papers at no cost, then the answer will be *yes*, as accurate reporting always helps readers make more informed decisions. However, when they have cognitive limitations to process the information, the answer will be *no*. Among the various possible frictions such as cost of communicating results, this paper will focus on an *aggregation friction*. Even though researchers know the details of statistical estimates, readers cannot fully process such details and thus only consider the binary (positive or negative) conclusions of each paper. This friction is documented in discussions of meta-analyses and some experiments⁴. The paper develops a model of communication between multiple researchers and one reader under aggregation frictions, and shows that the publication bias will emerge even when research is communicated optimally for readers. Various predictions of publication bias

³AEA RCT registry and clinical trial.gov. In economics, for example, the *American Economic Review* has banned stars in regression tables in 2016, and the *Journal of Development Economics* has promised result-independent reviews in 2018.

⁴While this assumption may appear too simplistic, even aggregation in major policy decisions may rely on such vote-counting. For example, an influential campaign that reached President Obama in 2013 had summarized 12,000 articles without consulting statistical details; it merely wrote "97 percent of climate papers stating a position on human-caused global warming agree that global warming is happening" (The Consensus Project 2014). Some experiments suggest dichotomous interpretation is common even among experts in statistics (McChane and Gal 2017). Due to the cost of processing information, to the limited expertise to understand subtleties, or to the paucity of memory to recall details, readers may only consult binary conclusions of each study to make their decisions.

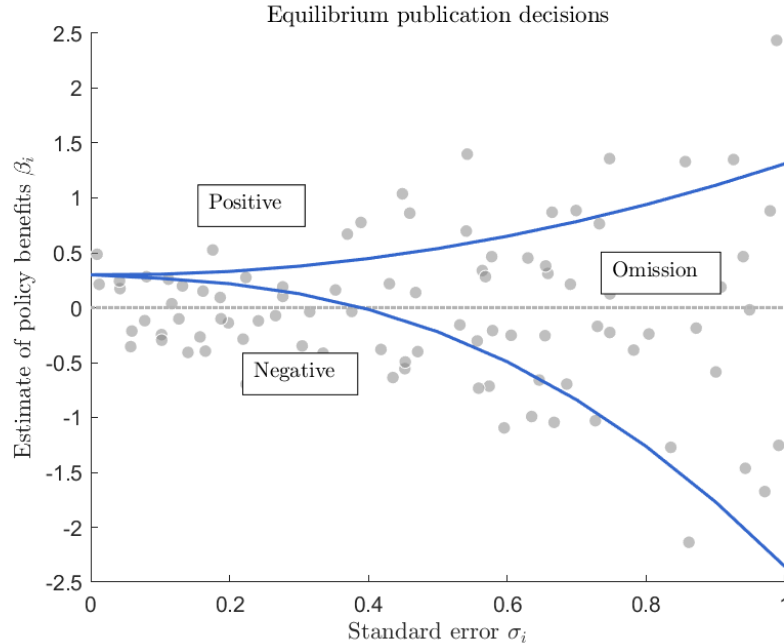


Figure 1-1: Equilibrium publication decision

Notes: Figure 1 is a “funnel plot” where a study i is plotted according to its coefficient estimate β_i on y -axis and its standard error σ_i on x -axis. The studies that have coefficient estimates and standard errors in the top region will be reported as positive; those in the bottom region will be reported as negative; those in the intermediate region will be omitted. Details of simulation environment are given in Section 2.4.2. Each dot is a simulated example rather than real-world data.

implied by this model can be summarized by a figure that plots the studies’ coefficient estimates and standard errors (Figure 1-1). It shows what sets of results will be published as either positive or negative results. We use this implication on conditional publication probabilities to develop a new, non-parametric *stem-based* bias correction method that provides robust meta-analysis estimate relative to existing methods.

The first set of theoretical results is that there will be an omission of non-extreme and imprecise estimates and also a bias towards publishing more positive results than negative ones. To see why omission occurs, note that researchers wish to convey as much information as possible but the aggregation friction permits them to convey only the signs, but not the strengths, of their estimates. In this context, omission of “weak” (i.e. non-extreme or imprecise) results can, in equilibrium, convey more information because when studies are reported, even readers with aggregation frictions can know that those studies are “strong” (i.e. either extreme or precise). To see why the omission will be biased, note that researchers publish results to provide useful

information that will improve the reader's decision, jointly with other researchers. But without knowing what others researchers will report, a researcher can only guess the results of other research when their own reports are consequential. Now, suppose that the reader uses a supermajoritarian rule: that is, she chooses to adopt a policy (e.g. prescribe a new drug) whenever there are strictly more positive than negative results. In this situation, a single positive report will induce the reader to adopt the policy when other studies report an equal number of positive and negative results, while a single negative report will induce her to reject the policy when other studies report one more positive than negative results. Given these asymmetric conditions, in equilibrium, researchers will be more cautious of reporting negative results than positive results, leading to bias by omission. We show that the supermajoritarian rule is optimal for the reader among alternative aggregation rules that the reader could adopt.

The second main theoretical result is that, if journals only publish studies with sufficiently small p -values, then even unbiased researchers may inflate some marginal estimates to turn them statistically significant. To see why such inflation emerges, note that the Bayesian posterior mean of a normal distribution divides each coefficient estimate by its variance. However, null hypothesis testings rely on t -statistics, which divide each coefficient estimate by its standard error. Due to this sharp contrast between the use of variance versus standard errors, neither the researchers nor their readers will wish to strictly be concerned with hypothesis testing in deciding whether or not to adopt the policy. In this way, even though p -values could be a reasonable rule of thumb for publishing results that might appropriately guides the readers' decision, there will be some room for inflation of results to improve reader welfare.

The third, additional result is that, even when the differences in underlying biases are small, the researchers will have large polarization of reporting rules to draw their binary conclusions. This amplification of small biases arises because the reporting decisions are strategic substitutes. That is, if another researcher reports positive results frequently, then a researcher would like to report positive results less frequently to offset the bias of another researcher, and this adjustment induces another researcher to report positive results even more often. This result suggests that the reporting decisions will be highly sensitive to the objectives, making it difficult to model parsimoniously.

The paper presents a range of evidence to suggest that the above three theoretical results are highly relevant in the published literature of both economics in other sciences, and has important implications for the meta-analysis literature. First, under

moderate assumptions, a range of reported evidence suggest all forms of publication bias are prevalent in the social sciences. Second, evidence derived in this paper shows that the communication model presented here can explain the pattern of omission more adequately than other models used in meta-analyses. Specifically, it suggests that noisy null results are likely to be omitted, whereas other models, for tractability, suggest either that any null results or any extremely negative results are likely to be so. Evidence from labor union effects (Doucouliagos et al. 2018), some of the largest meta-analytic data set in economics, shows that the prediction of the communication model presented in this paper holds, and thus, the bias correction methods that assume the parsimonious selection process may have misspecification problems.

Finally, this paper offers a solution for correcting publication bias in meta-analyses of the type described above. The challenge is that publication selection will depend on asymmetric and non-uniform omission, as well as inflation of marginal results, that cannot be parsimoniously modeled. Further, their exact process will depend on economic primitives unobservable to meta-analysts. Nonetheless, even without estimating the publication selection process, we can mitigate the bias by focusing on a subset of data that exhibits less bias under any selection models commonly used. Theory suggests that, across the parameters of communication models and many models, precise estimates are less biased than imprecise ones. Heuristically, more precise studies have more reliable estimates, and thus there is little reason not to publish. The stem-based method proposed in this paper derives the meta-analysis estimate from some n^* most precise studies, where n^* is chosen optimally. A simulation study with realistic parameter values suggests that the estimates from the stem-based method are significantly more robust than other methods (Hedges 1992, Duval and Tweedie 2001) because the stem-based method imposes much weaker assumptions on the publication selection process and underlying distributions than do other studies. Robustness is particularly crucial in meta-analyses, where the goal is to build a consensus among often contradictory results in the literature.

Related Literature. This paper relates to two sets of microeconomic theory literature on communication. It also relates to the statistics literature on bias correction methods.

First, this paper builds on and derives contrasting results from the canonical models of information aggregation and transmission. The results of (i) omission, (ii) bias, and (iii) amplification of small bias relate to the models of voting as information aggregation (Austen-Smith and Banks 1996, Feddersen and Pesendorfer 1996, 1997, 1998, 1999): (i) omission of inconclusive results is analogous to abstention of uninformed

voters; (ii) biased reporting is similar to jurors' bias to convict in order to counter-balance unanimity rules, but differs in that such biased reporting is socially optimal; (iii) amplification of small bias due to strategic substitution between researchers is an extension of the result of the non-partisan voters' vote against the bias of partisan voters. Moreover, the result that coarseness of message⁵ space leads to biased reporting stands in contrast with the models of information transmission (Crawford and Sobel 1982, Li, Rosen, and Suen 2001): whereas bias of senders leads to coarse messages due to the incentive constraints in their model, technological restriction of coarse message space leads to biased reporting due to pivotality conditions in this model.

Second, this paper also relates to the growing literature of microeconomic models of scientific communication⁶ and, specifically, publication bias. The result of *t*-statistics relates to the microeconomic decision models of statistical testing (Manski 2004, Tetenov 2016). The broad conclusion that publication bias needs not be socially detrimental is consistent with papers with various reasoning, such as incentives for endogenous information acquisition with biased researchers (Glaeser 2006, Libgober 2015, Henry and Ottaviani 2014), or limited number of studies readers may process (de Winter and Happee 2012) or journal space for publication (Frankel and Kasy 2018). In contrast, this paper derives the result even when information is exogenously given, even when researchers are unbiased, and even when there is no limit or cost of communication. This paper instead focuses the cognitive friction, which some papers (Suen 2004, Fryer and Jackson 2008, Blume and Board 2013) have shown as a possible reason for biased communication and decisions, and applies to the context of scientific communication.

Finally, this paper also contributes to the large and growing (Simonsohn et al. 2014, Bom and Rachinger 2018, Andrews and Kasy 2019) literature on correction methods for publication bias in meta-analyses. In contrast with the most commonly used methods (Hedges 1992, Duval and Tweedie 2000) that rely on specific assumptions on publication selection process and underlying distribution, the stem-based method depends on assumptions that hold across various assumptions made in lit-

⁵Coarseness is also a key friction in papers such as Dewan et al. (2015). The difference is that signals are more coarse than states in their paper whereas messages are more coarse than signals in this paper.

⁶In a broad way, the paper contributes to the empirical analyses of communication models that have been advanced specifically in the research on media in the real world (Gentzkow et al. 2016, Puglisi and Snyder 2016) or on hypothetical communication settings in the laboratories (Crawford 1998, Battaglini et al. 2010). This paper advances these empirical studies by obtaining a direct measure of bias in the real world data.

erature. The method extends approaches to focus on some arbitrary number of the most precise studies (Barth et al. 2013, Stanley et al. 2010) by providing a formal criteria and estimation methods to select the optimal number of studies to focus on.

The remainder of the paper is organized as follows: Section 2 presents the communication model and its related evidence; Section 3 develops and implements its empirical test; Section 4 proposes a bias correction method given this observation; and Section 5 concludes.

1.2 A Communication Model of Publication Bias

This Section presents a communication model between multiple researchers and a policy maker with the friction that the researchers can only communicate yes-or-no conclusions even though they are informed of their experiments' estimates. The analysis will show that the aggregation friction can provide explanation for various kinds of publication bias, such as (i) omission of insignificant results, (ii) inflation of marginally insignificant results, and (iii) amplification of small researchers' bias. Results on omission will also provide an empirical prediction on the meta-analysis data sets.

1.2.1 Set-up

The set-up is based on a static communication model between N senders, called researchers, and 1 receiver, called a policymaker. Researchers receive private unbiased signals of the treatment effect of policy, β_i , and its standard error, σ_i , and report their results, m_i . Given the reports from the researchers, the policymaker decides whether to implement the policy $a \in \{0, 1\}$.

The model's key element is an *aggregation friction*: even though researcher's private signals, $\{\beta_i, \sigma_i\}$, take continuous values, researchers can only convey a positive result, $m_i = 1$, or a negative result, $m_i = 0$, or do not report their study, $m_i = \emptyset$. Given the standard error independently drawn from some distribution $\sigma_i \sim G(\sigma)$, the treatment effect estimate has a normal distribution around the true benefit, b , so that $\beta_i \sim \mathcal{N}(b, \sigma_i^2)$. However, the policymaker only considers what results the researchers have reported to make their decision. Henceforth, let us denote the number of positive results by $n_1 = \sum_{i=1}^N \mathbf{1}(m_i = 1)$, and negative results by $n_0 = \sum_{i=1}^N \mathbf{1}(m_i = 0)$.

Both the researchers and the policymaker maximize the social welfare:

$$a(\mathbb{E}b - c), \tag{1.1}$$

where c is the cost of policy implementation.⁷ They have a common prior $b \sim \mathcal{N}(0, \sigma_b^2)$. The number of players, their pay-offs, priors, and the cost of policy implementation c are public information and common knowledge.

The timeline is as follows: first, researchers⁸ receive their own signals, and simultaneously decide whether to publish their binary conclusions. Then, the policymaker sees the reports and makes her policy decision using her posterior belief. Finally, the payoffs are realized.

1.2.2 Analyses

The analysis will focus on Perfect Bayesian Nash Equilibria (PBNE), the standard equilibrium concept in communication models. The strategy of a researcher i , $s_i \in S_i$, is a mapping from his own signal $\{\beta_i, \sigma_i\}$ to probability distribution over his message m_i $s_i : \mathbb{R} \times \mathbb{R}_+ \mapsto \Delta\{1, \emptyset, 0\}$. The strategy of policymaker, π , is a mapping from the messages to the probability distribution over binary policy action $a \in \{0, 1\}$: that is, $\pi : \{1, \emptyset, 0\}^N \mapsto \Delta\{0, 1\}$. The policymaker's belief over the researchers' strategy is denoted by $\mu \in \Delta\left(\{S_i\}_{i=1, \dots, N}\right)$.

Definition 1.1 : *An equilibrium is a tuple of strategies and belief, $\{s_1, \dots, s_N, \pi, \mu\}$ such that (i) researcher i ' strategy maximizes (3.1) given strategies of all other researchers and the policymaker, for all i ; (ii) policymaker's strategy maximizes (3.1) given strategies of researchers and belief; (iii) the policymaker's belief is consistent with Bayes' rule.*

As communication models always have multiple equilibria, including a babbling equilibrium, we introduce following criteria to focus on non-trivial and plausible equilibria:

⁷While this set-up may appear to assume no uncertainty in welfare when the policy is not implemented since it is fixed to be 0, it also represents such setting: suppose the welfare under policy is $u_1 \sim \mathcal{N}(\bar{u}_1, \sigma_{u_1}^2)$ and the welfare in the absence of policy is $u_0 \sim \mathcal{N}(\bar{u}_0, \sigma_{u_0}^2)$. Then, it is optimal to implement the policy if and only if $\bar{u}_1 - \bar{u}_0 \geq c$. Thus, defining $b \equiv \bar{u}_1 - \bar{u}_0$ and $\sigma_b^2 \equiv \sigma_{u_1}^2 + \sigma_{u_0}^2$ above will be an equivalent set-up.

⁸I suggest readers of this paper to think of researchers in this model not as individual authors in the real world, but as a collection of authors, referees, and editors who jointly make the publication decisions. In peer reviewed journals, individual researchers play both roles of authors and referees. Even if researchers may seek to maximize publications when they are authors, journals ask referees and editors to publish socially valuable information.

Definition 1.2 : An equilibrium is fully responsive if $\forall m_i, m'_i \exists m_{-i}$ such that $\pi(m_i, m_{-i}) \neq \pi(m'_i, m_{-i})$; and fully informative if $\mathbb{E}[b|m]$ is not constant analogously.

In any fully responsive and fully informative equilibria⁹, both the policymaker's and researchers' strategy will be monotone in their information, at least in a benchmark set-up¹⁰:

Lemma 1 (Monotonicity of equilibrium strategies). For $N=2$, for any c , and constant σ_i , any fully responsive and fully informative equilibrium has strategies that are monotone:

- (i) the policymaker's decision $\pi^*(m)$ is increasing in number of positive results (n_1) and decreasing in number of negative results (n_0) in the following sense. For every i

$$\begin{aligned} \pi^*(m_i, m_{-i}) > 0 &\Rightarrow [\forall m'_i \text{ s.t. } m'_i \succ m_i, \pi^*(m'_i, m_{-i}) = 1] \\ \pi^*(m_i, m_{-i}) < 1 &\Rightarrow [\forall m'_i \text{ s.t. } m'_i \prec m_i, \pi^*(m'_i, m_{-i}) = 0] \end{aligned}$$

, where messages are ordered by $1 \succ \emptyset \succ 0$ without loss of generality.

- (ii) each researcher i takes the threshold strategy: there exist $\underline{\beta}_i, \bar{\beta}_i \in (-\infty, \infty)$ such that

$$m_i^* = \begin{cases} 1 & \beta_i \geq \bar{\beta}_i \\ \emptyset & \text{if } \beta_i \in [\underline{\beta}_i, \bar{\beta}_i) \\ 0 & \beta_i < \underline{\beta}_i \end{cases}$$

Sketch of proof. By Bayes' rule, combined with the law of iterated expectation and monotonicity of the mean of a truncated normal distribution with respect to the mean of the underlying distribution. Appendix A1.3 contains a full proof. \square

Finally, to assess desirability and plausibility of particular equilibrium among non-trivial equilibria, we introduce the following definitions:

Definition 1.3 : An equilibrium is optimal if no other sets of strategies can attain strictly higher ex ante welfare; and locally stable if, for every neighborhood of the equilibrium, there exists a sub-neighborhood of the equilibrium from which a myopic and iterative adjustment of strategies stays

⁹Note that this definition differs from usual responsiveness and informativeness in that it requires all messages to be influential and informative. For example, the policy rule $a^* = 1 \Leftrightarrow n_1 = 2$ is not fully responsive in that the choice between $m = 0, \emptyset$ never alters the decision.

¹⁰As shown in Appendix A1.3, the monotonicity requires a constant standard error since the monotone likelihood ratio property needs not hold under normal distribution with unknown standard errors.

in that neighborhood. Formally, a myopic and iterative adjustment on the tuple $\mathcal{E}_t \equiv \{\pi_t(m), \bar{\beta}_{1,t}, \underline{\beta}_{1,t}, \bar{\beta}_{2,t}, \underline{\beta}_{2,t}\}$, is a dynamic process for $t = 1, 2, \dots$ in which, in each t , (i) the policymaker plays best response to $\{\bar{\beta}_{1,t-1}, \underline{\beta}_{1,t-1}, \bar{\beta}_{2,t-1}, \underline{\beta}_{2,t-1}\}$, (ii) researcher 1 plays best response to $\{\pi_t(m), \bar{\beta}_{2,t-1}, \underline{\beta}_{2,t-1}\}$, and (iii) researcher 2 does so to $\{\pi_t(m), \bar{\beta}_{1,t}, \underline{\beta}_{1,t}\}$. An equilibrium \mathcal{E} is locally stable if for every $\hat{d} > 0$, there exists $\bar{d} > 0$ such that, given any disturbance ϵ such that sup metric $d(\epsilon) < \bar{d}$, $d(\mathcal{E} - \mathcal{E}_\infty) < \hat{d}$; that is, the equilibrium stays in the neighborhood of the equilibrium.

It is standard to focus on the most informative equilibrium in sender-receiver games, and on the optimal equilibrium in common interest games. Local stability assures robustness to small deviations from the equilibrium strategies.

Henceforth, the analysis will combine analytical and numerical approaches to show that the main mechanisms play important roles in plausible settings that satisfy the above criteria. Analytical results will illustrate the logic behind the kinds of publication bias that arises from the model by focusing on a tractable environment and equilibrium with various symmetry properties. Analytical results will focus primarily on the case of $N = 2$,¹¹ $c = 0$, and often constant σ . Then, numerical results will show that the mechanism will play important roles in asymmetric environment that is more plausible yet analytically difficult to solve.

1.2.3 Main Result 1. Omission of Insignificant Results

The first main result is that, in the optimal and locally stable equilibrium, there will be an asymmetric omission of intermediate results such that the average estimates underlying published studies will have an upward bias.

Analytical Results

The following propositions will first show, in a symmetric environment with constant σ , there will be equilibria with asymmetric omission, no omission, and symmetric omission; and, second show that the equilibrium with asymmetric omission is both

¹¹Given that meta-analyses often include more than 2 studies, this set-up with $N = 2$ may appear restrictive. However, this simple setting is common in the committee decision-making literature to which this paper is closely related, such as Gilligan and Krehbiel 1989, Austen-Smith 1993, Krishna and Morgan 2001, and Hao, Rosen, and Suen 2001. Analytical characterization for $N \geq 3$ is difficult due to technical challenge of analytically evaluating multivariate normal's truncated mean. Instead, this paper takes numerical approach to settings with $N \geq 3$.

optimal and locally stable whereas other two kinds of equilibria are not. The relation to the information aggregation models in voting theory will be discussed.

Proposition 1.1 (Equilibrium with asymmetric omission). *Let $N = 2$, $c = 0$, and $\sigma_i = \sigma$. There exists an equilibrium with the following strategies. The policymaker's strategy is a supermajoritarian policy decision rule:*

$$\pi^* = \begin{cases} 1 & \text{if } n_1 > n_0 \\ 0 & \text{if } n_1 \leq n_0 \end{cases} \quad (1.2)$$

The researchers' strategies are identical to each other and characterized by the unique thresholds, $\bar{\beta}, \underline{\beta}$, that satisfy

$$\bar{\beta} > 0 > \underline{\beta} \text{ and } \bar{\beta} < -\underline{\beta} \quad (1.3)$$

so that there will be an upward bias in the estimates of the reported studies: $\mathbb{E}[\beta_i | m_i \neq \emptyset] > 0$.

Sketch of Proof: By a combination of information asymmetry among researchers and a message space that is coarser than a signal space. Suppose the policymaker adopts (1.2). Since a researcher does not know what another researcher will observe and report, he conditions his reporting decision on the state in which his own report will be pivotal and swing the policy decision in the direction of its conclusion. A positive result switches the policymaker from canceling to implementing the policy only when another researcher did not report his result as his signal had an intermediate value. Thus, the optimal threshold $\bar{\beta}$ satisfies the indifference condition, $\beta_i + \beta_{-i} = 0$, in expectation conditional on pivotality:

$$\bar{\beta} + \mathbb{E}[\beta_{-i} | \bar{\beta} > \beta_{-i} \geq \underline{\beta}, \beta_i = \bar{\beta}] = 0 \quad (1.4)$$

In contrast, a negative result leads the policymaker to cancel rather than to implement the policy only when his report is positive. Thus, the optimal threshold $\underline{\beta}$ satisfies:

$$\underline{\beta} + \mathbb{E}[\beta_{-i} | \beta_{-i} \geq \bar{\beta}, \beta_i = \underline{\beta}] = 0 \quad (1.5)$$

Because the binary conclusion cannot convey the strength and can only communicate the sign of the continuous signal each researcher receives, omission of intermediate results can better convey information than always reporting either positive or negative results. In addition, comparing the conditions (1.4) and (1.5), researchers are more cautious of reporting negative results than positive results and thus, given that results

are reported, the coefficients are on average upward biased: $\mathbb{E}[\beta_i|m_i \neq \emptyset] > \mathbb{E}[\beta_i]$. Finally, it is strictly optimal for the policymaker to follow the supermajoritarian rule (1.2) since he focuses on the average value conditional on $\{m_i\}$ whereas the researcher focused on the marginal value. Appendix A2.1 contains a formal proof. \square

There will also be an equilibrium with asymmetric omission that generates a downward omission when $c = 0$. However, as Appendix Figure B3 shows, the equilibrium in Proposition 1.1 attains a strictly higher welfare when $c \geq 0$. In this sense, the “positive” results are defined as the messages that alter the default decisions whereas “negative” results are the ones that suggest to maintain the default. The following propositions will now show that there are also equilibria without bias of underlying estimates:

Proposition 1.2 (Equilibria with symmetric or no omission). *Let $N = 2$ and $c = 0$, and $\sigma_i = \sigma$. There exist equilibria with the following strategies:*

- *Symmetric omission: the policymaker follows*

$$\pi^* = \begin{cases} 1 & n_1 > n_0 \\ \frac{1}{2} & \text{if } n_1 = n_0 \\ 0 & n_1 < n_0 \end{cases} \quad (1.6)$$

and the researchers' thresholds satisfy $\bar{\beta} > 0 > \underline{\beta}$ and $\underline{\beta} = -\bar{\beta}$ so that $\mathbb{E}[\beta_i|m_i \neq \emptyset] = 0$.

- *No omission: the policymaker follows*

$$\pi^* = \begin{cases} 1 & n_1 = 2 \\ \tilde{\pi} & \text{if } n_1 = 1, n_0 = 0 \\ 0 & n_1 \leq n_0, \end{cases} \quad (1.7)$$

where $\tilde{\pi} \in (0, 1]$ and the researchers' thresholds satisfy $\underline{\beta} = \bar{\beta} < 0$ so that $\mathbb{E}[\beta_i|m_i \neq \emptyset] = 0$.

Sketch of Proof: The equilibrium with symmetric omission exists because (i) the indifference conditions of researchers that determine $\bar{\beta}, \underline{\beta}$ will be symmetric to each other when $\pi = \frac{1}{2}$ when $n_1 = n_0$, and (ii) given symmetric thresholds, the policymaker will be indifferent between $a = 0, 1$ when $n_1 = n_0$. The equilibria with no omission exist because (i) the welfare gain from omission exists only when another researcher omits, and (ii) given that researchers always report $m_i = 0, 1$, the decision when there are omissions may be defined arbitrarily. \square

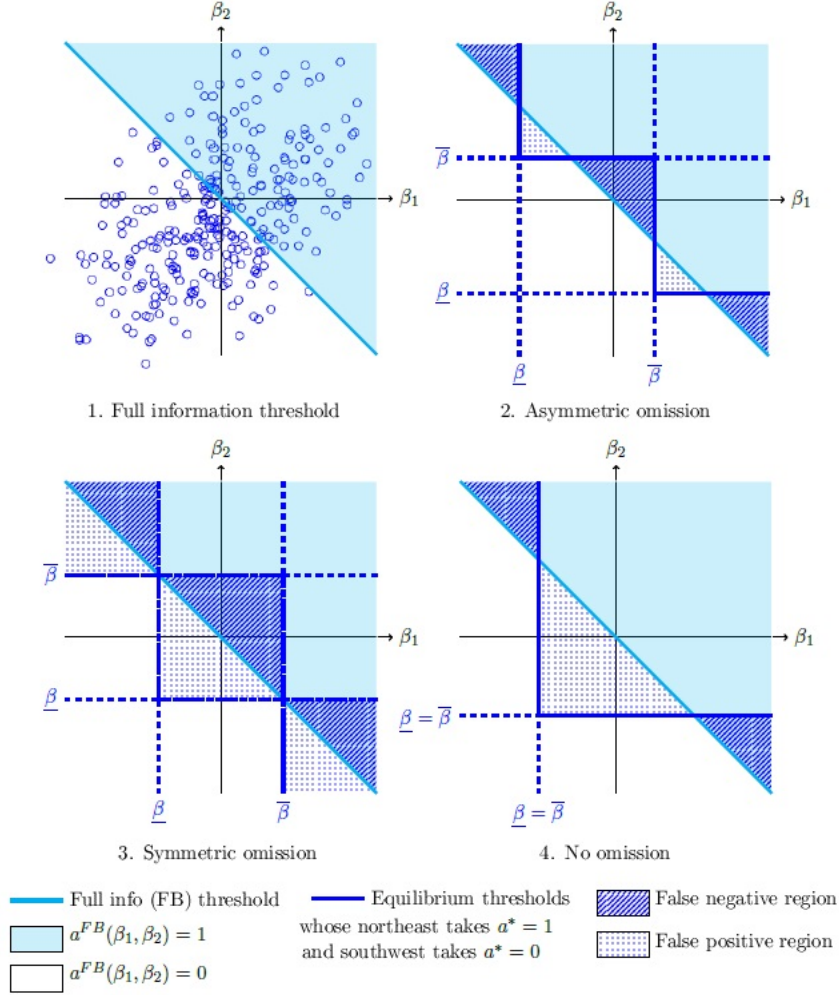


Figure 1-2: Equilibrium thresholds and policy decisions for $N = 2$, $c = 0$

Notes: Panel 1 plots the benchmark first-best policy implementation rule ($\beta_1 + \beta_2 \geq 0$) as given by the Bayes' rule and an example of the bivariate normal distribution of signal realizations $\{\beta_1, \beta_2\}$. Panel 2, 3, and 4 illustrate the equilibrium thresholds and policy decisions under three responsive and informative equilibria. The dotted line shows the thresholds for each equilibrium, and the policy is implemented if $\{\beta_1, \beta_2\}$ were in the region northwest to the solid line for Panel 1 and 2. For Panel 3, policy is implemented with $\frac{1}{2}$ probability in region surrounded by the dotted line. “False negative” (“False positive”) regions denote signal realizations such that the policy is (is not) implemented under the full information but is not (is) implemented in equilibrium. The figures' origin is $\{0, 0\}$.

While Proposition 1.2 shows that there are also fully informative and fully re-

sponsive equilibria without bias under reported studies, the following Proposition 1.3 shows that the equilibrium with bias of reported studies is more desirable and likely:

Proposition 1.3 (Optimality and local stability of equilibria). *The equilibrium with asymmetric omission as characterized in Proposition 1.1 is optimal and locally stable; the equilibria characterized in Proposition 1.2 are neither optimal nor locally stable.*

Sketch of Proof: The heuristic reasons for optimality and local stability are summarized by Figure 1-2. The equilibrium with asymmetric omission is optimal because its policy implementation rule as in (1.2) described in Panel 2 most closely approximates the first best threshold of $\beta_1 + \beta_2 = 0$ as depicted in Panel 1; it minimizes the probabilities of false positive and false negative errors that leads to welfare losses. It is also locally stable since the policymaker's decision is based on strict preference and the researchers' strategies are only moderate substitutes of one another. On the other hand, the equilibria with symmetric or no omission are neither optimal nor locally stable as small perturbation of researchers' thresholds and policymaker's strategy can improve the welfare and its subsequent iterative adjustment leads to a different equilibrium. For example, in the symmetric equilibrium (Panel 3), consider a small decrease in researchers' strategy, $\bar{\beta}' = \bar{\beta} - \Delta$ and $\underline{\beta}' = \underline{\beta} - \Delta$. This perturbation leads to a first order welfare increase because it increases $\mathbb{E}[b|n_1 > n_0]$ by quantity proportional to Δ but only has a second order welfare loss, and thus increases the total welfare. The symmetric equilibrium is also not locally stable because the policymaker now prefers to implement the supermajoritarian rule. Analogous argument applies to the equilibrium with no omission; and Appendix A2.3 contains a complete proof. \square

The concepts of optimality and local stability are closely related to each other because the model is a common interest game: if an equilibrium is not locally stable, then it cannot be optimal since every steps of iterative adjustment must be intended to improve welfare.

These results relate to the models of information aggregation and transmission. First, it builds on the models of voting as information aggregation (Austen-Smith and Banks 1996, Feddersen and Pesendorfer 1996, 1997, 1998, 1999) as the result regarding omission is analogous to the result that uninformed and unbiased voters abstain when there are other informed and unbiased voters (Feddersen and Pesendorfer 1996). The result regarding biased reporting echoes the result that unanimity rule, counter-intuitively, may increase the probability of false conviction if the jurors condition their votes on the states in which their votes are pivotal (Feddersen and Pesendorfer 1998).

The novel result of this paper is that the biased reporting is socially optimal whereas it was argued to be sub-optimal in their voting theory. This difference is due to the information coarsening: while these voting models have often assumed *binary* states, this paper assumes *continuous* states even though the messages can only be yes, no, or abstention.

Second, this paper relates to communication models that show that biases of senders result in coarse messages, both with one sender (Crawford and Sobel 1982) and multiple senders (Hao, Rosen, and Suen 2001). In contrast, this paper shows that the technological restriction of coarseness¹² leads to the biased reporting. In this sense, the relationship between conflict of interest and coarseness of information transmission may have causalities running in both ways.

Numerical Results

There are two results from the numerical simulation that are critical for understanding the asymmetric equilibrium characterized in Proposition 1.1. First, the probability of policy implementation is less than that in the environment where the estimates can be directly communicated: $\mathbb{P}(a = 1) \leq \frac{1}{2}$ (Appendix Figure B2). In this sense, the upward bias among the reported estimates is a way to mitigate the inherent conservativeness in supermajoritarian rule. Second, when $c > 0$, the welfare under the equilibrium with supermajoritarian rule, $\pi = 1 (n_1 > n_0)$, is higher than that with submajoritarian rule, $\pi = 1 (n_1 \geq n_0)$ (Appendix Figure B3). Therefore, while the equilibrium with $\pi = 1 (n_1 \geq n_0)$ also exists, this paper focuses on the case with $\pi = 1 (n_1 > n_0)$.

Evidence

A number of studies suggest that omission is prevalent by reporting a positive correlation between the coefficient magnitude and the standard error; on average, imprecisely estimated studies have higher coefficient values than precise studies¹³. In economics,

¹²This assumption is similar to some papers that examined the implication of communication frictions on biases, such as Suen 2004, Fryer and Jackson 2008, and Blume and Board 2013.

¹³This could be due to researchers omitting studies unless they are positive statistically significant. If the study is imprecise, then a large coefficient magnitude is needed whereas if the study is precise, then coefficient magnitude can be modest (Hedges 1992). Alternatively, this could also be due to researchers omitting studies when the coefficient values are low (Duval and Tweedie 2000, Copas and Li 1997). There are two formal tests that examine the presence of publication bias through examining the correlation between coefficient magnitude and study precision: an ordinal test that examines their rank correlations (Begg and Mezuemder 1994) and a cardinal test that examines the correlation by regression (Egger et al. 1997). Second, there is occasionally excess variance in the

important estimates such as the impact of minimum wage on employment (Card and Krueger, 1995), the return to schooling (Ashenfelter et al., 1999), and the intertemporal elasticity of substitution (Havránek, 2015), have evidence of a positive correlation. In environmental studies, estimates of the social cost of carbon, a key statistic for carbon tax policy, were shown to have the bias (Havránek et al., 2015). Moreover, the probability of omission around 30 percent is roughly consistent with some examples reported in Andrews and Kasy 2019.

1.2.4 Main Result 2. Inflation of Marginally Insignificant Results

The second main result is that, given heterogeneous standard errors, the equilibrium t -statistic threshold will not be constant across σ_i . This result has two empirical implications: (i) if journals apply a t -statistic threshold to publish positive results, then even unbiased researchers will inflate some marginally insignificant results; and (ii) there will be omission of imprecisely estimated null results.

Analytical Results

Analytical results show that the absolute value of t -statistics¹⁴ threshold will be increasing in σ_i . While the model set-up does not impose restrictions on the messages based on t -statistics, we can define the analogous t -statistics naturally in terms of the equilibrium threshold:

Definition 2 t -statistics: *Define the t -statistics thresholds, $\bar{t}_i(\sigma_i)$ and $\underline{t}_i(\sigma_i)$, as the ratio between the threshold coefficient and standard error: $\bar{t}_i(\sigma_i) = \frac{\bar{\beta}(\sigma_i)}{\sigma_i}$, $\underline{t}_i(\sigma_i) = \frac{\underline{\beta}(\sigma_i)}{\sigma_i}$.*

The following proposition claims that, in a unique symmetric equilibrium such that the two researchers apply the identical thresholds $\bar{t}(\sigma_i) = -\underline{t}(\sigma_i)$, the t -statistics thresholds will be increasing in σ_i so that precise studies will be more likely to be published than imprecise ones.

Proposition 1.2. (t -statistic threshold increasing in σ) *Suppose $N = 2$, $c = 0$, $\sigma_b = \infty$, and $\text{Supp}(G)$ is some interval in \mathbb{R}_{++} . There exists a unique symmetric*

estimates with an abundance of studies at the extreme values beyond significance thresholds and a scarcity of studies with intermediate coefficients (Stanley 2005).

¹⁴While it is formally a z -statistics since standard error is assumed to be known in the model, the paper uses the term t -statistics to be coherent with the way empirical studies usually conduct null hypothesis testing.

equilibrium such that the policymaker follows the mixed strategy π^* in (1.6), and the researchers adopt threshold strategies with cut-offs that depend on σ_i , as in Lemma 1.1. Then the t -statistics will be symmetric so that, $\bar{t}_i(\sigma_i) = -\underline{t}_i(\sigma_i) \equiv t(\sigma_i)$ for both $i = 1, 2$, and will be increasing in σ_i :

$$\frac{\partial t(\sigma_i)}{\partial \sigma_i} > 0$$

for every $\sigma_i \in \text{Supp}(G)$ for both researchers.

Sketch of Proof: By rearranging the indifference condition of researchers. By the Bayes' rule and law of iterated expectation, the researcher i 's indifference condition on $\bar{\beta}_i(\sigma_i)$ is

$$\mathbb{E} \left[\frac{\frac{\bar{\beta}_i(\sigma_i)}{\sigma_i^2} + \frac{\beta_j}{\sigma_j^2}}{\frac{1}{\sigma_i^2} + \frac{1}{\sigma_j^2}} \mid I_i, I_j \right] = 0, \quad (1.8)$$

where $I_i = \{\beta_i = \bar{\beta}_i(\sigma_i), \sigma_i\}$ is the information set of researcher i , and $I_j = \{\beta_j \in \text{Piv}(\sigma_j, \pi), \sigma_j\}$ is the information set of researcher j , where $\text{Piv}(\sigma_j, \pi)$ is the pivotality condition, and the expectation is taken over I_j . Reorganizing this condition (1.8), the threshold $\bar{t}(\sigma_i)$ must satisfy

$$\bar{t}(\sigma_i) = \frac{\bar{\beta}(\sigma_i)}{\sigma_i} = \sigma_i \frac{\mathbb{E} \left[-\frac{\beta_j}{\sigma_j^2 + \sigma_i^2} \mid I_i, I_j \right]}{\mathbb{E} \left[\frac{\sigma_j^2}{\sigma_j^2 + \sigma_i^2} \mid I_i, I_j \right]} \quad (1.9)$$

In this way, the t -statistics threshold is increasing in σ_i since $\mathbb{E} \left[-\frac{\beta_j}{\sigma_j^2 + \sigma_i^2} \mid I_i, I_j \right]$ is positive in equilibrium. Appendix A2.2 contains a complete proof, which focuses on the symmetric equilibrium that provides a tractable environment where the term, $\mathbb{E} \left[-\frac{\beta_j}{\sigma_j^2 + \sigma_i^2} \mid I_i, I_j \right] \times \mathbb{E} \left[\frac{\sigma_j^2}{\sigma_j^2 + \sigma_i^2} \mid I_i, I_j \right]^{-1}$, does not change substantively enough to alter this sign. Analogous results hold for $\underline{t}(\sigma_i)$. \square

The impossibility of equating the optimal thresholds with a constant t -statistics threshold arises from the difference in the use of standard errors between the Bayesian updating and the null hypothesis testing. In Bayesian updating, the coefficient is divided by the variance, σ^2 , since the weights on each coefficient must be proportional to its information that increases at rate n in the absence of study-specific effects ($\sigma_0 = 0$). In null hypothesis testing, the coefficient is divided by the standard errors, σ , since t -statistics normalize the convergence of distribution of β_i that occurs at rate \sqrt{n} . In this model, the thresholds are determined by approximating the Bayes rule, which stands in contrast with the focus on the p -value that measures how unlikely that a given observation occurs under the null hypothesis of zero effect.

This observation, while highlighting the contrast between t -statistics and optimal thresholds, renders support for t -statistics as a rule of thumb since the threshold $\bar{\beta}(\sigma)$ is increasing and $\underline{\beta}(\sigma)$ is decreasing. This result relates to a decision-theoretic and statistics literature that examines the optimality of null hypothesis testing as criteria for choosing alternative treatments (Manski 2004). A recent paper (Tetenov 2016) rationalized the t -statistics approach based on communication in the settings with commonly known value of standard error. This model considers the environment where the standard error is a private information, and shows that the equilibrium t -statistics may not be constant across σ_i .

Numerical Results

The analytical results have shown that, in the symmetric equilibrium with $N = 2$, $c = 0$, $\sigma_b = \infty$, the constant t -statistics threshold will be sub-optimal. While tractable, symmetric equilibria will not be locally stable and thus less plausible than the asymmetric equilibrium analogous to that characterized in Proposition 1.1. The numerical analysis henceforth will show that the key qualitative predictions will hold even under asymmetric equilibrium and even with $N \geq 3$, $c \geq 0$, $\sigma_b < \infty$. Moreover, it quantifies the bias, omission, and welfare gains from omission and inflation, and derives empirical predictions.

The equilibrium thresholds, $\bar{\beta}(\sigma)$, $\underline{\beta}(\sigma)$, of the asymmetric equilibrium in a plausible environment (Figure 1-3) show that the two analytical results hold in a more general set-up. First, $\bar{\beta}(\sigma)$, the threshold between sending positive message or not reporting the study, is strictly convex in σ . Thus, if academic communities impose a rule-of-thumb t -statics level to claim positive results, there could be some studies in the shaded region that might still be considered as a “positive” result even though it is marginally insignificant. Second, the omission occurs most importantly among the imprecisely estimated studies with intermediate coefficients. That is, when negative results are reported, they will be either precisely estimated null results or imprecisely estimated and yet extremely negative results. Since precisely estimated studies will be less subject to asymmetric omission, there will also be less bias among them.

A numerical simulation, presented in Appendix B1.5, shows a substantive welfare gain from allowing some inflation and frequency of omission can be substantive across a range of parameter values. Imposing a constant one-sided t -statistic, $\bar{\beta}(\sigma) = t\sigma$, with no omission, will lead to 3 ~ 50 percent of welfare loss relative to the environment

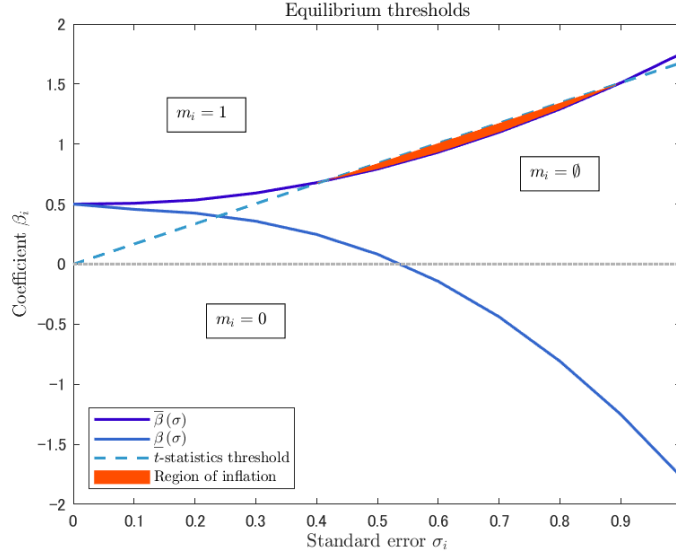


Figure 1-3: $\bar{\beta}(\sigma)$ and $\underline{\beta}(\sigma)$ thresholds

Notes: Figure 1-3 plots the $\bar{\beta}(\sigma)$, $\underline{\beta}(\sigma)$ thresholds under the prior standard deviation $\sigma_b = 0.7$, no study-specific effect $\sigma_0 = 0$, and distribution of standard errors such that, σ_i , that approximates the empirical data as shown in B1.1, and policy cost $c = \frac{1}{2}$. The darker solid line is $\bar{\beta}(\sigma)$, the lighter solid line is $\underline{\beta}(\sigma)$, and the dashed line represents a linear t -statistic threshold. Studies in the shaded region draw positive conclusions even though they are marginally statistically insignificant. The Figure 1-1 in Introduction shows the equilibrium thresholds in identical environment except $c = 0.3$.

in which estimates can be directly communicated, even when $t \in \mathbb{R}$ is chosen to minimize the welfare loss. Allowing for flexible equilibrium threshold can more than halve the welfare losses, leading to only 1 ~ 23 percent of welfare loss. The omission probability is roughly 7 percent among the most precisely estimated studies whereas it could be as large as 60 percent among the least precisely estimated studies. On average, omission probability is around 30 ~ 50 percents. Similarly, the bias is minimal and 0 ~ 20 percent of the underlying benefit distribution (σ_b) among the most precise studies whereas it could be very large among the least precise studies.

Numerical exploration also shows that the comparative static of thresholds with respect to N is ambiguous, and that the threshold for reporting negative results, $\underline{\beta}(\sigma)$, could be increasing in σ when c is high. Appendix B1.1 describes the details of the simulation set-up and procedure, and Appendices B1.2 and B1.3 contains a thorough discussion of these observations.

Evidence

There are various quantitative evidence of inflation and some qualitative evidence of omission that is heterogeneous across values of study precision.

Inflation: When the originally intended specification has a marginally insignificant t -statistic, researchers may inflate the statistical significance through the choice of specifications for outcome, control variables, and samples (Leamer 1978). If inflating t -statistic incurs search costs, then there will be an excess mass right above the threshold. In economics, Brodeur et al. (2016) argues that about 8% of results as statistically significant may be due to inflation. There is an excess mass right about the significance cut-off in sociology and political science, too (Gerber and Malhotra 2008a, 2008b). In psychology with lab experiments such that sample size can be adjusted subsequently, Simonsohn et al. (2014) reported density of p -values among the statistically significant tests were increasing in p -values and interpreted this as evidence of inflation.

Omission heterogeneity across study precision: There are some examples in which either precisely estimated null results and extremely negative results, while imprecisely estimated, are published. Some examples of precise null results¹⁵ include the large-scale clean cookstove study (Hanna et al. 2016), the air pollution regulation in Mexico (Davis 2008), and the community-based development programs (Casey et al. 2012). The examples of extreme negative results include the negative labor supply elasticities close to -1 among the taxi driver papers (Camerer 1997); the positive impact of inequality on economic growth (Forbes, 2000); the unexpected harmful effect of a therapeutic strategy on the cardiovascular events found in one trial (the Action to Control Cardiovascular Risk in Diabetes trial 2008). While these are only some examples, the empirical analyses in Section 3 will provide a formal evidence.

1.2.5 Additional Result. Amplification of Small Bias of a Researcher

The main results have shown that, even when researchers are completely unbiased, there will still be publication bias with omission and inflation. Nonetheless, in the real world, researchers' and policymakers' objectives are not completely aligned with

¹⁵There is one apparent counterexample: DEVTA study, the largest randomized trial that showed null effects of deworming and vitamin A supplementation on child mortality and health, was not published until 7 years after the data collection (Garner, 2013). While the delay required by the careful analysis of authors is extensive, that it was published in *the Lancet*, a top medical journal, is, in a way, reassuring of the academic journal's willingness to report precise negative results.

one another due to different information and interests regarding policies. This Section shows that there will be a large polarization of reporting rules among researchers even when a researcher's bias is small.

Analytical Results

Let us begin by introducing the *strategic multiplier* between researchers that quantifies on the strategic interdependence between them, keeping the policymaker's strategy fixed, given the researcher i 's objective, $a(\mathbb{E}b - c + d_i)$, so that d_i is his bias for policy implementation.

Definition 3 Strategic multiplier between researchers: *Define the strategic multipliers, $\bar{\zeta}, \underline{\zeta}$, as the ratio of the effect of small bias d_i of one researcher on the difference between thresholds of two researchers, between the environment with or without strategic effects, keeping the policymakers' strategy π^* fixed:*

$$\bar{\zeta} \equiv \frac{\partial(\bar{\beta}_i - \bar{\beta}_j)/\partial d_i}{\partial(\bar{\beta}_i - \bar{\beta}_j)/\partial d_i |_{\sigma_j = \sigma_j^*}} \quad \text{and} \quad \underline{\zeta} \equiv \frac{\partial(\underline{\beta}_i - \underline{\beta}_j)/\partial d_i}{\partial(\underline{\beta}_i - \underline{\beta}_j)/\partial d_i |_{\sigma_j = \sigma_j^*}}.$$

In a tractable case of symmetric equilibrium, the following proposition shows that the strategic multiplier is larger than 1; that is, the effect of small bias of one researcher will be amplified:

Proposition 1.3. (Amplification of Bias of a Researcher) *Suppose $N = 2$, $c = 0$, and $\sigma_i = \sigma$ for both $i = 1, 2$. In a symmetric equilibrium in Proposition 1.2, the strategic multiplier between researchers satisfies $\bar{\zeta} = \underline{\zeta} \equiv \zeta$, and*

$$\zeta = \frac{Var_{total}}{Var_{total} - Var_{truncated}}, \quad (1.10)$$

where $Var_{truncated} \equiv Var(\beta_i | \beta_i \leq \bar{\beta})$ and $Var_{total} \equiv Var(\beta_i) = \sigma^2 + \sigma_b^2$. Thus, $\zeta > 1$.

Sketch of Proof: By deriving the comparative statics with the researchers' indifference condition. Let us focus on the indifference condition for $\bar{\beta}_i$; the condition for $\underline{\beta}_i$ can be derived analogously. We consider the symmetric equilibrium with no bias at the beginning. The indifference condition of researcher i with bias, d_i , is

$$\bar{\beta}_i + \mathbb{E}[\beta_j | \beta_j \leq \bar{\beta}_j] = - \left(2 + \frac{\sigma^2}{\sigma_b^2} \right) d_i, \quad (1.11)$$

where $d_i = 0$ at first. This expression (1.11) already shows that $\bar{\beta}_i$ will be decreasing

in $\bar{\beta}_j$. The expression (1.10) is derived from the comparative statics of $\bar{\beta}_i$ on d_i with the expression (1.11). Since $Var_{total} > Var_{truncated}$ by definition of truncated distribution, $\zeta > 1$. Appendix 1.3 contains a complete proof. \square

In words, the two researchers' thresholds, $\bar{\beta}_i$ and $\bar{\beta}_j$, are strategic substitutes of one another: when a small increase in d_i shifts $\bar{\beta}_i$ downwards, $\bar{\beta}_j$ will be adjusted upwards to offset this effect, which then causes $\bar{\beta}_i$ to shift downwards even further, and so on. The strategic multiplier quantifies how the difference between $\bar{\beta}_i$ and $\bar{\beta}_j$ due to such repeated adjustments is larger than the case when there was only the first adjustment of $\bar{\beta}_i$, keeping $\bar{\beta}_j$ fixed.

The multiplier (1.10) shows, heuristically, that the strategic substitution effect is very large. If the equilibrium thresholds, $\bar{\beta}_i$ and $\bar{\beta}_j$, are low so that $Var_{truncated} = \frac{2}{3}$, then $\zeta = 3$: the equilibrium difference in reporting is three times larger than the primitive difference in objectives. When the conventional threshold, $\frac{\bar{\beta}}{\sigma} = 1.96$, is applied in the environment with zero true effect ($b = 0$) and small variation in true effects ($\sigma_b \simeq 0$), then $Var_{truncated} \simeq 0.88$, which suggests $\zeta \geq 8$: that is, the underlying difference in objectives is only 12 percent of the observed difference in reporting thresholds.

This amplification result builds on the results of information aggregation models (Feddersen and Pesendorfer 1996) that illustrates strategic substitution effects among voters. In the voting model, when there are partisan voters, independent voters vote against the bias of the partisan voters to offset their influence on electoral outcomes. In this model, when another researcher is biased in one direction, the researcher will bias her reporting in the opposite direction. The new result is that, because the voting model has considered binary decisions whereas this model considers continuous decisions of reporting thresholds, the original bias will be amplified in equilibrium.

Numerical Results

While the Proposition 1.3 focused on the analytically tractable case of symmetric equilibrium with $\bar{\beta} = -\beta$, the same effect of strategic substitution also exists in the asymmetric equilibrium in Proposition 1.1. A numerical simulation shows strategic substitution can have a quantitatively important influence not only on symmetric equilibrium but also on asymmetric equilibrium that is locally stable. Let us consider an example with 2 researchers, $c = 0$, and $\sigma_0 = \sigma = 1$. If neither researcher is biased, then the equilibrium threshold is $\bar{\beta}_i = 0.19$. If researcher i has bias $d_i = -0.1$ so that he has bias towards policy implementation and researcher j does not have bias, $d_j = 0$, then their thresholds will become $\bar{\beta}_i = -0.25$ and $\bar{\beta}_j = 0.5$. Note that, if there

were only 1 researcher, then the threshold for recommending policy only changes by -0.2 . Thus, the strategic multiplier is $\zeta \simeq 3$ in this example, consistent with the back-of-the-envelope calculation above.

Evidence

A large body of public health research has shown that industry funded research are more likely to have positive outcomes, and thus, interpreted this as a result of publication bias. A meta-analysis of 30 studies has found that industry-funded research is roughly four times more likely to have positive outcomes than the publicly funded research (Lexchin 2003). Given such evidence, it is common to consider that the pharmaceutical companies have large bias towards drug approval with little regards for patients' welfare (Goldcare 2010). This model's amplification result suggests that, however, caution is warranted when interpreting the difference in reporting decisions as quantitatively reflective of the underlying differences in the objectives. While research funded by industries will perhaps have some bias towards the outcomes favorable to the industry, the bias in objectives need not be so large to explain the strong associations between results and identities of funding sources.

1.2.6 Discussions of Key Assumptions

The analyses have shown that coarse aggregation can explain various kinds of publication bias. Overall, the discussions henceforth will show that the main conclusions are not highly sensitive to some auxiliary assumptions, and the main assumptions are standard in economics literature and relevant in the real world. The caveat must be in place if there is a reviewer who directly meta-analyze the results, or if the conflict of interest is large.

Sensitivity to Alternative Assumptions

The following discussions show the implications of (i) sequential reporting, (ii) conflict of interests, (iii) unknown number of researchers, and (iv) risk aversion:

- (i) sequential vs simultaneous reporting: the characterized equilibria will still remain as equilibria even if the reporting is sequential when there are 2 researchers, since the analysis of simultaneous reporting had researchers condition their reports on pivotality. This logic is analogous to Dekel and Piccione 2000. If

the later researcher observes the early researcher's estimate, then the later researcher can summarize both estimates through meta-analyses and full reporting of all estimates will be the optimal equilibrium.

- (ii) conflict vs consistency of interests: the Section 2.5 has shown that the omission and inflation results are robust to small conflict of interest. When there is a large conflict of interest such as when merely profit-maximizing pharmaceutical firms report studies, however, the incentive constraints will bind: rigid publication rules to eliminate publication bias will be optimal.
- (iii) unknown vs known number of researcher: the analysis has assumed that, N , the number of researchers is a public information. The numerical analysis in Appendix B1.3 shows that, in a plausible setting, the policymaker's optimal strategy is to implement the policy if and only if there are strictly more positive results than negative results. In this sense, the reader needs not know how many researchers there are to implement the optimal rule (1.2).
- (iv) risk aversion vs risk neutrality: when the payoff exhibits risk aversion, the study precision has benefits of reducing the uncertainty in addition to its role in determining weights of Bayesian updating. Nonetheless, small risk aversion does not alter the results since the objective is continuous in risk aversion parameter; by Taylor approximation of constant relative risk aversion preference γ implies the decision rule $\frac{(\mathbb{E}b)^{1-\gamma}}{1-\gamma} - \frac{\gamma}{2} \frac{Var(b)}{(\mathbb{E}b)^{1+\gamma}} \geq c$.

Validity of Main Assumptions

The results rely on the key assumptions that the message space is smaller than signal space, and that researchers can make contingent reasoning. The following discussions explore their validity:

- (i) large state and signal space vs limited action and message space: the critical assumption that drives the results is the distinction between space of states and signals that are continuous, and the space of action and messages that are discrete. When either assumptions are modified, then the omission with bias no longer arises. However, I argue that this set of assumptions is particularly appropriate for scientific communication: the information the researchers have are rich and complex whereas the messages they can convey will be limited and must be simple. Binary actions also apply in key applications such as whether to adopt a particular medicine or policy.

- (ii) PBNE and contingent reasoning: the implicit yet important assumption is that the senders of information condition the reporting decision on events in which their reports are pivotal. This logic is key to and common across models of information aggregation that have been applied in a number of settings, including general public's voting (Feddersen and Pesendorfer 1996), juror's voting (Feddersen and Pesendorfer 1998), opinion polls (Morgan and Stocken 2008), and demonstrations (Battaglini 2017). While such sophisticated reasoning may appear unrealistic and some lab experiments show individuals are unable to engage in such reasoning (Esponda and Vespa 2013), there is also evidence from both lab (Battaglini et al. 2010, Dickhaut 1995) and fields (Kawai and Watanabe 2013) that some people condition their voting decisions on others' decisions.

1.3 An Empirical Test of the Communication Model

The communication model has shown that, if aggregation frictions are a key reason of omission, then both precise null results and extremely negative results will be reported. This Section develops a new empirical test to compare this prediction against some other publication selection processes, and applies this to show it holds with an economics data set.

1.3.1 Various Models of Publication Bias

Various existing bias correction methods have assumed specific publication selection process. This sub-Section shows that the selection process based on the communication model makes different predictions than the two most commonly¹⁶ used methods assume.

Data Generating Process

The data generation process of the published estimates, $\{\beta_i, \sigma_i\}$, will be assumed to take three steps given various independence assumptions. First, the underlying effect,

¹⁶While there are also some other models, this paper focuses on the comparison with most commonly used models: as of December 2018, Duval and Tweedie (2000) that introduced trim-and-fill method is cited over 4,900 times, and Hedges and VIVEA (2000) that extends Hedges (1992) is cited over 2,000 times on Google Scholar. It is beyond the scope of this paper to fully explore the other selection models: Copas and Li (1997) (note that the working paper version had contained full discussions), Fafchamps and Labonne (2016), and Frankel and Kasy (2018).

$b_i \sim F$, and the study precision, $\sigma_i \sim G$, are independently¹⁷ determined. Second, the random error, $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$, is independently drawn and the coefficient, $\beta_i = b_i + \epsilon_i$, is determined. Third, the study is published with some probabilities, $P(\beta_i, \sigma_i)$, that depend on $\{\beta_i, \sigma_i\}$. We will denote $b_0 \equiv \int b_i dF$ and $\sigma_0^2 \equiv \int (b_i - b_0)^2 dF$ as the mean and variance of underlying effect. Moreover, let us denote $H(\beta_i)$ as the distribution of coefficient estimates given F and G . σ_0^2 measures heterogeneity of effects across studies, whereas σ_b^2 had measured heterogeneity of effects across policies.

Distinct Predictions of Various Models of Publication Bias


The communication model of this paper makes a distinct prediction on a form of publication selection, $P(\beta_i, \sigma_i)$, compared to the selection assumption behind the 2 most commonly used bias correction methods. Let us denote $\tilde{G}_\theta(\sigma)$ as the distribution of all standard errors conditional on the study being the non-positive results without any selection, and G_θ as its observed distribution with selection; let us also denote $\tilde{H}_0(\beta)$ as the distribution of coefficient estimates conditional on the study not rejecting the null hypothesis with threshold \bar{t} , and H_0 as its observed distribution with selection. The Figure 1-4 summarizes the distinct predictions.

(1) communication model-based selection: the model of this paper suggests that, under aggregation frictions, the imprecisely estimated results with coefficients with small absolute values will be omitted¹⁸. As has been discussed in Section 2.3 and illustrated in Figure 1-3, the omission probability shrinks to zero when the between-study heterogeneity, σ_0 , is small. Therefore, when negative results are published, they are either precise null results so that $G_\theta > \tilde{G}_\theta$, or imprecisely estimated but extremely negative results so that $H_0 > \tilde{H}_0$.

(2) uniform selection: the model behind ‘‘Hedges’’ bias correction method, proposed by Hedges 1992, suggests that the statistically insignificant results will be

¹⁷This independence assumption imposes that the studies with small vs large effects have identical true effects. This assumption is violated, for example, when the sample size affects the quality of treatment. Nonetheless, it is also assumed in other influential meta-analysis papers, such as Hedges 1992, Duval and Tweedie 2000, and Andrews and Kasy 2019.

¹⁸If the estimates can be directly communicated, then imprecise null results can be valuable (Abadie 2018).



(1) communication model: $\hat{G}_\theta > \tilde{G}_\theta$ and $\hat{H}_0 > \tilde{H}_0$ (2) uniform selection: $\hat{G}_\theta = \tilde{G}_\theta$ and $\hat{H}_0 > \tilde{H}_0$ (3) extremum selection: $\hat{G}_\theta > \tilde{G}_\theta$ and $\hat{H}_0 < \tilde{H}_0$

Figure 1-4: Selection models described in funnel plots

Notes: Using the funnel plot, the Figure 1-4 visualizes the (1) communication-based selection, (2) uniform selection, and (3) extremum selection models. The reported studies are depicted in dark dots, whereas the regions of omissions are approximated by the shaded area. The figures are generated by the author for an illustrative purpose, and do not reflect actual data.

uniformly less likely to be published than statistically significant results¹⁹:

$$\mathbb{P}(\text{study } i \text{ is reported}) = \begin{cases} \eta_1 & \text{if } \frac{|\beta_i|}{\sigma_i} \geq 1.96 \\ \eta_0 & \text{if } \frac{|\beta_i|}{\sigma_i} < 1.96, \end{cases} \quad (1.12)$$

where $\eta_1 > \eta_0 > 0$. This suggests that the results that are not positively significant are systematically unlikely to be published. Thus, conditional on being null results, the distribution is identical to the underlying distribution of null results: $\tilde{G}_\theta = G_\theta$. Nonetheless, published negative results will be likely to be extreme negative results since results with intermediate coefficients have low t -statistics: $H_0 > \tilde{H}_0$. With an assumption that F is normal, this model is commonly used in economics with maximum likelihood estimation to correct for publication bias. (Hedges, 1992, Ashenfelter et al. 1999, McCrary et al. 2016, Andrews and Kasy 2019). This selection is consistent with the setting in which the researchers select based only on statistical significance to make their publication decisions.

¹⁹While the t -statistic thresholds can be specified more flexibly, it is common to apply the conventional threshold of $\bar{t} = 1.96$ on both positive and negative signs; since the sample size is often small in meta-analyses, it is practically difficult to estimate a model with many cut-offs.

(3) extremum selection: the model behind “trim-and-fill” bias correction method, proposed by Duval and Tweedie 2000, suggests that the most *extreme* negative results will be omitted.

$$\mathbb{P}(\text{study } i \text{ is reported}) = \begin{cases} 1 & \text{if } \beta_i \geq \beta_{\min} \\ 0 & \text{if } \beta_i < \beta_{\min}, \end{cases} \quad (1.13)$$

where β_{\min} is some threshold. In common case where $\beta_{\min} < b_0$, there arises little selection among the most precise studies. Thus, null results are more likely to be reported when the standard error is small so that $G_\theta > \tilde{G}_\theta$. At the same time, the model also suggests that marginally insignificant negative results are not particularly more likely to be omitted since the selection is unrelated to the statistical significance milestones: $H_0 < \tilde{H}_0$. With an assumption that the underlying distribution of benefit, F , is symmetric, the trim-and-fill method imputes the most negative missing studies and computes the bias corrected estimate \hat{b}_0 . This reporting rule is consistent with the setting in which the researcher is biased towards positive results and do not hope to show extreme negative results, given a completely uninformed reader with improper uniform prior.

1.3.2 A Test to Distinguish the Various Models

Given that each model has distinct implications on the distribution of standard errors and coefficients of published studies, we develop an empirical test to examine them. The key obstacle is that the underlying distributions are unobserved. We (1) show that the underlying distribution can be predicted with some assumptions, and (2) describe the overview of the estimation and testing steps.

Assumptions

To estimate the underlying distribution of $\{\beta_i, \sigma_i\}$ without selection, we need some regions of the estimates that do not suffer from selection, and need ways to extrapolate from those regions to other regions with selection. To operationalize these requirements, we will use the following assumptions²⁰ that are parsimonious and common in meta-analyses:

²⁰These assumptions suggest that the bias due to inflation of marginal results is unimportant. With the limited sample sizes that is relevant for meta-analysis, it is infeasible to distinguish the inflation and omission; such test requires large sample size at the margin of statistical significance. While restrictive, this is an assumption applied in all other studies on bias correction.

A1. *Constant selection within statistically significant results:* $P(\beta_i, \sigma_i) = \bar{P} \in (0, 1)$ for any $\frac{|\beta_i|}{\sigma_i} \geq \bar{t}$.

A2. *Underlying effects with independent normal distribution:* $b_i \stackrel{iid}{\sim} F(b_i) = \Phi\left(\frac{b_i - b_0}{\sigma_0}\right)$

Under the extremum selection model, A1 will be satisfied so long as $\beta_{\min} < 0$; under the uniform selection, both A1 and A2 will be satisfied. Note that A1 does not require that the publication probability is 1 for statistically significant results.

Semi-parametric Estimation and Testing Steps

We estimate the underlying distributions semi-parametrically using the Assumptions A1 and A2, and compare them against the observed distribution with a Kolmogorov-Smirnov (KS)-type test. The estimation is non-parametric along the dimension of σ while it assumes normal distribution along the dimension of β . While the complete description and discussion are relegated to the Appendix B2.1, the following overview describes the three steps of estimation and testing:

1. estimate $\{\hat{b}_0, \hat{\sigma}_0\}$ by the stem-based bias correction method that is robust to various kinds of distribution, $F(b_i)$, and selection, $P(\beta_i, \sigma_i)$;
2. estimate (i) the distribution $G_\theta(\sigma|\hat{b}_0, \hat{\sigma}_0)$ using the studies such that $|t_i| \geq 1.96$, and (ii) the distribution $H_0(\sigma|\hat{b}_0, \hat{\sigma}_0)$ using the distribution $G_\theta(\sigma|\hat{b}_0, \hat{\sigma}_0)$ estimated using the studies such that $t_i \geq 1.96$;
3. estimate the KS statistic for each distribution, D^G and D^H :

$$D^G = \sup_{\sigma} \left\{ \hat{G}_\theta(\sigma) - \tilde{G}_\theta(\sigma|\hat{b}_0, \hat{\sigma}_0) \right\} \quad \text{and} \quad D^H = \sup_{\beta} \left\{ \hat{H}_0(\beta) - \tilde{H}_0(\beta|\hat{b}_0, \hat{\sigma}_0) \right\}$$

and associated one-sided p -values using the two-step bootstrap over estimates of b_0 and sampling of each study's σ_i and β_i .

1.3.3 Application

Using the test above, this sub-Section analyzes a meta-analysis data set with data selected from the papers that highlight the binary conclusions²¹. The result shows

²¹A previous version of this paper (Furukawa 2016) also analyzes the data of Intertemporal Elasticity of Substitution, which highlight the “binary conclusions” less since the quantity of interest is a continuous estimate.

that the communication model-based selection pattern fits the data more adequately than the other two models.

Data

The data come from the set of 106 studies (i) that are included in the total of 111 studies meta-analysis of labor union’s effect on firm productivity (Doucouliagos et al. 2018) and (ii) that have the binary conclusions from t -statistics that match with the conclusions highlighted in the original papers. How the labor union affects firms is a highly contested issue, with various evidence supporting both positive and negative views. As each paper contains many estimates from various specification, the analysis uses its median value. To focus on the coefficients underlying the highlighted conclusions, two independent readers examined the abstract, introduction, and conclusions of each paper and excluded some papers whose highlighted conclusions in the paper did not match the implication of t -statistics in Doucouliagos et al.’s data set²².

Results

The results suggest that the reporting patterns of null and negative results²³ are consistent with the communication model-based selection process, but not with other processes. Figure 1-5 visualizes the two results. First, observed null results tend to be more precise than predicted distribution of null results: $\hat{G}_0 > \tilde{G}_0$. Concretely, while only 20 percent of studies are predicted to have standard errors less than .08, above 70 percent of studies have standard errors smaller than 0.8 ($p = .000$). This pattern is not consistent with the uniform omission model, which suggest that the two distributions will be roughly equal with one another. Second, observed negative results, including null results and negative significant results, tend to be more negative than their predicted distribution: $\hat{H}_0 > \tilde{H}_0$. Concretely, while only roughly 14 percent of studies are expected to have coefficient less than -0.125, the observed distribution has over 29 percent of studies have such negative values ($p = .0311$). This pattern is

²²The detailed discussions on the inclusion criteria, as well as classification’s text evidence, are available in https://github.com/Chishio318/Data_publication_bias. The meta-analysis can include only one estimate from one study: when there are multiple estimates in one study, it is common to choose the estimates of median magnitude (for example, in Havránek 2015).

²³As Figure 1-2 suggests that no studies will be reported in an omission region, taken literally, the data may appear to contradict the communication model. However, existence of some studies in omission regions could be explained by dispersed beliefs across researchers that lead to different – even overlapping – thresholds of reporting positive vs negative results as illustrated in Section 2.5. The paper does not claim that aggregation friction is the only reason behind publication bias: instead, it only suggests that aggregation friction can explain some regularities of publication bias.

not consistent with the extremum omission model, which suggest that the reported studies will have more moderate coefficient values. Taken together, among the three models, the communication-based selection process is the only one that can account for the pattern of omission in this data set.

- (1) distribution of σ_i of null results: (2) distribution of β_i of non-positive results:

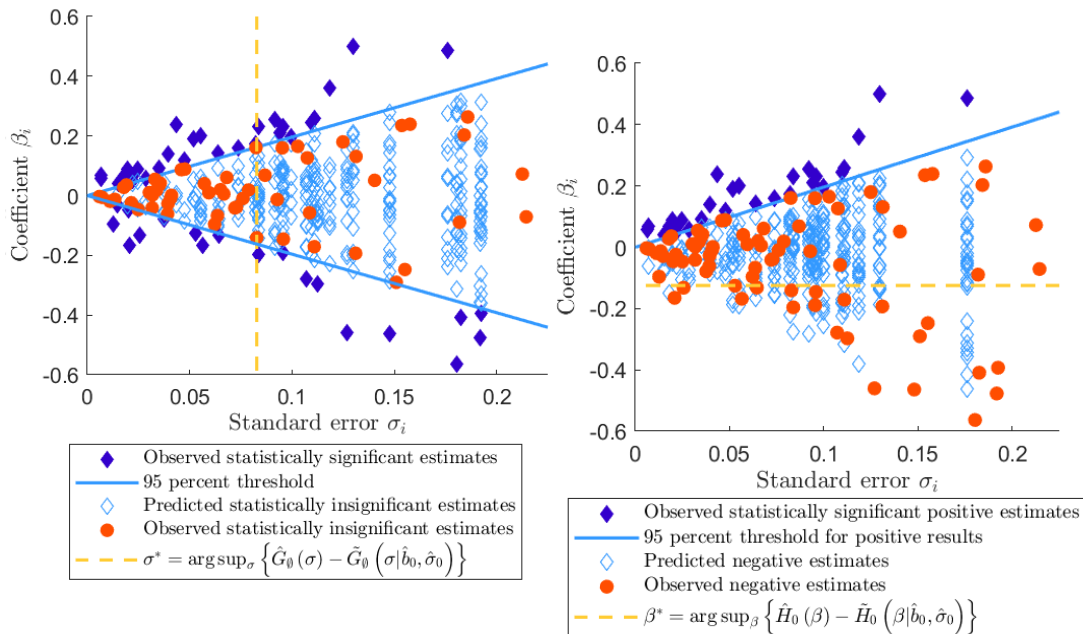


Figure 1-5: KS-type test illustrated in funnel plots

Notes: Figure 1-5 are the funnel plots that illustrate the KS-type test described in Section 3.2.2. The filled diamonds with dark blue are observed significant results; the empty diamonds are predicted non-positive results; and filled circle with orange are predicted non-positive results. The dashed lines represent the values at which the KS-type test is evaluated.

1.4 A New “Stem-based” Bias Correction Method


The communication model in Section 2 has suggested an alternative publication selection process, and the empirical analysis in Section 3 has shown its relevance in a real-world data set. Moreover, the model suggests, if aggregation friction is the important reason of publication selection, then the selection will depend on a number of economic primitives unobservable to meta-analysts. What could we do to alleviate the bias that arises from publication selection when aggregating various estimates?

This Section presents a new, non-parametric, fully data-dependent, and generally conservative bias correction method, to be called a “stem-based” bias correction method. The estimate uses the studies with highest precision, which correspond to the “stem” of the “funnel” plot, to estimate a bias corrected average effect. It has both theoretical and empirical merits over other existing methods: theoretically, the estimate is based on weaker assumptions on the publication selection process and the underlying distribution than other methods; empirically, the simulation shows that the estimate has adequate coverage probabilities across different publication selection processes.

1.4.1 Main Argument

The stem-based bias correction method uses some of the most precise studies because, across various selection models, precise studies suffer less from publication bias than imprecise studies. In the communication model of this paper, the most precise studies are omitted less often, as visualized in Figure 1-1. In the two most commonly used models described in Section 3.1.2²⁴, the following Proposition shows that the bias is decreasing in study precision, and that, under some conditions, the bias is zero as studies become infinitely precise.

²⁴While rigorous analysis is beyond the scope of this paper, this pattern also holds with other models such as Copas and Li (1997), Fafchamps and Labonne (2016), and both static and dynamic models of Frankel and Kasy (2018).



(1) communication model (2) uniform selection (3) extremum selection

Figure 1-6: Meta-analysis estimates in funnel plots

Notes: Using the funnel plot, the Figure 1-6 visualizes the (1) communication-based selection, (2) uniform selection, (3) extremum selection models. The reported studies are depicted in dark dots, whereas the regions of omissions are approximated by the shaded area. The thick red lines indicate the mean level of estimates at given values of standard errors. The figures are generated by the author for an illustrative purpose, and do not reflect actual data.

Proposition 2 (minimal bias among most precise studies across selection models).

Define the bias of studies with precision σ_i as $Bias(\sigma_i) \equiv \mathbb{E}[\beta_i | \sigma_i, \text{study } i \text{ reported}] - b_0$.

1. (Monotonicity) $Bias(\sigma_i)^2$ is increasing in σ_i for all σ_i under the extremum selection models, and for $\sigma_i \in [0, \underline{\sigma}]$ for some $\underline{\sigma} > 0$ under the uniform selection model.

2. (Limit) $\lim_{\sigma_i \rightarrow 0} Bias(\sigma_i)^2 = 0$ if $\sigma_0 = 0$ and threshold β_{\min} is sufficiently low under the extremum selection model, and always under the uniform selection model.

Sketch of Proof. By comparison of conditional across values of σ_i expectation given the bias selection model. Appendix A5 contains a proof. \square

That is, more precise studies are less subject to the publication selection; and moreover, the bias approaches zero as the studies become infinitely precise under some conditions. Heuristically, concerns for low statistical significance or extremely negative values become unimportant when studies are precise.

The stem-based method chooses the number of studies, n_{stem} , to include by op-

timizing over the bias-variance trade-off. While focusing only on the most precise studies give the least biased estimate, it also suffers from high variance. Therefore, the stem-based method includes n_{stem} to minimize the Mean Squared Error (MSE) of the estimate while ensuring that the assumed and implied variance are consistent with one another. Denoting the publication selection process as P as in Section 3.1.1, it strives to solve:

$$\min_n MSE \left(\hat{b}_0^n | \sigma_0 \right) = Var \left(\hat{b}_0^n, \sigma_0 \right) + Bias^2 \left(\hat{b}_0^n, b_0 \right) \text{ subject to } Var \left(b_i | \hat{b}_0^n, P \right) = \sigma_0^2 \quad (1.14)$$

However, this problem of minimizing the exact MSE requires the knowledge of b_0 , true mean, and P , publication selection process. Since solving this criteria is infeasible, the method instead solves its empirical analogue:

$$\min_n \tilde{MSE} \left(\hat{b}_0^n, \hat{b}_0, \sigma_0 \right) \text{ subject to } \hat{Var} \left(b_i | \hat{b}_0^n \right) = \sigma_0^2. \quad (1.15)$$

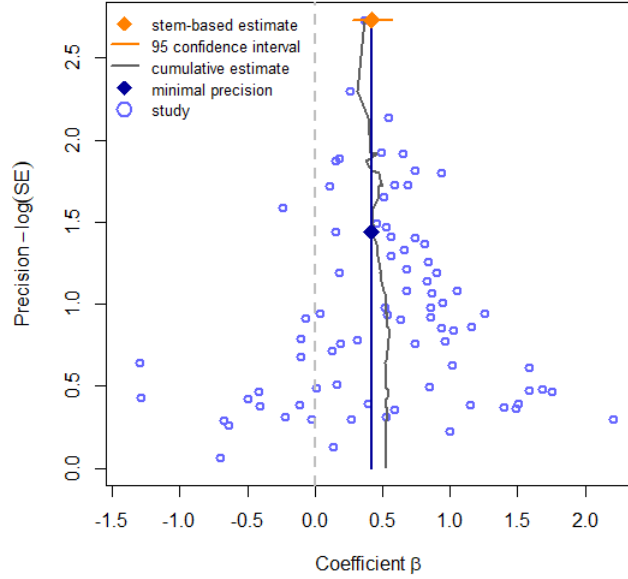
That is, the MSE is replaced by an unbiased estimate of its relevant component, $\tilde{MSE} \left(\hat{b}_0^n, \hat{b}_0, \sigma_0 \right)$, and the implied variance term is replaced by its empirical analogue, $\hat{Var} \left(b_i | \hat{b}_0^n \right)$.

The estimates of stem-based method can be visually represented with a funnel plot (Figure 1-7): with uniform selection generating 80 studies in this simulation, it is optimal to include 17 studies, which is roughly an average number of studies included in simulations. Including all studies leads to an upward bias, as indicated by theory. On the other hand, only a few most precise study can be noisy, both because of between-study heterogeneity and of insufficient high within-study heterogeneity due to low total sample size, so that inclusion of more studies leads to a smaller 95 confidence interval that covers 0.4, the true mean value.


1.4.2 Estimation

Given the objective (1.14), this sub-Section describes a particular approximation to (1.15) that this paper proposes, illustrates estimation steps, and discusses the assumptions necessary to ensure its reliability.

Figure 1-7: A visualization of stem-based method



(1) Funnel plot



(2) Mean Squared Error

Notes: Figure 1-7 are an illustration of (1) a funnel plot of stem-based bias correction method, and (2) the Mean Squared error criteria for choosing the n_{stem} , the optimal number of studies to include. The data comes from a simulation of 80 studies under the uniform selection model such that the number of included studies is 17. (1) The funnel plot, with y -axis denoting a measure of precision, describes the stem-based method. The orange diamond at the top indicates the stem-based estimate along with its 95 percent confidence interval. The connected line is the estimate with various $n_{stem} \in \{1, \dots, N\}$, indicating how aggregate estimates change. The diamond at the middle of the curve indicates minimal level of precision for the inclusion. Therefore, the stem-based estimate is given by the studies, represented by circle, whose precision are above this diamond. (2) The relevant components of Mean Squared Error is plotted, indicating that the $Bias^2$ is increasing while Var is decreasing in n_{stem} .

Estimation Steps

The stem-based method computes the estimates with the following inner and outer algorithms: the inner algorithm computes the bias corrected mean given an assumed value of σ_0^2 ; the outer algorithm computes the implied variance and ensure that it is consistent with its assumed value.

I. Inner algorithm: estimate $\hat{b}_{\text{stem}}, SE(\hat{b}_{\text{stem}}), n_{\text{stem}}$ given an assumed value of σ_0 .

1. rank and index studies in the ascending order of standard error so that $\sigma_1 \leq \sigma_2 \dots \leq \sigma_N$
2. for each $n = 2, \dots, N$, compute the relevant mean squared error, $M\tilde{S}E(n)$, as follows: given weights $w_i \equiv \frac{1}{\sigma_i^2 + \sigma_0^2}$,

$$M\tilde{S}E(n) = \frac{\sum_{i=2}^n \sum_{j \neq i}^n w_i w_j \beta_i \beta_j}{\sum_{i=2}^n \sum_{j \neq i}^n w_i w_j} - 2\beta_1 \frac{\sum_{i=2}^n w_i \beta_i}{\sum_{i=2}^n w_i}$$

3. compute the optimal number of included studies, n_{stem} : n_{stem} minimizes the relevant components of MSE, so that

$$n_{\text{stem}} \in \arg \min_n M\tilde{S}E(n)$$

Thus, the stem-based estimate is $\hat{b}_{\text{stem}} \equiv \frac{\sum_{i=1}^{n_{\text{stem}}} w_i \beta_i}{\sum_{i=1}^{n_{\text{stem}}} w_i}$, $SE(\hat{b}_{\text{stem}}) \equiv \frac{1}{\sqrt{\sum_{i=1}^{n_{\text{stem}}} w_i}}$. The estimation of \hat{b}_{stem} applies the inverse variance weights since they minimize the variance of the estimator given the sample.

II. Outer algorithm: search over values of σ_0^2 such that the implied $\hat{V}ar(b_i | \hat{b}_0^n)$ is consistent. Throughout, we adopt the formula of variance proposed by DerSimonian and Laird (1996)²⁵: given weights $w'_i \equiv \frac{1}{\sigma_i^2}$, $\hat{V}ar(\beta_i | \hat{b}_0^n) = \max\{\hat{V}ar(\beta_i | \hat{b}_0^n), 0\}$, where

$$\hat{V}ar(\beta_i | \hat{b}_0^n) = \frac{\sum_{i=1}^N w'_i (\beta_i - \hat{b}_0^n)^2 - (N-1)}{\sum_{i=1}^N w'_i - \frac{\sum_{i=1}^N w_i'^2}{\sum_{i=1}^N w'_i}}. \quad (1.16)$$

Here, $\hat{b}_0^n \equiv \frac{\sum_{i=1}^n w_i \beta_i}{\sum_{i=1}^n w_i}$ is the estimate based on n studies.

²⁵This formula is commonly used in non-parametric estimations of F . For example, trim-and-fill proposed by Duval and Tweedie (2000) also uses this originally. While there are some criticisms to this approach (Veroniki et al. 2015), it is left for future work to explore how between-study heterogeneity can be adequately estimated.

1. set two initial estimates of σ_0^2 by applying (1.16) to $\hat{b}_0^{\min} = \frac{\sum_{i=1}^N w_i' \beta_i}{\sum_{i=1}^N w_i'}$ and $\hat{b}_0^{\max} = \sum_{i=1}^N \beta_i$.
2. compute the implied stem-based estimates and their variance by applying (1.16)
3. iterate step 2 until it converges; if the limit of maximum and minimum disagree, then choose the maximum.

Additional Arguments

Turning an ideal problem (1.14) into a feasible problem (1.15) had required ways to approximate the knowledge of b_0 , true mean, and P , publication selection process. The method had applied non-parametric estimation techniques of unbiased Cross-Validation criteria to approximate b_0 ; and estimated σ_0^2 to give a conservative confidence interval of \hat{b}_{stem} given unknown P :

Unbiased Cross-Validation criteria for b_0 : we can replace the component, $MSE(n)$, by its relevant term, $M\tilde{S}E(n)$, since they differ only by a constant²⁶. The formula proposed in the inner algorithm provides an approximately unbiased estimate of $M\tilde{S}E(n)$ under some assumptions: if (A1) $\mathbb{E}\beta_1 \simeq b_0$ and (A2) $\mathbb{E}\hat{b}_{2,n} \simeq \mathbb{E}\hat{b}_{1,n}$, then

$$\begin{aligned} \mathbb{E}M\tilde{S}E(n) &= \mathbb{E} \frac{\sum_{i=2}^n \sum_{j \neq i}^i w_i w_j \beta_i \beta_j}{\sum_{i=2}^n \sum_{j \neq i}^i w_i w_j} - 2\mathbb{E}\beta_1 \frac{\sum_{i=2}^n w_i \beta_i}{\sum_{i=2}^n w_i} \\ &= \frac{\sum_{i=2}^n \sum_{j \neq i}^i w_i w_j \mathbb{E}\beta_i \mathbb{E}\beta_j}{\sum_{i=2}^n \sum_{j \neq i}^i w_i w_j} - 2\mathbb{E}\beta_1 \mathbb{E} \frac{\sum_{i=2}^n w_i \beta_i}{\sum_{i=2}^n w_i} \\ &\simeq \hat{b}_0^2 - 2b_0 \hat{b}_0 \end{aligned}$$

There are two statistical techniques involved in these steps: the first term computes the squared term by leaving one sample out in order to avoid the bias that arises due to the squared term²⁷. More importantly, the second term applies a Cross-Validation (CV) technique by replacing the true value of b_0 by its estimate. Since β_1 is the least biased estimate of b_0 , we apply the “leave-one-out” method in CV technique by splitting the sample into the most precise estimate that constitutes a testing set and all other estimates that constitute a training set.

²⁶To see this, we can expand the bias squared term:

$$M\tilde{S}E(\hat{b}_0) \equiv \hat{b}_0^2 - 2b_0 \hat{b}_0 = (\hat{b}_0 - b_0)^2 - b_0^2 = MSE(\hat{b}_0, b_0) - b_0^2.$$

²⁷The method requires at least $N = 3$ studies to compute the relevant MSE.

Equating $\widehat{Var}(b_i|\hat{b}_0^n) = \hat{\sigma}_0^2$: the outer algorithm likely leads to a large estimate of σ_0^2 for three reasons. (i) While the exact selection process, P , is unknown, the variance is overestimated when it is the intermediate results that is omitted as suggested by theory; (ii) the estimation of $\widehat{Var}(b_i|\hat{b}_0^n)$ uses the entire sample so as to avoid underestimating the variance with only few samples used in stem-based estimation; and (iii) when there are multiple values of σ_0^2 that are consistent with one another, the method uses a larger one. By choosing the specification such that the estimate of σ_0^2 is large, the method strives to make a conservative estimate of the 95 confidence interval for \hat{b}_{stem} .

Summary of Assumptions

In summary, for the stem-based method to generate a reliable estimate, we need that the MSE can be well-approximated, (A1) $\beta_1 \simeq b_0$ and (A2) $\hat{b}_{2,n} \simeq \hat{b}_{1,n}$, and that the variance implied is close to the true variance, $\widehat{Var}(b_i|\hat{b}_0^n) \simeq \sigma_0^2$. These conditions may not be satisfied when the underlying variance, σ_0^2 , is large since even the most precise studies may not approximate the true underlying mean. While the method imposes no assumptions on underlying distribution, F , and only monotonicity assumption on the selection process, P , it instead relies on these assumptions to mitigate publication bias.

One implicit and important assumption is that the studies' precision is correctly reported. An inflation of t -statistics through under-reporting the standard error, such as through choice of units of clustering, can compromise the reliability of this method. I recommend investigating the specifications of most precise studies in detail to avoid severe misreporting of study precision.

1.4.3 Assessment

Given the theoretical foundations and the assumptions made in estimation steps, how does stem-based bias correction method perform across various selection processes? The simulation henceforth shows that the stem-based correction method provides a more reliable estimate of confidence intervals than other commonly used methods in meta-analysis settings calibrated to plausible values²⁸.

²⁸There are new methods that have been developed, including regression-based approach of PET-PEESE (Stanley 2008), maximum likelihood approach (Andrews and Kasy 2018), selection model analogous to Heckman's two step process (Copas and Li 1997), other methods that focus on precise studies such as top10 (Stanley et al. 2010) and kink-based methods (Bomy and Rachinger 2018), and bias correction using only significant studies (Simonsohn et al. 2014). It is left for future work

Simulation set-up

This simulation will compute the coverage probabilities and interval lengths with a Monte Carlo experiment. The studies' standard errors, σ_i , is drawn from $G(\sigma)$ that approximates the implied distribution from the labor union data sets (Appendix B1.1). Concretely, $\hat{G}(\sigma)$ has the distribution of σ^2 that is χ^2 distribution with 2 degrees of freedom with support of $[0, 4]$ that is re-scaled such that $Supp(G) = [0, 1]$. The studies' coefficients are determined by $\beta_i = b_i + \epsilon_i$, where $b_i \sim \mathcal{N}(b_0, \sigma_0^2)$ and $\epsilon_i \sim \mathcal{N}(0, \sigma_i^2)$ and is independently drawn. $\sigma_0 = 0.3$ so that match the degree of heterogeneity in the labor union data set; $b_0 = 0.4$ so that the average effect size ($\frac{b_0}{\sqrt{2}\sigma_0} \simeq 0.94$) is large but reasonable. We consider the data with 30 and 80 published studies to investigate how sample size affects the reliability of estimates; these are the range of small and large meta-analysis data also used in other simulation studies (Duval and Tweedie 2000, Stanley and Doucouliagos 2014).

There will be three sets of data generating processes and four estimation methods²⁹, as presented in Table 1. We begin by simulating the data without selection (row 1) and the estimation method without any bias correction (columns (i) and (ii)). Then, row 2 presents the estimates with uniform selection model in which statistically insignificant results with $t = 1.96$ thresholds are reported with only 30 percent of the time ($\eta_1 = 1, \eta_0 = 0$), and columns (iii) and (iv) show the uniform MLE method (Hedges 1992, Hedges and Vevea 2005) that assumes this; row 3 presents the estimates with extremum selection model in which some very negative results are reported ($\hat{\beta} = -0.1$), and columns (v) and (vi) presents the trim-and-fill method that assumes such selection process. The selection parameters, η_1, η_0 , are based on the estimates from Andrews and Kasy 2019, and the parameter $\hat{\beta}$ is chosen so that the coverage probability with no correction is roughly equal between the two selection models. Finally, columns (vii) and (viii) presents estimation results using the stem-based method. In this way, with realistic parameter values, this simulation will assess not only how each method performs given the selection process that the method assumes, but also how each performs given the process that it does *not* assume.

to exhaustively investigate the relative merits and demerits of these methods.

²⁹Each estimation has utilized the canned command available in R. The trim-and-fill correction uses a package in `metafor` (Viechtbauer 2010), with between-study variance estimated using DerSimonian and Laird method as proposed in the original paper by Duval and Tweedie 2000. The uniform correction uses the package `weightr` (Coburn and Vevea 2017). Note that each algorithm had implementation difficulties due to non-convergence in trim-and-fill, and non-singularity of Hessian. In this simulation, each estimation method was evaluated with the data sets that do not have these estimation problems.

Results

The main result is that the confidence intervals based on the stem-based correction method are more reliable across various selection models than those based on other methods. With estimation with no correction (columns (i) and (ii)), the coverage probabilities are close 0.95 when there is no publication selection but are 0.26 when there is serious omission; with estimation with correction methods, the coverage probabilities are reasonable when their respective assumed selection process is correct, they can be low when it is different. With uniform MLE, it is roughly 0.76~0.88 given uniform selection model, but is 0.13~0.47 with extremum selection; with trim-and-fill method, the coverage probability is roughly 0.64~0.67 given extremum selection model, but is 0.43~0.69 with uniform selection; On the other hand, the stem-based estimates have coverage probabilities of above 0.78 across selection models.

The improvement of robustness of stem-based methods comes with the disadvantage of larger average interval lengths. The simulation underlying the Table 1-1 finds, on average, roughly $n_{stem}^* = 9$ studies for $N = 30$, and $n_{stem}^* = 15$ studies for $N = 80$ since the distribution $\hat{G}(\sigma)$ has high density of very precise studies.³⁰ Table 1-1 shows that the average interval length is roughly 1.5 to 2 times larger than the other methods that use all data points. Nonetheless, when a less permissive estimation methods such as stem-based methods reject the null hypotheses, one can be more confident that the conclusion is not driven by particular selection method that the method has imposed.

³⁰The number of included studies, n_{stem}^* , varies substantially across replication data sets within the same simulation environment. This heterogeneity of n_{stem}^* suggests the advantage of stem-based method relative to the rule-of-thumb approach that uses some fixed number or fraction of all studies. At the same time, there are many studies with only a few studies included. While n_{stem}^* may appear to indicate severity of publication selection since $n_{stem}^* = N$ in the absence of any selection, simulation indicates that the difference in n_{stem}^* between the data with or without selection is limited.

	no correction		uniform MLE		trim-and-fill		stem-based	
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)
no selection	0.94	0.93	0.89	0.78	0.85	0.83	0.95	0.93
	[0.25]	[0.40]	[0.29]	[0.42]	[0.23]	[0.38]	[0.55]	[0.64]
uniform	0.26	0.63	0.88	0.76	0.43	0.69	0.80	0.85
	[0.25]	[0.41]	[0.31]	[0.42]	[0.22]	[0.37]	[0.58]	[0.65]
extremum	0.26	0.62	0.13	0.47	0.64	0.67	0.78	0.79
	[0.21]	[0.33]	[0.20]	[0.34]	[0.20]	[0.31]	[0.43]	[0.48]
N	80	30	80	30	80	30	80	30

Table 1.1: Simulation of bias correction methods across various models

Notes. Table 1-1. reports the coverage probability and average interval length (noted in []) for the sample in which there are $N = 80$ studies and $N = 30$ studies. The simulation is based on a 1,000 replications of the data sets.

1.4.4 Final Remarks

The two most influential bias correction methods with high citations have relied on specific assumptions about the publication selection process, P , and the underlying distribution³¹, F . Various authors defend their own assumptions against each other: Duval and Tweedie (2000) justifies the extremum selection model by writing “A number of authors ... have pointed out that this simple p -value suppression scenario is rather simplistic since it fails to acknowledge the role of other criteria, such as size of study.” Simonsohn (2014) criticizes this approach and writes “In most fields, however, publication bias is governed by p -values rather than effect size.” The communication model of this paper suggests both criticisms are valid: while the selection process can be approximated by the constant t -statistics approach, study precisions also have important impact on publication decisions.

The stem-based bias correction method takes a different approach that uses the monotonicity property of various selection processes, and makes no assumptions on the underlying distribution unlike in other methods that assumed normality or symmetry. While there are other assumptions in estimation steps to perform well, the

³¹Even the “trim-and-fill” method’s assumption that F is symmetric can be problematic when economists hope to produce a meta-analysis estimates of elasticity, on which microeconomic theories impose sign restrictions.

numerical simulation shows that the method has more adequate coverage probabilities across a range of publication selection processes. In fact, there have been authors who have suggested to focus on some arbitrary number of most precise studies (Barth et al. 2013, Stanley et al. 2010). This paper builds on their ideas, proposes a formal theoretical justification of this approach, and develops an algorithm to choose an optimal number of most precise studies to include. In this way, the method can provide a meta-analysis tool that has merits to the researchers who believe in either processes of publication selection and who wish to build consensus among researchers who believe in different processes.

1.5 Conclusion

There are two thought experiments that question the common interpretation that (i) the publication bias must arise from the biased motives of journals and researchers, and that (ii) it will be socially optimal if journals publish all binary conclusions:

- (i) if readers prefer publication outlets with full reporting of all results, then a journal or a researcher can singularly announce that they will publish all results that they observe. Given the current technology of record keeping and replication, this statement can be verifiable. Then, the demand for such journals and researchers must increase, resulting in a higher demand that journals and researchers are seeking. Yet this deviation from the current communication equilibrium with publication bias is not observed today.
- (ii) if researchers report all results of null hypothesis testing and readers wish to use any drugs with positive effects but consider only the binary conclusions, then readers must use the drug even when only 3 percent of studies are positive and 97 percent of studies are negative. This is because, with a conventional null hypothesis testing, zero effect implies exactly 2.5 percent of positive and 97.5 percent of negative results; thus, when there are many studies, having positive results more than 2.5 percent of the time implies that the true underlying effect is positive. Yet ordinary readers will, I think, interpret 97 percent negative results not as an approval but as a disapproval of the drug.

While most discussions on publication bias have focused on biased incentives of researchers, the model of this paper, along with these thought experiments, suggest aggregation frictions may play important roles in understanding reasons why publication bias is prevalent and persistent.

Publication bias is commonly believed to contradict the unbiasedness of researchers, which has been put forth as a core ethos of science (Merton 1947). If information can be fully and costlessly communicated, then conveying all results, as they are, is what unbiased researchers must do. Yet this paper has shown that aggregation frictions can explain various kinds of publication bias. The casual expressions such as “exciting” vs “boring” results appear to suggest biases and irrationality among researchers. This paper is an attempt to provide a rational theory of “interesting results” – they are results whose binary conclusions can influence the decisions of the readers, when other results are collectively inconclusive.

The model also suggests that the publication selection process under aggregation frictions will not only differ from commonly assumed parsimonious processes, but also cannot have other parsimonious representations. This impossibility arises because (i) omission will be asymmetric between positive and negative results; (ii) inflation due to nonlinear thresholds will be difficult to address; and (iii) exact thresholds will depend critically on primitives – objectives and prior beliefs – that are unobservable to meta-analysts. Shared across commonly assumed processes and this model is the prediction that more precisely estimated studies suffer less from publication bias. This paper extends the existing methods that use arbitrary number of most precise studies (Stanley et al. 2010, Barth et al. 2013) by developing a formal method to choose an optimal number of studies to include. In this way, this paper provides a tool that has merits not only to meta-analysts who believe in different forms of publication bias, and but also to those who wish to build a consensus by relying not on contested assumptions but only on regularities common across them.

References

- Abadie, Alberto. 2018. "Statistical Non-Significance in Empirical Economics." *MIT Working Paper*
- Andrews, Isaiah, and Maximilian Kasy. 2019. "Identification of and Correction for Publication Bias" *American Economic Review*, forthcoming.
- Ashenfelter, Orley, Colm Harmon, and Hessel Oosterbeek. 1999. "A Review of Estimates of the Schooling/Earnings Relationship, with Tests for Publication Bias." *Labour Economics* 6 (4): 453–470.
- Austen-Smith, David. 1993. "Interested Experts and Policy Advice: Multiple Referrals under Open Rule." *Games and Economic Behavior* 5 (1): 3–43.
- Austen-Smith, David, and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review* 90 (1): 34–45.
- Barth, Jürgen, Thomas Munder, Heike Gerger, Eveline Nüesch, Sven Trelle, Hansjörg Znoj, Peter Jüni, and Pim Cuijpers. 2013. "Comparative Efficacy of Seven Psychotherapeutic Interventions for Patients with Depression: A Network Meta-Analysis." *PLOS Med* 10 (5): e1001454.
- Bastian, Hilda, Paul Glasziou, and Iain Chalmers. 2010. "Seventy-Five Trials and Eleven Systematic Reviews a Day: How Will We Ever Keep Up?" *PLoS Medicine* 7 (9): e1000326.
- Battaglini, Marco. 2017. "Public Protests and Policy Making." *The Quarterly Journal of Economics* 132 (1):485–549.
- Battaglini, Marco, Rebecca B Morton, and Thomas R. Palfrey. 2010. "The Swing Voter's Curse in the Laboratory." *The Review of Economic Studies* 77 (1): 61–89.
- Begg, Colin B., and Madhuchhanda Mazumdar. 1994. "Operating Characteristics of a Rank Correlation Test for Publication Bias." *Biometrics* 50 (4): 1088.
- Blanco-Perez, Cristina and Abel Brodeur. 2017. "Publication Bias and Editorial Statement on Negative Findings" Working Paper
- Blume, Andreas, and Oliver Board. 2013. "Language Barriers." *Econometrica* 81 (2): 781–812.
- Boccia, Stefania, Kenneth J. Rothman, Nikola Panic, Maria Elena Flacco, Annalisa Rosso, Roberta Pastorino, Lamberto Manzoli, et al. 2016. "Registration Practices for Observational Studies on ClinicalTrials.gov Indicated Low Adherence." *Journal of Clinical Epidemiology* 70 (February): 176–82.
- Boussageon, Rémy, Theodora Bejan-Angoulvant, Mitra Saadatian-Elahi, Sandrine

- Lafont, Claire Bergeonneau, Behrouz Kassai, Sylvie Erpeldinger, James M. Wright, François Gueyffier, and Catherine Cornu. 2011. "Effect of Intensive Glucose Lowering Treatment on All Cause Mortality, Cardiovascular Death, and Microvascular Events in Type 2 Diabetes: Meta-Analysis of Randomised Controlled Trials." *BMJ* 343 (July): d4169
- Bom, Pedro, and Heiko Rachinger. 2018. "A Kinked Meta-Regression Model for Publication Bias Correction" Working Paper.
- 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, Linda Babcock, George Loewenstein, and Richard Thaler. 1997. "Labor Supply of New York City Cabdrivers: One Day at a Time." *The Quarterly Journal of Economics* 112 (2): 407–41.
- Card, David, and Alan B. Krueger. 1995. "Time-Series Minimum-Wage Studies: A Meta-Analysis." *The American Economic Review* 85 (2): 238–243.
- Carey, Benedict. 2008. "Antidepressant Studies Unpublished." *The New York Times*, January 17. <http://www.nytimes.com/2008/01/17/health/17depress.html>.
- Casey, Katherine, Rachel Glennerster, and Edward Miguel. 2012. "Reshaping Institutions: Evidence on Aid Impacts Using a Preanalysis Plan." *The Quarterly Journal of Economics* 127 (4): 1755–1812.
- Chalmers I. 1990. "Underreporting Research Is Scientific Misconduct." *JAMA* 263 (10): 1405–8.
- Coburn, Kathleen, and Jack L. Vevea. 2017. "Package 'weightr'" *CRAN package*
- Copas, J. B., and H. G. Li. 1997. "Inference for Non-Random Samples." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 59 (1): 55–95.
- Christensen, Garret, and Edward Miguel. 2018 "Transparency, Reproducibility, and the Credibility of Economics Research." *Journal of Economic Literature*.
- Crawford, Vincent. 1998. "A Survey of Experiments on Communication via Cheap Talk." *Journal of Economic Theory* 78 (2): 286–98.
- Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.
- de Winter, Joost, and Riender Happee. 2013. "Why Selective Publication of Statistically Significant Results Can Be Effective." *PloS One* 8 (6): e66463.
- Davis, Lucas W. 2008. "The Effect of Driving Restrictions on Air Quality in Mexico City." *Journal of Political Economy* 116 (1): 38–81.
- Dekel, Eddie, and Michele Piccione. 2000. "Sequential Voting Procedures in Sym-

- metric Binary Elections.” *Journal of Political Economy* 108 (1): 34–55.
- De Long, J. Bradford, and Kevin Lang. 1992. “Are All Economic Hypotheses False?” *Journal of Political Economy* 100 (6): 1257–72.
- DerSimonian, Rebecca, and Nan Laird. 1986. “Meta-Analysis in Clinical Trials.” *Controlled Clinical Trials* 7 (3): 177–88.
- Dewan, Torun, Andrea Galeotti, Christian Ghiglino and Francesco Squintani. 2015. “Information Aggregation and Optimal Structure of the Executive.” *American Journal of Political Science*
- Dickhaut, John W., Kevin A. McCabe, and Arijit Mukherji. 1995. “An Experimental Study of Strategic Information Transmission.” *Economic Theory* 6 (3): 389–403.
- Doucouliafos, Christos, Richard Freeman, Patrice Laroche, and Tom Stanley. 2018. “How Credible is Trade Union Research? Forty Years of Evidence on the Monopoly-Voice Tradeoff.” *Industrial Relations: A Journal of Economy and Society*
- Doucouliafos, Chris, and Tom Stanley. 2013. “Are All Economic Facts Greatly Exaggerated? Theory Competition and Selectivity.” *Journal of Economic Surveys* 27 (2): 316–39.
- Doucouliafos, Christos, and Patrice Laroche. 2003. “What Do Unions Do to Productivity? A Meta-Analysis.” *Industrial Relations: A Journal of Economy and Society*
- Duval, Sue, and Richard Tweedie. 2000. “Trim and Fill: A Simple Funnel-Plot-based Method of Testing and Adjusting for Publication Bias in Meta-Analysis.” *Biometrics* 56 (2)
- Egger, Matthias, George Davey Smith, Martin Schneider, and Christoph Minder. 1997. “Bias in Meta-Analysis Detected by a Simple, Graphical Test.” *BMJ* 315 (7109): 629–634.
- Esponda, Ignacio, and Emanuel Vespa. 2014. “Hypothetical Thinking and Information Extraction in the Laboratory.” *American Economic Journal: Microeconomics*
- Every-Palmer, Susanna, and Jeremy Howick. 2014. “How Evidence-Based Medicine Is Failing due to Biased Trials and Selective Publication: EBM Fails due to Biased Trials and Selective Publication.” *Journal of Evaluation in Clinical Practice* 20 (6): 908–14.
- Fafchamps, Marcel, and Julien Labonne. 2017. “Using Split Samples to Improve Inference about Causal Effects.” *NBER Working Paper*.
- Feddersen, Timothy J. 2004. “Rational Choice Theory and the Paradox of Not Voting.” *Journal of Economic Perspectives* 18 (1): 99–112.
- Feddersen, Timothy J., and Wolfgang Pesendorfer. 1996. “The Swing Voter’s Curse.” *The American Economic Review*, 408–424.

- . 1997. “Voting Behavior and Information Aggregation in Elections With Private Information.” *Econometrica* 65 (5): 1029.
- . 1998. “Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts under Strategic Voting.” *American Political Science Review* 92 (1): 23–35.
- . 1999. “Abstention in Elections with Asymmetric Information and Diverse Preferences.” *American Political Science Review* 93 (2): 381–98.
- Forbes, Kristin J. 2000. “A Reassessment of the Relationship between Inequality and Growth.” *American Economic Review* 90 (4): 869–87.
- Foster, Andrew, Dean Karlan, and Ted Miguel 2018 “Registered Reports: Piloting a Pre-Results Review Process at the Journal of Development Economics” *Development Impact Guest Blogger*
- Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. “Publication Bias in the Social Sciences: Unlocking the File Drawer.” *Science* 345 (6203): 1502–5.
- Freeman, Richard. 1988. “Union Density and Economic Performance: An Analysis of US States.” *European Economic Review* 32 (2–3): 707–716.
- Fryer, Roland, and Matthew O. Jackson. 2008. “A Categorical Model of Cognition and Biased Decision Making.” *The BE Journal of Theoretical Economics* 8 (1).
- Fudenberg, Drew, and David K. Levine. 1998. “The Theory of Learning in Games” MIT Press.
- Furukawa, Chishio. 2016. “Unbiased Publication Bias: Theory and Evidence” *MIT Working Paper*.
- Garner, Paul, David Taylor-Robinson, and Harshpal Singh Sachdev. 2013. “DEVTA: Results from the Biggest Clinical Trial Ever.” *The Lancet* 381 (9876): 1439–41.
- Gentzkow, Matthew, Jesse Shapiro, and Daniel Stone 2016. “Chapter 14 - Media Bias in the Marketplace: Theory.” In *Handbook of Media Economics*, edited by Joel Waldfogel and David Strömberg Simon P. Anderson, 1:647–67.
- Gerber, Alan. 2008. “Do Statistical Reporting Standards Affect What Is Published? Publication Bias in Two Leading Political Science Journals.” *Quarterly Journal of Political Science* 3 (3): 313–26.
- Gerber, Alan S., and Neil Malhotra. 2008. “Publication Bias in Empirical Sociological Research: Do Arbitrary Significance Levels Distort Published Results?” *Sociological Methods & Research*.
- Gilligan, Thomas W., and Keith Krehbiel. 1987. “Collective Decisionmaking and Standing Committees: An Informational Rationale for Restrictive Amendment Procedures.” *Journal of Law, Economics, & Organization* 3 (2): 287–335.
- Glaeser, Edward L. 2008. “Researcher Incentives and Empirical Methods.” in *The*

- Foundations of Positive and Normative Economics, Andrew Caplin and Andrew Schotter
- Goldacre, Ben. 2010. “Bad Science: Quacks, Hacks, and Big Pharma Flacks.” New York: Farrar, Straus and Giroux.
- Hanna, Rema, Esther Duflo, and Michael Greenstone. 2016. “Up in Smoke: The Influence of Household Behavior on the Long-Run Impact of Improved Cooking Stoves.” *American Economic Journal: Economic Policy* 8 (1): 80–114.
- Havránek, Tomáš. 2015. “Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting.” *Journal of the European Economic Association* 13 (6): 1180–1204.
- Havránek, Tomáš, Zuzana Irsova, Karel Janda, and David Zilberman. 2015. “Selective Reporting and the Social Cost of Carbon.” *Energy Economics* 51 (September): 394–406.
- Hay, Michael, David W. Thomas, John L. Craighead, Celia Economides, and Jesse Rosenthal. 2014. “Clinical Development Success Rates for Investigational Drugs.” *Nature Biotechnology* 32 (1): 40–51.
- Hedges, Larry V. 1992. “Modeling Publication Selection Effects in Meta-Analysis.” *Statistical Science*, 246–255.
- Henry, Emeric. 2009. “Strategic Disclosure of Research Results: The Cost of Proving Your Honesty.” *The Economic Journal* 119 (539): 1036–64.
- Henry, Emeric, and Marco Ottaviani. 2014. “Research and the Approval Process.” mimeo.
- Hunter, John E., and Frank L. Schmidt. 2004. *Methods of Meta-Analysis: Correcting Error and Bias in Research Findings*. 2nd edition. Thousand Oaks, Calif: SAGE Publications, Inc.
- ICMJE. 2016. “Recommendations for the Conduct, Reporting, Editing, and Publication of Scholarly Work in Medical Journals.” <http://www.icmje.org/recommendations/>.
- Ioannidis, John P. A. 2005. “Why Most Published Research Findings Are False.” *PLoS Medicine* 2 (8): e124.
- John, K. Leslie, George Loewenstein, and Drazen Prelec. 2012. “Measuring the Prevalence of Questionable Research Practices With Incentives for Truth Telling.” *Psychological Science* 23 (5): 524–32.
- Kawai, Kei, and Yasutora Watanabe 2013. “Inferring Strategic Voting.” *American Economic Review*.
- Krishna, Vijay, and John Morgan. 2001. “A Model of Expertise.” *The Quarterly Journal of Economics* 116 (2): 747–75.

- Leamer, Edward E. 1978. *Specification Searches: Ad Hoc Inference with Nonexperimental Data*. 1 edition. New York: Wiley.
- Lexchin, Joel, Lisa A. Bero, Benjamin Djulbegovic, and Otavio Clark. 2003. "Pharmaceutical Industry Sponsorship and Research Outcome and Quality: Systematic Review." *BMJ* 326 (7400): 1167–70.
- Li, Hao, Sherwin Rosen, and Wing Suen. 2001. "Conflicts and Common Interests in Committees." *American Economic Review* 91 (5): 1478–97.
- Libgober, Jonathan. 2015. "False Positives in Scientific Research." Harvard University mimeo.
- Manski, Charles F. 2004. "Statistical Treatment Rules for Heterogeneous Populations." *Econometrica* 72 (4): 1221–46.
- McCrary, Justin, Garret Christensen, and Daniele Fanelli. 2016. "Conservative Tests under Satisficing Models of Publication Bias." Edited by Daniele Marinazzo. *PLOS ONE* 11 (2): e0149590.
- McShane, Blakeley B., and David Gal. 2017. "Statistical Significance and the Dichotomization of Evidence." *Journal of the American Statistical Association* 112 (519):885–95.
- Merton, Robert K. 1942 "The Normative Structure of Science", in the *Sociology of Science: Theoretical and Empirical Investigations*, Chicago: University of Chicago Press
- Miguel, Edward, Colin Camerer, Katherine Casey, Joshua Cohen, Kevin M. Esterling, Alan Gerber, Rachel Glennerster, et al. 2014. "Promoting Transparency in Social Science Research." *Science* 343 (6166): 30–31.
- Morgan, John, and Phillip C. Stocken. 2008. "Information Aggregation in Polls." *American Economic Review* 98 (3): 864–96.
- Munafo, M. R., I. J. Matheson, and J. Flint. 2007. "Association of the DRD2 Gene Taq1A Polymorphism and Alcoholism: A Meta-Analysis of Case-control Studies and Evidence of Publication Bias." *Molecular Psychiatry* 12 (5): 454–461.
- Newey, Whitney K., and Daniel McFadden. 1994. "Chapter 36 Large Sample Estimation and Hypothesis Testing." in *Handbook of Econometrics*, 4:2111–2245.
- Olken, Benjamin A. 2015. "Promises and Perils of Pre-Analysis Plans." *Journal of Economic Perspectives* 29 (3): 61–80.
- Powell Bingham. 1986. "American Voter Turnout in Comparative Perspective." *American Political Science Review*.
- Puglisi, Riccardo, and James M. Snyder Jr. 2016. "Chapter 15 - Empirical Studies of Media Bias." In *Handbook of Media Economics*, edited by Joel Waldfogel and David

- Strömberg Simon P. Anderson, 1:647–67.
- Reinhart, Alex. 2015. *Statistics Done Wrong: The Woefully Complete Guide*. 1 edition. San Francisco: No Starch Press.
- Rosenthal, Robert. 1979. “The File Drawer Problem and Tolerance for Null Results.” *Psychological Bulletin* 86 (3): 638.
- Sánchez-Pagés, Santiago, and Marc Vorsatz. 2007. “An Experimental Study of Truth-Telling in a Sender–receiver Game.” *Games and Economic Behavior* 61 (1): 86–112.
- Shapiro, Jesse M. 2016. “Special Interests and the Media: Theory and an Application to Climate Change.” *Journal of Public Economics*
- Siegfried, John J. 2012. “Minutes of the Meeting of the Executive Committee Chicago, IL January 5, 2012.” *American Economic Review* 102 (3): 645–52.
- Silverman, Bernard W. 1998. *Density Estimation for Statistics and Data Analysis*. Routledge.
- Simonsohn, Uri, Leif D. Nelson, and Joseph P. Simmons. 2014. “P-Curve: A Key to the File-Drawer.” *Journal of Experimental Psychology: General* 143 (2): 534–47.
- Stanley T.D. 2005 “Beyond Publication Bias” in “Meta-Regression Analysis: Issues of Publication Bias in Economics edited” by Collin Roberts and T. D. Stanley
- Stanley, T. D., Stephen B. Jarrell, and Hristos Doucouliagos. 2010. “Could It Be Better to Discard 90% of the Data? A Statistical Paradox.” *The American Statistician* 64 (1): 70–77.
- Suen, Wing. 2004. “The Self-Perpetuation of Biased Beliefs.” *The Economic Journal* 114 (495): 377–396
- Tetenov, Aleksey. 2016. “An Economic Theory of Statistical Testing.” *Working Paper, University of Bristol*
- The Consensus Project. 2014. “Why we need to talk about the scientific consensus on climate change”
<http://www.skepticalscience.com/news.php?f=why-we-need-to-talk-about-climate-consensus>
- Turner, Erick H., Annette M. Matthews, Eftihia Linardatos, Robert A. Tell, and Robert Rosenthal. 2008. “Selective Publication of Antidepressant Trials and Its Influence on Apparent Efficacy.” *New England Journal of Medicine* 358 (3): 252–60.
- van Assen, Marcel A. L. M., Robbie C. M. van Aert, Michèle B. Nuijten, and Jelte M. Wicherts. 2014. “Why Publishing Everything Is More Effective than Selective Publishing of Statistically Significant Results.” Edited by K. Brad Wray. *PLoS ONE* 9 (1): e84896.
- van Assen, Marcel A. L. M., Robbie C. M. van Aert, and Jelte M. Wicherts. 2015.

“Meta-Analysis Using Effect Size Distributions of Only Statistically Significant Studies.” *Psychological Methods* 20 (3): 293–309.

Vevea, Jack L., and Larry V. Hedges. 1995 “A General Linear Model for Estimating Effect Size in the Presence of Publication Bias.” *Psychometrika* 60, no. 3: 419–35.

Viechtbauer, Wolfgang. 2010. "Conducting Meta-Analyses in R with the metafor Package." *Journal of Statistical Software* Volume 36, Issue 3.

Wasserstein, Ronald L., and Nicole A. Lazar. 2016. “The ASA’s Statement on P-Values: Context, Process, and Purpose.” *The American Statistician* 70 (2): 129–33.

Ziliak, Stephen T., and Deirdre N. McCloskey. 2008. *The Cult of Statistical Significance: How the Standard Error Costs Us Jobs, Justice, and Lives*. 1st edition. Ann Arbor: University of Michigan Press.

Appendix A. Proofs

Appendix A presents the proofs of propositions and some additional analytical results. Appendix A1 presents some preliminary results to prepare for the main analyses; Appendix A2 presents proofs of propositions in Section 2.3; Appendix A3 presents the proof of Section 2.4; Appendix A4 presents the proof of Section 2.5; Appendix A5 presents an example in Section 2.5; Appendix A6 presents the proof of Section 4.1. For notational ease, let us denote $\pi(n_1, n_0)$ as the policymaker’s strategy given the number of positive and negative results.

A1. Preliminaries

We begin by proving three lemmas that will be relevant throughout the proofs: with normally distributed random variable, (1) conditional mean of will be increasing in the mean of its underlying distribution, (2) higher conditional mean implies that the likelihood ratio will be increasing in the mean of its underlying distribution, and (3) strategies will be monotone as in Lemma 1 in any fully responsive and fully informative equilibria. While these properties need not hold in general, normal distribution imposes sufficient structure to facilitate the analyses of the model.

A1.1. Monotonicity of Conditional Mean

Lemma A1 will show that the conditional mean of normally distributed random variable with any arbitrary condition will equal the ratio of conditional variance to total

variance. This is a generalization of the proof for truncated normal distribution by Alecos Papadopoulos (2013).

Lemma A1. Derivative of conditional mean with respect to unconditional mean. *Given any $s_m(\beta) \equiv \mathbb{P}(s = 1|\beta)$, the derivative of conditional mean $\mathbb{E}[\beta s | s_m(\beta)]$ of $\beta \sim \mathcal{N}(\mu, \sigma^2)$ with respect to its mean μ satisfies*

$$\frac{\partial \mathbb{E}[\beta s | s_m(\beta)]}{\partial \mu} = \frac{Var_m}{\sigma^2}, \quad (1.17)$$

where $Var_m \equiv \mathbb{E}\{\beta - \mathbb{E}[\beta s | s_m(\beta)]\}^2$ is the variance of the random variable conditional on the message.

Proof. By applying the property of density of normal distribution. We will first express the conditional mean, then take the derivative by the chain rule, and finally reorganize the expression to see that (1.17) holds for any $s(\beta)$.

First, by definition, we have

$$\mathbb{E}[\beta s | s_m(\beta)] = \frac{f_1(\mu)}{f_2(\mu)},$$

where $f_1(\mu) \equiv \int \beta s_m(\beta) \phi\left(\frac{\beta - \mu}{\sigma}\right) d\beta$, $f_2(\mu) \equiv \int s_m(\beta) \phi\left(\frac{\beta - \mu}{\sigma}\right) d\beta$, and $\phi(\cdot)$ is the density of standard normal distribution.

Second, we can apply the chain rule to obtain

$$\frac{\partial \mathbb{E}[\beta s | s_m(\beta)]}{\partial \mu} = \frac{f_1'(\mu) f_2(\mu) - f_1(\mu) f_2'(\mu)}{[f_2(\mu)]^2}, \quad (1.18)$$

where

$$f_1'(\mu) = -\frac{1}{\sigma} \int \beta s_m(\beta) \phi'\left(\frac{\beta - \mu}{\sigma}\right) d\beta \text{ and } f_2'(\mu) = -\frac{1}{\sigma} \int s_m(\beta) \phi'\left(\frac{\beta - \mu}{\sigma}\right) d\beta.$$

Third, using the property of normal density that $\phi'\left(\frac{\beta - \mu}{\sigma}\right) = -\frac{\beta - \mu}{\sigma} \phi\left(\frac{\beta - \mu}{\sigma}\right)$, we can reorganize them as

$$\begin{aligned} f_1'(\mu) &= \frac{1}{\sigma^2} \int \beta s_m(\beta) (\beta - \mu) \phi\left(\frac{\beta - \mu}{\sigma}\right) d\beta = \frac{f_3(\mu) - \mu f_1(\mu)}{\sigma^2} \\ f_2'(\mu) &= \frac{1}{\sigma^2} \int s_m(\beta) (\beta - \mu) \phi\left(\frac{\beta - \mu}{\sigma}\right) d\beta = \frac{f_1(\mu) - \mu f_2(\mu)}{\sigma^2} \end{aligned}$$

where $f_3(\mu) \equiv \int \beta^2 s_m(\beta) \phi\left(\frac{\beta-\mu}{\sigma}\right) d\beta$. By substituting into the condition (1.18),

$$\begin{aligned}
\frac{\partial \mathbb{E}[\beta s | s_m(\beta)]}{\partial \mu} &= \frac{1}{\sigma^2} \frac{(f_3 - \mu f_1) f_2 - f_1 (f_1 - \mu f_2)}{f_2^2} \\
&= \frac{1}{\sigma^2} \frac{f_3 f_2 - f_1^2}{f_2^2} \\
&= \frac{1}{\sigma^2} \left[\frac{f_3}{f_2} - \left(\frac{f_1}{f_2} \right)^2 \right] \\
&= \frac{1}{\sigma^2} \left\{ \mathbb{E}[\beta^2 s_m | s(\beta)] - (\mathbb{E}[\beta s_m | s(\beta)])^2 \right\} \\
&= \frac{Var_m}{\sigma^2}
\end{aligned}$$

where the last line followed by the definition of variance. □

A1.2. Monotonicity of Mean Likelihood Ratios

Lemma A2. will show that the messages with higher conditional mean will also be more likely to be sent when the mean of underlying distribution increases. This monotonicity of likelihood ratio is not equivalent to the standard Monotone Likelihood Ratio Property of normal distribution with known variance, since the standard statement is concerned with each value whereas the following lemma addresses the average value. This property is key to deriving the Lemma 1 monotonicity of equilibrium strategies: the analogue of Lemma 1 will not hold when the standard errors are heterogeneous and unknown.

Lemma A2. Equivalence of change in likelihood ratio and mean ranking.

Consider two strategies, $s_m(\beta)$ and $s_{\tilde{m}}(\beta)$, given $\beta \sim \mathcal{N}(\mu, \sigma^2)$ and the associated likelihood ratio of each messages, $LR(\mu) \equiv \frac{\mathbb{P}(m|\mu)}{\mathbb{P}(\tilde{m}|\mu)}$. Then, $LR'(\mu) > 0$ if and only if $\mathbb{E}[\beta s | s_m(\beta)] > \mathbb{E}[\beta s | s_{\tilde{m}}(\beta)]$.

Proof. By definition,

$$LR(\mu) \equiv \frac{\int s_m(\beta) \phi\left(\frac{\beta-\mu}{\sigma}\right) d\beta}{\int s_{\tilde{m}}(\beta) \phi\left(\frac{\beta-\mu}{\sigma}\right) d\beta}$$

By chain rule and the property of normal density that $\phi' \left(\frac{\beta - \mu}{\sigma} \right) = -\frac{\beta - \mu}{\sigma} \phi \left(\frac{\beta - \mu}{\sigma} \right)$,

$$\begin{aligned}
LR'(\mu) &\equiv -\frac{1}{\sigma} \frac{\int s_m \phi' \times \int s_{\tilde{m}} \phi - \int s_m \phi \times \int s_{\tilde{m}} \phi'}{\left(\int s_{\tilde{m}} \phi \right)^2} \\
&= \frac{\int [\beta - \mu] s_m \phi \times \int s_{\tilde{m}} \phi - \int s_m \phi \times \int [\beta - \mu] s_{\tilde{m}} \phi}{\left(\sigma \int s_{\tilde{m}} \phi \right)^2} \\
&= \frac{\int \beta s_m \phi \times \int s_{\tilde{m}} \phi - \int s_m \phi \times \int \beta s_{\tilde{m}} \phi}{\left(\sigma \int s_{\tilde{m}} \phi \right)^2}
\end{aligned}$$

Therefore,

$$\begin{aligned}
LR'(\mu) > 0 &\Leftrightarrow \int \beta s_m \phi \times \int s_{\tilde{m}} \phi > \int s_m \phi \times \int \beta s_{\tilde{m}} \phi \\
&\Leftrightarrow \frac{\int \beta s_m \phi}{\int s_m \phi} > \frac{\int \beta s_{\tilde{m}} \phi}{\int s_{\tilde{m}} \phi} \\
&\Leftrightarrow \mathbb{E}[\beta s | s_m(\beta)] > \mathbb{E}[\beta s | s_{\tilde{m}}(\beta)]
\end{aligned}$$

where the last line followed by the definition of conditional mean. \square

A1.3. Monotonicity of Equilibrium Strategies

Lemma 1 in Section 2.2 claims that, for any c and $\sigma_i = \sigma$, the strategies will be monotone if the equilibrium is fully responsive and fully informative: (i) researchers will apply threshold strategies and (ii) the policymaker's probability of policy implementation will be increasing in positive results and decreasing in negative results.

Proof. Since the result of monotonicity of policymaker's strategy will be used for that of researcher's strategy, we will first derive the result of policymaker's and then that of researchers'.

(i) Policymaker's strategy: suppose $\pi^*(n_1, n_0) > 0$ for some n_1, n_0 . Then, by the policymaker's optimization condition, $\mathbb{E}[b | n_1, n_0] \geq c$. By full informativeness and Bayes' rule, $\mathbb{E}[b | n_1 + k, n_0] > \mathbb{E}[b | n_1, n_0]$ for $k \in \{1, 2\}$ and thus, $\mathbb{E}[b | n_1 + k, n_0] \geq c$. By the policymaker's optimization, $\pi^*(n_1 + k, n_0) > 0$. Analogous argument holds for $\pi^*(n_1, n_0) < 1 \Rightarrow \pi^*(n_1, n_0 + k) = 0$.

(ii) Researchers' strategies: by the Bayes' rule, full responsiveness, and domain of signals. The proof consists of three steps: the first step organizes the indifference conditions, and the second step shows the existence of solution, and the third step shows the uniqueness.

Step 1. indifference conditions: given any policymaker's strategy $\pi(m_i, m_{-i})$ and

another researcher's strategy $s(\beta_{-i})$, the expected welfare of reporting message m_i given the signal β_i can be written as

$$W(m_i, \beta_i) = \sum_{m_{-i}} \pi(m_i, m_{-i}) \mathbb{P}(m_{-i}|\beta_i) \left\{ \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} [\beta_i + \mathbb{E}[\beta_{-i}|m_{-i}, \beta_i]] - c \right\}$$

by the objective (3.1).

At some thresholds, $\bar{\beta}$ and $\underline{\beta}$, in which the researcher will be willing to switch the messages, the indifference conditions $W(1, \bar{\beta}) = W(\emptyset, \bar{\beta})$ and $W(0, \underline{\beta}) = W(\emptyset, \underline{\beta})$ must be satisfied. Rewriting, for each threshold, the conditions are,

$$I(1, \emptyset, \bar{\beta}) = 0 \text{ and } I(\emptyset, 0, \underline{\beta}) = 0,$$

where

$$I(m_i, m'_i, \beta_i) \equiv \sum_{m_{-i}} p(m_i, m'_i|m_{-i}) q(m_{-i}|\beta_i) r(\beta_i|m_{-i}),$$

where

$$\begin{aligned} p(m_i, m'_i|m_{-i}) &\equiv \pi(m_i, m_{-i}) - \pi(m'_i, m_{-i}) \\ q(m_{-i}|\beta_i) &\equiv \mathbb{P}(m_{-i}|\beta_i) \\ r(\beta_i|m_{-i}) &\equiv \beta_i + \mathbb{E}[\beta_{-i}|m_{-i}, \beta_i] - c \left(2 + \frac{\sigma^2}{\sigma_b^2} \right) \end{aligned}$$

By full responsiveness, there must exist some m_{-i}, m'_{-i} such that $\pi(1, m_{-i}) > \pi(\emptyset, m_{-i})$ and $\pi(\emptyset, m'_{-i}) > \pi(0, m'_{-i})$. Thus, these conditions are not vacuous. The meaning of messages is, without loss of generality, assigned to be consistent with the set-up.

Step 2. existence: for all m_{-i} , another researcher's strategy, $s(\beta_{-i})$, by Lemma A1, there exist some β'_i such that $r(\beta'_i|m_{-i}) < 0$ and some other β''_i such that $r(\beta''_i|m_{-i}) > 0$. Since $p(1, \emptyset|m_{-i}) > 0$ and $q(m_{-i}|\beta_i) > 0$ and $q(\cdot)$ and $r(\cdot)$ are continuous functions of β_i , there must exist some $\bar{\beta}$ and $\underline{\beta}$ that satisfy the indifference condition by intermediate value theorem.

Step 3. uniqueness: to show that there is a unique value that satisfies an indifference condition, we first show that $p(m_i, m'_i|m_{-i}) = 0$ for at least 1 m_{-i} and then show $\partial I(m_i, m'_i, \beta_i) / \partial \beta_i > 0$ when evaluated at $I(m_i, m'_i, \beta_i) = 0$ so that there is a unique value of β_i that satisfies this.

- $p(m_i, m'_i|m_{-i}) = 0$ for at least 1 m_{-i} : first, note that $\pi(1, 1) = 1$ and $\pi(0, 0) = 0$. To see why, suppose $\pi(1, 1) < 1$. Then by policymaker's optimization,

$\pi(1, \emptyset) = \pi(\emptyset, 1) = 0$, which then implies $\pi(0, m_{-i}) = 0$ for all m_{-i} and $\pi(m_i, 0) = 0$ for all m_i , contradicting full responsiveness. Second, we consider three cases of $\pi(\emptyset, \emptyset)$:

- when $\pi(\emptyset, \emptyset) = 1$, $\pi(1, \emptyset) = \pi(\emptyset, 1) = 1$ by policymaker's monotonicity. Thus, $p(1, \emptyset|\emptyset) = p(1, \emptyset|1) = 0$. Moreover, by full responsiveness for another researcher, either $\{\pi(0, \emptyset), \pi(0, 1)\} = \{0, \bar{\pi}\}$ with $\bar{\pi} > 0$ or $\{\pi(0, \emptyset), \pi(0, 1)\} = \{\bar{\pi}, 1\}$ with $\bar{\pi} < 1$. If former, $p(\emptyset, 0|0) = 0$ another researcher and if later, $p(\emptyset, 0|1) = 0$ for the researcher himself. In this way, $p(m_i, m'_i|m_{-i}) = 0$ for at least 1 m_{-i} for both $\{m_i, m'_i\} \in \{\{1, \emptyset\}, \{\emptyset, 0\}\}$ in any fully responsive equilibria.
 - when $\pi(\emptyset, \emptyset) = 0$, a symmetric argument analogous to above applies.
 - when $\pi(\emptyset, \emptyset) = \bar{\pi}$ for $\bar{\pi} \in (0, 1)$, $\pi(1, \emptyset) = \pi(\emptyset, 1) = 1$ and $\pi(0, \emptyset) = \pi(\emptyset, 0) = 0$ by policymaker's monotonicity. Thus, $p(1, \emptyset|1) = 0$ and $p(\emptyset, 0|0) = 0$.
- $\partial I(m_i, m'_i, \beta_i) / \partial \beta_i > 0$ at $I(m_i, m'_i, \beta_i) = 0$: for each $\{m_i, m'_i\} \in \{\{1, \emptyset\}, \{\emptyset, 0\}\}$, let us consider two cases:

- when $p(m_i, m'_i|m_{-i}) = 0$ for 2 values of m_{-i} : denoting m_{-i}^* as the value such that $p(m_i, m'_i|m_{-i}^*) > 0$, the indifference condition is $p(m_i, m'_i|m_{-i}^*) r(\beta_i|m_{-i}^*) = 0$. $r(\beta_i|m_{-i})$ is strictly increasing.
- when $p(m_i, m'_i|m_{-i}) = 0$ for only 1 value of m_{-i} : denoting m_{-i}^*, m_{-i}^{**} as the value such that $p(m_i, m'_i|m_{-i}) > 0$,

$$p(m_i, m'_i|m_{-i}^*) \tilde{q}(m_{-i}^*|\beta_i) r(\beta_i|m_{-i}^*) + p(m_i, m'_i|m_{-i}^{**}) \tilde{q}(m_{-i}^{**}|\beta_i) r(\beta_i|m_{-i}^{**}) = 0,$$

where $\tilde{q}(m_{-i}|\beta_i) \equiv \frac{q(m_{-i}|\beta_i)}{q(m_{-i}^*|\beta_i) + q(m_{-i}^{**}|\beta_i)}$ is the normalized probability. The derivative of the indifference condition with respect to β_i is

$$p(m_i, m'_i|m_{-i}^*) \tilde{q}(m_{-i}^*|\beta_i) r'(\beta_i|m_{-i}^*) + p(m_i, m'_i|m_{-i}^{**}) \tilde{q}(m_{-i}^{**}|\beta_i) r'(\beta_i|m_{-i}^{**}) \\ + p(m_i, m'_i|m_{-i}^*) \tilde{q}'(m_{-i}^*|\beta_i) r(\beta_i|m_{-i}^*) + p(m_i, m'_i|m_{-i}^{**}) \tilde{q}'(m_{-i}^{**}|\beta_i) r(\beta_i|m_{-i}^{**})$$

* by Lemma A1, $r'(\beta_i|m_{-i}^*) > 0$ and $r'(\beta_i|m_{-i}^{**}) > 0$.

* by Lemma A1 and full responsiveness, $r(\beta_i|m_{-i}^*) < 0$ and $r(\beta_i|m_{-i}^{**}) > 0$ must hold at the indifference condition $I(m_i, m'_i, \beta_i) = 0$ since all other terms are positive (the meaning of m_{-i}^*, m_{-i}^{**} is without loss of

generality.) By Lemma A2, $\tilde{q}'(m_{-i}^*|\beta_i) < 0$ and $\tilde{q}'(m_{-i}^{**}|\beta_i) > 0$: higher mean implies higher relative likelihood of message m_{-i}^{**} sent by another researcher.

Since all terms are thus positive, $\partial I(m_i, m'_i, \beta_i) / \partial \beta_i > 0$ at $I(m_i, m'_i, \beta_i) = 0$.

Since $p(m_i, m'_i | m_{-i}) = 0$ at least for 1 value of m_{-i} , we saw that the indifference condition must be increasing in β_i when it is satisfied so that the solution will be unique.

□

A2. Proofs of 2.3 Omission of Insignificant Results

This sub-Section presents the proofs of propositions in Section 2.2. A2.1 proves Proposition 1.1; A2.2 proves Proposition 1.2; A2.3 proves Proposition 1.3.

A2.1. Proof of Proposition 1.1

Proposition 1.1 claims that there exists an equilibrium in which the policymaker adopts a supermajoritarian decision rule and the researchers apply a threshold that is asymmetric such that the estimates underlying reported studies will have an upward bias. The proof will show first that the policymaker's strategy is a part of the equilibrium, and second the researchers' strategies are also a part of the equilibrium.

(i) Policymaker's strategy: If the decision rule (1.2) is consistent with policymaker's optimization, we need that $\mathbb{E}[b | n_1 > n_0] \geq 0$ and $\mathbb{E}[b | n_1 \leq n_0] \leq 0$ given thresholds (1.3). This holds because the researchers' strategy must satisfy the indifference condition at the margin whereas the policymaker assess whether the condition holds on average. To see why $\mathbb{E}[b | n_1 > n_0] \geq 0$, note that $\mathbb{E}[b | n_1 = 1, n_0 = 0] \geq 0$ since

$$\begin{aligned} \mathbb{E}[b | m_1 = 1, m_2 = \emptyset] &= \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} \mathbb{E}[\beta_1 + \mathbb{E}[\beta_2 | \bar{\beta} > \beta_2 \geq \underline{\beta}, \beta_1] | \beta_1 \geq \bar{\beta}] \\ &> \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} \{ \bar{\beta} + \mathbb{E}[\beta_2 | \bar{\beta} > \beta_2 \geq \underline{\beta}, \beta_1 = \bar{\beta}] \} \\ &= 0, \end{aligned}$$

where the last equality holds due to the researchers' indifference condition. Analogous arguments also hold for $\mathbb{E}[b \mid n_1 \leq n_0] \leq 0$, showing (1.2) is an equilibrium.

(ii) Researchers' strategy: given the supermajoritarian voting rule in (1.2), the equilibrium thresholds will (1) be unique and symmetric between researchers, (2) lead to some omissions ($\bar{\beta} \geq 0 > \underline{\beta}$), (3) be asymmetric between the thresholds for positive vs negative results ($\bar{\beta} < -\underline{\beta}$), so that together will (4) have an upward bias of reported studies $\mathbb{E}[\beta_i \mid m_i \neq \emptyset] > 0$. The following proof shows these results in turn.

- (1) uniqueness of symmetric solution: given π^* as in (1.2), the thresholds will be unique since the researchers' strategies are moderate strategic substitutes of one another. When the best response function satisfies $\frac{\partial \bar{\beta}_i(\bar{\beta}_j)}{\partial \bar{\beta}_j} \in (-1, 0)$, there can be at most one value that satisfies the equilibrium conditions and will be symmetric between researchers so that $\bar{\beta}_i = \bar{\beta}_j = \bar{\beta}$ and $\underline{\beta}_i = \underline{\beta}_j = \underline{\beta}$. The following Lemma A3 shows $\frac{\partial \bar{\beta}_i(\bar{\beta}_j)}{\partial \bar{\beta}_j} \in (-1, 0)$:

Lemma A3. (Moderate Strategic Substitution). *Define $\bar{\beta}_i(\bar{\beta}_j)$ as the best response to some threshold $\bar{\beta}_j$ that satisfies the equilibrium conditions:*

$$\begin{aligned}\bar{\beta}_i + \mathbb{E}[\beta_j \mid \beta_j > \bar{\beta}_j \geq \underline{\beta}_j, \beta_i = \bar{\beta}_i] &= 0 \\ \underline{\beta}_j + \mathbb{E}[\beta_i \mid \beta_i \geq \bar{\beta}_i, \beta_j = \underline{\beta}_j] &= 0\end{aligned}$$

Then

$$-1 < \frac{\partial \bar{\beta}_i(\bar{\beta}_j)}{\partial \bar{\beta}_j} < 0$$

Proof. By totally differentiating the equilibrium conditions. Writing $\bar{K}_j = \frac{\partial \mathbb{E}[\beta_j \mid \bar{\beta}_j > \beta_j \geq \underline{\beta}_j, \beta_i = \bar{\beta}_i]}{\partial \bar{\beta}_j}$, $\underline{K}_j = \frac{\partial \mathbb{E}[\beta_j \mid \bar{\beta}_j > \beta_j \geq \underline{\beta}_j, \beta_i = \bar{\beta}_i]}{\partial \beta_j}$, $\bar{K}_i = \frac{\partial \mathbb{E}[\beta_j \mid \bar{\beta}_j > \beta_j \geq \underline{\beta}_j, \beta_i = \bar{\beta}_i]}{\partial \bar{\beta}_i}$, and $\underline{L}_j = \frac{\partial \mathbb{E}[\beta_i \mid \beta_i \geq \bar{\beta}_i, \beta_j = \underline{\beta}_j]}{\partial \bar{\beta}_j}$ and $\bar{L}_i = \frac{\partial \mathbb{E}[\beta_i \mid \beta_i \geq \bar{\beta}_i, \beta_j = \underline{\beta}_j]}{\partial \bar{\beta}_i}$, the system of derivatives satisfy:

$$\begin{bmatrix} 1 + \bar{K}_i & \underline{K}_j \\ \bar{L}_i & 1 + \underline{L}_j \end{bmatrix} \begin{bmatrix} \frac{\partial \bar{\beta}_i}{\partial \bar{\beta}_j} \\ \frac{\partial \underline{\beta}_j}{\partial \bar{\beta}_j} \end{bmatrix} = \begin{bmatrix} -\bar{K}_j \\ 0 \end{bmatrix}$$

Inverting the matrix, we have

$$\begin{bmatrix} \frac{\partial \bar{\beta}_i}{\partial \bar{\beta}_j} \\ \frac{\partial \underline{\beta}_j}{\partial \bar{\beta}_j} \end{bmatrix} = \frac{1}{(1 + \bar{K}_i)(1 + \underline{L}_j) - \underline{K}_j \bar{L}_i} \begin{bmatrix} 1 + \underline{L}_j & -\bar{L}_i \\ -\underline{K}_j & 1 + \bar{K}_i \end{bmatrix} \begin{bmatrix} -\bar{K}_j \\ 0 \end{bmatrix}$$

Thus,

$$\frac{\partial \bar{\beta}_i}{\partial \bar{\beta}_j} = \frac{-\bar{K}_j (1 + \underline{L}_j)}{(1 + \bar{K}_i) (1 + \underline{L}_j) - \underline{K}_j \bar{L}_i}$$

By the definition of truncated distribution, $\bar{K}_i + \bar{K}_j + \underline{K}_j = 1$ and $\bar{L}_i + \underline{L}_j = 1$. Moreover, all terms, $\bar{K}_j, \underline{K}_j, \bar{K}_i, \underline{L}_j, \bar{L}_i$, are positive and less than 1 by Lemma A1. Thus,

$$\begin{aligned} -\frac{\partial \bar{\beta}_i}{\partial \bar{\beta}_j} &< -1 \text{ since } \bar{K}_j (1 + \underline{L}_j) < (1 + \bar{K}_i) (1 + \underline{L}_j) - \underline{K}_j \bar{L}_i \Leftrightarrow (1 - \bar{K}_i - \bar{K}_j) (1 - \underline{L}_j) < \\ &\quad (1 + \bar{K}_i - \bar{K}_j) (1 + \underline{L}_j) \\ -\frac{\partial \bar{\beta}_i}{\partial \bar{\beta}_j} &< 0 \text{ since } (1 + \bar{K}_i) (1 + \underline{L}_j) > \underline{K}_j \bar{L}_i. \quad \blacksquare \end{aligned}$$

Given that $\bar{\beta}_i = \bar{\beta}_j = \bar{\beta}$ and $\underline{\beta}_i = \underline{\beta}_j = \underline{\beta}$, we will be able to substitute the threshold values to derive the results.

- (2) omission $\bar{\beta} > 0 > \underline{\beta}$: towards contradiction, suppose $\bar{\beta} \leq 0$. By the indifference condition (1.4) and by (1) $\bar{\beta}_i = \bar{\beta}_j = \bar{\beta}$, $\bar{\beta} = -\mathbb{E} [\beta_{-i} | \beta_{-i} \in [\underline{\beta}, \bar{\beta}], \beta_i = \bar{\beta}] > 0$ since $\mathbb{E} [\beta_{-i} | \beta_{-i} < \bar{\beta}] < 0$ regardless of other conditions. Because this contradicts the assumption, $\bar{\beta} \geq 0$. Substituting this into (1.5), $\underline{\beta} = -\mathbb{E} [\beta_{-i} | \beta_{-i} > \bar{\beta}, \beta_i = \underline{\beta}] < 0$.
- (3) asymmetry $\bar{\beta} < -\underline{\beta}$: towards contradiction, suppose $\bar{\beta} \geq -\underline{\beta}$ given $\bar{\beta} \geq 0$. Then, $\bar{\beta} = -\mathbb{E} [\beta_{-i} | \beta_{-i} \in [\underline{\beta}, \bar{\beta}), \beta_i = \bar{\beta}] < 0$ because the combinations of conditions $\beta_{-i} \in [\underline{\beta}, \bar{\beta})$ by (1) and $\beta_i = \bar{\beta}$ imply $\mathbb{E} [\beta_{-i} | \beta_{-i} \in [\underline{\beta}, \bar{\beta}), \beta_i = \bar{\beta}] > 0$. Since this contradicts the assumption, $\bar{\beta} < -\underline{\beta}$ must hold.
- (5) bias of estimates underlying reported studies $\mathbb{E} [\beta_i | m_i \neq \emptyset] > 0$: the following algebraic argument formally shows that the asymmetry in (1.3) leads to the upward bias:

$$\begin{aligned} &\mathbb{E} [\beta_i | m_i \neq \emptyset] \\ &= \mathbb{P} [m_i = 1 | m_i \neq \emptyset] \mathbb{E} [\beta_i | m_i = 1] + \mathbb{P} [m_i = 0 | m_i \neq \emptyset] \mathbb{E} [\beta_i | m_i = 0] \\ &= \frac{1 - \Phi(\bar{\beta})}{1 - \Phi(\bar{\beta}) + \Phi(\underline{\beta})} \sqrt{\sigma^2 + \sigma_b^2} \frac{\phi(\bar{\beta})}{1 - \Phi(\bar{\beta})} - \frac{\Phi(\underline{\beta})}{1 - \Phi(\bar{\beta}) + \Phi(\underline{\beta})} \sqrt{\sigma^2 + \sigma_b^2} \frac{\phi(\underline{\beta})}{\Phi(\underline{\beta})} \\ &= \sqrt{\sigma^2 + \sigma_b^2} \frac{\phi(\bar{\beta}) - \phi(\underline{\beta})}{1 - \Phi(\bar{\beta}) + \Phi(\underline{\beta})} > 0 \end{aligned}$$

Since both the policymaker and researchers' strategies satisfy the indifference conditions given the strategy of one another, and beliefs are consistent with the Bayes' rule, the strategies in Proposition 1.1. constitutes an equilibrium. \square

A2.2. Proof of Proposition 1.2

Proposition 1.2 claims that there are both (1) an equilibrium with symmetric omission with policymaker's decision rule, $\pi(n_0 = n_1) = \frac{1}{2}$, and (2) an equilibrium with no omission with policymaker's decision rule, $\pi(n_1 = 2) = 1$ and $\pi(n_1 = 1, n_0 = 0) \in (0, 1]$. We will prove this for the one with (1) symmetric omission, and then with (2) no omission.

(1) Proof for equilibrium with symmetric omission. The policymaker's strategy is an equilibrium because, given $n_0 = n_1$, the researcher will be indifferent between implementing or not implementing the policy. The researchers' strategies will constitute an equilibrium because, given $\pi(n_0 = n_1) = \frac{1}{2}$, the criteria for the thresholds $\bar{\beta}$ and $\underline{\beta}$ will be symmetric with one another.

(i) Policymaker's strategy: For the policymaker's strategies to be optimal, it is necessary that $\mathbb{E}[b | n_1 = n_0] = 0$; moreover, this condition is also sufficient due to the monotonicity as in Lemma 1. We consider two cases, $n_1 = n_0 = 0$ and $n_1 = n_0 = 1$ in turn given the researchers' strategies such that $\bar{\beta} = -\underline{\beta}$:

- Case of $n_1 = n_0 = 0$: by Bayes' rule,

$$\mathbb{E}[b | n_1 = n_0 = 0] = \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} \mathbb{E}[\beta_1 + \beta_2 | \beta_1, \beta_2 \in [\underline{\beta}, \bar{\beta}]]$$

- When $\beta_1 = \Delta$,

$$\mathbb{E}[\beta_1 + \beta_2 | \beta_1 = \Delta, \beta_2 \in [\underline{\beta}, \bar{\beta}]] = \phi\left(\frac{\Delta}{\sigma}\right) \int_{\underline{\beta}}^{\bar{\beta}} (\Delta + \beta_2) \phi\left(\frac{\beta_2 - \rho\Delta}{\sigma}\right) d\beta_2$$

- When $\beta_1 = -\Delta$,

$$\mathbb{E}[\beta_1 + \beta_2 | \beta_1 = -\Delta, \beta_2 \in [\underline{\beta}, \bar{\beta}]] = \phi\left(-\frac{\Delta}{\sigma}\right) \int_{\underline{\beta}}^{\bar{\beta}} (-\Delta + \beta_2) \phi\left(\frac{\beta_2 + \rho\Delta}{\sigma}\right) d\beta_2$$

- Since β_1 is symmetrically distributed, $\phi\left(\frac{\Delta}{\sigma}\right) = \phi\left(-\frac{\Delta}{\sigma}\right)$. Using the two expressions above, we have

$$\mathbb{E}[\beta_1 + \beta_2 | \beta_1 = \Delta, \beta_2 \in [\underline{\beta}, \bar{\beta}]] + \mathbb{E}[\beta_1 + \beta_2 | \beta_1 = -\Delta, \beta_2 \in [\underline{\beta}, \bar{\beta}]] = 0$$

for the following two reasons by $\bar{\beta} = -\underline{\beta}$:

* on the term multiplied by Δ ,

$$\begin{aligned}
& \int_{\underline{\beta}}^{\bar{\beta}} \Delta \left[\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] d\beta_2 \\
&= \Delta \left\{ \left[\Phi \left(\frac{\bar{\beta} - \rho\Delta}{\bar{\sigma}} \right) - \Phi \left(\frac{\underline{\beta} - \rho\Delta}{\bar{\sigma}} \right) \right] - \left[\Phi \left(\frac{\bar{\beta} + \rho\Delta}{\bar{\sigma}} \right) - \Phi \left(\frac{\underline{\beta} + \rho\Delta}{\bar{\sigma}} \right) \right] \right\} \\
&= \Delta \left\{ \left[\Phi \left(\frac{\bar{\beta} - \rho\Delta}{\bar{\sigma}} \right) - \Phi \left(\frac{\underline{\beta} - \rho\Delta}{\bar{\sigma}} \right) \right] - \left[\Phi \left(-\frac{\underline{\beta} - \rho\Delta}{\bar{\sigma}} \right) - \Phi \left(-\frac{\bar{\beta} - \rho\Delta}{\bar{\sigma}} \right) \right] \right\} \\
&= \Delta \left\{ \left[\Phi \left(\frac{\bar{\beta} - \rho\Delta}{\bar{\sigma}} \right) - \Phi \left(\frac{\underline{\beta} - \rho\Delta}{\bar{\sigma}} \right) \right] - \left[1 - \Phi \left(\frac{\underline{\beta} - \rho\Delta}{\bar{\sigma}} \right) - \left[1 - \Phi \left(\frac{\bar{\beta} - \rho\Delta}{\bar{\sigma}} \right) \right] \right] \right\} \\
&= 0
\end{aligned}$$

* on the term multiplied by β_2 ,

$$\begin{aligned}
& \int_{\underline{\beta}}^{\bar{\beta}} \beta_2 \left[\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] d\beta_2 \\
&= \int_0^{\bar{\beta}} \beta_2 \left[\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] d\beta_2 \\
&\quad + \int_{\underline{\beta}}^0 \beta_2 \left[\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] d\beta_2 \\
&= \int_0^{\bar{\beta}} \beta_2 \left[\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] d\beta_2 \\
&\quad - \int_0^{\bar{\beta}} \beta_2 \left[\phi \left(\frac{-\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{-\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] d\beta_2 \\
&= \int_0^{\bar{\beta}} \beta_2 \left\{ \left[\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] \right. \\
&\quad \left. - \left[\phi \left(\frac{-\beta_2 - \rho\Delta}{\bar{\sigma}} \right) - \phi \left(\frac{-\beta_2 + \rho\Delta}{\bar{\sigma}} \right) \right] \right\} d\beta_2 \\
&= 0
\end{aligned}$$

where the last line followed by $\phi \left(\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right) = \phi \left(-\frac{\beta_2 + \rho\Delta}{\bar{\sigma}} \right)$ and $\phi \left(\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right) = \phi \left(-\frac{\beta_2 - \rho\Delta}{\bar{\sigma}} \right)$.

- Case of $n_1 = n_0 = 1$: by Bayes' rule, without loss of generality, let us consider

$m_1 = 1$ and $m_2 = 0$.

$$\mathbb{E}[b \mid n_1 = n_0 = 1] = \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} \mathbb{E}[\beta_1 + \beta_2 \mid \beta_1 \geq \bar{\beta}, \beta_2 \leq \underline{\beta}]$$

Note that we can express $\mathbb{E}[\beta_1 + \beta_2 \mid \beta_1 \geq \bar{\beta}, \beta_2 \leq \underline{\beta}]$ as

$$\int_{-\infty}^{\underline{\beta}} \int_{\bar{\beta}}^{-\beta_2} [\beta_1 + \beta_2] \tilde{\phi}(\beta_1, \beta_2) d\beta_1 d\beta_2 + \int_{\bar{\beta}}^{\infty} \int_{-\beta_1}^{\underline{\beta}} [\beta_1 + \beta_2] \tilde{\phi}(\beta_1, \beta_2) d\beta_2 d\beta_1$$

By the change of variable using the symmetry of distribution,

$$\int_{-\infty}^{\underline{\beta}} \int_{\bar{\beta}}^{-\beta_2} [\beta_1 + \beta_2] \tilde{\phi}(\beta_1, \beta_2) d\beta_1 d\beta_2 + \int_{\bar{\beta}}^{\infty} \int_{-\beta_2}^{\underline{\beta}} [\beta_1 + \beta_2] \tilde{\phi}(\beta_1, \beta_2) d\beta_1 d\beta_2$$

At each $\beta_2 = \Delta$,

$$\begin{aligned} & \int_{\bar{\beta}}^{-\Delta} [\beta_1 + \Delta] \phi\left(\frac{\beta_1 - \rho\Delta}{\bar{\sigma}}\right) d\beta_1 + \int_{-\Delta}^{\underline{\beta}} [\Delta + \beta_1] \phi\left(\frac{\beta_1 - \rho\Delta}{\bar{\sigma}}\right) d\beta_1 \\ &= \int_{\bar{\beta}}^{-\Delta} \{[\beta_1 + \Delta] - [\Delta + \beta_1]\} \phi\left(\frac{\beta_1 - \rho\Delta}{\bar{\sigma}}\right) d\beta_1 \\ &= 0 \end{aligned}$$

Thus, taking together, $\mathbb{E}[\beta_1 + \beta_2 \mid \beta_1 \geq \bar{\beta}, \beta_2 \leq \underline{\beta}] = 0$, satisfying the policymaker's indifference condition. \blacksquare

(ii) Researchers' strategy: given the policymaker's strategy, the researchers' indifference conditions are given by

$$\begin{aligned} \bar{\beta} + \frac{1}{2} \sum_{m_{-i} \in \{0, \emptyset\}} \{ \mathbb{P}(m_{-i} \mid \bar{\beta}, m_{-i} \in \{0, \emptyset\}) \mathbb{E}[\beta_{-i} \mid m_{-i}, \bar{\beta}] \} &= 0 \\ \underline{\beta} + \frac{1}{2} \sum_{m_{-i} \in \{0, 1\}} \{ \mathbb{P}(m_{-i} \mid \underline{\beta}, m_{-i} \in \{0, 1\}) \mathbb{E}[\beta_{-i} \mid m_{-i}, \underline{\beta}] \} &= 0 \end{aligned}$$

Applying the formula of truncated normal distribution,

$$\begin{aligned} \bar{\beta}_i + \frac{1}{2} \mathbb{E}[\beta_{-i} \mid \beta_{-i} \leq \bar{\beta}_{-i}, \bar{\beta}_i] &= 0 \\ \underline{\beta}_i + \frac{1}{2} \mathbb{E}[\beta_{-i} \mid \beta_{-i} \geq \underline{\beta}_{-i}, \underline{\beta}_i] &= 0 \end{aligned}$$

Note that when $\bar{\beta}_i = -\underline{\beta}_i$ and $\bar{\beta}_{-i} = -\underline{\beta}_{-i}$, these conditions are equivalent to each

other. Moreover, the solution $\bar{\beta}_i$ is strictly decreasing in $\bar{\beta}_{-i}$. Combining, there exists a unique solution $\bar{\beta}_i = -\underline{\beta}_i = \bar{\beta}_{-i} = -\underline{\beta}_{-i}$ that satisfies the researchers' indifference conditions. ■

(2) Proof for equilibrium with no omission. the policymaker's strategy will be a part of the equilibrium by an immediate implication of researchers' indifference condition; the researchers strategy $\bar{\beta}, \underline{\beta}$ will be determined by the identical conditions, leading to $\bar{\beta} = \underline{\beta}$.

(i) Policymaker's strategy: given that the researcher will be indifferent at the switching point, $\bar{\beta} = \underline{\beta}$, the policymaker will also be indifferent between implementing or not implementing the policy since the policymaker knows $\beta_i = \bar{\beta} = \underline{\beta}$ if $m_i = \emptyset$ (even though $m_i = \emptyset$ occurs with probability zero). Since the decisions when $m_i \neq \emptyset$ for both i can be given by the monotonicity of Lemma 1, the policy rule (1.7) is a part of the equilibrium. ■

(ii) Researchers' strategy: suppose that another researcher follows $\bar{\beta} = \underline{\beta}$ and the policymaker adopts the policy rule as (1.7). Then, regardless of one's own signal, $\mathbb{P}(m_{-i} = \emptyset | \beta_i) = 0$. Therefore, one can write the indifference conditions as

$$\mathbb{P}(m_{-i} = 1 | \bar{\beta}_i) (1 - \tilde{\pi}) \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} [\mathbb{E} [\beta_{-i} | \beta_{-i} \geq \bar{\beta}_{-i}, \bar{\beta}_i] + \bar{\beta}_i] = 0 \quad (1.19)$$

$$\mathbb{P}(m_{-i} = 1 | \underline{\beta}_i) \tilde{\pi} \frac{\frac{1}{\sigma^2}}{\frac{1}{\sigma_b^2} + \frac{2}{\sigma^2}} [\mathbb{E} [\beta_{-i} | \beta_{-i} \geq \bar{\beta}_{-i}, \underline{\beta}_i] + \underline{\beta}_i] = 0 \quad (1.20)$$

Since these conditions are proportional to each other, researcher i 's optimal strategy has $\bar{\beta}_i = \underline{\beta}_i$. □

A2.3. Proof of Proposition 1.3

Proposition 1.3 claims that the asymmetric equilibrium is locally stable whereas equilibria with symmetric or no omission are not; the former is also optimal whereas later are not. We will first prove the results of local stability, and then that of optimality.

Proof of local stability: we will focus on the concept of local stability in Definition 1.3, which is adopted from Definition 6.1 in Chapter 1 of Fudenberg and Levine (1998) with a particular order of adjustment. The equilibria satisfying local stability are more plausible to emerge than those without local stability since small perturbation of strategies likely occur in the real world.

We consider a perturbation of equilibrium with monotone strategies, $\mathcal{E} \equiv \{\pi(n), \bar{\beta}_1, \underline{\beta}_1, \bar{\beta}_2, \underline{\beta}_2\}$,

and consider the distance between two equilibria, $\mathcal{E}, \tilde{\mathcal{E}}$, as $d(\mathcal{E} - \tilde{\mathcal{E}}) \equiv \max_s \{|\epsilon_s|\}$, where $\epsilon \equiv \mathcal{E} - \tilde{\mathcal{E}}$. While one could consider a richer perturbation on researchers' strategies as the mapping from the signals $\beta_i \times \sigma_i \in \mathbb{R}^2$ into probability distribution over messages, this definition is intuitive and analytically tractable. Moreover, Lemma 1 has shown that all fully responsive and fully informative equilibria will take this form. We first consider the asymmetric equilibrium in Proposition 1.1, and then analyze the other equilibria in Proposition 1.2.

(1) *Asymmetric equilibrium in Proposition 1.1. is locally stable:* Let us denote the perturbation of researcher $i = 1, 2$'s strategies by the set of perturbations, $\{\bar{\epsilon}_i, \underline{\epsilon}_i\}$ so that $\bar{\beta}_{i,0} = \bar{\beta} + \bar{\epsilon}_i, \underline{\beta}_{i,0} = \underline{\beta} + \underline{\epsilon}_i$. Without loss of generality, suppose that the researcher 2 receives a larger perturbation so that $\max\{|\bar{\epsilon}_2|, |\underline{\epsilon}_2|\} \geq \max\{|\bar{\epsilon}_1|, |\underline{\epsilon}_1|\}$.

The proof takes four steps: first, we observe that the policymaker's strategy does not change; second, consider researcher 1's adjustment in $t = 1$; third, consider researcher 2's adjustment in $t = 1$; and finally argue that these results show the local stability of the asymmetric equilibrium.

Step 1. policymaker's strategy: in $t = 1$, even with small perturbation of researchers' strategy, the policymaker's strategy will not change since it relied on strict preference. Thus, the strategy (1.2) will continue to be played.

Step 2. researcher 1's strategy: given the supermajoritarian rule (1.2) and the researcher 2's initial strategy $\bar{\beta}_{2,0}, \underline{\beta}_{2,0}$, the researcher 1's strategy will satisfy

$$\begin{aligned} \left| \bar{\beta}_{1,1}(\bar{\beta}_{2,0}, \underline{\beta}_{2,0}) - \bar{\beta} \right| &< \max\{|\bar{\epsilon}_2|, |\underline{\epsilon}_2|\} \\ \left| \underline{\beta}_{1,1}(\bar{\beta}_{2,0}) - \underline{\beta} \right| &< |\bar{\epsilon}_2| \end{aligned}$$

by the property of derivative of the truncated normal distribution.

Step 3. researcher 2's strategy: given the supermajoritarian rule (1.2) and the researcher 1's strategy after adjustment $\bar{\beta}_{1,1}, \underline{\beta}_{1,1}$, the researcher 2's strategy will satisfy

$$\begin{aligned} \left| \bar{\beta}_{2,1}(\bar{\beta}_{1,1}, \underline{\beta}_{1,1}) - \bar{\beta} \right| &< \max\{|\bar{\epsilon}_2|, |\underline{\epsilon}_2|\} \\ \left| \underline{\beta}_{2,1}(\bar{\beta}_{1,1}) - \underline{\beta} \right| &< |\bar{\epsilon}_2| \end{aligned}$$

by the property of derivative of the truncated normal distribution.

Step 4. relating to the definition of local stability: for equilibrium \mathcal{E} to be locally stable, we need for every $\hat{d} > 0$, there exist some \bar{d} such that

$$d(\mathcal{E} - \mathcal{E}_0) < \bar{d} \Rightarrow d(\mathcal{E} - \mathcal{E}_\infty) < \hat{d}.$$

By Step 2 and 3, we know that $d(\mathcal{E} - \mathcal{E}_1) < \max\{|\bar{\epsilon}_2|, |\epsilon_2|\} = d(\mathcal{E} - \mathcal{E}_0)$. Iterating the adjustment ad infinity, we have $d(\mathcal{E} - \mathcal{E}_\infty) < d(\mathcal{E} - \mathcal{E}_0)$. Thus, setting any $\bar{d} \leq \hat{d}$ can satisfy the condition.

(2) *Equilibria in Proposition 1.2. are not locally stable:*

Lemma A3: *any equilibrium such that $\pi^*(n_1, n_0) \in (0, 1)$ for some n_1, n_0 is not locally stable.*

Proof. This is because the policymaker must be exactly indifferent between whether or not implementing the policy; that is, the posterior belief $\mathbb{E}[b|n_1, n_0] = 0$ must hold for such n_1, n_0 . However, even with a small perturbation of some thresholds $\{\bar{\beta}_i, \underline{\beta}_i\}$, the policymaker will have $\mathbb{E}[b|n_1, n_0] \neq 0$ so that his optimal strategy in $t = 1$ will be either $\pi^*(n_1, n_0) \in \{0, 1\}$ for that n_1, n_0 . Since the modification in policymaker's strategy is large, the researchers' strategies will not converge back to the original strategies. ■

On the equilibrium with no omission such that $\tilde{\pi} = 1$, we can consider how a small perturbation of another researcher's strategy, $\bar{\beta}_i + \Delta$, makes the probability of omission to be strictly positive; that is, with such perturbation, the iterative adjustment to examine the local stability will lead to the asymmetric equilibrium characterized in Proposition 1.1. □

Proof of optimality: we will first show that the equilibria characterized in Proposition 1.2 are not optimal by using the relationship to the concept of local stability; then show that the equilibrium in Proposition 1.1 is optimal by examining all other possible equilibria.

(1) *Equilibria in Proposition 1.2 are not optimal:* since the model is a common interest game, we have the following close relationship between local stability and optimality:

Lemma A3: *if an equilibrium \mathcal{E} is optimal, then it is locally stable.*

Proof. Let us write the welfare attained in the equilibrium $\mathcal{E} \equiv \{\pi(n), \bar{\beta}_1, \underline{\beta}_1, \bar{\beta}_2, \underline{\beta}_2\}$ as $W(\pi(n), \bar{\beta}_1, \underline{\beta}_1, \bar{\beta}_2, \underline{\beta}_2)$. By policymaker's optimization, $\frac{\partial W}{\partial \pi}|_{\pi=1} \geq 0$, $\frac{\partial W}{\partial \pi}|_{\pi=0} \leq 0$, $\frac{\partial W}{\partial \pi}|_{\pi \in (0,1)} = 0$ (where the derivative at the boundaries are either or left or right derivatives) and by researchers' optimization, $\frac{\partial W}{\partial \beta_i} = \frac{\partial W}{\partial \bar{\beta}_i} = 0$ when evaluated at the equilibrium since W is continuous in each element. If an equilibrium is optimal, then it is locally stable since W must be locally concave at each local maximum. ■

Using the contrapositive of Lemma A3, if an equilibrium is not locally stable, then it is not optimal. Since the equilibria in Proposition 1.2 are not locally stable, it is not optimal.

(2) *Asymmetric equilibrium in Proposition 1.1 is optimal:* while optimality implies

local stability, the converse does not hold. The proof consists of arguing that (i) any optimal equilibria must be fully responsive, and (ii) there are only two fully responsive and locally stable equilibria: one with $\pi^*(m) = 1 (n_1 \geq n_0)$ and another with $\pi^*(m) = 1 (n_1 > n_0)$. Then, since these two equilibria are symmetric to one another, they attain the identical level of welfare.

- (i) if an equilibrium is not fully responsive or not fully informative, then it is not optimal. Suppose there exists m_i, m'_i such that $\pi(m_i, m_{-i}) = \pi(m'_i, m_{-i})$ or $\mathbb{E}[b|m_i, m_{-i}] = \mathbb{E}[b|m'_i, m_{-i}]$ for all m_{-i} . Then, by changing the reporting strategies of m_i, m'_i , the researcher i can better convey the private information β_i to the policymaker. Since this is a common interest game, this strictly improves welfare.
- (ii) the only fully responsive equilibria that are also locally stable have $\pi^*(m) = 1 (n_1 \geq n_0)$ or $\pi^*(m) = 1 (n_1 > n_0)$. By Lemma A4, local stability requires $\pi(m_i, m_{-i}) \in \{0, 1\}$ for all m_i, m_{-i} . Moreover, by Lemma 1, $\pi(m_i, m_{-i})$ will be monotone.

- if $\pi(m_i, m_{-i}) = 0$ for all $m_i \in \{\emptyset, 0\}$ for either i , then it is not fully responsive. Thus, the fully responsive equilibria with minimum number of $\pi = 1$ is $\pi^*(m) = 1 (n_1 > n_0)$.
- suppose $\pi^*(m) = 1$ if $m_i = 1$ and $\{m_i, m_{-i}\} = \{\emptyset, 1\}$. Then for another researcher, whether $m_{-i} = \emptyset$ or 0 does not make any difference. Suppose $\pi^*(m) = 1$ if $m_i \neq 0$ for both researchers. Then, for either researcher, whether $m_{-i} = \emptyset$ or 1 does not make any difference. Thus, the fully responsive equilibria with the second minimum number of $\pi = 1$ is $\pi^*(m) = 1 (n_1 \geq n_0)$.

Since we can consider $\pi = 0$ symmetrically, we have considered all equilibria with monotone strategy of researchers and $\pi(m_i, m_{-i}) \in \{0, 1\}$ for all m_i, m_{-i} . Verifying that $\pi^*(m) = 1 (n_1 > n_0)$ and $\pi^*(m) = 1 (n_1 \geq n_0)$ satisfies full responsiveness, we have that there can be at most two equilibria.

Note that the two equilibria $\pi^*(m) = 1 (n_1 > n_0)$ and $\pi^*(m) = 1 (n_1 \geq n_0)$ are symmetric to one another when $c = 0$, and thus, attains identical welfare. Since the welfare under $\pi^*(m) = 1 (n_1 \geq n_0)$ is not strictly higher than that under $\pi^*(m) = 1 (n_1 > n_0)$, the equilibrium characterized in Proposition 1.1 is an optimal equilibrium. Since an optimal set of strategies must constitute an equilibrium, the equilibrium that attains weakly higher welfare than any other equilibria must also be optimal. \square

A3. Proof of 2.4 Inflation of Marginally Insignificant Results

This sub-Section presents the proof of Proposition 1.2 in Section 2.4.

Proof of Proposition 1.2 The Proposition 1.2 claims that there exists a unique symmetric equilibrium (researchers' thresholds are identical with one another, and their thresholds are symmetric so that $\bar{\beta}(\sigma) = -\underline{\beta}(\sigma)$), and in that equilibrium the absolute value of the t-statistics must be increasing in σ_i . The proof proceeds in three steps: first, we express and simplify the researchers' indifference conditions assuming the policymakers' strategy in Proposition 1.2; second, we show the existence of the solution and characterize that solution; and finally, we verify that the researchers' and policymakers' strategies constitute an equilibrium.

Step 1. researchers' indifference conditions with heterogeneous σ_i : we express the indifference conditions by extending the expression of thresholds derived in Proposition 1.2(i) in three sub-steps.

First, by researcher i 's optimization, the posterior belief on expected benefit must equal zero at the thresholds at every $\sigma_i \in \text{Supp}(\sigma)$: by Bayes' rule and the law of iterated expectations,

$$\int \frac{\frac{\bar{\beta}_i}{\sigma_i^2} + \frac{\mathbb{E}[\beta_j | \sigma_j, \beta_j \in \text{Piv}(\sigma_j), \pi, \beta_i = \bar{\beta}_i]}{\sigma_j^2}}{\frac{1}{\sigma_b^2} + \frac{1}{\sigma_i^2} + \frac{1}{\sigma_j^2}} g(\sigma_j) d\sigma_j = 0, \quad (1.21)$$

where the expectation is taken over another researcher's signals $\{\sigma_j, \beta_j\}$ and the policymaker's strategy π . $\text{Piv}(\sigma_j)$ is a set of values of β_j such that the researcher i 's message can alter the policymaker's decision.

Second, we rearrange the condition (1.21) by (i) assumption of policymaker's strategy, (ii) improper prior assumption ($\sigma_b = \infty$), (iii) change of variables from $\bar{\beta}_i$ to $t(\sigma_i)$, and (iv) re-expressing the inverse Mills ratio:

- (i) as shown in Section A2.2, if policymaker adopts the strategy in symmetric equilibrium (1.6), then $\mathbb{E}[\beta_j | \sigma_j, \beta_j \in \text{Piv}(\sigma_j, \pi), \bar{\beta}_i] = \rho_{ij} \bar{\beta}_i - \sigma_{ij} \frac{\phi(\cdot)}{\Phi(\cdot)}$, where $\rho_{ij} = \frac{\sigma_b^2}{\sqrt{\sigma_i^2 + \sigma_b^2} \sqrt{\sigma_j^2 + \sigma_b^2}}$ is the correlation coefficient, $\sigma_{ij} = \sqrt{\sigma_i^2 + \sigma_b^2} \sqrt{\sigma_j^2 + \sigma_b^2} - \sigma_b^2$ is the standard deviation of β_j conditional on β_i , and the argument of inverse Mills ratio is $\frac{\bar{\beta}_j(\sigma_j) - \rho_{ij} \bar{\beta}_i}{\sigma_{ij}}$ with $\bar{\beta}_j(\sigma_j)$ denoting the researcher j 's threshold conditional on σ_j .

- (ii) by the assumption $\sigma_b = \infty$, $\frac{1}{\sigma_b^2} = 0$, $\rho_{ij} = 1$, and $\sigma_{ij} = \sqrt{\frac{\sigma_i^2 + \sigma_j^2}{2}}$. Therefore, the

indifference condition (1.21) is equivalent to

$$2\bar{\beta}_i - \frac{\sigma_i}{\sqrt{2}} \int \sqrt{\frac{\sigma_i^2}{\sigma_i^2 + \sigma_j^2}} \frac{\phi(\cdot)}{\Phi(\cdot)} g(\sigma_j) d\sigma_j = 0, \quad (1.22)$$

where the argument of $\phi(\cdot)$ and $\Phi(\cdot)$ is $\frac{\bar{\beta}_j(\sigma_j) - \bar{\beta}_i}{\sqrt{\frac{\sigma_i^2 + \sigma_j^2}{2}}}$.

(iii) dividing the condition (1.22) by σ_i and writing $t_i = \frac{\bar{\beta}_i(\sigma_i)}{\sigma_i}$ for brevity,

$$t_i = \frac{1}{4} \int \sqrt{\frac{2}{1 + \sigma_r^2}} \frac{\phi(t_j \sigma_r - t_i)}{\Phi(t_j \sigma_r - t_i)} g(\sigma_j) d\sigma_j, \quad (1.23)$$

where $\sigma_r \equiv \frac{\sigma_j}{\sigma_i}$.

(iv) we can express the condition (1.23) as a conditional mean of a truncated standard normal distribution using some hypothetical random variables: Noting that, for some random variable, $\tau_1 \sim \mathcal{N}\left(t_i, \frac{1 + \sigma_r^2}{2}\right)$, $\mathbb{E}[\tau_1 | \tau_1 \leq t_j \sigma_r] = t_i - \sqrt{\frac{1 + \sigma_r^2}{2}} \frac{\phi(t_j \sigma_r - t_i)}{\Phi(t_j \sigma_r - t_i)}$,

$$\begin{aligned} \frac{\phi(t_j \sigma_r - t_i)}{\Phi(t_j \sigma_r - t_i)} &= \sqrt{\frac{2}{1 + \sigma_r^2}} \{t_i - \mathbb{E}[\tau_1 | \tau_1 \leq t_j \sigma_r]\}, \\ &= \sqrt{\frac{2}{1 + \sigma_r^2}} \{-\mathbb{E}[\tau_2 | \tau_2 \leq t_j \sigma_r - t_i]\}, \text{ where } \tau_2 \sim \mathcal{N}\left(0, \frac{1 + \sigma_r^2}{2}\right) \\ &= \sqrt{\frac{2}{1 + \sigma_r^2}} \mathbb{E}[\tau_2 | \tau_2 \geq t_i - t_j \sigma_r] \\ &= \sqrt{\frac{2}{1 + \sigma_r^2}} \mathbb{E}\left[\sqrt{\frac{1 + \sigma_r^2}{2}} \tau_3 \mid \sqrt{\frac{1 + \sigma_r^2}{2}} \tau_3 \geq t_i - t_j \sigma_r\right], \text{ where } \tau_3 \sim \mathcal{N}(0, 1) \\ &= \mathbb{E}\left[\tau_3 \mid \tau_3 \geq \sqrt{\frac{2}{1 + \sigma_r^2}} (t_i - t_j \sigma_r)\right] \end{aligned}$$

Combining with (1.23), the indifference condition is

$$t_i = \frac{1}{4} \int \sqrt{\frac{2}{1 + \sigma_r^2}} \mathbb{E}\left[\tau \mid \tau \geq \sqrt{\frac{2}{1 + \sigma_r^2}} (t_i - t_j \sigma_r)\right] g(\sigma_j) d\sigma_j, \quad (1.24)$$

where $\tau \sim \mathcal{N}(0, 1)$.

Step 2. characterization: using the formula (1.25), we can show that $\frac{\partial t(\sigma_i)}{\partial \sigma_i} > 0$.
 Writing $K(\sigma_r) = \sqrt{\frac{2}{1+\sigma_r^2}}$, $L(\sigma_r) = \mathbb{E} \left[\tau | \tau \geq \sqrt{\frac{2}{1+\sigma_r^2}} (t_i - t_j \sigma_r) \right]$,

$$\frac{\partial t_i(\sigma_i)}{\partial \sigma_i} = \frac{\partial \sigma_r}{\partial \sigma_i} \times \frac{1}{4} \int [K(\sigma_r) L'(\sigma_r) + K'(\sigma_r) L(\sigma_r)] g(\sigma_j) d\sigma_j \quad (1.25)$$

by the chain rule. Note that $L'(\sigma_r) < 0$ since

$$\begin{aligned} \frac{\partial \sqrt{\frac{2}{1+\sigma_r^2}} (t_i - t_j \sigma_r)}{\partial \sigma_r} &= -t_j \sqrt{\frac{1+\sigma_r^2}{2}} - (t_i - t_j \sigma_r) \left(\frac{1}{2}\right)^{\frac{3}{2}} (1+\sigma_r^2)^{-\frac{1}{2}} \\ &= -\frac{1}{2} \sqrt{\frac{1+\sigma_r^2}{2}} \left[2t_j + \frac{t_i - t_j \sigma_r}{1+\sigma_r^2} \right] \\ &= -\frac{1}{2} \sqrt{\frac{1+\sigma_r^2}{2}} \left\{ \frac{t_i + 2t_j [(1-\sigma_r)^2 + \sigma_r]}{1+\sigma_r^2} \right\} < 0 \end{aligned}$$

Together with observations that $\frac{\partial \sigma_r}{\partial \sigma_i} < 0$, $K(\sigma_r) > 0$, and $L(\sigma_r) > 0$ and $K'(\sigma_r) < 0$, $\frac{\partial t_i(\sigma_i)}{\partial \sigma_i} > 0$ holds at every σ_i , as stated by Proposition 1.2.

Step 3. existence and uniqueness of researchers' strategies in symmetric equilibrium:
 to show that the equilibrium characterization is meaningful, we will show the existence of such threshold $t(\sigma_i | \{t_j\})$ defined by (1.24). We will apply the contraction mapping theorem, which also shows that the thresholds will be unique.

The functional equation $t(t_j)$ (a simplification of notation $t(\sigma_i | \{t_j\})$) is said to be a contraction mapping if there exists some constant $k \in (0, 1)$ such that $d\{t(\tau_0), t(\tau_1)\} \leq kd\{\tau_0, \tau_1\}$ for any τ_0, τ_1 . Here, let us define the distance between two strategies, $\tau_0(\sigma)$ and $\tau_1(\sigma)$, as the sup metric:

$$d\{\tau_0, \tau_1\} = \sup_{\sigma} |\tau_0(\sigma) - \tau_1(\sigma)|.$$

We first show that $t(\sigma_i | \{t_j\})$ is a contraction mapping when τ_0, τ_1 are differentiable with respect to σ in three sub-steps, and then apply the contraction mapping theorem.

- (i) Sub-step 1: we first show that we can consider a corresponding totally ordered two functions, rather than considering directly the arbitrary functions that may not be totally ordered.

Sub-Lemma 1. *Define $\bar{\tau}(\sigma) \equiv \max\{\tau_0(\sigma), \tau_1(\sigma)\}$ and $\underline{\tau}(\sigma) \equiv \min\{\tau_0(\sigma), \tau_1(\sigma)\}$. For any k , if $d\{t(\bar{\tau}), t(\underline{\tau})\} \leq kd\{\bar{\tau}, \underline{\tau}\}$, then $d\{t(\tau_0), t(\tau_1)\} \leq kd\{\tau_0, \tau_1\}$.*

Proof. The expression (1.24) shows that $t(\tau)$ is strictly decreasing in τ so that

$$t(\sigma|\{\bar{\tau}\}) \leq t(\sigma|\{\tau_s\}) \leq t(\sigma|\{\underline{\tau}\})$$

for both $s = 0, 1$, for all σ . Thus, $d\{t(\bar{\tau}), t(\underline{\tau})\} \geq d\{t(\tau_0), t(\tau_1)\}$. Since $\tau_0(\sigma)$ and $\tau_1(\sigma)$ are assumed to be differentiable, it is continuous and attains the maximum and minimum. Since $d\{\bar{\tau}, \underline{\tau}\} = d\{\tau_0, \tau_1\}$, the Sub-Lemma 1 holds. ■

- (ii) Sub-step 2: we then show that we can consider the derivative of function to prove that the two functions satisfy the condition to be a contraction map.

Sub-Lemma 2. *If there exists some k such that*

$$\frac{\partial d\{t(\tau + \delta), t(\tau)\}}{\partial \delta} < k$$

for any τ and for any δ , then the function t satisfies $d\{t(\bar{\tau}), t(\underline{\tau})\} \leq kd\{\bar{\tau}, \underline{\tau}\}$ for any totally ordered $\{\bar{\tau}, \underline{\tau}\}$.

Proof. Given any $\{\bar{\tau}, \underline{\tau}\}$ that is totally ordered (i.e. $\bar{\tau}(\sigma) > \underline{\tau}(\sigma)$ for all σ), let $\delta \equiv \sup_{\sigma} \{\bar{\tau}(\sigma) - \underline{\tau}(\sigma)\}$. Then

$$\begin{aligned} d\{t(\bar{\tau}), t(\underline{\tau})\} &\leq d\{t(\underline{\tau} + \delta), t(\underline{\tau})\} \\ &= \underbrace{d\{t(\underline{\tau}), t(\underline{\tau})\}}_{=0} + \int_0^{\delta} \underbrace{\frac{\partial d\{t(\tau + \tilde{\delta}), t(\tau)\}}{\partial \tilde{\delta}}}_{\leq k} d\tilde{\delta} \\ &\leq k\delta \\ &= kd\{\bar{\tau}, \underline{\tau}\} \end{aligned}$$

The first line follows because t is strictly decreasing in τ ; the second line follows from the fundamental theorem of calculus; the third line follows by assumption of Sub-Lemma 2; and the fourth line by definition of δ . ■

- (iii) Sub-step 3: using the expression (1.24), we derive the expression of the bound on the derivative.

Sub-Lemma 3. *For any τ and for any δ , with t defined as (1.24),*

$$\frac{\partial d\{t(\tau + \delta), t(\tau)\}}{\partial \delta} < \frac{1}{2}$$

Proof. There are three steps to prove this: first, we note that analyzing $\frac{\partial}{\partial \delta} t(\sigma^*(\delta) | \{\tau + \delta\})$ will be sufficient:

$$\begin{aligned} \frac{\partial d\{t(\tau + \delta), t(\tau)\}}{\partial \delta} &= \frac{\partial}{\partial \delta} [t(\sigma^*(\delta) | \{\tau\}) - t(\sigma^*(\delta) | \{\tau + \delta\})], \\ &\quad \text{where } \sigma^*(\delta) \equiv \arg \max_{\sigma} \{t(\sigma | \{\tau + \delta\}) - t(\sigma | \{\tau\})\} \\ &= -\frac{\partial}{\partial \delta} t(\sigma^*(\delta) | \{\tau + \delta\}) \\ &\quad + \frac{\partial}{\partial \sigma} [t(\sigma^*(\delta) | \{\tau\}) - t(\sigma^*(\delta) | \{\tau + \delta\})] |_{\sigma=\sigma^*(\delta)} \times \frac{\partial \sigma^*(\delta)}{\partial \delta} \\ &= -\frac{\partial}{\partial \delta} t(\sigma^*(\delta) | \{\tau + \delta\}) \end{aligned}$$

where the first equality followed by definition (the maximum exists since $\text{Supp}(G)$ is closed and bounded), the second equality followed by the chain rule, and the third equality followed by the envelope theorem given differentiability of t .

Second, writing $\Sigma \equiv \sqrt{\frac{2}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2}} \left\{ t(\sigma^* | \{\tau + \delta\}) - [\tau(\sigma_j) + \delta] \frac{\sigma^*}{\sigma_j} \right\}$ for notational ease, we can derive the bound on $-\frac{\partial}{\partial \delta} t(\sigma^*(\delta) | \{\tau + \delta\})$: given (1.24),

$$\frac{\partial}{\partial \delta} t(\sigma^* | \{\tau + \delta\}) = \frac{1}{4} \int \sqrt{\frac{2}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2}} \frac{\partial \mathbb{E}[\tilde{\tau} | \tilde{\tau} \geq \Sigma]}{\partial \Sigma} \times \frac{\partial \Sigma}{\partial \delta} g(\sigma_j) d\sigma_j \quad (1.26)$$

By the chain rule, $\frac{\partial \Sigma}{\partial \delta} = \sqrt{\frac{2}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2}} \left\{ \frac{\partial}{\partial \delta} t(\sigma^* | \{\tau + \delta\}) - \frac{\sigma^*}{\sigma_j} \right\}$. Thus, rearranging (1.26),

$$-\frac{\partial}{\partial \delta} t(\sigma^* | \{\tau + \delta\}) = \frac{\int \frac{\frac{\sigma^*}{\sigma_j}}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2} \frac{\partial \mathbb{E}[\tau | \tau \geq \Sigma]}{\partial \Sigma} g(\sigma_j) d\sigma_j}{2 - \int \frac{1}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2} \frac{\partial \mathbb{E}[\tau | \tau \geq \Sigma]}{\partial \Sigma} g(\sigma_j) d\sigma_j}$$

Note that $\frac{\partial \mathbb{E}[\tau | \tau \geq \Sigma]}{\partial \Sigma} < 1$ because it is a derivative with respect to a truncated normal distribution. Given $\frac{1}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2} < 1$, we note that $2 - \int \frac{1}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2} \frac{\partial \mathbb{E}[\tau | \tau \geq \Sigma]}{\partial \Sigma} g(\sigma_j) d\sigma_j >$

1. Combining,

$$-\frac{\partial}{\partial \delta} t(\sigma^* | \{\tau + \delta\}) < \int \frac{\frac{\sigma^*}{\sigma_j}}{1 + \left(\frac{\sigma^*}{\sigma_j}\right)^2} g(\sigma_j) d\sigma_j.$$

Third, note that $\frac{\frac{\sigma_j^*}{\sigma_j}}{1 + \left(\frac{\sigma_j^*}{\sigma_j}\right)^2} \leq \frac{1}{2}$ for any σ^* and σ_j since

$$0 \leq (\sigma_i - \sigma_j)^2 \Rightarrow 2\sigma_i\sigma_j \leq \sigma_i^2 + \sigma_j^2 \Rightarrow \frac{\sigma_i\sigma_j}{\sigma_i^2 + \sigma_j^2} = \frac{\frac{\sigma_j^*}{\sigma_j}}{1 + \left(\frac{\sigma_j^*}{\sigma_j}\right)^2} \leq \frac{1}{2}.$$

Combining the three steps, we conclude that the Sub-Lemma 3 holds. \blacksquare

Since the expression (1.24) implies that the equilibrium threshold must be differentiable with respect to σ , we do not have to consider functions $\tau(\sigma)$ that is not differentiable. Since (i) $k = \frac{1}{2} < 1$ and (ii) the space of continuous functions is a complete metric space under sup metric, we can apply the contraction mapping theorem to claim that the function $t(\sigma)$ satisfying (1.24) exists and is unique.

Step 4. verifying the policymaker's indifference condition: To show that the policymaker will be willing to follow the strategy in the symmetric equilibrium (1.6), we need to show $\mathbb{E}[b|n_1 = n_0] = 0$. $\underline{\beta}(\sigma) = -\overline{\beta}(\sigma)$ holds at every σ by the uniqueness of the threshold that satisfies the indifference condition. By the proof of existence of symmetric equilibrium in A2.2, $\mathbb{E}[b|n_1 = n_0, \sigma] = 0$ for all σ . Thus,

$$\mathbb{E}[b|n_1 = n_0] = \int \mathbb{E}[b|n_1 = n_0, \sigma] g(\sigma) d\sigma = 0.$$

By combining Steps 1 to 4, the result $\frac{\partial t(\sigma_i)}{\partial \sigma_i} > 0$ holds in the unique symmetric equilibrium. \square

A4. Proof of 2.5 Amplification of Small Bias of a Researcher

This sub-Section proves the Proposition 1.3 in Section 2.5.

A3 Proof of Proposition 1.3 The Proposition 1.3 provides an expression for the strategic multiplier between researchers, and claims that it will be greater than 1. The proof consists of two steps: first, we derive comparative statics in equilibrium; and second, derive the multiplier and show that it is greater than 1. Note that the results for $\underline{\beta}_i$ can be derived analogously.

Step 1. comparative static with researchers' indifference conditions: as derived in Appendix [], in a symmetric equilibrium with $d_i = d_j = 0$, the indifference conditions

are given by

$$\bar{\beta}_i + \mathbb{E} [\beta_j | \beta_j \leq \bar{\beta}_j, \beta_i = \bar{\beta}_i] = - \left(2 + \frac{\sigma^2}{\sigma_b^2} \right) d_i \quad (1.27)$$

$$\bar{\beta}_j + \mathbb{E} [\beta_i | \beta_i \leq \bar{\beta}_i, \beta_j = \bar{\beta}_j] = - \left(2 + \frac{\sigma^2}{\sigma_b^2} \right) d_j \quad (1.28)$$

Totally differentiating the indifference conditions with respect to d_i , we have

$$\begin{bmatrix} 1 & \frac{Var_{truncated}}{Var_{total}} \\ \frac{Var_{truncated}}{Var_{total}} & 1 \end{bmatrix} \begin{bmatrix} \partial \bar{\beta}_i / \partial d_i \\ \partial \bar{\beta}_j / \partial d_i \end{bmatrix} = - \begin{bmatrix} 2 + \frac{\sigma^2}{\sigma_b^2} \\ 0 \end{bmatrix}$$

since $\frac{\mathbb{E}[\beta_j | \beta_j \leq \bar{\beta}_j, \beta_i = \bar{\beta}_i]}{\partial \bar{\beta}_i} = \rho \frac{Var_{truncated}}{\sigma^2}$ with $\rho = \frac{\sigma^2}{Var_{total}}$, and $\bar{\beta}_i = \bar{\beta}_j$ in the symmetric equilibrium. Rearranging,

$$\begin{bmatrix} \partial \bar{\beta}_i / \partial d_i \\ \partial \bar{\beta}_j / \partial d_i \end{bmatrix} = - \frac{1}{1 - \left(\frac{Var_{truncated}}{Var_{total}} \right)^2} \begin{bmatrix} 1 & -\frac{Var_{truncated}}{Var_{total}} \\ -\frac{Var_{truncated}}{Var_{total}} & 1 \end{bmatrix} \begin{bmatrix} 2 + \frac{\sigma^2}{\sigma_b^2} \\ 0 \end{bmatrix}$$

Therefore, we have

$$\frac{\partial (\bar{\beta}_i - \bar{\beta}_j)}{\partial d_i} = - \frac{2 + \frac{\sigma^2}{\sigma_b^2}}{1 - \frac{Var_{truncated}}{Var_{total}}} \quad (1.29)$$

Step 2. deriving and interpreting the strategic multiplier: using the expression (1.29), we can derive the expression of the multiplier, and show that it will always be greater than 1.

- expression: in the absence of strategic effects,

$$\frac{\partial (\bar{\beta}_i - \bar{\beta}_j)}{\partial d_i} \Big|_{\sigma_j = \sigma_j^*} = \frac{\partial \bar{\beta}_i}{\partial d_i} \Big|_{\sigma_j = \sigma_j^*} - \frac{\partial \bar{\beta}_j}{\partial d_i} \Big|_{\sigma_j = \sigma_j^*} = - \left(2 + \frac{\sigma^2}{\sigma_b^2} \right) - 0 = - \left(2 + \frac{\sigma^2}{\sigma_b^2} \right).$$

Thus, the multiplier is $\zeta = \frac{1}{1 - \frac{Var_{truncated}}{Var_{total}}}$.

- interpretation: by the definition of truncated distribution, $\frac{Var_{truncated}}{Var_{total}} \in (0, 1)$. Thus, $\zeta \in (1, +\infty)$.

□

A5. Proof of 4.1 A New “Stem-based” Bias Correction Method

This sub-Section contains the proof of Proposition 2 in Section 4.1, concerning the properties of bias used for the stem-based bias correction method.

A3.1. Proof of Proposition 2.

Proposition 2 claims that the bias squared is increasing in the standard error of the studies under some conditions, both for the extremum and uniform selection models. We will prove the result for the extremum selection, and then for the uniform selection. For notational ease, let us henceforth write $\bar{\sigma} = \sqrt{\sigma_0^2 + \sigma_i^2}$.

Proof of bias under extremum selection. We derive the monotonicity and limit results from the definition of truncated normal distribution. For notational ease, we the true mean, b_0 , to zero.

(i) Monotonicity: let us write

$$Bias(\sigma_i) = -\bar{\sigma} \frac{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}{\int_{\beta_{\min}}^{\infty} \phi\left(\frac{\beta}{\bar{\sigma}}\right) d\beta} = -\bar{\sigma} \left[\underbrace{\int_{\beta_{\min}}^{|\beta_{\min}|} \frac{\phi\left(\frac{\beta}{\bar{\sigma}}\right)}{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)} d\beta}_{=0} + \int_{|\beta_{\min}|}^{\infty} \frac{\phi\left(\frac{\beta}{\bar{\sigma}}\right)}{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)} d\beta \right]^{-1}$$

where the second line considered the case when $\beta_{\min} < 0$. By the chain rule, $Sign\left(\frac{\partial |Bias(\bar{\sigma})|}{\partial \bar{\sigma}}\right) = Sign(D)$, where

$$\begin{aligned} D &= \int_{|\beta_{\min}|}^{\infty} \frac{\phi\left(\frac{\beta}{\bar{\sigma}}\right)}{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)} d\beta - \bar{\sigma} \frac{\partial}{\partial \bar{\sigma}} \int_{|\beta_{\min}|}^{\infty} \frac{\phi\left(\frac{\beta}{\bar{\sigma}}\right)}{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)} d\beta \\ &= \int_{|\beta_{\min}|}^{\infty} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) - 2 \frac{\beta^2 - \beta_{\min}^2}{\bar{\sigma}^2} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta \\ &= \int_{|\beta_{\min}|}^{\infty} \left[1 + 2 \frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right] \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta \\ &= \int_{|\beta_{\min}|}^{\infty} \left[1 + 2 \frac{\beta_{\min}^2}{\bar{\sigma}^2}\right] \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta + \int_{|\beta_{\min}|}^{\infty} -2 \frac{\beta^2}{\bar{\sigma}^2} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta \\ &= |\beta_{\min}| \\ &> 0, \end{aligned}$$

where the last line followed by the integration by parts³². Thus,

$$\frac{\partial Bias^2(\sigma_i)}{\partial \sigma_i} = 2Bias(\bar{\sigma}) \frac{\partial Bias(\bar{\sigma})}{\partial \bar{\sigma}} \frac{\partial \bar{\sigma}}{\partial \sigma_i} > 0.$$

(ii) Limit: let us consider the two cases in turn while considering the original expression $Bias(\sigma_i) = -\bar{\sigma} \frac{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}{1-\Phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}$: by using the L'Hopital's rule wherever applicable,

- $\beta_{\min} < 0$:

$$\lim_{\bar{\sigma} \rightarrow 0} Bias(\bar{\sigma}) = -\lim_{\bar{\sigma} \rightarrow 0} \bar{\sigma} \frac{\lim_{\bar{\sigma} \rightarrow 0} \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}{\lim_{\bar{\sigma} \rightarrow 0} [1 - \Phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)]} = -0 \times \frac{0}{1} = 0$$

- $\beta_{\min} = 0$:

$$\lim_{\bar{\sigma} \rightarrow 0} Bias(\bar{\sigma}) = -\lim_{\bar{\sigma} \rightarrow 0} \bar{\sigma} \frac{\lim_{\bar{\sigma} \rightarrow 0} \phi(0)}{\lim_{\bar{\sigma} \rightarrow 0} [1 - \Phi(0)]} = -0 \times \frac{\phi(0)}{\frac{1}{2}} = 0$$

- $\beta_{\min} \geq 0$:

$$\begin{aligned} \lim_{\bar{\sigma} \rightarrow 0} Bias(\bar{\sigma}) &= -\frac{\lim_{\bar{\sigma} \rightarrow 0} \bar{\sigma} \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}{\lim_{\bar{\sigma} \rightarrow 0} [1 - \Phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)]} \\ &= -\frac{\lim_{\bar{\sigma} \rightarrow 0} \left[\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right) - \frac{\beta_{\min}}{\bar{\sigma}^2} \bar{\sigma} \phi'\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)\right]}{\lim_{\bar{\sigma} \rightarrow 0} \left[\frac{\beta_{\min}}{\bar{\sigma}^2} \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)\right]} \\ &= -\frac{\lim_{\bar{\sigma} \rightarrow 0} \left[\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right) + \left(\frac{\beta_{\min}}{\bar{\sigma}}\right)^2 \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)\right]}{\lim_{\bar{\sigma} \rightarrow 0} \left[\frac{\beta_{\min}}{\bar{\sigma}^2} \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)\right]} \\ &= -\left[\lim_{\bar{\sigma} \rightarrow 0} \frac{\bar{\sigma}^2 \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}{\beta_{\min} \phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)} + \lim_{\bar{\sigma} \rightarrow 0} \beta_{\min} \frac{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)}{\phi\left(\frac{\beta_{\min}}{\bar{\sigma}}\right)} \right] \\ &= \beta_{\min}, \end{aligned}$$

³²Concretely, we can write:

$$\begin{aligned} \int_{|\beta_{\min}|}^{\infty} -2 \frac{\beta^2}{\bar{\sigma}^2} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta &= \int_{|\beta_{\min}|}^{\infty} \beta \times \frac{\partial}{\partial \beta} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta \\ &= \beta \times \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) \Big|_{|\beta_{\min}|}^{\infty} - \int_{|\beta_{\min}|}^{\infty} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta \\ &= |\beta_{\min}| - \int_{|\beta_{\min}|}^{\infty} \exp\left(\frac{\beta_{\min}^2 - \beta^2}{\bar{\sigma}^2}\right) d\beta \end{aligned}$$

where the third line followed by the property of normal density that $\phi'(x) = -x\phi(x)$.

Thus, the most precise study is unbiased as $\sigma_i \rightarrow 0$ when $\sigma_0 = 0$ if and only if $\beta_{\min} \leq 0$. \blacksquare

Proof of bias under uniform selection. We derive the monotonicity and limit results from the definition of truncated normal distribution: we can write the bias as

Denoting $\bar{\beta} = \frac{t\sigma_i - b_0}{\bar{\sigma}}$ and $\underline{\beta} = \frac{-t\sigma_i - b_0}{\bar{\sigma}}$, we can write the bias as

$$\begin{aligned} \text{Bias}(\bar{\sigma}) &= \bar{\sigma} \frac{\eta_1 [-\phi(\underline{\beta}) + \phi(\bar{\beta})] - \eta_0 [\phi(\bar{\beta}) - \phi(\underline{\beta})]}{\eta_1 [\Phi(\underline{\beta}) + 1 - \Phi(\bar{\beta})] + \eta_0 [\Phi(\bar{\beta}) - \Phi(\underline{\beta})]} \\ &= -\bar{\sigma}^2 \frac{\partial \ln \{ \eta_1 + (\eta_1 - \eta_0) \Delta\Phi(b_0) \}}{\partial b_0}, \end{aligned}$$

where $\Delta\Phi(b_0) = \Phi(\bar{\beta}) - \Phi(\underline{\beta}) > 0$.

(i) Monotonicity: let us write (from the expression above, we can write

$$\text{Bias}(\bar{\sigma}) = \bar{\sigma} \frac{(\eta_1 - \eta_0)}{\eta_1 + (\eta_1 - \eta_0) \Delta\Phi(b_0)} K(\sigma_i),$$

where $K(\sigma_i) = \phi(\bar{\beta}) - \phi(\underline{\beta})$

- By an assumption $\eta_1 - \eta_0 > 0$, we know that $\frac{(\eta_1 - \eta_0)}{\eta_1 + (\eta_1 - \eta_0) \Delta\Phi(b_0)}$ is increasing in σ_i .
- We show that there exists some range $[0, \underline{\sigma}]$ such that $K(\sigma_i)$ will be increasing in σ_i :

$$\frac{\partial K(\sigma_i)}{\partial \sigma_i} = \phi'(\bar{\beta}) \frac{\partial \bar{\beta}(\sigma_i)}{\partial \sigma_i} - \phi'(\underline{\beta}) \frac{\partial \underline{\beta}(\sigma_i)}{\partial \sigma_i}$$

By the definitions above,

$$\begin{aligned} \frac{\partial \bar{\beta}(\sigma_i)}{\partial \sigma_i} &= \frac{t\bar{\sigma} - (t\sigma_i - b_0) \times \frac{1}{2}\bar{\sigma}^{-1} \times 2\sigma_i}{\bar{\sigma}^2} = \frac{t}{\bar{\sigma}} - \frac{(t\sigma_i - b_0)\sigma_i}{\bar{\sigma}^3} \\ \frac{\partial \underline{\beta}(\sigma_i)}{\partial \sigma_i} &= \frac{-t\bar{\sigma} - (-t\sigma_i - b_0) \times \frac{1}{2}\bar{\sigma}^{-1} \times 2\sigma_i}{\bar{\sigma}^2} = -\frac{t}{\bar{\sigma}} + \frac{(t\sigma_i + b_0)\sigma_i}{\bar{\sigma}^3} \end{aligned}$$

Substituting,

$$\begin{aligned}\frac{\partial K(\sigma_i)}{\partial \sigma_i} &= \phi'(\bar{\beta}) \left[\frac{t}{\bar{\sigma}} - \frac{(t\sigma_i - b_0)\sigma_i}{\bar{\sigma}^3} \right] - \phi'(\underline{\beta}) \left[-\frac{t}{\bar{\sigma}} + \frac{(t\sigma_i + b_0)\sigma_i}{\bar{\sigma}^3} \right] \\ &= \frac{1}{\bar{\sigma}} \left\{ t [\phi'(\underline{\beta}) - \phi'(\bar{\beta})] + t \frac{\sigma_i}{\bar{\sigma}} [\phi'(\underline{\beta}) \underline{\beta} - \phi'(\bar{\beta}) \bar{\beta}] \right\}\end{aligned}$$

When σ_i is small, term $[\phi'(\underline{\beta}) - \phi'(\bar{\beta})]$ determines the sign of $\frac{\partial K(\sigma_i)}{\partial \sigma_i}$. Since $\bar{\beta} > \underline{\beta}$, $\phi'(\underline{\beta}) > \phi'(\bar{\beta})$, and thus, $\frac{\partial K(\sigma_i)}{\partial \sigma_i} > 0$. On the other hand, When σ_i is large, the term, $\phi'(\underline{\beta}) \underline{\beta} - \phi'(\bar{\beta}) \bar{\beta} = \phi(\bar{\beta}) \bar{\beta}^2 - \phi(\underline{\beta}) \underline{\beta}^2$ will be important, and can be negative since $\phi(\underline{\beta}) > \phi(\bar{\beta})$ when the thresholds are at the tail of normal distribution.

(ii) Limit: since the cumulative distribution function is continuously differentiable, we can analyze by distributing the limit:

$$\begin{aligned}\lim_{\sigma_i \rightarrow 0} Bias(\sigma_i) &= - \lim_{\sigma_i \rightarrow 0} \bar{\sigma}^2 \times \frac{\partial \ln \{ \eta_1 + (\eta_1 - \eta_0) \lim_{\sigma_i \rightarrow 0} \Delta \Phi(b_0) \}}{\partial b_0} \\ &= -\sigma_0^2 \times \frac{\partial \ln \{ \eta_1 + (\eta_1 - \eta_0) [\Phi(\frac{-b_0}{\bar{\sigma}}) - \Phi(\frac{-b_0}{\bar{\sigma}})] \}}{\partial b_0} \\ &= 0\end{aligned}$$

Thus, $\lim_{\sigma_i \rightarrow 0} Bias(\sigma_i) = 0$ for any parameter values. ■

□

Appendix B. Supplementary Numerical Discussions

Appendix A has provided various analytical proofs. Due to limited analytical tractability, however, this paper has extensively employed numerical approach. Appendix B provides details of numerical simulations and presents some additional results: B1 will illustrate equilibrium thresholds under the general environments, and B2 describes details of empirical tests.

B1. Thresholds under General Environments

This Section describes the simulation of thresholds, $\bar{\beta}(\sigma_i)$ and $\underline{\beta}(\sigma_i)$, in a more general environment and in an equilibrium that is the main focus of the analysis. The environment is more general since the simulation can consider settings with $N \geq 3$, $c > 0$, $\sigma_b < \infty$, and heterogeneous values of σ_i . While Proposition 2.1 concerning

the thresholds under heterogeneous σ_i , for analytical tractability, focused on the symmetric equilibrium such that $\bar{\beta}(\sigma_i) = -\underline{\beta}(\sigma_i)$, the numerical analysis can explore the properties of the asymmetric equilibrium with $\bar{\beta}(\sigma_i) < -\underline{\beta}(\sigma_i)$.

This analysis will show when the analytical results are robust to alternative environments. B1.1 will first describe the overview of simulation algorithm; B1.2 shows some additional results regarding omission; B1.3 shows that the threshold $\underline{\beta}(\sigma)$ need not be concave when c is high; B1.4 explores the implication of N , the number of researchers, on the thresholds; B1.5 summarizes the magnitude of omission, bias, and welfare consequences of various reporting rules.

B1.1 Simulation Step Overview

We compute the equilibrium thresholds, $\bar{\beta}(\sigma_i), \underline{\beta}(\sigma_i)$, that are symmetric between N researchers, given primitive environments' parameters such as threshold policy effectiveness, c , and number of researchers, N , as well as underlying variance, σ_b^2 , and between-study heterogeneity, σ_0^2 . By discretizing the support of standard errors to $\{\sigma_1, \sigma_2, \dots, \sigma_S\}$ with $\sigma_1 \leq \sigma_2 \leq \dots \leq \sigma_S$, the equilibrium thresholds at each standard error, $\bar{\beta}(\sigma_i), \underline{\beta}(\sigma_i)$, is given by a system of $2 \times S$ equations:

$$\bar{\beta}(\sigma_i) = \frac{\sigma_i^2}{\mathbb{E}^{\sigma_{-i}} \left[\frac{1}{\sum_{j=0}^N \sigma_j^{-2}} \middle| Piv_1 \right]} \left\{ c - \mathbb{E}^{\sigma_{-i}} \left[\frac{\mathbb{E}^{m_{-i}} \left[\frac{\mathbb{E}^{\beta_{-i}} \sum \frac{\beta_{-i}}{\sigma_{-i}^2} \middle| Piv_1 \right]}{\sum_{j=0}^N \sigma_j^{-2}} \middle| Piv_1 \right] \right\} \quad (1.30)$$

$$\underline{\beta}(\sigma_i) = \frac{\sigma_i^2}{\mathbb{E}^{\sigma_{-i}} \left[\frac{1}{\sum_{j=0}^N \sigma_j^{-2}} \middle| Piv_0 \right]} \left\{ c - \mathbb{E}^{\sigma_{-i}} \left[\frac{\mathbb{E}^{m_{-i}} \left[\frac{\mathbb{E}^{\beta_{-i}} \sum \frac{\beta_{-i}}{\sigma_{-i}^2} \middle| Piv_0 \right]}{\sum_{j=0}^N \sigma_j^{-2}} \middle| Piv_0 \right] \right\} \quad (1.31)$$

where the Piv_m for $m \in \{0, 1\}$ denotes the other's message realization such that the researchers' switch between \emptyset and m is pivotal. Concretely, denoting the number of others' positive results and negative results as n'_1, n'_0 respectively, Piv_1 is $n'_1 = n'_0$ and Piv_0 is $n'_1 = n'_0 + 1$ in the asymmetric equilibrium with supermajoritarian rule (1.2).

The algorithm solves the above system of $2 \times S$ equations with $2 \times S$ unknowns iteratively by inner and outer loops. The inner loop computes $\mathbb{E}^{m_{-i}} [\cdot | Piv_m]$ for every combination of m_{-i} in Piv_m ; the outer loop computes $\mathbb{E}^{\sigma_{-i}} [\cdot]$ for every $\sigma_{-i} \in \{\sigma_1, \sigma_2, \dots, \sigma_S\}^N$. Since there is no analytical solution of mean of correlated multivariate normal distribution, $\mathbb{E}^{\beta_{-i}} [\cdot]$, the algorithm used numerical integration with rejection sampling. The iterative adjustment takes the estimated thresholds under $N - 1$ researchers as an input conjecture, $\beta_{conjecture}$, and computes the updated thresh-

olds, β_{new} , by $\beta_{new} = \Delta\beta_{sol} + (1 - \Delta)\beta_{conjecture}$, where Δ is a step of adjustment, looping over every σ_i . The initial values for $N = 2$ are some linear functions $\bar{\beta}(\sigma) = A\sigma + c$, $\underline{\beta}(\sigma) = -A\sigma + c$; but the thresholds are not sensitive to the choice of $A > 0$. The algorithm stops when the updates, $|\beta_{new} - \beta_{conjecture}|$, are smaller than some tolerance level.

For a sufficiently large S that permits fine grid for $Supp(\sigma)$, the computational time increases exponentially as N increases. This is because dimensions of the inputs into computation increase exponentially: the weights on probabilities given message realizations take S^N dimensions and the message realizations take $\sum_{k=0}^{\lfloor N/2 \rfloor} \binom{N}{k} \times (N - 1)$ dimensions to compute. Moreover, we have set $\Delta = 0.5$ and tolerance level to be 0.05. For the simulation with heterogeneous priors, we chose $\{\sigma_1, \sigma_2, \dots, \sigma_S\} = \{0.1, 0.2, \dots, 1\}$ so that $S = 10$. Due to limitations of feasibility, the simulation with heterogeneous thresholds compute only up to $N = 4$.

To approximate some real-world settings with reasonable algorithms, we choose a distribution of $G(\sigma)$ close to the distribution of σ in the labor union data set (Doucouliagos et al. 2018). Since the observed distribution of σ in the data set is the distribution with publication selection, we impute the underlying distribution with the positive significant results from the example of labor union ($G(\sigma) = \frac{1}{C} \sum_i \left(1 - \Phi\left(\frac{1.96\sigma_i - \hat{b}_0}{\sqrt{\sigma_i^2 - \hat{\sigma}_0^2}}\right)\right)^{-1} \mathbb{1}(\sigma \geq \sigma_i)$, where $C = \sum_i \left(1 - \Phi\left(\frac{1.96\sigma_i - \hat{b}_0}{\sqrt{\sigma_i^2 - \hat{\sigma}_0^2}}\right)\right)^{-1}$ is the normalizing constant, and $\{\hat{b}_0, \hat{\sigma}_0^2\}$ are estimated with the stem-based method. The largest standard error is normalized to be 1. The figure shows that χ^2 distribution with 2 degrees of freedom with support $[0, 4]$ approximates the empirical distribution of variance, σ^2 , reasonably.




Figure B1: Approximation of empirical distribution

Notes: Figure B1 plots the imputed empirical distribution of variance, σ^2 against χ^2 distribution with 2 degrees of freedom with support of $[0, 4]$ normalized to support of $[0, 1]$.

The simulation henceforth will incorporate between-study heterogeneity, σ_0^2 , on equilibrium thresholds. When there is study-specific effects on underlying benefits, $b_i = b + \zeta_i$ with $\zeta_i \sim \mathcal{N}(0, \sigma_0^2)$, the estimates are generated by $\beta_i = b_i + \epsilon_i = b + \zeta_i + \epsilon_i$. Thus, given the estimated standard error σ_i due to the sampling variance, the true variance, σ_i^{2*} , satisfies $\sigma_i^{2*} = \sigma_i^2 + \sigma_0^2$. As the formula shows, we can consider these heterogeneities by shifting the values of inverse variance weights used in Bayesian updating.

B1.2 Numerical results on omission

The following two figures show that $\mathbb{P}(a = 1) < \frac{1}{2}$ under supermajoritarian rule, and that the welfare attained under supermajoritarian rule is higher than those under submajoritarian rule.




Figure B2: Likelihood of policy implementation

Notes: Figure B2 plots the probability of policy implementation $\mathbb{P}(a = 1)$ for $N = 2$, $c = 0$, $\sigma_b = \frac{1}{3}$, and various values of standard error of signal, σ , in the equilibrium characterized by Proposition 1. It shows that, relative to the policy implementation probability under communication of estimates, $\mathbb{P}(a = 1) = \frac{1}{2}$, policy is slightly less likely to be implemented. This is primarily due to the conservative rule of supermajoritarian voting rule $a^* = 1 \Leftrightarrow n_1 > n_0$, largely mitigated by the thresholds $\bar{\beta}, \underline{\beta}$ that lead to the upward bias of the estimates that underlie reported studies.




Figure B3: Optimality of supermajoritarian voting rule

Notes: Figure B3 plots the welfare (measured as a fraction of benchmark case with communication of estimates) under a supermajoritarian rule ($a^* = 1 \Leftrightarrow n_1 > n_0$) and a submajoritarian rule ($a^* = 1 \Leftrightarrow n_1 \geq n_0$) for $N = 2$, $\sigma_b = \sigma = 1$, and various values of $c \geq 0$. It shows that the supermajoritarian rule attains higher welfare than the submajoritarian rule for $c > 0$, and identical welfare for $c = 0$. The supermajoritarian rule is better than the submajoritarian rule especially when c is high and thus there is large relative welfare loss.

B1.3 Shape of $\underline{\beta}(\sigma)$ under high $c > 0$

Simulations show that, while Proposition 2.1 suggested that $\underline{\beta}(\sigma_i)$ will be concave in the symmetric equilibrium with $c = 0$, it can be convex in the asymmetric equilibrium with large $c > 0$ and supermajoritarian voting rule. Figure B2 illustrates this in a setting with $N = 2$, $c = 2$, $\sigma_b = \frac{3}{4}$, $\sigma_0 = 0$ and distribution $G(\sigma)$ that approximates an empirical distribution as discussed in Section B1.1. This result suggests that the pattern of omission and inflation may be very different between positive and negative results.




Figure B4: Example of convex $\underline{\beta}(\sigma)$

Notes: Figure B4 plots the thresholds $\bar{\beta}(\sigma)$ and $\underline{\beta}(\sigma)$ for $N = 2$, $c = 2$, $\sigma_b = \frac{3}{4}$, $\sigma_0 = 0$ and distribution $G(\sigma)$ that approximates an empirical distribution as discussed in Section B1.1.

This result is driven by the prior belief of effectiveness, $\mathbb{E}b$, that is more conservative than the target effectiveness, c . By Bayes' rule, when the estimates can be

conveyed, it is optimal to implement the policy if and only if the average signal $\hat{\beta}$ satisfies $\hat{\beta} \geq \frac{c - \mathbb{E}b}{\sigma_b^2} \frac{\sigma_i^2}{N}$: that is, whenever the prior is conservative so that $c > \mathbb{E}b$, the required level of average signal $\hat{\beta}$ is convex in σ_i^2 . When c is large, this force can dominate the effect of less information as characterized in Proposition 2.1, turning the threshold $\underline{\beta}(\sigma_i)$ to be convex rather than concave.

This result clarifies that the omission is not driven by the uninformedness of researchers per se, but by the lack of strong belief in whether it is optimal to implement the policy. In the most stark example, the omission probability approaches zero as the researchers' signals become imprecise ($\sigma_i \rightarrow \infty$). This observation clarifies the discussion of informedness and abstention in voting theories (Feddersen and Pesendorfer 1999): the lack of information needs not arise from lack of signal, but can also arise from lack of strong prior belief.

B1.4 Implications of Many Researchers (High N)

Due to the analytical tractability, the propositions did not examine the implications of many researchers on the equilibria. This sub-Section shows that the supermajoritarian rule holds even with high N , and analyzes how the omission and bias change given high N .

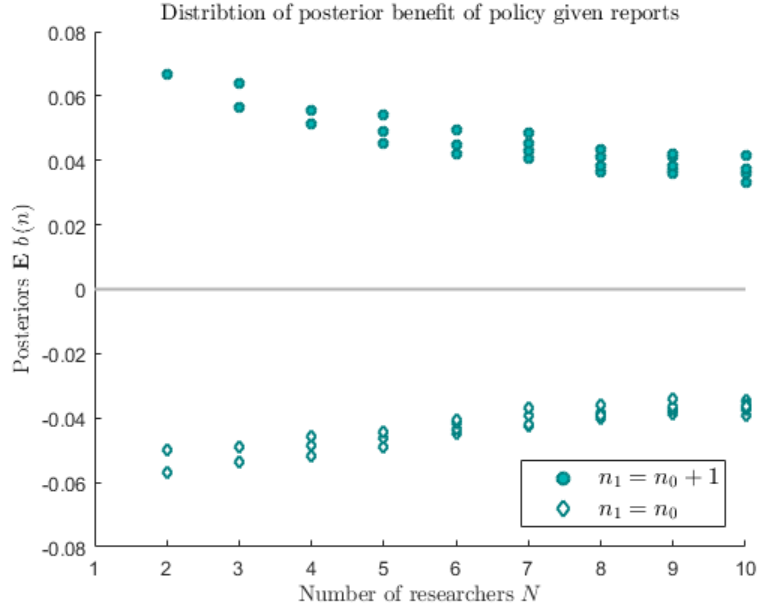


Figure B5: consistency of policymakers’ posterior beliefs

Notes: Figure B5 plots the distribution of posterior belief $\mathbb{E}b(n)$ in the equilibrium of supermajoritarian policy rule ($a^* = 1 \Leftrightarrow n_1 \geq n_0$) with $\sigma_b = \sigma = 0$ and $c = 0$ for $N = 2, \dots, 10$. It shows that, when there are marginally more positive results than negative results ($n_1 = n_0 + 1$), then the posterior belief is positive; conversely, when there are equal numbers of positive and negative results ($n_1 = n_0$), then the posterior belief is negative. Since the posterior beliefs are monotone in the number of positive and negative results, this confirms that the conjectured supermajoritarian policy rule ($a^* = 1 \Leftrightarrow n_1 \geq n_0$) is consistent with the belief and utility maximization of the policymaker. This result of consistency of supermajoritarian policy rule holds for various values of c . Here, the analysis restricts to the case of constant σ due to the computational feasibility.

First, the policymaker’s supermajoritarian rule ($a^* = 1 \Leftrightarrow n_1 > n_0$) holds in a communication equilibrium even for $N = 3, \dots, 10$ (Figure B5). This suggests that even if the policymaker does not know the number of underlying studies, they can still compare the number of reported positive vs negative results to make the decisions they would have taken with knowledge of N . While this result may rely on risk neutrality as will be discussed in sensitivity analysis, it suggests that the assumption that the policymaker knows N needs not be critical.³³

³³The set-up needs to maintain the assumption that the researchers know how many other researchers exist on the same subject. First, this appears to be closer to actual scientific practice than the assumption that the readers also know number of researchers. Second, the Figure B4 demon-

Second, the researchers' omission probability gradually increases as N increases. The Figure B6 depicts the example equilibrium thresholds for $N = 1, \dots, 4$ keeping all other environment constant when (A) σ_b is high and σ_0 is high, (B) σ_b is low and σ_0 is low, and (C) σ_b is high and σ_0 is low. In all cases, the probability of omission conditional on study precision, $\mathbb{P}(m_i = \emptyset | \sigma_i)$, increases with N for any values of σ_i . This is because, as N increases, the total information owned by other researchers rises and leaving the decisions to others' papers becomes more desirable.³⁴

Nevertheless, the bias on the coefficients that underlie reported studies may increase or decrease as increase in N may shift the thresholds in either directions.³⁵ This is because there are three channels through which N alters the thresholds. As Figure B6(A) shows, the effect of changing pivotality condition shifts $\underline{\beta}$ upwards and $\overline{\beta}$ in ambiguous directions, potentially mitigating the bias as N increases.³⁶ As Figure B6(B) shows, the decreasing effect of conservative prior shifts both $\overline{\beta}$ and $\underline{\beta}$ downwards. When there are more researchers, each researcher needs less extreme signals to overturn the default decision. As Figure B6(C) suggests, there is also the effect of equilibrium thresholds adjustment that shifts both $\overline{\beta}$ and $\underline{\beta}$ in directions that offset the effect of the first two effects. For example, $\underline{\beta}$ may shift downwards due to upward shift in $\overline{\beta}$ especially among noisy studies, increasing the bias. These considerations jointly determine the conditions under which the increase in number of researchers, N , may increase or decrease the bias.

strates that the thresholds do not change qualitatively with N and the changes are not quantitatively large: thus, even if there were uncertainty in N from researchers' perspective, they may still be able to choose approximately optimal thresholds.

³⁴This result is consistent with the literature on voting theory that shows that the abstention probability increases as the number of voters increases.

³⁵This inquiry relates to the literature on media that explores the impact of market competition on media bias, and finds that the higher number of competing senders may increase the bias arising from taste and decrease bias arising from reputation motives (Gentzkow et al., 2016).

³⁶Consider, for example, the pivotality conditions for $N = 2$ and $N = 3$. $\underline{\beta}(N = 2)$ is lower than $\underline{\beta}(N = 3)$ because it satisfies the indifference condition (1.30) when one researcher receives high signal as opposed to when one receives high signal and another receives intermediate signal. $\overline{\beta}(N = 2)$ may be lower or higher than $\overline{\beta}(N = 3)$ because it satisfies the indifference condition (1.31) when only one another researcher receives intermediate negative signal as opposed to two researchers receiving intermediate negative signals or one receiving positive and another receiving negative signals.




Figure B6. $\bar{\beta}(\sigma), \underline{\beta}(\sigma)$ thresholds under various σ_b and σ_0

Notes: Figure B6 plots the $\bar{\beta}(\sigma), \underline{\beta}(\sigma)$ thresholds under $c = 2$ and distribution $G(\sigma)$ that approximates an empirical distribution as discussed in Section B1.1, and varying N, σ_b and σ_0 . In each figure, the convex solid lines are $\bar{\beta}(\sigma)$ and concave solid lines are $\underline{\beta}(\sigma)$. There is only one threshold for $N = 1$ since there is no benefit of omission when there is only one researcher.

B1.5 Quantifying bias, omission, and welfare

The discussions of bias, omission, and welfare in Section 2.4.2 were based on a simulation under various parameters. This Appendix describes the simulation more fully. The results are presented in Table B1.

Table B1: Bias, omission, and welfare

		(1)	(2)	(3)	(4)	(5)
		Baseline	High σ_b	High σ_0	High c	High N
(A) Bias:	i. overall	0.42	0.27	0.22	-0.78	0.38
	ii. $\sigma_i = 0.1$	0.06	0.19	0.00	-0.31	0.04
	iii. $\sigma_i = 1$	0.93	0.40	0.73	-1.06	0.92
(B) $\mathbb{P}(\text{omission})$:	i. overall	0.49	0.46	0.33	0.54	0.53
	ii. $\sigma_i = 0.1$	0.24	0.41	0.07	0.36	0.26
	iii. $\sigma_i = 1$	0.69	0.53	0.62	0.65	0.77
(C) Welfare:	i. unrestricted	0.95	0.95	0.99	0.77	0.93
	ii. restricted	0.85	0.85	0.97	0.46	0.84
Specification changes from baseline		-	$\sigma_b = 1$	$\sigma_0 = 1$	$c = .05$	$N = 3$

Notes: Table B1 summarizes (A) bias $\mathbb{E}[\beta_i | m_i \neq \emptyset, \sigma_i]$, (B) omission probability $\mathbb{P}(m_i = \emptyset | \sigma_i)$, and (C) ex-ante welfare $\mathbb{E}[a^*(m^*(\beta_i, \sigma_i)) [\mathbb{E}b(\beta_i, \sigma_i) - c]]$ under various settings. In (A) and (B), i. overall values are expected values unconditional on σ_i realization; ii focuses on the most precise studies ($\sigma_i = 0.1$); and iii focuses on the least precise studies ($\sigma_i = 1$). In (C), welfare is computed as a fraction of full information welfare. i. unrestricted refers to the environment without linear t -statistics rule whereas ii. restricted refers to the hypothetical setting in which no omission is allowed and the threshold is restricted to follow $\bar{\beta}(\sigma) = t\sigma$. Here, t is computed to be the optimal value. Since there is no omission by the exogenous restriction, bias and omission probability are both zero. The baseline specification applies $\sigma_b = 0.25$, $\sigma_0 = 0.1$, $c = 0$, and $N = 2$. Columns (2)-(4) modifies this environment for each specification as presented. The simulation is based on heterogeneous $G(\sigma)$ that approximates an empirical distribution as discussed in Section B1.1.

In many specifications, Panel (A) suggests sizable upward bias in unweighted³⁷

³⁷Note that this is different from the meta-analysis estimates with Bayes' rule, which puts higher weight on the more precise studies. Many meta-analysis studies often discuss these unweighted estimates.

average estimates. Consistent with Figures 2, these biases are driven by noisy studies. In contrast, most precise studies have very small biases. Note, for large $c > 0$, there can be downward bias because the studies whose coefficients near c are omitted. Quantitatively, despite its symmetric set-up with $c = 0$, among the least precise studies, the model can explain the bias greater than one standard deviation of the underlying distribution of benefits.

Panel (B) shows that the average omission rate can be as high as roughly 30 to 50 percents. While most precise studies have only 7 percent omission rate, most noisy studies may be omitted roughly 70 to 80 percents.

Finally, the model can quantify welfare gains from permitting publication bias relative to imposing restriction that every binary conclusion of null hypothesis testing needs to be reported. As discussed in the introduction, coarse communication can be largely welfare-reducing, and even with sophisticated readers who compute the posterior correctly and flexible adjustment of t -statistics, there is roughly 3 to 30 percent welfare loss relative to the full information benchmark. The gain from allowing for some omission and inflation is to roughly halve these welfare losses to 1 to 12 percents. This analysis demonstrates that the welfare consequences can be quantitatively important.

B2. Estimation and Testing Steps

Section 3.2.2 has described the overview of the estimation and testing steps of the semi-parametric Kolmogorov-Smirnov (KS)-type test. This Appendix Section adds additional description and discussion of computing the p -values in this test, and provides some results supplementary to Section 3.3.2.

The computation of KS statistics requires estimating the two theoretical distributions: $G_\theta(\sigma)$, the distribution of standard errors of null results, and $H_0(\beta)$, the distribution of coefficients of negative results. The estimates use the stem-based estimates of $\{\hat{b}_0, \hat{\sigma}_0\}$ and apply

$$\tilde{G}_\theta(\sigma|\hat{b}_0, \hat{\sigma}_0) = \frac{1}{C} \sum_{i||\beta_i| \geq \bar{t}\sigma_i} \frac{\Phi\left(\frac{\bar{t}\sigma_i - \hat{b}_0}{\hat{\sigma}}\right) - \Phi\left(\frac{-\bar{t}\sigma_i - \hat{b}_0}{\hat{\sigma}}\right)}{\left[1 - \Phi\left(\frac{\bar{t}\sigma_i - \hat{b}_0}{\hat{\sigma}}\right)\right] \mathbf{1}(\beta_i \geq \bar{t}\sigma_i) + \Phi\left(\frac{\bar{t}\sigma_i - \hat{b}_0}{\hat{\sigma}}\right) \mathbf{1}(\beta_i < \bar{t}\sigma_i)} \mathbf{1}(\sigma \geq \sigma_i) \quad (1.32)$$

$$\tilde{H}_0\left(\beta|\hat{b}_0, \hat{\sigma}_0\right) = \frac{1}{n_0} \sum_{i|\beta_i \geq \bar{\tau}\sigma_i} \min \left\{ \frac{\Phi\left(\frac{\beta - \hat{b}_0}{\bar{\sigma}}\right)}{\Phi\left(\frac{\bar{\tau}\sigma_i - \hat{b}_0}{\bar{\sigma}}\right)}, 1 \right\}, \quad (1.33)$$

$$\text{where } C = \sum_{i|\beta_i \geq \bar{\tau}\sigma_i} \frac{\Phi\left(\frac{\bar{\tau}\sigma_i - \hat{b}_0}{\sqrt{\sigma_i^2 + \hat{\sigma}_0^2}}\right) - \Phi\left(\frac{-\bar{\tau}\sigma_i - \hat{b}_0}{\sqrt{\sigma_i^2 + \hat{\sigma}_0^2}}\right)}{\left[1 - \Phi\left(\frac{\bar{\tau}\sigma_i - \hat{b}_0}{\sqrt{\sigma_i^2 + \hat{\sigma}_0^2}}\right)\right] \mathbb{1}(\beta_i \geq \bar{\tau}\sigma_i) + \Phi\left(\frac{\bar{\tau}\sigma_i - \hat{b}_0}{\sqrt{\sigma_i^2 + \hat{\sigma}_0^2}}\right) \mathbb{1}(\beta_i < \bar{\tau}\sigma_i)} \text{ and } \bar{\sigma} = \sqrt{\sigma_i^2 + \sigma_0^2}.$$

We can understand these formula by considering how many null or negative studies in some intervals of parameters there must have been in order to have the number of observed positive or significant studies. For example, let us consider some interval with length $\Delta > 0$ that has n_1 positive studies. If the mean and variance of underlying normal distribution is given by $\{b_0, \sigma_0\}$, then in expectation, there must have been $\frac{\Phi\left(\frac{\bar{\tau}\sigma_i - b_0}{\bar{\sigma}}\right) - \Phi\left(\frac{-\bar{\tau}\sigma_i - b_0}{\bar{\sigma}}\right)}{1 - \Phi\left(\frac{\bar{\tau}\sigma_i - b_0}{\bar{\sigma}}\right)} \times n_1$ null results. The formulas (1.32) and (1.33) are constructed with this logic.

***p*-values:** we wish to compute the probability of observing the discrepancy between observed ($\hat{G}_\theta(\sigma)$ and $\hat{H}_0(\beta)$) and predicted distributions, defined by the KS-type statistics, $D^G = \sup \left\{ \hat{G}_\theta(\sigma) - \tilde{G}_\theta(\sigma|\hat{b}_0, \hat{\sigma}_0) \right\}$ and $D^H = \sup \left\{ \hat{H}_0(\beta) - \tilde{H}_0(\beta|\hat{b}_0, \hat{\sigma}_0) \right\}$. We cannot apply the standard KS tests since they compare either one theoretical and one empirical, or two empirical distributions; here, the predicted distribution contains uncertainties not only in studies used for estimation but also in the estimates of parameters $\{\hat{b}_0, \hat{\sigma}_0\}$; ignoring the uncertainty in two-step estimation (Newey and McFadden 1994) may underestimate the *p*-values.

The *p*-value of this test equals the average *p*-values given each value of $\{b_0, \sigma_0\}$ simulated given errors in their estimates. This is because the overall *p*-value is defined as the probability that the maximum difference between the empirical and predicted distributions is at least as large as the observed difference. For each draw of $\{b_0, \sigma_0\}$ and resultant predicted distribution, the algorithm applies the inverse cumulative distribution function method to generate a simulated distribution with sample size n_0 . Then, the *p*-value given each value of $\{b_0, \sigma_0\}$ is the fraction of simulated estimates such that their KS statistic is at least as large as D^G and D^H respectively. The test computes one-sided *p*-values, and repeats the simulation until the estimated *p*-value converges.

The bootstrap estimates are appropriate since the parameters $\{b_0, \sigma_0\}$ are not the extremum statistics of the distribution. Since the stem-based method treats as a nuisance parameter, the estimation employs the bootstrap method to obtain the distribution of $\{\hat{b}_0, \hat{\sigma}_0\}$ estimates. Since each study has equal level of information

on the distribution of study-specific effects, b_i , each study has equal weight in the bootstrap method. The stem-based method suggests $\{\hat{b}_0, \hat{\sigma}_0\} = \{-0.02, 0.05\}$ in this data set. In addition, the KS-type statistics are $D^G = \hat{G}_\theta(.08) - \tilde{G}_\theta(.08) = 0.40$ and $D^H = \hat{H}_0(-.13) - \tilde{H}_0(-.13) = 0.15$.

Chapter 2

Power Laws in Superspreading Events: Evidence from Coronavirus Outbreaks and Implications for SIR Models

2.1 Introduction

On March 10th, 2020, choir members were gathered for their rehearsal in Washington. While they were all cautious to keep distance from one another and nobody was coughing, three weeks later, 52 members had COVID-19, and two passed away. There are numerous similar anecdotes worldwide.¹ Many studies have shown that the average basic reproduction number (\mathcal{R}_0) is around 2.5-3.0 for this coronavirus (e.g. [Liu et al., 2020](#)), but 75% of infected cases do not pass on to any others ([Nishiura et al., 2020](#)). The superspreading events (SSEs), wherein a few primary cases infect an extraordinarily large number of others, are responsible for the high average number. As SSEs were also prominent in SARS and MERS before COVID-19, epidemiology research has long sought to understand them (e.g. [Shen et al., 2004](#)). In particular, various parametric distributions of infection rates have been proposed, and their variances have been estimated in many epidemics under an assumption that they exist (e.g. [Lloyd-Smith et al., 2005](#)). On the other hand, stochastic Susceptible-Infectious-Recovered (SIR) models have shown that, as long as the infected population is moderately large, the idiosyncratic uncertainties of SSEs will cancel out each other ([Roberts](#)

¹See Table 2.5.2 in Appendix for a list of several examples.

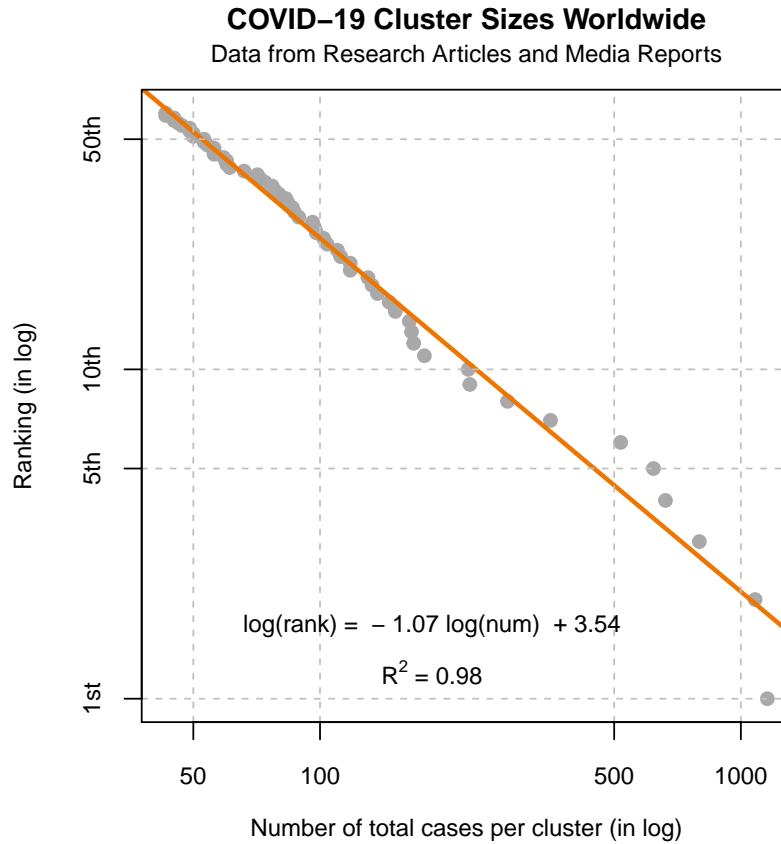
et al., 2015). That is, following the Central Limit Theorem (CLT), stochastic models quickly converge to their deterministic counterparts, and become largely predictable. From this perspective, the dispersion of SSEs is unimportant in itself, but is useful only to the extent it can help target lockdown policies to focus on SSEs to efficiently reduce the average rates \mathcal{R}_0 (Endo et al., 2020).

In this paper, we extend this research by closely examining the distribution of infection rates, and rethinking how its dispersion influences the uncertainties of aggregate dynamics. Using data from SARS, MERS, and COVID-19 from around the world, we provide consistent evidence that SSEs follow a power law, or Pareto, distribution with fat tails, or infinite variance. That is, the true variance of infection rates cannot be empirically estimated as any estimate will be an underestimate however large it may be. When the CLT assumption of finite variance does not hold, many theoretical and statistical implications of epidemiology models will require rethinking. Theoretically, even when the infected population is large, the idiosyncratic uncertainties in SSEs will persist and lead to large aggregate uncertainties. Statistically, the standard estimate of the average reproduction number (\mathcal{R}_0) may be far from its true mean, and the standard errors will understate the true uncertainty. Because the infected population for COVID-19 is already large, our findings have immediate implications for statistical inference and current policy.

We begin with evidence. Figure 2-1 plots the largest clusters reported worldwide for COVID-19 from data gathered by Leclerc et al. (2020). If a random variable follows a power law distribution with an exponent α , then the log of its scale (e.g. a US navy vessel had 1,156 cases tested positive) and the log of its severity rank (e.g. that navy case ranked 1st in severity) will have a linear relationship, with its slope indicating $-\alpha$. Figure 2-1 shows a fine fit of the power law distribution.² Moreover, the slope is very close to 1, indicating a significant fatness of the tail to the extent that is analogous to natural disasters such as earthquakes (Gutenberg and Richter, 1954) that are infrequent but can be extreme³. While data collection through media reports may be biased towards extreme cases, analogous relationships consistently hold for other SARS, MERS, and COVID-19 data based on surveillance data, with

²In Appendix 2.5, we also estimate the exponent with a small sample bias correction proposed by Gabaix and Ibragimov (2011), which shows the exponent is 1.16, and the R^2 is 0.98. With maximum likelihood estimation, the exponent is 1.01. When using the Kolmogorov-Smirnov test (Clauset et al., 2009), the p-value given $\alpha = 1.01$ is 0.75, failing to reject the null hypothesis that the empirical observation arises from the power law distribution. On the other hand, the p-value given $\alpha = 2$ is 0.000, rejecting the null hypothesis that the distribution is observed from power law distribution with a finite variance.

³The power law distribution with $\alpha = 1$ is called the Zipf's law.



Source: CMMID COVID-19 Working Group online database (Leclerc et al., 2020)

Figure 2-1: Log cluster size vs log rank for COVID-19 worldwide

Notes: Figure 2-1 plots the number of total cases per cluster (in log) and their ranks (in log) for COVID-19, last updated on June 3rd. It fits a linear regression for the clusters with size larger than 40. The data are collected by the Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group (Leclerc et al., 2020).

exponents often indicating fat tails. Note that other distributions, including the negative binomial distributions commonly applied in epidemiology research, cannot predict these relationships, and significantly underestimate the risks of extremely severe SSEs.

Using fat-tailed power law distributions, we show that stochastic SIR models predict substantial uncertainties in aggregate epidemiological outcomes. Concretely, we consider a stochastic model with a population of one million, whereby a thousand people are initially infected, and apply epidemiological parameters adopted from the literature. We consider effects of tails of distribution while keeping the average rate (\mathcal{R}_0) constant. Under thin-tailed distributions, such as the estimated negative bino-

mial distribution or power law distribution with $\alpha = 2$, the epidemiological outcomes will be essentially predictable. However, under fat-tailed distributions close to those estimated in the COVID-19 data worldwide ($\alpha = 1.1$), there will be immense variations in all outcomes. For example, the peak infection rate is on average 14%, but its 90th percentile is 31% while its 10th percentile is 4%. Under thin-tailed distribution such as negative binomial distribution, the average, 90th percentile and 10th percentile of the peak infection is all concentrated at 26-27%, generating largely deterministic outcomes.

While our primary focus was on the effect on aggregate uncertainty, we also find important effects on average outcomes. In particular, under a fat-tailed distribution, the cumulative and peak infection, as well as the herd immunity threshold, will be lower, and the timing of outbreak will come later than those under a thin-tailed distribution, *on average*. For example, the average herd immunity threshold is 66% with thin-tailed distribution, it is 39% with a fat-tailed distribution. These observations suggest that the increase in aggregate uncertainty over \mathcal{R}_0 has effects analogous to a decrease in average \mathcal{R}_0 . This relationship arises because the average future infection will be a *concave* function of today's infection rate: because of concavity, mean preserving spread will lower the average level. In particular, today's higher infection rate has two countering effects: while it increases the future infection, it also decreases the susceptible population, which decreases it. We provide theoretical interpretations for each outcome by examining the effect of mean-preserving spread of \mathcal{R}_0 in analytical results derived in deterministic models.

Our findings have critical implications for the design of lockdown policies to minimize the social costs of infection. Here, we study lockdown policies that target SSEs. We assume that the maximum size of infection rate can be limited to a particular threshold (e.g., 50, 100, or 1000 per day) with some probabilities by banning large gatherings. Because both the uncertainty and mean of the infection rate in the fat-tailed distribution are driven by the tail events, such policies substantially lower the uncertainty and improve the average outcomes. Because the cost of such policy⁴ is difficult to estimate reliably, we do not compute the cost-effectiveness of such policy. Nonetheless, we believe this is an important consideration in the current debates on how to re-open the economy while mitigating the risks of subsequent waves.

Finally, we also show the implications of a fat-tailed distributions for the estimation of the average infection rate. Under such a distribution with small sample

⁴For example, it is prohibitively costly to shut down daycare, but it is less costly to prevent a large concert.

sizes, the sample mean yields estimates that are far from the true mean and standard errors that are too small. To address such possibility, it will be helpful to estimate the power law exponent. If the estimate indicates a thin-tailed distribution, then one can be confident with the sample mean estimate. If it indicates a fat-tailed distribution, then one must be aware that there is much uncertainty in the estimate not captured by its confidence interval. While such fat-tailed distributions cause notoriously difficult estimation problems, we explore a “plug-in” method that uses the estimated exponent. Such estimators generate median estimates closer to the true mean with adequate confidence intervals that reflect the substantial risk of SSEs.

Related Literature. First, our paper belongs to a large literature on stochastic epidemiological models. The deterministic SIR model was initiated by [Kermack and McKendrick \(1927\)](#), and later, [Bartlett \(1949\)](#) and [Kendall \(1956\)](#) developed stochastic SIR models (see [Britton \(2010, 2018\)](#), [Britton et al. \(2015\)](#) for surveys). The traditional view of the stochastic SIR model is that while useful when the number of infected is small, once the infected population is moderately large, it behaves similarly to the deterministic model due to the CLT. [Britton \(2010\)](#) writes “Once a large number of individuals have been infected, the epidemic process may be approximated by the deterministic counter-part.” [Roberts \(2017\)](#) also considers an SIR model with small fluctuations of epidemiological parameters, but shows that deterministic models approximate its average reasonably. Here, we consider large aggregate fluctuations arising from idiosyncratic shocks and show that even the average deviates significantly from predictions of deterministic models. There are recent applications of stochastic SIR models that study the very beginning of COVID-19 outbreaks when the number of infection is small (for example, [Abbott et al. \(2020\)](#), [Karako et al. \(2020\)](#), [Simha et al. \(2020\)](#) and [Bardina et al. \(2020\)](#)). However, the major modeling effort has been to use deterministic models based on the common justification above. Our point is that when the distribution is fat-tailed, which we found an empirical support for, the CLT no longer applies, and hence the stochastic model behaves qualitatively differently from its deterministic counterpart even with a large number of infected individuals.

Second, the empirical importance of SSEs is widely recognized in the epidemiological literature before COVID-19 ([Lloyd-Smith et al., 2005](#); [Galvani and May, 2005](#)) and for COVID-19 ([Frieden and Lee, 2020](#); [Endo et al., 2020](#)). These papers fit the parametric distribution that is by construction thin-tailed, such as negative binomial distribution. It has been common to estimate “the dispersion parameter k ” of the negative binomial distribution. We argue that the fat-tailed distribution provides a

better fit to the empirical distribution of SSEs, in which a tail parameter, α , parsimoniously captures the fatness of the tail. A recent contribution by [Cooper et al. \(2019\)](#) consider Pareto rule in the context of malaria transmission, but they nonetheless estimate the dispersion with finite variance for the entire infections.

Third, our paper also relates to studies that incorporate heterogeneity into SIR models, incorporating differences in individual characteristics or community structures. Several recent papers point out that the permanent heterogeneity in individual infection rates lower the herd immunity threshold ([Gomes et al., 2020](#); [Hébert-Dufresne et al., 2020](#); [Britton et al., 2020](#)). Although we obtain a similar result, our underlying mechanisms are distinct from theirs. In our model, there is no ex-ante heterogeneity across individuals, and thus their mechanism is not present. [Zhang et al. \(2013\)](#) and [Szabó \(2020\)](#) consider a model in which individuals have heterogeneous infection rates that follow power laws in scale-free networks, but their heterogeneity is permanent (i.e. due to individual characteristics). Instead, what matters for us is the aggregate fluctuations in \mathcal{R}_0 (i.e. due to idiosyncratic variations in environments), which their models do not exhibit. Some recent papers emphasize the importance of age-dependent heterogeneity and its implications for lockdown policies ([Acemoglu et al., 2020](#); [Davies et al., 2020](#); [Gollier, 2020](#); [Rampini, 2020](#); [Glover et al., 2020](#); [Brotherhood et al., 2020](#)). We emphasize another dimension of targeting: targeting toward large social gatherings, and this policy reduces the uncertainty regarding various epidemiological outcomes. [Roberts \(2013\)](#) analyzes a deterministic model in which basic reproduction number is estimated with noise, and derives probability distributions over epidemiological outcomes due to the uncertainty of the estimates.

Finally, it is well-known that many variables follow a power law distribution. These include the city size ([Zipf, 1949](#)), the firm size ([Axtell, 2001](#)), income ([Atkinson et al., 2011](#)), wealth ([Kleiber and Kotz, 2003](#)), consumption ([Toda and Walsh, 2015](#)) and even the size of the earthquakes ([Gutenberg and Richter, 1954](#)), the moon craters and solar flares ([Newman, 2005](#)). Regarding COVID-19, [Beare and Toda \(2020\)](#) document that the cumulative number of infected population across cities and countries is closely approximated by a power law distribution. They then argue that the standard SIR model is able to explain the fact. We document that the infection at the individual level follows a power law. We are also partly inspired by economics literature which argue that the fat-tailed distribution in firm-size has an important consequence for the macroeconomics dynamics, originated by [Gabaix \(2011\)](#). We follow the similar route in documenting that the SSEs are well approximated by a power law distribution and arguing that such empirical regularities have important

consequences for the epidemiological dynamics.

Roadmap. The rest of the paper is organized as follows. Section 2.2 documents evidence that the distribution of SSEs follows power law. Section 2.3 embed the evidence into an otherwise standard SIR models to demonstrate its implications for the epidemiological dynamics. Section 2.4 studies estimation of the reproduction numbers under fat-tailed distribution. Section 2.5 concludes by discussing what our results imply for ongoing COVID-19 pandemic.

2.2 Evidence

We present evidence from SARS, MERS, and COVID-19 that the SSEs follow power law distributions. Moreover, our estimates suggest the distributions are often fat-tailed, with critical implications for the probabilities of extreme SSEs. Evidence also suggests a potential role of policies in reducing the tail distributions.

2.2.1 Statistical model

Let us define the SSEs and their distribution. Following the notations of [Lloyd-Smith et al. \(2005\)](#), let $z_{it} \in \{0, 1, 2, \dots\}$ denote the number of secondary cases⁵ an infected individual i has at time t . Then, given some threshold \underline{Z} , an individual i is said to have caused SSE at time t if $z_{it} \geq \underline{Z}$. To make the estimation flexible, suppose the distribution for non-SSEs, $z_{it} < \underline{Z}$, needs not follow the same distribution as those for SSEs.

In this paper, we consider a power law (or Pareto) distribution on the distribution of SSE. Denoting its exponent by α , the countercumulative distribution is

$$\mathbb{P}(z_{it} \geq Z) = \pi (Z/\underline{Z})^{-\alpha} \quad \text{for } Z \geq \underline{Z}, \quad (2.1)$$

where π is the probability of SSEs. Notably, its mean and variance may not exist when α is sufficiently low: while its mean is $\frac{\alpha}{\alpha-1}\underline{Z}$ if $\alpha > 1$, it is ∞ if $\alpha \leq 1$. While its variance is $\frac{\alpha}{(\alpha-1)^2(\alpha-2)}\underline{Z}^2$ if $\alpha > 2$, it is ∞ if $\alpha \leq 2$. In this paper, we formally call a distribution to be fat-tailed if $\alpha < 2$ so that they have infinite variance. While non-existence of mean and variance may appear pathological, a number of socioeconomic and natural phenomenon such as city sizes ($\alpha \approx 1$), income ($\alpha \approx 2$), and earthquake energy ($\alpha \approx 1$) have tails well-approximated by this distribution as reviewed in the

⁵Note that the number of “secondary” cases include only direct transmissions and exclude indirect transmissions. This is how the COVID-19 data in Figure 2-1 were also collected ([Leclerc et al., 2020](#)).

Introduction. One concrete example⁶ that can explain a power law distribution is due to the result in [Beare and Toda \(2019\)](#): suppose each participant can invite some others with some probability. Conditional on inviting, the number of people each participant invites follows some distributions such as log-normal distribution. Then, the resulting distribution of all participants follows a power law.

This characteristics stands in contrast with the standard assumption in epidemiology literature that the full distribution of z_{it} follows a negative binomial (or Pascal) distribution⁷ with finite mean and variance. The negative binomial distribution has been estimated to fit the data better than Poisson or geometric distribution for SARS ([Lloyd-Smith et al., 2005](#)), and given its theoretical bases from branching model (e.g. [Gay et al., 2004](#)), it has been a standard distributional assumption in the epidemiology literature (e.g. [Nishiura et al., 2017](#)).

2.2.2 Data

This paper uses five datasets of recent coronavirus outbreaks for examining the distribution of SSEs: COVID-19 data from (i) around the world, (ii) Japan, and (iii) India, and (iv) SARS data, (v) MERS data.

(i) COVID-19 data from around the world: this dataset contains clusters of infections found by a systematic review of academic articles and media reports, conducted by the Centre of the Mathematical Modelling of Infectious Diseases COVID-19 Working Group ([Leclerc et al., 2020](#)). The data are restricted to first generation of cases, and do not include subsequent cases from the infections. The data are continuously updated, and in this draft, we have used the data downloaded on June 3rd. There were a total of 227 clusters recorded.

(ii) COVID-19 data from Japan: this dataset contains a number of secondary cases of 110 COVID-19 patients across 11 clusters in Japan until February 26th, 2020, reported in [Nishiura et al. \(2020\)](#). This survey was commissioned by the Ministry of Health, Labor, and Welfare of Japan to identify high risk transmission cases.

(iii) COVID-19 data from India: this dataset contains the state-level data

⁶Another theoretical reason why this distribution could be relevant for airborne diseases is that the number of connections in social networks often follow a power law ([Barabasi and Frangos, 2014](#)).

⁷Denoting its mean by R and dispersion parameter by k , the distribution is

$$\mathbb{P}(z_{it} \geq Z) = 1 - \sum_{z=0}^Z \frac{\Gamma(z+k)}{z! \Gamma(k)} \left(\frac{\mathcal{R}}{k}\right)^z \left(1 + \frac{\mathcal{R}}{k}\right)^{-(z+k)}$$

The variance of this distribution is $\mathcal{R} \left(1 + \frac{\mathcal{R}}{k}\right)$. The distribution nests Poisson distribution (as $k \rightarrow \infty$) and geometric distribution (when $k = 1$.)

collected by the Ministry of Health and Family Welfare, and individual data collected by covid19india.org.⁸ We use the data downloaded on May 31st.

(iv) SARS from around the world: this dataset contains 15 incidents of SSEs from SARS in 2003 that occurred in Hong Kong, Beijing, Singapore, and Toronto, as gathered by [Lloyd-Smith et al. \(2005\)](#)⁹ through a review of 6 papers. The rate of community transmission was not generally high so that, for example, the infections with unknown route were only about 10 percent in the case of Beijing. The data consist of SSEs, defined by epidemiologists ([Shen et al., 2004](#)) as the cases with more than 8 secondary cases. For Singapore and Beijing, the contact-tracing data is available from [Hsu et al. \(2003\)](#) and [Shen et al. \(2004\)](#), respectively. When compare the fit to the negative binomial distribution, we compare the fit of power law to that of negative binomial using these contact tracing data.

(v) MERS from around the world: this dataset contains MERS clusters reported up to August 31, 2013. The cases are classified as clusters when they are linked epidemiologically. The data come from three published studies were used in [Kucharski and Althaus \(2015\)](#). Total of 116 clusters are recorded.

We use multiple data sets in order to examine the robustness of findings.¹⁰ Having multiple data sets can address each other’s weaknesses in data. While data based on media reports is broad, they may be skewed to capture extreme events; in contrast, data based on contact tracing may be reliable, but are restricted to small population. By using both, we can complement each data’s weaknesses.

2.2.3 Estimation

The datasets report cumulative number of secondary cases, either $\sum_i z_{it}$ (when a particular event may have had multiple primary cases) or $\sum_t z_{it}$ (when an individual infects many others through multiple events over time). Denoting these cumulative numbers by Z , we consider this distribution for some $Z \geq \underline{Z}^*$. As discussed in Appendix 2.5, we can interpret the estimates of this tail distribution as approximately the per-period and individual tail distribution and therefore map directly to the pa-

⁸<https://www.kaggle.com/sudalairajkumar/covid19-in-india>. covid19india.org is a volunteer-based organization that collects information from municipalities.

⁹Even though [Lloyd-Smith et al. \(2005\)](#) had analyzed 6 other infectious diseases, SARS was the only one with sufficient sample sizes to permit reliable statistical analyses.

¹⁰the infectious diseases considered here share some commonalities as SARS-CoV that causes SARS, MERS-CoV that causes MERS, and SARS-CoV-2 that causes COVID-19 are human coronaviruses transmitted through the air. They have some differences in terms of transmissibility, severity, fatality, and vulnerable groups ([Petrosillo et al., 2020](#)). But overall, as they are transmitted through the air, they are similar compared to other infectious diseases.

parameter of the SIR model in the next section. The thresholds for inclusion, \underline{Z}^* , will be chosen to match the threshold for SSEs when possible, but also adjust for the sample size. For COVID-19 in the world, we apply $\underline{Z} = 40$ to focus on the tail of the SSE distribution. For SARS, we apply $\underline{Z} = 8$ as formally defined (Shen et al., 2004). For other samples, we apply $\underline{Z} = 2$ because the sample size is limited.

To assess whether the distribution of Z follows the power law, we adopt the regression-based approach that is transparent and commonly used. If Z follows power law distribution, then by (2.1), the log of Z and the log of its underlying rank have a linear relationship: $\log \text{rank}(Z) = -\alpha \log Z + \log(N\pi\underline{Z}^\alpha)$. This is because, when there are N individuals, the expected ranking of a realized value Z is $\mathbb{E} \text{rank}(Z) \simeq \mathbb{P}(z \geq Z)N$ for moderately large N . Thus, when N is large, we obtain a consistent estimate of α by the following regression:

$$\log \text{rank}(Z) = -\alpha \log Z + \log(N\pi\underline{Z}^\alpha) + \varepsilon \quad (2.2)$$

When N is not large, however, the estimate will exhibit a downward bias because log is a concave function and thus $\mathbb{E} \log \text{rank}(Z) < \log \mathbb{E} \text{rank}(Z)$. While we present the analysis according to (2.2) in Figures 2-1 and 2-2 for expositional clarity, we also report the estimates with small sample bias correction proposed by Gabaix and Ibragimov (2011) in Appendix 2.5.¹¹ We also estimate using the maximum likelihood in Appendix 2.5. Note that when there are ties (e.g. second and third largest had 10 infections), we assigned different values to each observation (e.g. assigning rank of 2 and 3 to each observation).

Next, we also compare the extent to which a power law distribution can approximate the distribution of SSEs adequately relative to the negative binomial distribution. First, we plot what the predicted log-log relationship in (2.2) would be given the estimated parameters of negative binomial distribution.¹² Second, to quantify the predictive accuracy, we compute the ratio of likelihood of observing the actual data.

¹¹Their approach is to turn the dependent variable into $\log [\text{rank}(Z) - \frac{1}{2}]$ instead of $\log [\text{rank}(Z)]$. We examine the performance of their bias correction method through a estimating regression given random variables generated from power law distributions. While their bias correction almost eliminates bias when N is moderately large, it has an upward bias of α whereas the equation (2.2) has a downward bias. The magnitude of bias is similar when $N = 10$ or $N = 15$. Thus, our preferred approach is to refer to both methods for robustness.

¹²This approach stands in contrast with a common practice to plot the probability mass functions. Unlike such approaches where differences in tail densities are invisible since it is very close to zero, this approach highlights the differences in tail densities.

2.2.4 Results

Our analysis shows that the power law finely approximates the distribution of SSEs. Figure 2-1 visualizes this for COVID-19 from across the world, and Figure 2-2 for SARS, MERS, and COVID-19 in Japan and India. Their R^2 range between 0.93 and 0.99, suggesting high levels of fit to the data. Because our focus is on upper-tail distribution, Figure 2-1 truncates below at the cluster size 40, Figure 2-2 truncates at 8 for SARS and at 2 for MERS and COVID-19 in India and Japan. Figure 2.5.1 in Appendix presents a version of Figure 2-1 truncated below at 20.

In addition, the estimates of regression (2.2) suggest that the power law exponent, α , is below 2 and even close to 1. Table 2.1 summarizes the main findings. The estimated exponents near 1 suggest that extreme SSEs are not uncommon. For COVID-19 in Japan and India, the estimated exponents are larger than 1 but often below 2. Since applying the threshold of $\underline{Z}^* = 2$ is arguably too low, we must interpret out-of-sample extrapolation from these estimates with caution. When higher thresholds are applied, the estimated exponents tend to be higher. For example, when applying the threshold of $\underline{Z}^* = 8$ as in SARS 2003 to COVID-19 in India, the estimated exponent is 1.85 or 2.25. This pattern is already visible in Figure 2-2. Table 2.5.1 in Appendix 2.5 presents results using bias correction technique of [Gabaix and Ibragimov \(2011\)](#) as well as maximum likelihood. The results are very similar.

Notably, the estimated exponent of India is higher than those of other data. There are two possible explanations. First, the lockdown policies in India have been implemented strictly relative to moderate approaches in Japan and some other parts of the world during the outbreaks. By discouraging and prohibiting large-scale gatherings, sometimes by police enforcement, they may have been successful at targeting SSEs. Second, contact tracing to ensure data reliability may have been more difficult in India until end of May than in Japan until end of February.¹³ While missing values will not generate any biases if the attritions were proportional to the number of infections, large gatherings may have dropped more than in Japan where the SSEs were found through contact tracing. Nonetheless, these estimates suggest that various environments and policies could decrease the risks of the extreme SSEs. This observation motivates our policy simulations to target SSEs.

¹³Concretely, there were only 248 cases of more than one secondary infections reported in the data among 27,890 primary cases in the data from India. That is, only 0.8 percents of primary cases were reported to have infected more than one persons. In contrast, there were 27 cases with more than one secondary infections among 110 primary cases in Japan. That is, 25 percent of primary cases were infectious. This difference in ration likely reflects the data collection quality than actual infection dynamics.

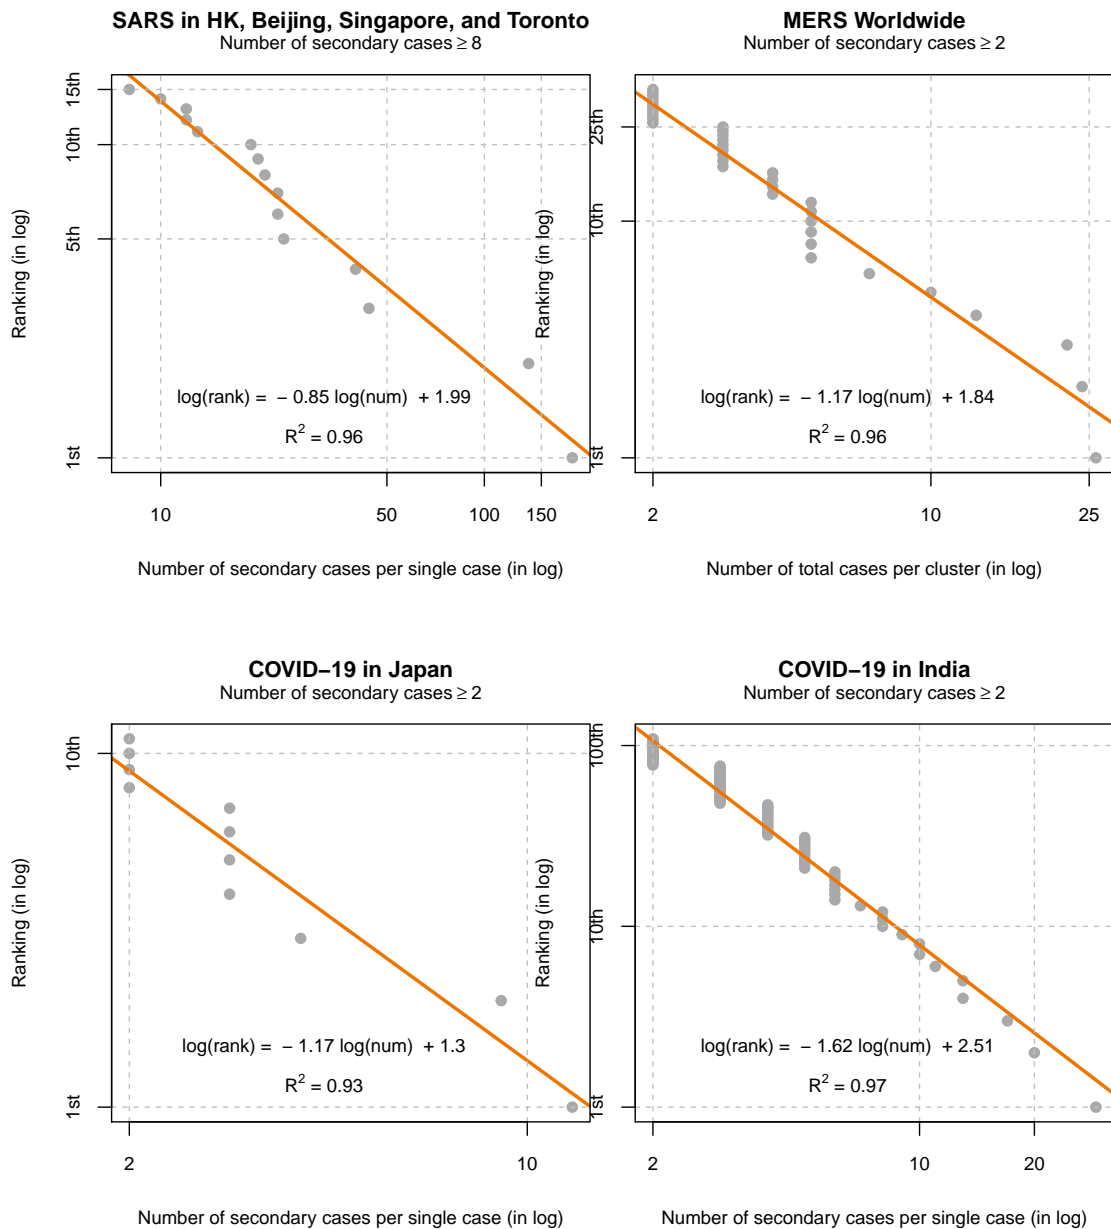


Figure 2-2: Log size vs log rank for COVID-19

Notes: Figure 2-2 plots the number of total cases per cluster (in log) and their ranks (in log) for MERS, and the number of total cases per cluster (in log) and their ranks (in log) for SARS and COVID-19 in Japan and India. The data for SARS are from [Lloyd-Smith et al. \(2005\)](#), and focus on SSEs defined to be the primary cases that have infected more than 8 secondary cases. The data for MERS come from [Kucharski and Althaus \(2015\)](#). The data for Japan comes from periods before February 26, 2020, reported in [Nishiura et al. \(2020\)](#). The data for India are until May 31, 2020, reported by the Ministry of Health and Family Welfare, and [covid19india.org](#). The plots are restricted to be the cases larger than 2.

	COVID-19			SARS			MERS
	World	Japan	India	World	Singapore	Beijing	World
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\hat{\alpha}$	1.07	1.17	1.62	0.85	0.75	0.75	1.17
	(0.04)	(0.10)	(0.03)	(0.06)	(0.08)	(0.06)	(0.07)
\underline{Z}	40	2	2	8	2	2	2
Obs.	60	11	109	15	19	8	36
R ²	0.98	0.93	0.97	0.96	0.91	0.94	0.96
$\log_{10} LR$	-	11.39	-	-	19.51	8.04	40.89

Table 2.1: Estimates of power law exponent ($\hat{\alpha}$) and their fit with data

Notes: Table 2.1 summarizes the estimates of power law exponent ($\hat{\alpha}$) given as the coefficient of regression of log of number of infections (or size of clusters) on the log of their rankings. Heteroskedasticity-robust standard errors are reported in the parenthesis. \underline{Z} denote the threshold number of infection to be included. $\log_{10}(LR)$ denotes “likelihood ratios”, expressed in the log with base 10, of probability of observing this realized data with power law distributions relative to that with estimated negative binomial distributions. Columns (1)-(3) report estimates for COVID-19; columns (4)-(6) for SARS, and column (7) for MERS.

Next, we compare the assumption of power law distribution relative to that of a negative binomial distribution. Figure 2-3 shows that the negative binomial distributions would predict that the extreme SSEs will be fewer than the observed distribution: while it predicts the overall probability of SSEs accurately, they suggest that, when they occur, they will not be too extreme in magnitude. Table 1 reports the relative likelihood, in logs, of observing the data given the estimated parameters. It shows that, under the estimated power law distribution relative to the estimated negative binomial distribution, it is $10^8 - 10^{20}$ times more likely to observe the SARS data (10^{40} times more for MERS, and 10^{11} times more COVID-19 data in Japan). Such large differences emerge because the negative binomial distribution, given its implicit assumption of finite variance, suggests that the extreme SSEs are also extremely rare when estimated with entire data sets¹⁴. If our objective is to predict the overall incidents of infections parsimoniously, then negative binomial distribution is well-validated and theoretically founded (Lloyd-Smith et al., 2005).¹⁵ However, if

¹⁴For example, the binomial distribution estimate suggests an incidence of 185 cases (residential infection in Hong Kong) only has a chance of 9.5×10^{-10} occurring for any single primary case.

¹⁵Since the power law distribution is fitted only to SSEs, estimated power law distribution may

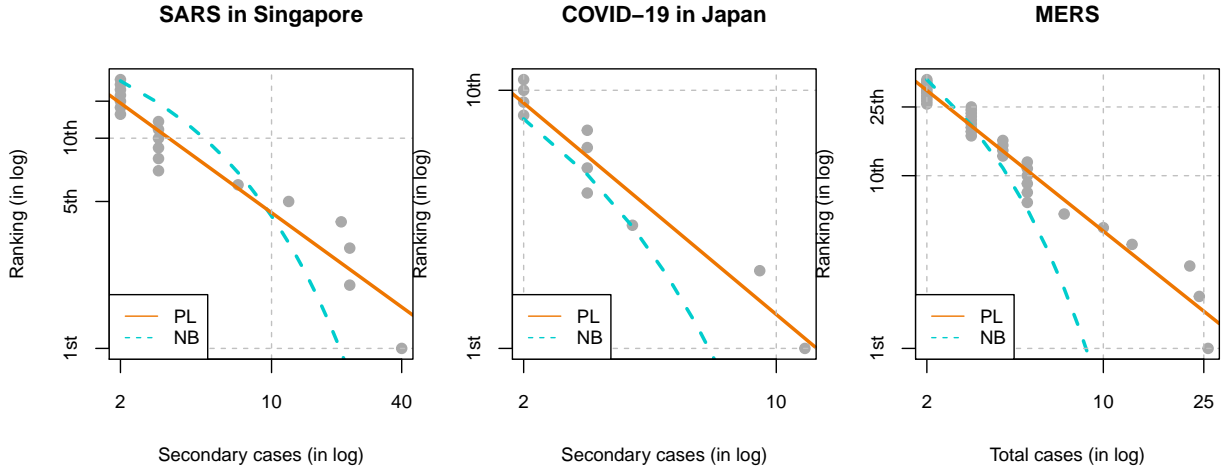


Figure 2-3: Comparison of power law and negative binomial distributions

Notes: Figure 2-3 plots the predicted ranking of infection cases given the estimated negative binomial (NB) distribution, in addition to the log-log plots and estimated power law (PL) distributions. The negative binomial distribution is parameterized by (R, k) , where R is mean and k is the dispersion parameter with the variance being $R(1 + R/k)$. The estimates for SARS Singapore come from our own estimates using the maximum likelihood ($R = 0.88, k = 0.09$); MERS come from the world ($R = 0.47, k = 0.26$) estimated in [Kucharski and Althaus \(2015\)](#); and COVID-19 in Japan were from our own estimates using the maximum likelihood ($R = 0.56, k = 0.21$). The estimates of Singapore is slightly different from [Lloyd-Smith et al. \(2005\)](#) because we pool all the samples.

our goal is to estimate the risks of extreme SSEs accurately, then using only two parameters with finite variance to estimate together with the entire distribution may be infeasible.

These distributional assumptions have critical implications for the prediction of the extreme SSEs. Table 2.2 presents what magnitude top 1%, top 5%, and top 10% among SSEs will be given each estimates of the distribution. Given the estimates of the negative binomial distribution, even the top 1% of SSEs above 8 cases will be around the magnitude of 19-53. However, given a range of estimates from power law distribution, the top 1% could be as large as 569. Thus, it is no longer surprising that the largest reported case for COVID-19 will be over 1,000 people. In contrast,

fit the data better than the estimated negative binomial distribution that was meant to fit the entire data set. Rather than making such comparison, this estimation is intended to illustrate the magnitude of difference between the two distributional assumptions. Because of significant missing values for the low number of infections in the COVID-19 from across the world and India, we will not use the data sets for estimation of negative binomial distributions.

	Power Law					Negative Binomial		
	$\alpha = 1.08$	$\alpha = 1.1$	$\alpha = 1.2$	$\alpha = 1.5$	$\alpha = 2$	SARS	MERS	COVID-19
1%	569	526	371	172	80	44	18	19
5%	128	122	97	59	36	31	15	15
10%	67	65	55	37	25	25	13	14

Table 2.2: Probabilities of extreme SSEs under each distribution

Notes: Table 2.2 shows the size of secondary cases at each quantile, top 1 percentile, 5 percentile, and 10 percentile, given each distributions. The negative binomial distribution’s estimates for SARS are from Singapore, for COVID-19 are from Japan, and for MARS is from around the world.

such incidents have vanishingly low chance under binomial distributions. Since the SSEs are rare, researchers will have to make inference about their distribution based some parametric methods. Scrutinizing such distributional assumptions along with the estimation of parameters themselves will be crucial in accurate prediction of risks of extreme SSEs.

2.3 Theory

Motivated by the evidence, we extend an otherwise standard stochastic SIR model with a fat-tailed SSEs. Unlike with thin-tailed distributions, we show that idiosyncratic risks of SSEs induce aggregate uncertainties even when the infected population is large. We further show that the resulting uncertainties in infection rates have important implications for average epidemiological outcomes. Impacts of lockdown policies that target SSEs are discussed.

2.3.1 Stochastic SIR model with fat-tailed distribution

Suppose there are $i = 1, \dots, N$ individuals, living in periods $t = 1, 2, \dots$. Infected individuals pass on and recover from infection in heterogeneous and uncertain ways. Let β_{it} denote the number of new infection in others an infected individual i makes at time t . Let $\gamma_{it} \in \{0, 1\}$ denote the recovery/removal, where a person recovers ($\gamma_{it} = 1$) with probability $\gamma \in [0, 1]$. Note that, whereas z_{it} in Section 2.2 was a stochastic analogue of “effective” reproduction number, β_{it} here is such analogue of “basic reproduction number.” Assuming enough mixing in the population, these two models are related by $z_{it} = \beta_{it} \frac{S_t}{N}$, where S_t is a number of susceptible individuals in

the population.

This model departs from other stochastic SIR models only mildly: we consider a fat-tailed, instead of thin-tailed, distribution of infection rates. Based on the evidence, we consider a power law distribution of β_{it} : its countercumulative distribution is given by

$$\mathbb{P}(\beta_{it} \geq \beta) = \pi(\beta/\underline{\beta})^{-\alpha}$$

for the exponent α and a normalizing constant $\underline{\beta}$, and $\pi \in [0, 1]$ is the probability that $\beta \geq \underline{\beta}$. Note that the estimated exponent α can be mapped to this model, as discussed in Appendix 2.5. If we assume β_{it} is distributed according to exponential distribution or negative binomial distribution, we obtain a class of stochastic SIR models commonly studied in the epidemiological literature (see Britton (2010, 2018) for surveys). We will compare the evolution dynamics under this power law distribution against those under negative binomial distribution as commonly assumed, keeping the average basic reproduction number the same. To numerically implement this, we will introduce normalization to the distributions.

The evolution dynamics is described by the following system of stochastic difference equations. Writing the total number of infected and recovered/removed populations by I_t and R_t , we have

$$S_{t+1} - S_t = - \sum_{i=1}^{I_t} \beta_{it} \frac{S_t}{N} \tag{2.3}$$

$$I_{t+1} - I_t = \sum_{i=1}^{I_t} \beta_{it} \frac{S_t}{N} - \sum_{i=1}^{I_t} \gamma_{it} \tag{2.4}$$

$$R_{t+1} - R_t = \sum_{i=1}^{I_t} \gamma_{it}. \tag{2.5}$$

This system is a discrete-time and finite-population analogue of the continuous-time and continuous-population differential equation SIR models.

Parametrization: we parametrize the model as follows. The purpose of simulation is a proof of concept, rather than to provide realistic numbers. We take the length of time to be one week. We set the sum of the recovery and the death rate per day is 1/18 following Wang et al. (2020), so that $\gamma = 7/18$. The total population is set to $N = 10^5$, and initially infected population is 1% of the total population. As a benchmark case, we set $\alpha = 1.1$, which is in line with the estimates for the COVID-19

Parameter	Description	Value	Source
A. Common parameters			
γ	recovery & death rate	7/18	(Wang et al., 2020)
N	total population	10^5	
I_0	initially infected populatoion	10^3	1% of population
$\mathcal{R}_0 \equiv \mathbb{E}[\beta_{it}]/\gamma$	mean basic reproduction number	2.5	(Remuzzi and Remuzzi, 2020)
B. Power law			
π	probability of infecting	0.25	(Nishiura et al., 2020)
α	tail parameter	{1.08, 1.1, 1.2, 1.5, 2}	
C. Negative binomial			
k	overdispersion parameter	0.16	(Lloyd-Smith et al., 2005)

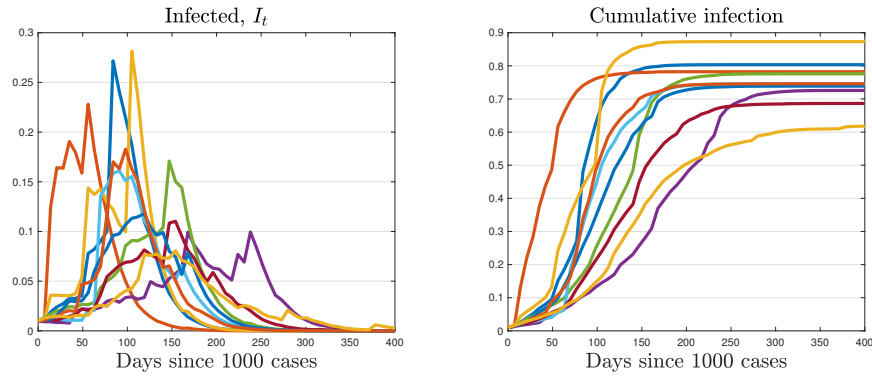
Table 2.3: Parameter values

data worldwide, but we explore several other parametrization, $\alpha \in \{1.08, 1.2, 1.5, 2\}$. As documented in Nishiura et al. (2020), 75% of people did not infect others. We therefore set $\pi = 0.25$. This number is also in line with the evidence from SARS reported in Lloyd-Smith et al. (2005), in which 73% of cases were barely infectious. We choose $\underline{\beta}$, which controls the mean of β_{it} , so that the expected $\mathcal{R}_0 \equiv \mathbb{E}\beta_{it}/\gamma$ per day is 2.5, corresponding to the middle of the estimates obtained in Remuzzi and Remuzzi (2020). This leads us to choose $\underline{\beta} = 0.354$ in the case of $\alpha = 1.1$.


We will contrast the above model to a model in which β_{it} is distributed according to negative binomial, $\beta_{it}/\gamma \sim \text{negative binomial}(\mathcal{R}_0, k)$. The mean of this distribution is $\mathbb{E}\beta_{it}/\gamma = \mathcal{R}_0$, ensuring that it has the same mean basic reproduction number as in the power law case, and the variance is $\mathcal{R}_0(1 + \mathcal{R}_0/k)$. The smaller values of k indicate greater heterogeneity (larger variance). We use the estimates of SARS by Lloyd-Smith et al. (2005), $k = 0.16$. The mean is set to the same value as power law case, $\mathcal{R}_0 = 2.5$,

2.3.2 Effects of fat-tailed distribution on uncertainty

Figure 2-4a shows 10 sample paths of infected population generated through the simulation of the model with $\alpha = 1.1$. One can immediately see that even though all the simulation start from the same initial conditions under the same parameters, there is enormous uncertainty in the timing of the outbreak of the disease spread, the maximum number of infected, and the final number of susceptible population. The timing of outbreak is mainly determined by when SSEs occur. To illustrate the



(a) Power law ($\alpha = 1.1$)



(b) Negative binomial

Figure 2-4: Ten sample paths from simulation

Note: Figure 2-4 plots 10 sample path of the number of infected population from simulation, in which we draw $\{\beta_{it}, \gamma_{it}\}$ randomly every period in an i.i.d. manner. Figure 2-4a plots the case with power law distribution, and Figure 2-4b plots the case with negative binomial distribution.

importance of a fat-tailed distribution, Figure 2-4b shows the same sample path but with a thin-tailed negative binomial distribution. In this case, as already 1,000 people are infected in the initial period, the CLT implies the aggregate variance is very small and the model is largely deterministic. This is consistent with Britton (2018). Britton (2018) shows that when the total population is as large as 1,000 or 10,000, the model quickly converges to the deterministic counterpart.

Figure 2-5 compares the entire distribution of the number of cumulative infection (top-left), the herd immunity threshold (top-right), the peak number of infected (bottom-left), and the days it takes to infect 5% of population (bottom-right). The herd immunity threshold is defined as the cumulative number of infected at which the number of infected people is at its peak. The histogram contrast the case with

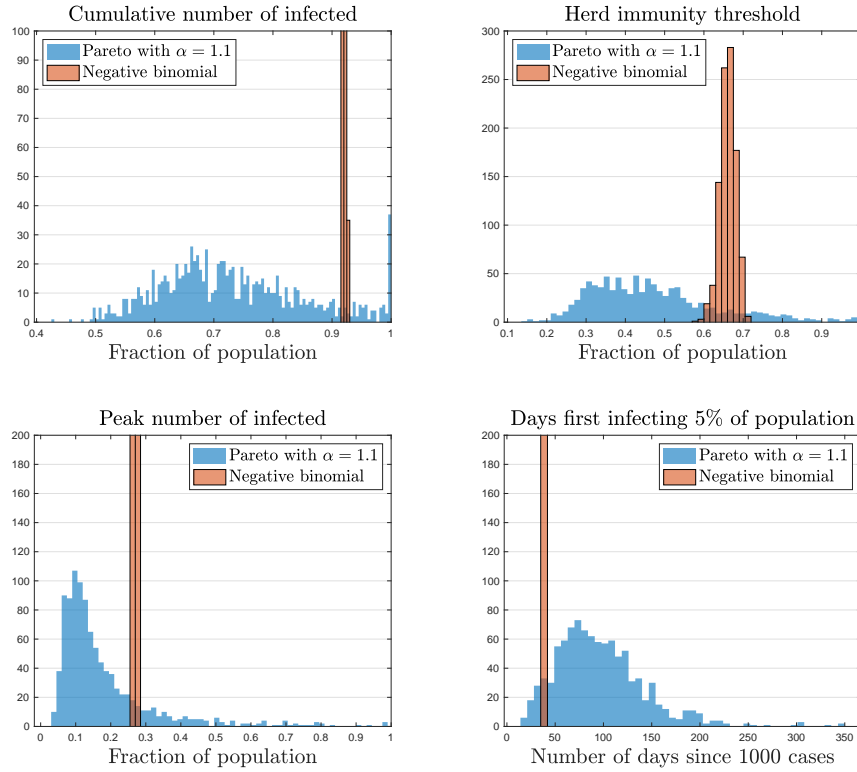


Figure 2-5: Histogram from 1000 simulation

Note: Figure 2-5 plots the histogram from 1000 simulations, in which we draw $\{\beta_{it}, \gamma_{it}\}$ randomly every period in an i.i.d. manner. The cumulative number of infected is S_T , where we take $T = 204$ weeks. The herd immunity threshold is given by the cumulative number of infected, at which the infection is at the peak. Formally, S_{t^*} where $t^* = \arg \max_t I_t$. The peak number of infected is $\max_t I_t$.

power law distribution with $\alpha = 1.1$ to the case with negative binomial distribution. It is again visible that uncertainty remains in all outcomes when the distribution of infection rate is fat-tailed. For example, the cumulative infection varies from 65% to 100% in the power law case, while the almost all simulation is concentrated around 92% in the case of negative binomial distribution.

Table 2.4 further shows the summary statistics for the epidemiological outcomes for various power law tail parameters, α , as well as for negative binomial distribution. With fat-tails, i.e. α close to one, the range between 90th percentile and 10th percentile for all statistics is wide, but this range is substantially slower as the tail becomes thinner (α close to 2). For example, when $\alpha = 1.08$ the peak infection rate

can vary from 6% to 32% as we move from 10th percentile to 90th percentile. In contrast, when $\alpha = 2$, the peak infection rate is concentrated at 26–27%. Moreover, when $\alpha = 2$, the model behaves similarly to the model with negative binomial distribution because the CLT applies to both cases.

2.3.3 Effects of fat-tailed distribution on average

While our primary focus was the effect on the uncertainty of epidemiological outcomes, Figure 2-5 also shows significant effects on the mean. In particular, fat-tailed distribution also lowers cumulative infection, the herd immunity threshold, the peak infection, and delays the time it takes to infect 5% of population, *on average*. Why could such effects emerge?

To understand these effects, we consider a deterministic SIR model with continuous time and continuum of population. In such a textbook model, we consider the effect of small uncertainties (i.e. mean-preserving spread) in \mathcal{R}_0 . Such theoretical inquiry can shed light on the effect because the implication of fat-tailed distribution is essentially to introduce time-varying fluctuation in aggregate \mathcal{R}_0 . We can thus examine how the outcome changes by \mathcal{R}_0 , and invoke Jensen’s inequality to interpret the results.¹⁶

1. **Effect on cumulative infection:** note that the cumulatively infected population is given by $1 - S_\infty/N$, where S_∞ is the ultimate susceptible population as $t \rightarrow \infty$. Taking the standard derivation, S_∞ satisfies the following equation:¹⁷

$$\log(S_\infty/N) = -\mathcal{R}_0(1 - S_\infty/N) \tag{2.6}$$

In Appendix 2.5, we prove that S_∞ is a convex function of \mathcal{R}_0 if $\mathcal{R}_0 > 1.125$, which is likely to be met in SARS or COVID-19.¹⁸ Thus, the cumulative infection is concave in \mathcal{R}_0 , and the mean-preserving spread in \mathcal{R}_0 lowers the cumulative infection.

2. **Effect on herd immunity threshold:** denoting the number of recovered/removed and infected population by R , the infection will stabilize when $\mathcal{R}_0 \left(\frac{N-R}{N} \right) = 1$.

¹⁶This assumes that \mathcal{R}_0 is drawn at time 0, and stay constant thereafter for each simulation. This exercise is not exactly the same as our original SIR model because there \mathcal{R}_0 fluctuates over time within a simulation. Thus this is for providing intuition, rather than a proof.

¹⁷Here, we set the initially recovered population to zero, $R_0 = 0$.

¹⁸Numerically, we did not find any counterexample even when $\mathcal{R}_0 \in [1, 1.125]$.

Rearranging this condition, the herd immunity threshold, R^* is given by

$$\frac{R^*}{N} = 1 - \frac{1}{\mathcal{R}_0}, \quad (2.7)$$

where $\mathcal{R}_0 \equiv \beta/\gamma$. Since R^* is concave in \mathcal{R}_0 , the mean-preserving spread in \mathcal{R}_0 lowers the herd immunity threshold.

3. **Effect on timing of outbreak:** let us consider the time t^* when some threshold of outbreak $(\frac{I}{N})^*$ is reached. Supposing $S/N \approx 1$ at the beginning of outbreak, t^* satisfies

$$\left(\frac{I}{N}\right)^* \approx \frac{I_0}{N} \exp\left(\frac{1}{\gamma}(\mathcal{R}_0 - 1)t^*\right) \quad (2.8)$$

Thus, t^* is convex in \mathcal{R}_0 , and the mean-preserving spread in \mathcal{R}_0 delays the timing of the outbreak.

4. **Effect on peak infection rate:** the peak infection rate, denoted by $\frac{I^{\max}}{N}$, satisfies

$$\frac{I^{\max}}{N} = 1 - \frac{1}{\mathcal{R}_0} - \frac{1}{\mathcal{R}_0} \log(\mathcal{R}_0 S_0), \quad (2.9)$$

where S_0 is initial susceptible population. We show in the Appendix that (2.9) implies that the peak infection, I^{\max}/N , is a concave function of \mathcal{R}_0 if and only if $\mathcal{R}_0 \geq \frac{1}{S_0} \exp(0.5)$. If we let $S_0 \approx 1$, this implies $\mathcal{R}_0 \geq \exp(0.5) \approx 1.65$. This explains why we found a reduction in peak infection rate, as we have assumed $\mathcal{R}_0 = 2.5$. Loosely speaking, since the peak infection rate is bounded above by one, it has to be concave for sufficiently high \mathcal{R}_0 .

Overall, we have found that the increase in the uncertainty over \mathcal{R}_0 has effects similar to a decrease in the level of \mathcal{R}_0 . This is because the aggregate fluctuations in \mathcal{R}_0 introduce negative correlation between the future infection and the future susceptible population. High value of today's $\mathcal{R}_0 \equiv \mathbb{E} \frac{\beta_{it}}{\gamma}$ increases tomorrow's infected population, I_{t+1} , and decreases tomorrow's susceptible population, S_{t+1} . That is, $Cov(S_{t+1}, I_{t+1}) < 0$. Because the new infection tomorrow is a realization of β_{t+1} multiplied by the two (that is, $\beta_{t+1} I_{t+1} \frac{S_{t+1}}{N}$) this negative correlation reduces the spread of the virus in the future on average, endogenously reducing the magnitude of the outbreak.

This interpretation also highlights the importance of intertemporal correlation of infection rates, $Cov(\beta_t, \beta_{t+1})$. When some individuals participate in events at

infection-prone environments more frequently than others, the correlation will be positive. Such effects can lead to a sequence of clusters and an extremely rapid rise in infections (Cooper et al., 2019) that overwhelm the negative correlation between S_{t+1} and I_{t+1} highlighted above. On the other hand, when infections take place at residential environments (e.g. residential compound in Hong Kong for SARS, and dormitory in Singapore for COVID-19), then the infected person will be less likely to live in another residential location to spread the virus. In this case, the correlation will be negative. In this way, considering the correlation of infection rates across periods will be crucial.

Note that the mechanism we identified on herd immunity thresholds is distinct from the ones described in Gomes et al. (2020); Hébert-Dufresne et al. (2020); Britton et al. (2020). They note that when population has permanently heterogenous activity rate, which captures both the probability of infecting and being infected, the herd immunity can be achieved with lower threshold level of susceptible. They explain this because majority of “active” population becomes infected faster than the remaining population. Our mechanism does not hinge on the permanent heterogeneity in population, which could have been captured by $Cov(\beta_{it}, \beta_{it+1}) = 1$. The fat-tailed distribution in infection rate alone creates reduction in the required herd immunity rate in expectation.

2.3.4 Lockdown policy targeted at SSEs

How could the policymaker design the mitigation policies effectively if the distribution of infection rates is fat-tailed? Here, we concentrate our analysis on lockdown policy. Unlike the traditionally analyzed lockdown policy, we consider a policy that particularly targets SSEs. Specifically we assume that the policy can impose an upper bound on $\beta_{it} \leq \bar{\beta}$ with probability ϕ . The probability ϕ is meant to capture some imperfection in enforcements or impossibility in closing some facilities such as hospitals and daycare¹⁹. Here, we set $\phi = 0.5$. For tractability, we assume that the government implements targeted lockdown policies for entire periods. We experiment with $\bar{\beta}$ for various values: 1000 cases per day, 100 cases per day, and 50 cases per day.

While Table 2.5.3 in Appendix presents results in detail, we briefly summarize the main results here. First, the policy reduces the mean of the peak infection rate if and only if the distribution features fatter tails. Second, the targeted lockdown policy is

¹⁹Note that, even though the theoretical variance is infinite, the realized variance in numerical simulations will always be finite. Therefore, such stochastic reductions can still reduce the simulated variance even though the theoretical variance remains infinite.

effective in reducing the volatility of the peak infection rate in the case that such risks exist in the first place. For example, consider the case with $\alpha = 1.1$. Moving from no policy to the upper-limit of 100 cases reduces the 90th percentile of peak infection from 31% to 17%.²⁰ In contrast, when $\alpha = 2$ or with negative binomial distribution, the policy has virtually no effect. Therefore the policy is particularly effective in mitigating the upward risk of overwhelming the medical capacity. This highlights that while the fat-tailed distribution induces the aggregate risk in the epidemiological dynamics, the government can partly remedy this by appropriately targeting the lockdown policy.

We conclude this section by discussing several modeling assumptions. First, we have assumed that $\{\beta_{it}\}$ is independently and identically distributed across individuals and over time. This may not be empirically true. For example, a person who was infected in a big party is more likely to go to a party in the next period. This introduces ex ante heterogeneities as discussed in (Gomes et al., 2020; Hébert-Dufresne et al., 2020; Britton et al., 2020), generating positive correlation in $\{\beta_{it}\}$ along the social network. Or, a person who tends to be a superspreader may be more likely to be a superspreader in the next period. This induces a positive correlation in $\{\beta_{it}\}$ over time. If the resulting cascading effect were large, then the average effects on the epidemiological outcomes we have found may be overturned. Second, we have exogenously imposed power law distributions without fully exploring underlying data generation mechanisms behind them. The natural next step is to provide a model in which individual infection rate follows a power law. We believe SIR models with social networks along the line of Pastor-Satorras and Vespignani (2001), Moreno et al. (2002), Castellano and Pastor-Satorras (2010), May and Lloyd (2001), Zhang et al. (2013), Gutin et al. (2020), and Akbarpour et al. (2020) are promising avenue to generate endogenous power law in individual infection rates.

2.4 Estimation methods

We began with the evidence that SSEs follow a power law distribution with fat tails in many settings, and showed that such distributions substantively change the pre-

²⁰We may be concerned that the unbounded support of power law distribution is unrealistic; at the extreme case, one cannot infect more than 8 billion people since that will exceed the world population. Imposing some upperbound on the distribution of infection rate will be equivalent to imposing a lockdown policy with perfect implementation ($\phi = 1$). As shown in the results of lockdown policy, imposing such upperbounds can significantly reduce the volatility relative to the unbounded case, and nonetheless, some uncertainties will persist and remain much larger than the predictions of negative binomial distributions.

dictions of SIR models. In this Section, we discuss the implications of power law distributions for estimating the effective reproduction number.

2.4.1 Limitations of sample means

Estimation of average reproduction numbers (\mathcal{R}_t) has been the chief focus of empirical epidemiology research (e.g. [Becker and Britton, 1999](#)). Our estimates across five different data sets suggest that the exponent satisfies $\alpha \in (1, 2)$ in many occasions: that is, the infection rates have a finite mean but an infinite variance. Since the mean exists, by the Law of Large Numbers, the sample mean estimates (see e.g. [Nishiura, 2007](#)) that have been used in the epidemiology research will be consistent (i.e. converge to the true mean asymptotically) and also unbiased (i.e. its expectation equals the true mean with finite samples.)

Due to the infinite variance property, however, the sample mean will converge very slowly to the true mean because the classical CLT requires finite variance. Formally, while the convergence occurs at a rate \sqrt{N} for distributions with finite variance, or thin tails, it occurs only at a rate $N^{1-\frac{1}{\alpha}}$ for the power law distributions with fat tails, $\alpha \in (1, 2)$ ([Gabaix, 2011](#)).²¹ Under distributions with infinite variance, or fat tails, the sample mean estimates could be far from the true mean with reasonable sample sizes, and their estimated 95 confidence intervals will be too tight. Figure 2-6 plots a Monte Carlo simulation of sample mean's convergence property. For thin-tailed distributions such as the negative binomial distribution or the power law distribution with $\alpha = 2$, even though the convergence is slow due to their very large variance, they still converge to the true mean reasonably under a few 1,000 observations. In contrast, with fat-tailed distributions such as power law distribution with $\alpha = 1.1$ or $\alpha = 1.2$, the sample mean will remain far from the true mean. Their sample mean estimates behave very differently as the sample size increases. Every so often, some extraordinarily high values occur that significantly raises the sample mean and its standard errors. When such extreme values are not occurring, the sample means gradually decrease. With thin tails, such extreme values are rare enough not to cause such sudden increase in sample means; however, with fat tails, the extreme values are not so rare.

²¹For $\alpha = 1$ exactly, the convergence will occur at rate $\ln N$.

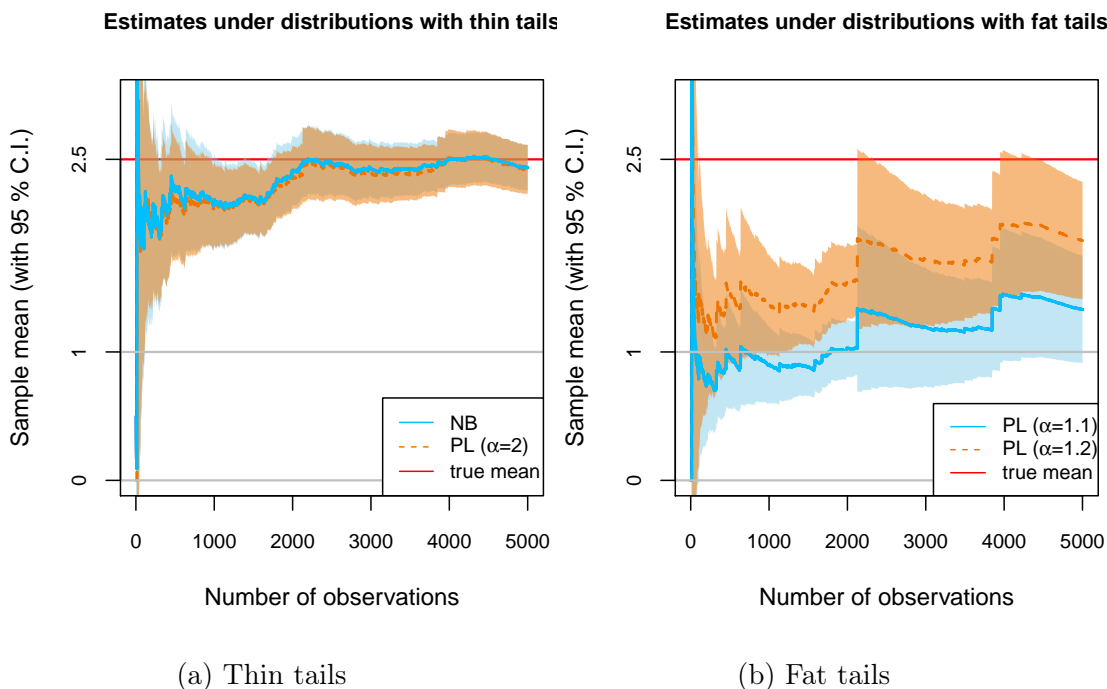


Figure 2-6: An example of sample mean estimates

Notes: Figure 2-6 depicts an example of sample mean estimates for thin-tailed and fat-tailed distributions. The draws of observations are simulated through the inverse-CDF method, where the identical uniform random variable is applied so that the sample means are comparable across four different distributions. All distributions are normalized to have the mean of 2.5. The negative binomial (NB) distribution has the dispersion parameter $k = 0.16$ taken from (Lloyd-Smith et al., 2005). The range of power law (PL) parameters is also taken from the empirical estimates.

2.4.2 Using power law exponents to improve inference

What methods could address the concerns that the sample mean may be empirically unstable? One approach may be to exclude some realizations as an outlier, and focus on subsamples without extreme values²². However, such analysis will neglect major source of risks even though extreme "outlier" SSEs may fit the power law distributions as shown in Figure 2-1. While estimating the mean of distributions with rare but extreme values has been notoriously difficult²³, there are some approaches to address

²²In Japan, the case of over 620 infections in the cruise ship Diamond Princess was excluded from all other analyses.

²³Consider, for example, a binary distribution of infection rates such that one infects N others with $1/N$ probability, and 0 others with $1 - 1/N$ probability. In this case, the true mean $R_t = 1$. Suppose a statistician observes 10 infected cases for each estimation. If N were 1,000, then with 99 ($\approx 0.999^{10}$)

this formally.

With power law distributions, the estimates of exponent have information that can improve the estimation of the mean. Figure 2-7 shows that the exponents α can be estimated adequately with reasonable sample sizes.²⁴ If $\alpha > 2$, as may be the case for the India under strict lockdown, then one can have more confidence in the reliability of sample mean estimates. However, if $\alpha < 2$, the sample mean may substantially differ from the true mean. At the least, one can be aware of the possibility.

One transparent approach is a “plug-in” method: to estimate the exponent $\hat{\alpha}$, and plug into the formula of the mean $\frac{\hat{\alpha}}{\hat{\alpha}-1}\underline{Z}$. This method yields a valid 95 confidence intervals (C.I.) of the median²⁵ since the estimated $\hat{\alpha}$ has valid confidence intervals.²⁶ Figure 2-7 shows the estimation results for the same data with $\alpha = 1.1, 1.2$ as shown in Figure 2-6. First, while the sample mean in Figure 2-6 had substantially underestimated the mean, this estimated median is close to the true mean. Second, while the sample mean estimation imposed symmetry between lower and upper bounds of 95 percent confidence intervals, this estimate reflects the skewness of uncertainties: upward risks are much higher than downward risks because of the possibility of extreme events. Third, the standard errors are much larger, reflecting the inherent uncertainties given the limited sample sizes.²⁷ Fourth, the estimates are more stable and robust to the extreme values²⁸ than the sample mean estimates that have sudden jumps in the estimates after the extreme values.

Table 2.5 demonstrates the validity of the “plug-in” method through a simulation experiment. The table shows the comparison of the probability that the constructed 95% C.I. covers the true mean using the 1,000 Monte-Carlo simulation. When the

percent chance, nobody becomes infected so that $\hat{R}_t = 0$, and the estimates’ confidence interval will be $[0, 0]$. But with less than 1 percent chance when any infection occurs, \hat{R}_t will be larger than 100. Thus, the 95 percent confidence interval contains the true mean in less than 1 percent of the time. To the best of our knowledge, there is no techniques that can help us completely avoid this problem given the fundamental constraint of small sample size.

²⁴The standard errors are computed by the maximum likelihood approach, as the linear regressions are known to underestimate the standard errors (see [Gabaix and Ibragimov, 2011](#)).

²⁵Note that the estimate corresponds to the median estimate because $\frac{\hat{\alpha}}{\hat{\alpha}-1}$ is a non-linear transformation of $\hat{\alpha}$.

²⁶To be more formal, the correct C.I. will be to consider the uncertainties with the mean of observations below \underline{Z} . To focus on the uncertainty from upper tail, we construct the 95 percent C.I. from that of the estimate of α here.

²⁷When the number of observations is less than 1000, the estimated confidence interval of α contains values less than 1.0, turning the upper bound of the mean to be ∞ . This does not mean that a correct expectation is ∞ infections in the near future, but that there is serious upward risks in infection rates.

²⁸This is because the estimation through log-likelihood will take the log of the realized value, instead of its level.

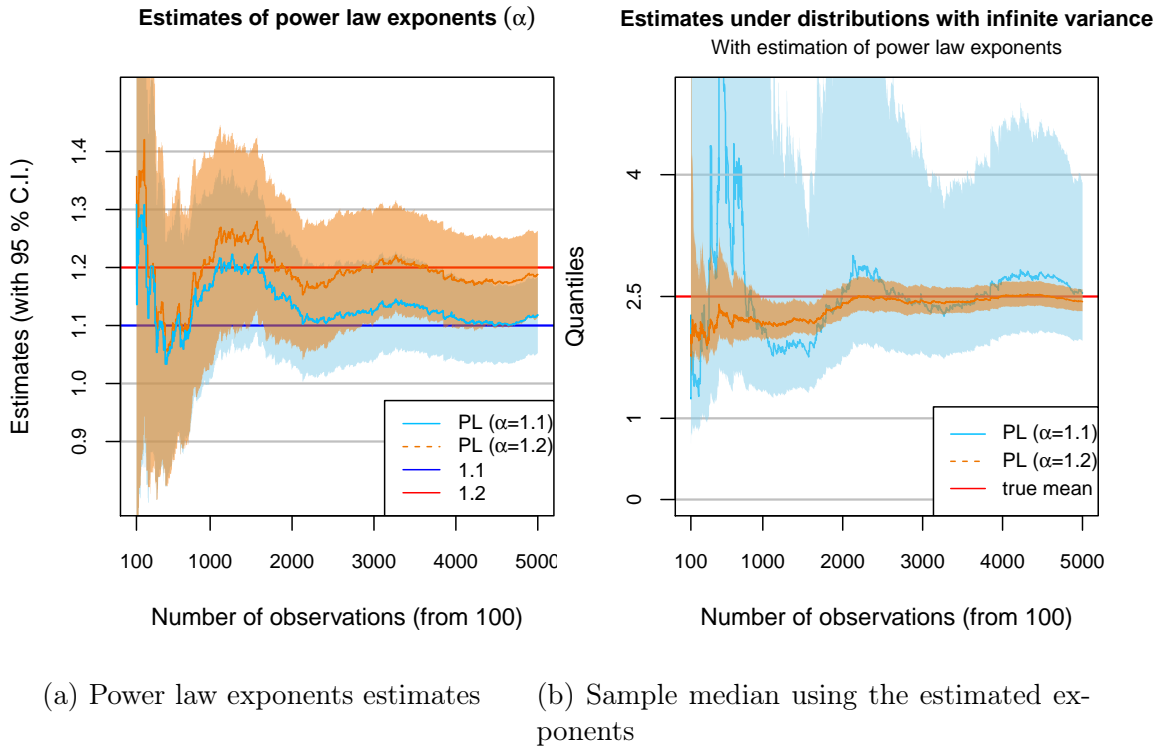


Figure 2-7: An example of “plug-in” estimates

Notes: Figure 2-7 plots the estimates of power law exponents and the resulting estimates of sample median, using the same data as in Figure 2-6. Note that while the number of observations contains all observations, the data points contributing to the estimates are only above some thresholds: only less than 25 percents of the data contribute to the estimation of the exponents.

estimate is unbiased and has correct standard errors, this coverage probability is 95%. When the power law exponent is close to one, the traditional “sample means” approach has the C.I. that covers the true mean only with 20-40% for all sample sizes. By contrast the “plug-in” method covers the true estimates close to 95%. As the tail becomes thinner toward $\alpha = 2$, the difference between the two tends to disappear, with “sample mean” approach performing better some times. When the underlying distribution has fat-tails, however, estimation using the plug-in method is preferred.

While the C.I. in the plug-in method has adequate coverage probabilities, it is often very large and possibly infinite. Figure 2-7 visualizes this. This large C.I. occurs especially when $\alpha \simeq 1$ because the mean of a power law distribution is proportional to $\frac{\alpha}{1-\alpha}$. How could the policymakers plan their efforts do given such large uncertainty in \mathcal{R}_0 ? Given the theoretical results in Section 3 that the epidemiological dynamics

will be largely uncertain even when $\alpha \simeq 1$ is perfectly known, we argue that applying the estimated \mathcal{R}_0 into a deterministic SIR model will not lead to a reliable prediction. Instead of focusing on the mean, it will be more adequate and feasible to focus on the distribution of near-future infection outcomes. For example, using the estimated power law distribution, policymakers can compute the distribution of the future infection rate. The following analogy might be useful: in planning for natural disasters such as hurricanes and earthquakes, policymakers will not rely on the estimates of average rainfall or average seismic activity in the future; instead, they consider the probabilities of some extreme events, and propose plans contingent on realizations. Similar kinds of planning may be also constructive regarding preparation for future infection outbreaks.

To overcome data limitations, epidemiologists have developed a number of sophisticated methods such as backcalculation assuming Poisson distribution (Becker et al., 1991), and ways to account for imported cases. There are also a number of methods developed to account for fat-tailed distributions (see e.g. Stoyanov et al., 2010, for a survey), such as tail tempering (Kim et al., 2008) and separating the data into subgroups (Toda and Walsh, 2015). In the future, it will be important to examine what power law distributions will imply about existing epidemiological methods, and how statistical techniques such as plug-in methods can be combined with epidemiological techniques to allow more reliable estimation of risks.

2.5 Conclusion: implications for COVID-19 pandemic

Most research on infection dynamics has focused on deterministic SIR models, and have estimated its key statistics, the *average* reproduction number (\mathcal{R}_0). In contrast, some researchers have concentrated on SSEs, and estimated the *dispersion* of infection rates using negative binomial distributions. Nonetheless, stochastic SIR models based on estimated distributions have predicted that idiosyncratic uncertainties in SSEs would vanish when the infected population is large, and thus, the epidemiological dynamics will be largely predictable. In this paper, we have documented evidence from SARS, MERS, and COVID-19 that SSEs actually follow a power law distribution with the exponent $\alpha \in (1, 2)$: that is, their distributions have infinite variance, or fat tails. Our stochastic SIR model with these fat-tailed distributions have shown that idiosyncratic uncertainties in SSEs will persist even when the infected population is large, inducing major unpredictability in aggregate infection dynamics.

Since the currently infected population is estimated to be around 3 million in

the COVID-19 pandemic,²⁹ our analysis has immediate implications for policies of today. For statistical inference, the aggregate unpredictability suggests caution is warranted on drawing inferences about underlying epidemiological conditions from observed infection outcomes. First, large geographic variations in infections may be driven mostly by idiosyncratic factors, and not by fundamental socioeconomic factors. While many looked for underlying differences in public health practices to explain the variations, our model shows that these variations may be more adequately explained by the presence of a few, idiosyncratic SSEs. Second, existing stochastic models would suggest that, keeping the distribution of infection rates and pathological environments constant, recent infection trends can predict the future well. In contrast, our analysis shows that even when the average number of new infections may seem to have stabilized at a low level in recent weeks, subsequent waves can suddenly arrive in the future.

Such uncertainties in outbreak timing and magnitude introduce substantial socio-economic difficulties, and measures to assess and mitigate such risks will be invaluable. The death rate is shown to increase when the medical capacity binds. Thus, reducing uncertainties can reduce average fatality. Furthermore, uncertainties can severely deter necessary investments and impede planning for reallocation and recovery from the pandemic shocks. To assess such risks, we can estimate the tail distributions to improve our inference on the average number. To address such risks, social distancing policies and individual efforts can focus on large physical gatherings in infection-prone environments. Our estimates suggest, like earthquakes, infection dynamics will be largely unpredictable. But unlike earthquakes, they are a consequence of social decisions, and efforts to reduce SSEs can significantly mitigate the uncertainty the society faces as a whole.

²⁹According to [worldometers.info](https://www.worldometers.info), the cumulative infection worldwide is 7 million, among which 4 million have already recovered or died, as of June 9, 2020.

	Power law					Negative binomial
	$\alpha = 1.08$	$\alpha = 1.1$	$\alpha = 1.2$	$\alpha = 1.5$	$\alpha = 2$	
1. Cumulative infected						
mean	60%	73%	89%	92%	92%	92%
90th percentile	85%	91%	95%	93%	92%	92%
50th percentile	59%	71%	88%	92%	92%	92%
10th percentile	39%	59%	84%	91%	92%	92%
2. Herd immunity threshold						
mean	39%	49%	62%	65%	66%	66%
90th percentile	65%	75%	78%	71%	69%	69%
50th percentile	35%	45%	59%	65%	66%	66%
10th percentile	17%	29%	51%	60%	62%	64%
3. Peak infection						
mean	14%	18%	25%	27%	27%	27%
90th percentile	31%	34%	36%	29%	28%	27%
50th percentile	9%	13%	22%	26%	27%	27%
10th percentile	4%	7%	18%	25%	26%	26%
4. Days infecting 5%						
mean (days)	137	93	47	37	35	35
90th percentile	252	147	56	42	35	35
50th percentile	119	84	49	35	35	35
10th percentile	49	42	35	35	35	35

Table 2.4: Summary statistics for epidemiological outcomes

Note: Table 2.4 shows the summary statistics from 1000 simulations for five different tail parameters for the case of power law distribution, and for the negative binomial distribution.

	$\alpha = 1.08$	$\alpha = 1.1$	$\alpha = 1.2$	$\alpha = 1.5$	$\alpha = 2$
1. $N = 100$					
Sample means	21%	26%	42%	74%	89%
Plug-in	98%	98%	98%	94%	87%
2. $N = 500$					
Sample means	24%	29%	45%	78%	90%
Plug-in	98%	98%	95%	94%	84%
3. $N = 1000$					
Sample means	24%	26%	48%	78%	92%
Plug-in	97%	97%	93%	93%	86%

Table 2.5: Coverage probability of 95% confidence interval

Note: Thable 2.5 reports the probability that the 95% confidence interval, constructed in two different ways, covers the true value in 1000 simulation. “Sample means” is simply uses the sample mean. “Using power laws uses” first estimates the Pareto exponent using the maximum likelihood, and then convert it to the mean estimates.

Bibliography

- ABBOTT, S., J. HELLEWELL, J. MUNDAY, CMMID nCoV WORKING GROUP, AND S. FUNK (2020): “The Transmissibility of Novel Coronavirus in the Early Stages of the 2019-20 Outbreak in Wuhan: Exploring Initial Point-Source Exposure Sizes and Durations Using Scenario Analysis,” *Wellcome Open Research*, 5, 17.
- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. D. WHINSTON (2020): “Optimal Targeted Lockdowns in a Multi-Group SIR Model,” Working Paper 27102, National Bureau of Economic Research.
- AKBARPOUR, M., C. COOK, A. MARZUOLI, S. MONGEY, A. NAGARAJ, M. SACCAROLA, P. TEBALDI, S. VASSERMAN, AND H. YANG (2020): “Socioeconomic Network Heterogeneity and Pandemic Policy Response,” 77.
- ATKINSON, A. B., T. PIKETTY, AND E. SAEZ (2011): “Top Incomes in the Long Run of History,” *Journal of Economic Literature*, 49, 3–71.
- AXTELL, R. L. (2001): “Zipf Distribution of U.S. Firm Sizes,” *Science*, 293, 1818–1820.
- BARABASI, A.-L. AND J. FRANGOS (2014): *Linked: How Everything Is Connected to Everything Else and What It Means for Business, Science, and Everyday Life*, Basic Books.
- BARDINA, X., M. FERRANTE, AND C. ROVIRA (2020): “A Stochastic Epidemic Model of COVID-19 Disease,” *arXiv:2005.02859 [q-bio]*.
- BARTLETT, M. S. (1949): “Some Evolutionary Stochastic Processes,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 11, 211–229.
- BEARE, B. K. AND A. A. TODA (2019): “Geometrically Stopped Markovian Random Growth Processes and Pareto Tails,” *arXiv:1712.01431 [econ, math, stat]*.

- (2020): “On the Emergence of a Power Law in the Distribution of COVID-19 Cases,” *arXiv:2004.12772 [physics, q-bio]*.
- BECKER, N. G. AND T. BRITTON (1999): “Statistical Studies of Infectious Disease Incidence,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61, 287–307.
- BECKER, N. G., L. F. WATSON, AND J. B. CARLIN (1991): “A Method of Non-Parametric Back-Projection and Its Application to Aids Data,” *Statistics in Medicine*, 10, 1527–1542.
- BRITTON, T. (2010): “Stochastic Epidemic Models: A Survey,” *Mathematical Biosciences*, 225, 24–35.
- (2018): “Basic Stochastic Transmission Models and Their Inference,” *arXiv:1801.09594 [stat]*.
- BRITTON, T., T. HOUSE, A. L. LLOYD, D. MOLLISON, S. RILEY, AND P. TRAPMAN (2015): “Five Challenges for Stochastic Epidemic Models Involving Global Transmission,” *Epidemics*, 10, 54–57.
- BRITTON, T., P. TRAPMAN, AND F. G. BALL (2020): “The Disease-Induced Herd Immunity Level for Covid-19 Is Substantially Lower than the Classical Herd Immunity Level,” Preprint, Infectious Diseases (except HIV/AIDS).
- BROTHERHOOD, L., P. KIRCHER, C. SANTOS, AND M. TERTILT (2020): “An Economic Model of the Covid-19 Epidemic: The Importance of Testing and Age-Specific Policies,” SSRN Scholarly Paper ID 3594329, Social Science Research Network, Rochester, NY.
- CASTELLANO, C. AND R. PASTOR-SATORRAS (2010): “Thresholds for Epidemic Spreading in Networks,” *Physical Review Letters*, 105, 218701.
- CLAUSET, A., C. R. SHALIZI, AND M. E. J. NEWMAN (2009): “Power-Law Distributions in Empirical Data,” *SIAM Review*, 51, 661–703.
- COOPER, L., S. Y. KANG, D. BISANZIO, K. MAXWELL, I. RODRIGUEZ-BARRAQUER, B. GREENHOUSE, C. DRAKELEY, E. ARINAITWE, S. G. STAEDKE, P. W. GETHING, P. ECKHOFF, R. C. REINER, S. I. HAY, G. DORSEY, M. R. KAMYA, S. W. LINDSAY, B. T. GRENFELL, AND D. L. SMITH (2019): “Pareto

- Rules for Malaria Super-Spreaders and Super-Spreading,” *Nature Communications*, 10, 3939.
- DAVIES, N. G., P. KLEPAC, Y. LIU, K. PREM, M. JIT, C. C.-. WORKING GROUP, AND R. M. EGGO (2020): “Age-Dependent Effects in the Transmission and Control of COVID-19 Epidemics,” *medRxiv*, 2020.03.24.20043018.
- ENDO, A., CENTRE FOR THE MATHEMATICAL MODELLING OF INFECTIOUS DISEASES COVID-19 WORKING GROUP, S. ABBOTT, A. J. KUCHARSKI, AND S. FUNK (2020): “Estimating the Overdispersion in COVID-19 Transmission Using Outbreak Sizes Outside China,” *Wellcome Open Research*, 5, 67.
- FRIEDEN, T. R. AND C. T. LEE (2020): “Identifying and Interrupting Superspreading Events-Implications for Control of Severe Acute Respiratory Syndrome Coronavirus 2,” *Emerging Infectious Diseases*, 26, 1059–1066.
- GABAIX, X. (2009): “Power Laws in Economics and Finance,” *Annual Review of Economics*, 1, 255–294.
- (2011): “The Granular Origins of Aggregate Fluctuations,” *Econometrica*, 79, 733–772.
- GABAIX, X. AND R. IBRAGIMOV (2011): “Rank - 1 / 2: A Simple Way to Improve the OLS Estimation of Tail Exponents,” *Journal of Business & Economic Statistics*, 29, 24–39.
- GALVANI, A. P. AND R. M. MAY (2005): “Dimensions of Superspreading,” *Nature*, 438, 293–295.
- GAY, N. J., G. DE SERRES, C. P. FARRINGTON, S. B. REDD, AND M. J. PAPANIA (2004): “Assessment of the Status of Measles Elimination from Reported Outbreaks: United States, 1997-1999,” *The Journal of Infectious Diseases*, 189 Suppl 1, S36–42.
- GLOVER, A., J. HEATHCOTE, D. KRUEGER, AND J.-V. RÍOS-RULL (2020): “Health versus Wealth: On the Distributional Effects of Controlling a Pandemic,” *CEPR Covid Economics: Vetted and Real-Time Papers*, 6, 22–64.
- GOLLIER, C. (2020): “Cost-Benefit Analysis of Age-Specific Deconfinement Strategies,” *CEPR Covid Economics: Vetted and Real-Time Papers*, 24, 1–31.

- GOMES, M. G. M., R. M. CORDER, J. G. KING, K. E. LANGWIG, C. SOUTO-MAIOR, J. CARNEIRO, G. GONCALVES, C. PENHA-GONCALVES, M. U. FERREIRA, AND R. AGUAS (2020): “Individual Variation in Susceptibility or Exposure to SARS-CoV-2 Lowers the Herd Immunity Threshold,” *medRxiv*, 2020.04.27.20081893.
- GUTENBERG, B. AND C. RICHTER (1954): *Seismicity Of The Earth And Associated Phenomena*.
- GUTIN, G., T. HIRANO, S.-H. HWANG, P. R. NEARY, AND A. A. TODA (2020): “The Effect of Social Distancing on the Reach of an Epidemic in Social Networks,” *arXiv:2005.03067 [physics, q-bio]*.
- HÉBERT-DUFRESNE, L., B. M. ALTHOUSE, S. V. SCARPINO, AND A. ALLARD (2020): “Beyond R_0 : Heterogeneity in Secondary Infections and Probabilistic Epidemic Forecasting,” *arXiv:2002.04004 [physics, q-bio]*.
- HSU, L.-Y., C.-C. LEE, J. A. GREEN, B. ANG, N. I. PATON, L. LEE, J. S. VILLACIAN, P.-L. LIM, A. EARNEST, AND Y.-S. LEO (2003): “Severe Acute Respiratory Syndrome (SARS) in Singapore: Clinical Features of Index Patient and Initial Contacts,” *Emerging infectious diseases*, 9, 713.
- JESSEN, A. H. AND T. MIKOSCH (2006): “Regularly Varying Functions,” *Publications de l’Institut Mathématique*, 80(94), 171–192.
- KARAKO, K., P. SONG, Y. CHEN, AND W. TANG (2020): “Analysis of COVID-19 Infection Spread in Japan Based on Stochastic Transition Model,” *BioScience Trends*, advpub.
- KENDALL, D. G. (1956): “Deterministic and Stochastic Epidemics in Closed Populations,” in *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 4: Contributions to Biology and Problems of Health*, The Regents of the University of California.
- KERMACK, W. O. AND A. G. MCKENDRICK (1927): “A Contribution to the Mathematical Theory of Epidemics,” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115, 700–721.
- KIM, Y. S., S. T. RACHEV, M. L. BIANCHI, AND F. J. FABOZZI (2008): “Financial Market Models with Lévy Processes and Time-Varying Volatility,” *Journal of Banking & Finance*, 32, 1363–1378.

- KLEIBER, C. AND S. KOTZ (2003): *Statistical Size Distributions in Economics and Actuarial Sciences*, John Wiley & Sons.
- KUCHARSKI, A. J. AND C. L. ALTHAUS (2015): “The Role of Superspreading in Middle East Respiratory Syndrome Coronavirus (MERS-CoV) Transmission,” *Eurosurveillance*, 20, 21167.
- LECLERC, Q. J., N. M. FULLER, L. E. KNIGHT, CMMID COVID-19 WORKING GROUP, S. FUNK, AND G. M. KNIGHT (2020): “What Settings Have Been Linked to SARS-CoV-2 Transmission Clusters?” *Wellcome Open Research*, 5, 83.
- LIU, Y., A. A. GAYLE, A. WILDER-SMITH, AND J. ROCKLÖV (2020): “The Reproductive Number of COVID-19 Is Higher Compared to SARS Coronavirus,” *Journal of Travel Medicine*, 27.
- LLOYD-SMITH, J. O., S. J. SCHREIBER, P. E. KOPP, AND W. M. GETZ (2005): “Superspreading and the Effect of Individual Variation on Disease Emergence,” *Nature*, 438, 355–359.
- MAY, R. M. AND A. L. LLOYD (2001): “Infection Dynamics on Scale-Free Networks,” *Physical Review E*, 64, 066112.
- MORENO, Y., R. PASTOR-SATORRAS, AND A. VESPIGNANI (2002): “Epidemic Outbreaks in Complex Heterogeneous Networks,” *The European Physical Journal B - Condensed Matter and Complex Systems*, 26, 521–529.
- NEWMAN, M. E. J. (2005): “Power Laws, Pareto Distributions and Zipf’s Law,” *Contemporary Physics*, 46, 323–351.
- NISHIURA, H. (2007): “Time Variations in the Transmissibility of Pandemic Influenza in Prussia, Germany, from 1918–19,” *Theoretical Biology & Medical Modelling*, 4.
- NISHIURA, H., K. MIZUMOTO, AND Y. ASAI (2017): “Assessing the Transmission Dynamics of Measles in Japan, 2016,” *Epidemics*, 20, 67–72.
- NISHIURA, H., H. OSHITANI, T. KOBAYASHI, T. SAITO, T. SUNAGAWA, T. MATSUI, T. WAKITA, M. C.-. R. TEAM, AND M. SUZUKI (2020): “Closed Environments Facilitate Secondary Transmission of Coronavirus Disease 2019 (COVID-19),” *medRxiv*, 2020.02.28.20029272.
- PASTOR-SATORRAS, R. AND A. VESPIGNANI (2001): “Epidemic Spreading in Scale-Free Networks,” *Physical Review Letters*, 86, 3200–3203.

- PETROSILLO, N., G. VICECONTE, O. ERGONUL, G. IPPOLITO, AND E. PETERSEN (2020): “COVID-19, SARS and MERS: Are They Closely Related?” *Clinical Microbiology and Infection: The Official Publication of the European Society of Clinical Microbiology and Infectious Diseases*, 26, 729–734.
- RAMPINI, A. A. (2020): “Sequential Lifting of COVID-19 Interventions with Population Heterogeneity,” Working Paper 27063, National Bureau of Economic Research.
- REMUZZI, A. AND G. REMUZZI (2020): “COVID-19 and Italy: What Next?” *The Lancet*, 395, 1225–1228.
- ROBERTS, M., V. ANDREASEN, A. LLOYD, AND L. PELLIS (2015): “Nine Challenges for Deterministic Epidemic Models,” *Epidemics*, 10, 49–53.
- ROBERTS, M. G. (2013): “Epidemic Models with Uncertainty in the Reproduction Number,” *Journal of Mathematical Biology*, 66, 1463–1474.
- (2017): “An Epidemic Model with Noisy Parameters,” *Mathematical Biosciences*, 287, 36–41.
- SHEN, Z., F. NING, W. ZHOU, X. HE, C. LIN, D. P. CHIN, Z. ZHU, AND A. SCHUCHAT (2004): “Superspreading SARS Events, Beijing, 2003,” *Emerging Infectious Diseases*, 10, 256–260.
- SIMHA, A., R. V. PRASAD, AND S. NARAYANA (2020): “A Simple Stochastic SIR Model for COVID 19 Infection Dynamics for Karnataka: Learning from Europe,” *arXiv:2003.11920 [math, q-bio]*.
- STOYANOV, S. V., S. RACHEV, B. RACHEVA-IOTOVA, AND F. J. FABOZZI (2010): “Fat-Tailed Models for Risk Estimation,” SSRN Scholarly Paper ID 1729040, Social Science Research Network, Rochester, NY.
- SZABÓ, G. M. (2020): “Propagation and Mitigation of Epidemics in a Scale-Free Network,” *arXiv:2004.00067 [physics, q-bio]*.
- TODA, A. A. AND K. WALSH (2015): “The Double Power Law in Consumption and Implications for Testing Euler Equations,” *Journal of Political Economy*, 123, 1177–1200.
- WANG, H., Z. WANG, Y. DONG, R. CHANG, C. XU, X. YU, S. ZHANG, L. TSAMLAG, M. SHANG, J. HUANG, Y. WANG, G. XU, T. SHEN, X. ZHANG, AND

Y. CAI (2020): “Phase-Adjusted Estimation of the Number of Coronavirus Disease 2019 Cases in Wuhan, China,” *Cell Discovery*, 6, 1–8.

ZHANG, H., Z.-H. GUAN, T. LI, X.-H. ZHANG, AND D.-X. ZHANG (2013): “A Stochastic SIR Epidemic on Scale-Free Network with Community Structure,” *Physica A: Statistical Mechanics and its Applications*, 392, 974–981.

ZIPF, G. K. (1949): *Human Behavior and the Principle of Least Effort*, Human Behavior and the Principle of Least Effort, Oxford, England: Addison-Wesley Press.

Appendix

A Empirical Appendix

A.1 Relating empirical distribution of Z to theoretical distribution of β_{it}

In this paper, we have used the estimates from the data to simulate the evolution dynamics of the epidemiological model. The key step in our argument is that the tail distribution of $\sum_i z_{it}$ or $\sum_t z_{it}$, the *cumulative* “effective” number of infections, is equivalent to the tail distribution of β_{it} , the *individual and per-period* “basic” number of infection. However, in general, this needs not hold: for example, even if β_{it} were normally distributed (i.e. thin tailed), Z may follow a t -distribution (i.e. fat-tailed). Under what conditions is our interpretation about the relationship between distribution of Z and distribution of β_i valid? Are they plausible in the settings of the coronaviruses?

To clarify this question, let us lay out a model. Formally, Z is a *mixture distribution* of the *weighted sum* of β_{it} . Here, we provide notations for $\sum_t z_{it}$ but the identical argument will also apply to $\sum_i z_{it}$. Specifically, suppose i stays infected for \bar{t} periods, and let the probability mass be $\delta(\bar{t})$. In the case of exponential decay as in the SIR model, $\delta(\bar{t}) = \gamma^{\bar{t}}$. Denoting the countercumulative distribution of Z_i by Φ , and that of β_{it} by F , we have

$$\Phi(Z_i) = \sum_{\bar{t}=1}^{\infty} \delta(\bar{t}) G_{\bar{t}} \left(\sum_{t=1}^{\bar{t}} \frac{S_t}{N} \beta_{it} \right), \quad \beta_{it} \sim F,$$

where $G_{\bar{t}}$ denotes the distribution of $\sum_{t=1}^{\bar{t}} \frac{S_t}{N} \beta_{it}$.

A.1.1 Empirical evidence on causes of SSEs

First, we may be concerned that, even if Φ is a power law distribution, F may not be a power law distribution. A counterexample is that a geometric Brownian motion with stochastic stopping time that follows exponential distribution can also generate power law distributions of the tail (Beare and Toda, 2020). That is, the tail property of Φ needs not be due to tails of F : for $\sum_t z_{it}$, it could also be due to some individuals

staying infectious for an extremely long periods. For $\sum_i z_{it}$, it could also be due to some events having extremely high number of infected primary cases.

While we acknowledge such possibilities, we argue that for superspreaders or SSEs of the coronaviruses, the main mechanism of extremely high number of cumulative infection is primarily due to some extreme events at particular time t . Let us be concrete. If the counterexample's reasoning were true for $\sum_t z_{it}$, then a superspreader is someone who goes, for example, to a restaurant and infect two other people at time t , and then goes to a shopping mall and infects three other people at time $t + 1$, and then goes to meet her two friends and infect them, and so on. However, this interpretation is inconsistent with numerous anecdotes. Instead, a superspreader infects many people because he attends a SSE that has infection-prone environment at a particular time t . Conferences, parties, religious gatherings, and sports gyms are a particular place that can infect many at the same time. Moreover, [Nishiura et al. \(2020\)](#) paper whose data we use has identified particular environment that has caused SSEs. This interpretation is important because, if the extremely high cumulative number of infection were due to some staying infectious for a long time or some events having extremely high number of primary cases, then our model's prediction of sudden outbreak due to SSE is no longer a valid prediction.

A1.2 Theoretical analysis on interpretation of exponents

Second, we may be concerned that the exponent of $\Phi(Z_i)$ may be different than the exponent of $F(\beta_{i\tau})$, even if both have tails that follow power laws. We use two steps to show that this is not a concern:

- (i) if a random variable has a power law distribution with exponent α , then its weighted sum also has a tail distribution that follows a power law with exponent α (see e.g. [Jessen and Mikosch \(2006\)](#) or [Gabaix \(2009\)](#)). Thus, neither summation over multiple periods nor the weights of $\frac{S_t}{N}$ will change this.
- (ii) the tail property of distribution can be examined by considering $\alpha_F(Z) = \frac{f(Z)}{f(cZ)}$ for some $c \neq 1$ and taking its limit. In particular, if F has a power law distribution, then $\alpha_F(Z) = c^\alpha$.³⁰ Denoting the probability mass of $G_{\bar{t}}(\cdot)$ by $g_{\bar{t}}(\cdot)$, and the normalizing constant of each \bar{t} by $A_{\bar{t}}$,

$$\lim_{Z \rightarrow \infty} \alpha_\Phi(Z) = \frac{\sum_{\bar{t}=1}^{\infty} \delta(\bar{t}) \lim_{Z \rightarrow \infty} g_{\bar{t}}(Z)}{\sum_{\bar{t}=1}^{\infty} \delta(\bar{t}) \lim_{Z \rightarrow \infty} g_{\bar{t}}(cZ)} = \frac{\sum_{\bar{t}=1}^{\infty} \delta(\bar{t}) A_{\bar{t}} Z^{-\alpha}}{\sum_{\bar{t}=1}^{\infty} \delta(\bar{t}) A_{\bar{t}} (cZ)^{-\alpha}} = c^\alpha.$$

³⁰This capture the essence of power laws – that whatever the value of Z , its frequency and frequency of cZ has the same ratio.

Thus, the exponent of $\Phi(Z_i)$ will be identical to the exponent of $F(\beta_{i\tau})$ asymptotically.

This discussion suggests that whenever possible, it is desirable to take the estimates from the tail end of the distribution instead of using moderate values of Z . For the COVID-19 from the world, the distributions are estimated from the very extreme tail. But when the sample size of SSEs is limited, choice of how many observations to include thus faces a bias-variance trade-off. Nonetheless, as many statistical theories are based on asymptotic results, these arguments show that it is theoretically founded to interpret the exponent of $\Phi(Z_i)$ as the exponent of $F(\beta_{i\tau})$, at least given the data available.

A2 Robustness

We present several robustness checks on our empirical results.

Figure 2-1 with a different cut-off

In Figure 2-1, we truncated the size of cluster from below at 40. Figure 2.5.1 instead show results with a cut-off of 20. The fit is worse at the lower tail of the distribution, which suggests that the lower tail may not be approximated by power law distribution. This is a common feature among many examples. However, what matters for the existence of variance is the upper tail distribution, we do not think this is a concern. Moreover, given that the data partly come from media reports, the clusters of small sizes likely suffer from omission due to lack of media coverage.

Robustness of power law exponents estimates


Gabaix and Ibragimov (2011) show that an estimate of 2.2 is biased in a small sample and propose a simple bias correction method that replace the dependent variable with $\ln(\text{rank} - 1/2)$. Panel A of Table 2.5.1 show the results with this bias correction method. The results are broadly very similar to our baseline results in Table 2.1.

Panel B of Table 2.5.1 conduct another robustness check, where we estimate using the maximum likelihood. Again, the point estimates are overall similar to the baseline results, although standard errors are larger.

A.3 Additional Tables and Figures

Table 2.5.2 shows several examples of superspreading events during COVID-19 pandemic.

COVID-19 Cluster Sizes Worldwide Data from Research Articles and Media Reports



Source: CMMID COVID-19 Working Group online database (Leclerc et al., 2020)

Figure 2.5.1: Log size vs log rank for Superspreading Events in SARS 2003

Notes: Figure 2.5.1 plots the number of total cases per cluster (in log) and their ranks (in log) for COVID-19, last updated on June 3rd. It fits a linear regression for the clusters with size larger than 30. The data are collected by the Centre for the Mathematical Modelling of Infectious Diseases COVID-19 Working Group (Leclerc et al., 2020).

Panel A. Bisas corrected regression estimates							
	COVID-19			SARS			MERS
	World	Japan	India	World	Singapore	Beijing	World
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\hat{\alpha}$	1.16	1.45	1.70	1.02	0.86	0.96	1.29
	(0.07)	(0.16)	(0.06)	(0.10)	(0.12)	(0.10)	(0.11)
\underline{Z}	40	2	2	8	2	2	2
Obs.	60	11	109	15	19	8	36
R ²	0.97	0.93	0.96	0.95	0.89	0.93	0.95
$\log_{10} LR$	-	11.73	-	-	19.92	8.05	41.19

Panel B. Maximum likelihood estimates							
	COVID-19			SARS			MERS
	World	Japan	India	World	Singapore	Beijing	World
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$\hat{\alpha}$	1.01	1.96	1.71	0.89	1.21	0.87	1.49
	(0.13)	(0.59)	(0.16)	(0.23)	(0.28)	(0.31)	(0.25)
\underline{Z}	40	2	2	8	2	2	2
Obs.	60	11	109	15	19	8	36
$\log_{10} LR$	-	11.93	-	-	20.34	8.07	46.93

Table 2.5.1: Estimates of power law exponent: robustness

Notes: Table 2.5.1 summarizes two robustness check exercises of power law exponent ($\hat{\alpha}$). Panel A. bias corrected estimates take $\log(\text{rank} - \frac{1}{2})$ as the dependent variable. This is a small sample bias correction proposed by (Gabaix and Ibragimov, 2011). Heteroskedasticity-robust standard errors are reported in the parenthesis. Panel B. presents the maximum likelihood estimates. Standard errors are reported in the parenthesis. In both panels, $\log_{10}(LR)$ denotes “likelihood ratios”, expressed in the log with base 10, of probability of observing this realized data with power law distributions relative to that with estimated negative binomial distributions. Columns (1)-(3) report estimates for COVID-19; columns (4)-(6) for SARS, and column (7) for MERS.

Major super-spraeding evernts	Confirmed cases	Date
Choir practice in Washington, the US	52	03/10
Conference in Boston, the US	89	02/26
Religious gathering in Daegu, South Korea	49	02/19
Religious gathering in Frankfurt, Germany	49	02/19
Wedding ceremony in New Zealand	76	03/21
Prison in IL, the US	351	04/23
Food processing plant in Ghana	533	05/11
Dormitory in Singapore	797	04/09

Table 2.5.2: Examples of superspreading events

Noes: Table 2.5.2 summarizes some examples of superspreading events, their dates and the number of confirmed cases for COVID-19. *Source:* [COVID-19 settings of transmission - database](#) (accessed, June 4, 2020)

B. Theory Appendix

Proof that S_∞ is convex in \mathcal{R}_0 if $\mathcal{R}_0 > \frac{9}{8(1-R_0)}$

We show that S_∞ is a concave function in \mathcal{R}_0 . Recall that S_∞ is a solution to

$$\log S_\infty = -\mathcal{R}_0(1 - S_\infty).$$

By the implicit function theorem,

$$\begin{aligned} \frac{dS_\infty}{d\mathcal{R}_0} &= -\frac{1}{\left(\frac{1}{S_\infty} - \mathcal{R}_0\right)}(1 - S_\infty) \\ &< 0. \end{aligned}$$

because $S_\infty < 1/\mathcal{R}_0$. Applying the implicit function theorem again,

$$\underbrace{\left(\frac{1}{S_\infty} - \mathcal{R}_0\right)}_{>0} \frac{d^2 S_\infty}{d\mathcal{R}_0^2} = \underbrace{\frac{dS_\infty}{d\mathcal{R}_0}}_{<0} \left(2 - \frac{1/S_\infty - 1}{1 - \mathcal{R}_0 S_\infty}\right).$$

It remains to show that $\left(2 - \frac{1/S_\infty - 1}{1 - \mathcal{R}_0 S_\infty}\right) < 0$. We can rewrite this as

$$f(S_0) \equiv 2\mathcal{R}_0 S_\infty^2 - 3S_\infty + 1 > 0.$$

Note that $f(\cdot)$ is minimized at $S_\infty^* = \frac{3}{4\mathcal{R}_0}$. The minimum value is

$$\min_{S_0} f(S_\infty) = -\frac{9}{8\mathcal{R}_0} + 1.$$

Therefore $f(S_\infty) > 0$ for all S_∞ if and only if $\mathcal{R}_0 > \frac{9}{8}$. This implies that when $\mathcal{R}_0 > \frac{9}{8}$, S_∞ is a concave function of \mathcal{R}_0 .

2.5.1 Proof that I^{\max} is concave in \mathcal{R}_0 if and only if $\mathcal{R}_0 > \frac{1}{S_0} \exp(0.5)$

Recall that the peak infection rate is given by

$$I^{\max}/N = 1 - \frac{1}{\mathcal{R}_0} - \frac{1}{\mathcal{R}_0} \log(\mathcal{R}_0 S_0).$$

The derivative is

$$\frac{dI^{\max}/N}{d\mathcal{R}_0} = \frac{1}{(\mathcal{R}_0)^2} \log(\mathcal{R}_0 S_0).$$

The second derivative is

$$\frac{d^2(I^{\max}/N)}{d\mathcal{R}_0^2} = \frac{1}{(\mathcal{R}_0)^3} (1 - 2 \log(\mathcal{R}_0 S_0)),$$

which is negative if and only if $\mathcal{R}_0 > \frac{1}{S_0} \exp(0.5)$.

Results for targeted lockdown policy experiment

Table 2.5.3 shows the simulation results with lockdown policies targeted at SSEs. $\bar{\beta}$ is the daily upperbound of infection rates due to policies, and we consider cases of $\bar{\beta} = 1000, 100, 50$. As already discussed in the main text, when the distribution is fat-tailed, the targeted policy is not only effective in reducing the mean of the peak infection rate, but also its volatility (the interval between 90 percentile and 10 percentile).

	Power law					Negative binomial
	$\alpha = 1.08$	$\alpha = 1.1$	$\alpha = 1.2$	$\alpha = 1.5$	$\alpha = 2$	
1. $\bar{\beta}$: 1000 cases per day						
mean	11%	15%	23%	27%	27%	27%
90th percentile	19%	23%	29%	29%	28%	27%
50th percentile	8%	12%	21%	26%	27%	27%
10th percentile	4%	7%	17%	25%	26%	26%
3. $\bar{\beta}$: 100 cases per day						
mean	9%	12%	20%	26%	27%	27%
90th percentile	17%	20%	26%	27%	28%	27%
50th percentile	5%	8%	18%	26%	27%	27%
10th percentile	3%	5%	16%	24%	26%	26%
3. $\bar{\beta}$: 50 cases per day						
mean	8%	11%	19%	26%	27%	27%
90th percentile	14%	19%	26%	27%	28%	27%
50th percentile	4%	8%	17%	25%	27%	27%
10th percentile	2%	5%	14%	24%	26%	26%

Table 2.5.3: Peak infection under targeted lockdown policy

Note: Table 2.5.3 shows the summary statistics for peak infection rates from 1000 simulations with various policy parameters $\bar{\beta}$, where $\bar{\beta}$ is the upperbound on the infection imposed by the policy.

Chapter 3

A Model of Reflection and Meditation in Experimentation under Imperfect Recall

3.1 Introduction

Meditation, a process of regulating and focusing attention on the present – is an ancient practice that has received renewed interest in the modern era. For two thousand years, meditation has been embraced by numerous philosophical and religious traditions as a means of improving people’s well-being. Given the fast-paced nature of modern life, many people are now turning to meditation as a means to slow down, reflect, and restore their psychological health. According to the National Health Interview Survey, in the five years from 2012 to 2017, the proportion of American adults who have meditated increased from 4.1 to 14.2 percent (Tainya et al. 2018). There is now a nationwide initiative to teach meditation in K-12 schools (Ryan 2012), and even the US military and Congress have adopted the practice for their members. This renewed interest has accompanied scientific research, with the number of journal articles on the topic increasing from just 50 to over 500 from 1997 to 2017. However, the actual techniques of meditative practices, such as breathing, physical exercise and mantras, vary widely across practitioners, and this is reflected in research studies. Thus, to draw an externally valid inference from the evidence with different practices and subjects, a precise theory regarding its effects is needed.¹

¹Since the time of the initial thesis submission, there were some changes made to allow for completion of the proof. The main change is the change in the signal structure: Self 2 now receives an additional signal about the Self 1’s information, and the two Selves’ signals are no longer assumed

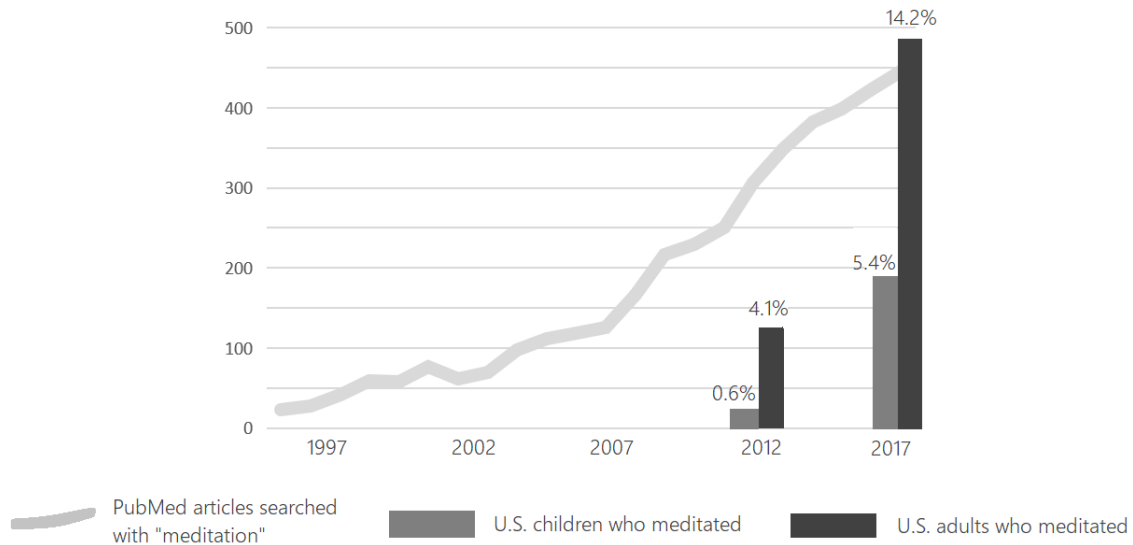
Across histories and cultures, teachers of meditation have surmised that the psychological benefits of meditation arise primarily from reducing the effects of negative thoughts. Jon Kabat-Zinn, the founder of mindfulness meditation, writes that “mindfulness is the awareness that arises when we *non-judgmentally* pay attention in the present moment” (Kabat-Zinn 2009). A Buddhist scholar writes that Bodhidharma, the founder of Zen meditation, aspired “to realize wisdom and transparency. While ordinary kind of knowledge helps us by informing our comparison of alternatives through our memories, wisdom is based on clearing our mind to be *empty*” (Yanagida 1974). While insightful, these teachings have naturally faced much suspicion. First, they appear paradoxical because it contradicts expected utility models that any knowledge will be welfare-enhancing because it can better inform the decisionmaker in their comparison of alternatives (Blackwell 1953). Second, their theories are ancient and thus based on introspection, a method of inquiry prone to subjective assessment and logical vagueness. Meanwhile, as an alternative, some economists rely only on neurological and physiological understanding. While this approach is based on scientific measurement and provides explanations for pharmaceutical treatments, it remains incomplete because the focus of meditation is on thoughts and beliefs rather than neurological reactions in the brain.

To provide logical rigor and derive testable empirical implications, this paper proposes a formal model of meditation that builds on the insights of the empty mind hypothesis. Our model is based on a standard two-period model of active learning (Rothschild 1974; see Sobel 2000 for review) augmented with two kinds of human memory imperfections. First, when experimenters learn from outcomes, they have an imperfect recall of past actions and information because the memories fade over time. Second, despite this imperfect recall, experimenters have persistent prior belief over one’s type due to “autobiographical memories.” In this setting, meditation is modeled as increasing the weights experimenters put on current signals rather than prior beliefs to form expectations about each other. This model shows that meditation improves experimentation decisions by diffusing the impact of autobiographical memories².

to be bounded. The additional changes were that the cost of effort is assumed to be constant, and when there is no effort, no information regarding the underlying types is generated. Due to these changes, the change in information called ‘meditation’ is now referred to as “reflection over past information” in the main text.

²Some may disagree with the interpretation of “empty mind” as the diffused prior. Instead, one may interpret this as a removal of negative thoughts, as negative thoughts are the cause of negative emotions. However, such selective suppression of negative thoughts, or positive illusion, is contrary

Figure 3.1.1: A rapid increase in evidence and practice of meditation



Notes: Figure 3.1.1 describes the number of articles with “meditation” and percentage of Americans who answer they have implemented meditation at least once in the past year in National Health Interview Survey.

Under these memory imperfections, the experimentation decision and belief updating will exhibit biases towards autobiographical memory. Suppose an experimenter has a autobiographical memory with low confidence. Then, even if she *knows* it is worth trying, she may not experiment in equilibrium. On the other hand, an experimenter with high confidence may continue even if she *knows* it is better to let go. Given the first-period outcomes, the inference will be expected to yield biases towards the autobiographical memory on average. Even if the experimenter with low confidence succeeds, she will likely attribute this to luck rather than to her past efforts. When the experimenter with high confidence experiences a failure, she will likely attribute to efforts instead of luck. In this way, unlike in the standard experimentation models of perfect memory, the level of confidence of the experimenter will no longer be a sufficient statistic that determines the behaviors (Gittins 1979).

While these biases emerge in an imperfect recall framework, they are alleviated and

to what many meditation practitioners teach. Others may then interpret this as a removal of any thoughts. However, taken literally, this will lead to inaction. In contrast, meditation is shown to make people more actively engaged by observing their conditions more clearly. While a no-mental-content-at-all state may be advanced by some practitioners, many popular meditations aim at lack of judgment.

even asymptotically eliminated by meditation. This effect arises because the memory imperfection has introduced a *coordination problem* between decisions over multiple periods. When attention is unstable across periods, each decision and inference must rely on long-term memory to achieve coordination. In contrast, when attention is stable, long-term memory will be uninformative regarding another Self’s action and, by avoiding its influence, the decisions will converge to the perfect recall optimum. This result is notable in that it stands in contrast with the ordinary understanding that coordination requires one to be informed about another’s action; instead, it is *lack* of information that leads to unbiasedness.

To sum up the discussions thus far, the proposed model shows that embedding memories and attention within a formal model of active learning can illuminate the role of meditation in “emptying” one’s mind. However, there are two limitations to this approach. First, the model does not include physiological responses. This omission is a modeling choice, as the model is intended to illustrate the main mechanism with minimal assumptions, rather than to include all real-world elements³. Second, the model appears to differ from other models economists have developed regarding psychological distress. At the bottom line, theories are intended to clarify mappings from assumptions to testable hypotheses, so their relevance must be examined based on evidence when introspections vary. However, such disagreements exist and persist in psychology, and difference from dominant models does not necessarily imply deficiencies of new models. In the remainder of this section, I contrast the proposed model against two other approaches.

First, one approach used by other economists have attributed states of psychological distress to biased beliefs, that is, the difference between subjective and objective probability assessments. They then suggest that a lack of experimentation, and a resulting belief-behavior feedback loop, is the sustaining mechanism of distress (de Quidt and Haushofer 2017). This conceptualization captures an important aspect of distress, for in Behavioral Therapy, clients are asked to deliberately experiment with new actions to overcome this feedback loop⁴. However, while important, there are two implications of this approach that are inconsistent. First, if psychological distress is defined as the difference between internal belief and external reality, then distress would be eliminated by a worsening external reality that comes to match a pessimistic internal belief. However, a common sense would suggest otherwise. The

³It is left for future work to incorporate physiological symptoms into the model. For example, one could assume that the cost of second-period effort is a function of the emotional response to the first-period outcome, say, because the emotional response affects the sleep one can have.

⁴Since my model extends the standard experimentation model, this mechanism can be present.

proposed model, by contrast, attributes a state of distress to one’s relationship with oneself, avoiding such predictions. Second, this approach suggests that the role of psychotherapy is to persuade one to think more positively. However, many professional cognitive therapists argue that this is a misconception of their service. In contrast, the proposed model is consistent with cognitive therapy practitioners who suggest that the role of psychotherapy to let go of judgmental thoughts and focus on the present, rather than merely trying to think more positively.

A second alternative framework is to consider affect as a state that shapes utility functions and belief updating. For example, when the decisionmaker is in an “angry” state, she thinks negatively of others. Meditation shrinks her amygdala, which shifts her state to be “calm,” and thereby changes her thoughts. As discussed in Section 4, the proposed model can also be interpreted similarly, as meditation changes the non-standard components of the objective function and inference process. However, one advantage of the proposed model is that this is an implication rather than an assumption, and mechanisms are illustrated through its derivation. Moreover, there is also a fundamental difference rooted in clinical psychology debates. Like this prominent alternative approach, many psychologists had previously believed that emotional states precede thoughts (e.g. “I think badly of others because I am angry.”) Since the 1960s, however, cognitive therapies were developed by Ellias and Beck (e.g. Beck 1967) based on their observation that there are preconscious thoughts that precede emotional states (e.g. “I feel angry because I, perhaps without awareness, perceive others to be bad.”) Meditation shares with these cognitive therapies the common objective of bringing these thoughts into awareness and to scrutinize them, and thereby improve the emotional states⁵. Evidence suggests this approach is highly effective.

Relation to the Literature. This paper draws upon and informs the literature on (i) information economics, (ii) behavioral economics, and (iii) organizational economics.

First, in information economics, while information acquisition models with perfect memory suggest that the level of confidence determines experimentation decisions, this paper extends the analysis to shows that attention and memory structures will also be key when memory is imperfect. In standard models (Rothschild 1974; see Sobel 2000 for review), the celebrated result is sufficiency of confidence level (e.g. Gitten’s index in “multi-armed bandit” models). Here, we extend this model with imperfect recall of past actions (Kuhn 1953, Harsanyi and Selton 1988, Piccione and

⁵This stands in contrast with, but can be consistent with, other pharmaceutical and physiological techniques that directly improve the emotional states.

Rubinstein 1997, Aumann et al. 1997), noisy signals (Bagwell 1995), and global game perturbations of preferences (Carlsson and van Damme 1993, Frankel et al. 2003) that arise due to noisy cognition (See Khaw, Li, and Woodford 2017 for a similar approach). The key implication is that without perfect recall, confidence no longer uniquely determines experimentation decisions, and that attention and memory structures will be critical instead. This is due to externalities of information acquisition (e.g. Bolton and Harris 1999, Keller et al. 2005; see Hörner and Skrzypacz 2016 for review). However, rather than incentives, coordination will be the key friction.

Second, this study contributes to behavioral economics, where various papers have examined the role of information processing imperfections in passive learning models. For example, biases in learning have been explained by quasi-Bayesian models (Rabin and Schrag 1999), motivated beliefs, and model mis-specification (Heidius et al. 2018⁶, He 2019). Wilson (2014) examines the role of coarse memories on the biases of eventual learning, while de Silveira and Woodford (2019) show that when the memory capacity of the decision-maker is limited, there will be an over-reaction to new information in both beliefs and actions. By contrast, this paper examines information processing imperfections within an active learning model, showing that biases in higher order beliefs result in biases in the first order belief in later periods, even under Bayes' rule.

Finally, this paper also contributes to the theory of “teams” in organizational economics, providing new results on the role of organizational memory and communication. Teams are groups of decisionmakers who share a common interest but who hold private information (Marschak and Radner 1972). More recently, Dessain and Santos 2006 ask when it is worthwhile to improve communication among workers when they face a trade-off between adaptation (i.e. responding to private information) and coordination (i.e. ensuring that decisions are well-coordinated with one another.) As multiple selves also share a common interest and yet face different information, the model presented in this paper also addresses this adaptation-coordination trade-off in the context of information acquisition (e.g. generations of workers deciding whether to invest in a risky project.) To maximize adaptation, it is optimal to reduce shared organizational memory, and allow workers to communicate with each other. However, if the firm's objective is to induce a high level of output, then it is optimal to instill positive organizational memory, and restrict communication among workers so that

⁶Even though their model has agents to choose actions to endogenously generate signals, the agents choose actions myopically without experimentation motives. In this sense, the strategic reasoning of this paper is absent in their model.

they will exert effort even when their private signals suggest the situation is difficult. By introducing a global game framework (Morris and Shin 2001), this paper analyzes the new role of shared organizational memory beyond what is captured by existing models.

3.2 Set-up

This Section introduces the set-up of the model and shows that the unique equilibrium has a threshold-form. The set-up extends a standard 2-period experimentation model with imperfect recall over first period action and information. The set-up is thus formulated as a model of multiple Selves since either decision-maker cannot control another's action.

3.2.1 Environment

There are two Selves, Self 1 and Self 2, each making effort decision in respective period $t = 1, 2$. Self 1 chooses her action in period 1, $a_1 \in \{0, 1\}$ to produce the output $y_1 \in \{0, 1\}$. Self 2 has an *imperfect recall over action*: she observes a noisy signal, $z \in \{0, 1\}$, of action a_1 , that is correct with probability $r_a \in [\frac{1}{2}, 1]$. After observing y_1 and z , Self 2 chooses her action in period 2, $a_2 \in \{0, 1\}$ to produce $y_2 \in \{0, 1\}$. The two Selves share a common underlying payoff:

$$\mathbb{E} [y_1 - ca_1 + y_2 - ca_2], \quad (3.1)$$

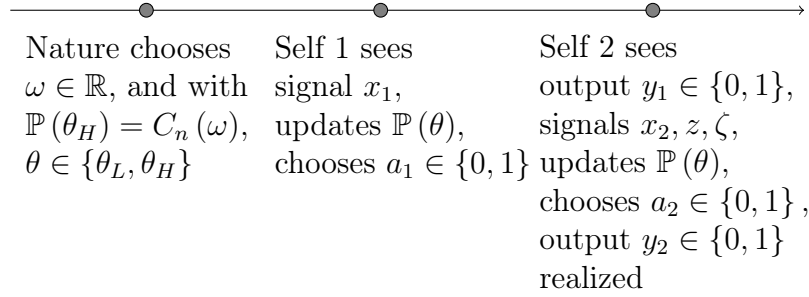
where $c > 0$ is the cost of effort $a_t = 1$.

There are also two types that the Selves learn about through noisy signals using the Bayes' rule. First, there is a productivity type, $\theta \in \{\theta_H, \theta_L\}$, and the probability of being a high type is $p \equiv \mathbb{P}(\theta_H) \in (0, 1)$. Effort raises the probability of high outcome more when the type is high: $\pi_H > \pi_L > \pi_0$, where $\pi_\theta = \mathbb{P}(y_t = 1 | a_t = 1, \theta) \in (0, 1)$ and $\pi_0 = \mathbb{P}(y_t = 1 | a_t = 0) \in (0, 1)$. The first-period effort, $a_1 = 1$, is called an *experimentation* since its outcome can inform the second-period action about the type θ . The Selves *actively* learn about θ by Self 1's experiments, and Self 2's inference is based on her recalled history of the first period. Let $\pi_H - \pi_0 > c$ and $\pi_L - \pi_0 < c$ so that their confidence is consequential for their effort decision.

Second, there is a potential type, $\omega \sim \mathcal{N}(\bar{\omega}_0, \sigma_\omega^2)$. Let us denote a censoring function $C_n : \mathbb{R} \mapsto [0, 1]$ that censors the values outside of a unit interval⁷. The

⁷That is, $C_n(l) = l$ if $l \in (0, 1)$, $C_n(l) = 0$ if $l \leq 0$, and $C_n(l) = 1$ if $l \geq 1$

Figure 3.2.2: Time line



Notes: Figure 3.2.2 describes the time line of this model. Here, $\mathbb{P}(\theta)$ denotes the average belief $p \equiv \mathbb{P}(\theta_H)$ conditional on signals.

potential type determines the distribution of the distribution of productivity type, p , by the censoring function: $p = C_n(\omega)$. Each Self *passively* learns about this potential type, ω , at the beginning of each period t through noisy signals, x_t . The second-period information includes both first-period signal, x_1 , and additional information, ξ :

$$x_1 = \omega + \varepsilon_1 \text{ and } x_2 = \frac{x_1 + \xi}{2}, \text{ where } \xi = \omega + \varepsilon_2 \quad (3.2)$$

and $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. However, Self 2 has an *imperfect recall over information*: while she observes x_2 perfectly, she cannot distinguish x_1 and ξ and instead observes an additional signal, $\zeta = x_1 + \varepsilon_0$, where $\varepsilon_0 \sim \mathcal{N}(0, [(1 - r_\varepsilon) \sigma_\varepsilon]^2)$. That is, while Self 2 remembers what Self 1 knew, she imperfectly recalls what Self 1 did not know. The parameters and the distributions are common knowledge.

The time line is as follows (Figure 3.2.2): first, Nature chooses the potential type, ω , and chooses the type θ according to the distribution $p = C_n(\omega)$; second, the Self 1 observes x_1 and makes the effort decision, $a_1 \in \{0, 1\}$; third, the Self 2 observes outcome y_1 , signal of distributon type x_2 , and signal of first-period action z and information ζ , and updates her belief $\mathbb{P}(\theta)$. Then he chooses $a_2 \in \{0, 1\}$ given her posterior belief; and finally, the output y_2 is realized.

Interpretation. To interpret this set-up concretely, let us consider college students who study for a linear algebra class. They choose how much to study, a_t , for the mid-term exam ($t = 1$) and the final exam ($t = 2$). The students can pass the exam, y_t , with chances that depends on both their efforts, a_t , and their specific preparedness

for the exam θ . When mid-terms are returned months after the exam, they have a noisy memory of how much they had studied for the mid-term ($r_a < 1$) and what they were thinking then ($r_\varepsilon < 1$). They can engage in reflection practices, such as journaling of study efforts, to improve these recall precisions.

The students do not know their specific preparedness for the exam, θ , but infer it from results of past exams and their readiness for the class ω . The students also do not know ω , either, but infer this from their knowledge of general ability for math, $\{\bar{\omega}_0, \sigma_\omega^2\}$, as well as information, x_t , they have from classes. The signals, x_t , represent, for example, whether they can do problem sets or they understand lectures easily. The students know their general ability, $\{\bar{\omega}_0, \sigma_\omega^2\}$, from their personal histories such as in high schools, such as whether their parents or teachers discouraged or pushed them from pursuing math.

3.2.2 Equilibrium

This sub-Section defines an equilibrium and shows that the equilibrium will take a threshold form.

Perfect Bayesian Nash Equilibria

This paper will analyze the Perfect Bayesian Nash Equilibrium, the standard equilibrium concept to analyze dynamic games of incomplete information, and henceforth call them as *equilibrium*. Let the action space be $\mathcal{A}_t \equiv \{0, 1\}$; the outcome space be $\mathcal{Y}_t \equiv \{0, 1\}$; and the interim type space be $\mathcal{X}_t \equiv \mathbb{R}$. The *Self 1's strategy*, s_1 , is a mapping from the interim type space, \mathcal{X}_1 , to a distribution over action space, \mathcal{A}_1 ; that is $s_1 : \mathcal{X}_1 \mapsto \Delta(\mathcal{A}_1)$. The *Self 2's strategy*, s_2 , is a mapping from the interim type space, \mathcal{X}_2 , the first-period outcome space, \mathcal{Y}_1 , the noisy signal of first-period action in the space, $\mathcal{Z} = \mathcal{A}_1$, the noisy signal of first-period information in the space, $\hat{\mathcal{X}}_1 = \mathcal{X}_1$, to a distribution over own action space, \mathcal{A}_2 ; that is $s_2 : \mathcal{X}_2 \times \mathcal{Y}_1 \times \mathcal{Z} \times \hat{\mathcal{X}}_1 \mapsto \Delta(\mathcal{A}_2)$. Denote the Self 2's prior belief over the Self 1's strategy by $\mu \equiv \mathbb{P}(s_1)$.

Definition 1 Equilibrium *An equilibrium is a tuple of strategies and beliefs $\{s_1, s_2, \mu\}$ such that (i) all strategies maximize the objective (3.1) given the strategies of each other; (ii) beliefs are consistent with the Bayes' rule; (iii) off-equilibrium strategies also maximize the objective (3.1).*

Unlike some other analyses of multiple equilibria, the results herein will not be driven by “implausible” off-path beliefs and are robust to refinements such as strategic stability (Kohlberg and Mertens 1986) and intuitive criterion (Kreps and Cho 1987).

Preliminary Analyses

To provide the benchmark, we first consider the setting with perfect recall over actions. While there is an imperfect recall over information, the equilibrium is identical to the standard setting without such imperfection. Henceforth, let us denote the period- t interim belief conditional on signal x_t by $\bar{p}_t(x_t)$.

Proposition 1. Equilibrium under Perfect Recall over Action and Imperfect Recall over Information ($r_a = 1, r_\varepsilon \leq 1$). *Self 1 and Self 2 exert their effort if and only if they are sufficiently confident:*

$$a_1^{**} = \begin{cases} 1 & \text{if } \bar{p}_1(x_1) \geq \bar{\Pi}_1 \\ 0 & \text{if o.w.} \end{cases} \quad \text{and} \quad a_2^{**} = \begin{cases} 1 & \text{if } \bar{p}_2(x_2) \geq \bar{\Pi}_2(z, y_1) \\ 0 & \text{if o.w.} \end{cases},$$

where the thresholds, $\{\bar{\Pi}_1, \bar{\Pi}_2(z, y_1)\}$, satisfy

$$0 < \bar{\Pi}_2(1, 1) < \bar{\Pi}_2(0, y) < \bar{\Pi}_2(1, 0) < 1. \quad (3.3)$$

and $\bar{\Pi}_1 < \bar{\Pi}_2(0, y)$, and are independent of r_ε and $\bar{\omega}_0$.

Sketch of Proof. By the backward induction. The ordering (3.3) suggests, for intermediate levels of interim belief, $\bar{p}_2(x_2) \in (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(1, 0))$, the second-period effort depends on the histories $\{a_1, y_1\}$ while for the extreme levels, the effort level only depends on the belief. In particular, for the moderately low level of confidence $\bar{p}_2(x_2) \in (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(0, y))$, the Self 2 exerts effort if the Self 1 has experimented and succeeded; for the moderately high level of confidence $\bar{p}_2(x_2) \in (\bar{\Pi}_2(0, y), \bar{\Pi}_2(1, 0))$, the Self 2 exerts effort unless the Self 1 has experimented and failed. In this way, Self 1's effort has an informational value to improve Self 2's decision. Thus, the confidence level required to exert the first-period effort, $\bar{\Pi}_1$, is less than the "myopic" level, $\bar{\Pi}_2(0, y)$, required to exert the second-period effort in the absence of additional information. Appendix A2.1 contains a complete proof. \square

The Proposition 1 shows that both Selves exert effort if and only if they are sufficiently confident given their information. For example, the students with low confidence will not study for the final unless they have additional confidence from their experience of studying hard for the mid-term and performing well; on the other hand, those with high confidence will study for the final unless they have an experience of trying hard for the mid-term and yet not succeeding. Some students with moderate level of confidence may thus study hard for the mid-term, not only to perform well

in the mid-term itself, but also to learn how much to study for the final. This result that the experimentation decision depends on the confidence level is a fundamental results in the models of experimentation, and often employed as the Gittins' index in the context of multi-armed bandit models.

To characterize the role of imperfect recall, let us first show that the unique equilibrium will take a threshold-form:

Lemma 1. Equilibrium Thresholds under Imperfect Recall over Action and Information ($r_a < 1, r_\varepsilon < 1$). *There exists some $\bar{r}_\varepsilon < 1$ such that for any $r_\varepsilon \in [\bar{r}_\varepsilon, 1)$, there is a unique equilibrium with the following threshold-form strategies: Self 1 chooses*

$$a_1^* = \begin{cases} 1 & \text{if } \bar{p}_1(x_1) \geq \Pi_1^r \\ 0 & \text{if o.w.} \end{cases}$$

Self 2's choice will depend on the histories:

- *When $y_1 = 1$: for each z , for $\bar{p}_2(x_2) \in (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(0, y))$*

$$a_2^* = \begin{cases} 1 & \text{if } \zeta \geq \zeta_{z1}^r(x_2) \\ 0 & \text{if o.w.;} \end{cases} \quad (3.4)$$

and $a_2^ = 1$ if $\bar{p}_2(x_2) > \bar{\Pi}_2(0, y)$ and $a_2^* = 0$ if $\bar{p}_2(x_2) \leq \bar{\Pi}_2(1, 1)$.*

- *When $y_1 = 0$: for each z , for $\bar{p}_2(x_2) \in (\bar{\Pi}_2(0, y), \bar{\Pi}_2(1, 0))$*

$$a_2^* = \begin{cases} 1 & \text{if } \zeta < \zeta_{z0}^r(x_2) \\ 0 & \text{if o.w.;} \end{cases} \quad (3.5)$$

and $a_2^ = 1$ if $\bar{p}_2(x_2) > \bar{\Pi}_2(1, 1)$ and $a_2^* = 0$ if $\bar{p}_2(x_2) \leq \bar{\Pi}_2(0, y)$.*

The thresholds, $\{\Pi_1^r, \zeta_{zy}^r(x_2)\}$, depend on r_a, r_ε , and \bar{w}_0 .

Sketch of Proof. By the logic of uniqueness of threshold-form equilibrium in “global games” (Carlsson and van Damme 1993)⁸. There are two steps in this proof. The first step shows that any rationalizable strategies must be bounded by some threshold-form strategies by the iterative elimination of strictly dominated strategies. In the

⁸Note that, unlike some parts of the global game literature, this paper does not use this technique as an equilibrium selection criterion. Instead, this paper takes the noisy signals as a characteristic of imperfect memory, as suggested by psychology and neuroscience literature, and studies the role of stability of the signals in the cases away from the limit.

-the term “essentially” is added to account for strategies at the cut-off

extreme levels of confidence such that Self 2's choices are independent of histories, the Self 2 has dominant strategies. From these regions, contagion restricts rationalizable strategies in the intermediate levels of confidence where the choices depend on recalled history of the Self 1's action. The second step shows that Self 1's such concerns of imperfect recall is small enough when information is recalled sufficiently accurately so that these bounds contain essentially a unique threshold, constituting the unique equilibrium. Appendix A3.2.3 provides a complete proof. \square

The Lemma 1 shows that the decisions under the imperfect recall remain heuristic despite the complexity of information structure. Self 1 experiments if she is sufficiently confident, but this threshold now depends on her expectation of Self 2's recall. In turn, Self 2 who experiences a surprising outcome – those with low initial confidence but with success, and those with high initial confidence but with failure – changes her second period effort if she is sufficiently sure that Self 1 has experimented. The subsequent analysis will illustrate how this imperfect recall leads to biases in decisions and inference.

3.3 Main Analyses

This Section first shows that, under imperfect recall, the threshold of experimentation will exhibit biases towards the prior belief. Then, reflection is shown to reduce this bias by alleviating biases in inference. The underlying mechanism of coordination problem is also illustrated.

3.3.1 Behaviors and Beliefs

Let us begin by considering the implications of imperfect memories on biases in experimentation and inference. By using the notation similar to the definition, let $\mu(x_2, \zeta) \equiv \mathbb{P}(a_1 = 1 | x_2, \zeta)$ be Self 2's belief that $a_1 = 1$ conditional on signals x_2, ζ . Let $H(\mu | \bar{p}_1, \bar{\omega}_0)$ denote the distribution of the Self 2's belief given Self 1 with the belief \bar{p}_1 and prior $\bar{\omega}_0$.

Proposition 2.1. Equilibrium with Bias to Conform to Prior Beliefs under Imperfect Recall over Action and Information ($r_a < 1, r_\varepsilon < 1$). *Given $r_\varepsilon \in [\bar{r}_\varepsilon, 1)$,*

$$\frac{\partial \Pi_1^r(\bar{\omega}_0)}{\partial \bar{\omega}_0} < 0.$$

There exists a unique prior ω_0^ such that $\Pi_1^r(\omega_0^*) = \bar{\Pi}_1$ such that the following holds:*

- (i) **Prior with Low Confidence:** when $\bar{\omega}_0 < \omega_0^*$, there will be an underexperimentation ($a_1^* = 0$ while $a_1^{**} = 1$) for $\bar{p}(x_1) \in [\bar{\Pi}_1, \Pi_1^r]$. At the threshold, the higher order expectation for experimentation is low: $H(\mu|\Pi_1^r, \bar{\omega}_0) > H(\mu|\bar{\Pi}_1, \omega_0^*)$.
- (ii) **Prior with High Confidence:** when $\bar{\omega}_0 > \omega_0^*$, there will be an overexperimentation ($a_1^* = 1$ while $a_1^{**} = 0$) for $\bar{p}(x_1) \in [\Pi_1^r, \bar{\Pi}_1]$. At the threshold, the higher order expectation for experimentation is high: $H(\mu|\Pi_1^r, \bar{\omega}_0) < H(\mu|\bar{\Pi}_1, \omega_0^*)$.

Sketch of Proof. When the Self 2 has an imperfect recall over Self 1's action, the Self 2 relies on her signal, x_2 , to infer Self 1's action. Since Self 2's information is uncertain from the perspective of Self 1, Self 1 will in turn use her signal, x_1 , to infer what Self 2 will think about what she thinks. Since their signals are noisy, Self 1 expects Self 2 to expect Self 1 to observe information closer to the prior. We can heuristically focus on the mean: Self 1 expects Self 2 to observe

$$\mathbb{E}[x_2|x_1] = \frac{x_1}{2} + \frac{\alpha x_1 + (1 - \alpha)\bar{\omega}_0}{2}, \text{ where } \alpha \equiv \frac{\sigma_\omega^2}{\sigma_\varepsilon^2 + \sigma_\omega^2}.$$

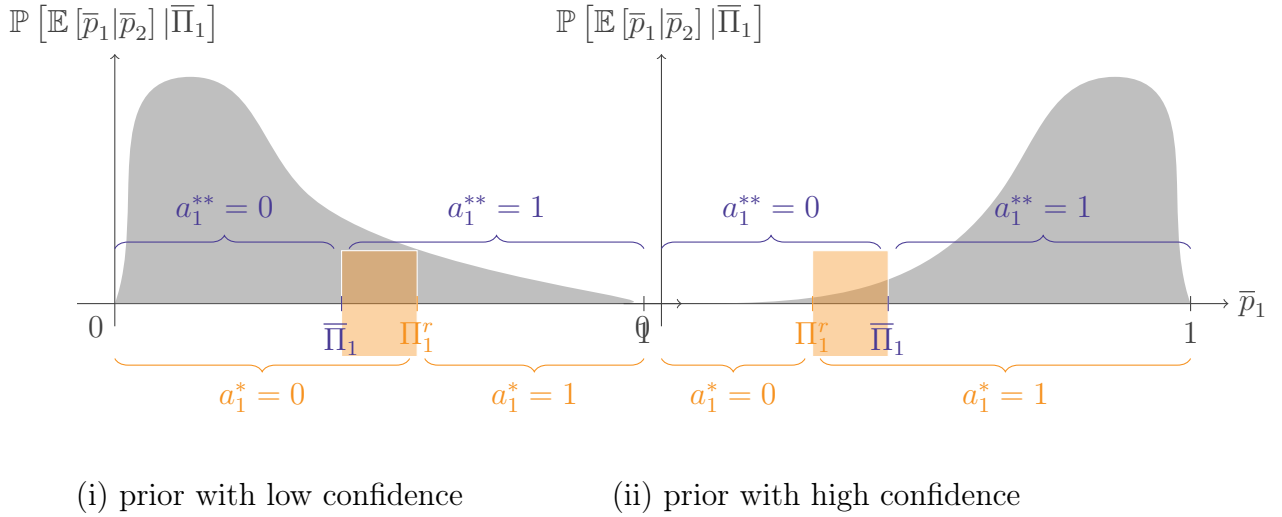
Further, substituting Self 2's expectation over Self 1's signal conditional on the information $\{x_2, \zeta\}$

$$\mathbb{E}[\mathbb{E}[x_1|x_2, \zeta]|x_1] = \gamma(1 - \alpha)\bar{\omega}_0 + [1 - \gamma(1 - \alpha)]x_1, \text{ where } \gamma \equiv \frac{(1 - r_\varepsilon)^2}{(1 - r_\varepsilon)^2 + 1}.$$

More precisely, the distribution of μ , Self 2's expectation over Self 1's action, is shifted towards the prior belief as stated in (i) and (ii).

Recall that the experimentation was a decision to improve future payoff by providing more information at the cost of today's payoff. Thus, the experimentation decision depends critically on the future decision of whether to incorporate the information from today's outcome. In this way, experimentation decision exhibits a coordination problem between Self 1 and Self 2 under the imperfect recall: when Self 1 thinks she is expected not to experiment due to the prior with low confidence, she does not; when Self 1 thinks she is expected to experiment, she does so. Therefore, the the minimum level of confidence to experiment in equilibrium, Π_1^r , is lower when the prior level, $\bar{\omega}_0$, is higher conditional on the interim belief. A complete proof is given in the Appendix A3.3.1. \square

Figure 3.3.3: Experimentation with bias towards prior belief



Notes: Figure 3.3.3 show the equilibrium experimentation decision vs compared to the perfect recall benchmark.

The Figure 3.3.3 visualizes the skewness of higher order beliefs and resulting biases in experimentation decision. When the prior confidence level is low, at the benchmark threshold $\bar{\Pi}_1$, Self 1 expects Self 2 to expect her to have observed signals with the downward skewness. To conform towards this expectation, Self 1 chooses $a_1^* = 0$ more frequently. The opposite bias occurs for the prior with high confidence. Note that, when the prior level is low, it is more common to have the interim belief, \bar{p}_2 , in the region $(\bar{\Pi}_2(1,1), \bar{\Pi}_2(0,y))$ so that the outcome of success, $y_1 = 1$, is consequential. In contrast, when the prior level is high, the outcome of failure, $y_1 = 0$, becomes consequential. This observation allows us to interpret this proposition with the following examples:

Let us consider some students from the disadvantaged backgrounds. For example, female students in the fields of math and sciences are often handicapped relative to their male peers. In the standard expected utility theory, they do not give a try to challenging classes either because (i) they are not confident about their ability, or (ii) rewards to success are low for them. But in this model, one could be just as confident in her capacity as other male students based on objective assessments, and there may even be strong rewards from success. However, one merely thinks that the efforts and success in areas of math and sciences are not what they expect from themselves. Even if successful in the mid-term exam, they tend to attribute that success to mere

luck, not to effort, and thus unlikely to keep exerting efforts for the final. Given such prospect, they never try in the first place.

Consider some other students with an elite background, who were always taught to work hard in areas they were expected to perform. In the standard theory, they exert effort either (i) because they are confident about their ability, or (ii) because their rewards to success are high. But in this model, some could in fact have a low confidence and think that there is little benefit to performing well in the class, and yet they keep putting efforts because that is what they expect of themselves. If they perform poorly, they tend to think of themselves as incompetent, even if the result may be due to mere lack of efforts. To avoid such judgment towards oneself in the future, they keep working hard while knowing it is better for themselves to take it easy.

While there are important works that have explored the implication of conformity bias, this model differs from them in both mechanisms and implications. Other models have directly modified the Bayes' rule to put heavier weights on prior beliefs relative to the signals (See e.g. Rabin and Schrag 1999), explored their implications in passive learning models. In contrast, this model maintains the Bayes' rule but focuses on the role of higher-order beliefs in the active learning setting. In their model, the decisionmakers think their decisions are optimal. Here, the decisionmaker knows that their decision is in some ways biased relative to what the first-order belief suggests is the optimal action under the perfect recall benchmark.

Note that even though the deviation of the thresholds from the perfect recall benchmark is called "bias" here, these thresholds are still optimal given the imperfect recall. This is an immediate consequence of common interests assumption in (3.1) and the uniqueness result in Lemma 1. In this sense, even though the two Selves are not acting according to their first-order belief, they are still acting "optimally."

3.3.2 Role of Reflection and Meditation

While the imperfect recall over actions can introduce biases, changing the recall over information can eliminate this bias asymptotically, while keeping the imperfect recall over actions.

Proposition 2.2. Role of Reflection over Information. *Given the information structure, define*

$$m \equiv \frac{(1 - r_\varepsilon) \sigma_\varepsilon}{\sigma_\omega}$$

Then,

$$\lim_{m \rightarrow 0} \Pi_1^r = \bar{\Pi}_1$$

$$\lim_{m \rightarrow 0} \mathbb{E} [\mu(x_2, \zeta)] = \frac{1}{2} \forall x_1, x_2$$

Sketch of Proof. By turning the Self 1's expectation over Self 2's expectation over Self 1's action to be agnostic. As the recall over information becomes precise, Self 1 now expects Self 2 to understand his information perfectly: since $\lim_{m \rightarrow 0} \gamma(1 - \alpha) = 0$,

$$\lim_{r_\varepsilon \rightarrow 1} \mathbb{E} [\mathbb{E} [x_1 | x_2, \zeta] | x_1] = x_1 \forall x_1, x_2$$

Since the Self 2 no longer relies on her information but instead uses the noisy information ζ to make an inference regarding μ , Self 1 expects Self 2 to have an agnostic expectation over her behaviors. Thus, Self 1 does not have biases in her experimentation decision and follows her own signal. A proof with the version of $r_\varepsilon \rightarrow 1$ is given in the Appendix A3.3.1. \square

Formally, this limit optimality result arises because *risk dominance* (Harsanyi and Selton 1988) implies the optimality in common interest games. Risk dominance is an equilibrium selection criteria in 2×2 games in which each player is assumed to face the largest uncertainty about another player's action, that is, she puts probability $\frac{1}{2}$ to either actions. Due to miscoordination payoffs, risk dominance will differ from payoff dominance in general. However, when the players have the common interests, the miscoordination payoffs cancel exactly with one another.

Such maximum strategic uncertainty arises even when the limit of $r_\varepsilon \rightarrow 1$ is considered. As Self 2 has almost perfect recall over Self 1's action, the strategic risks vanish to 0 as $r_\varepsilon \rightarrow 1$ in all information away from the cut-off. However, at the cut-off, the strategic uncertainty not only remains but also is maximized as it converges to $\frac{1}{2}$. As we will see in Section x, the equilibrium outcome differs starkly from when $r_\varepsilon = 1$.

The key to interpreting this result is not the decrease in uncertainty, but the *relative weights* between the signal ζ and the prior ω_0 that Self 2 uses to infer about Self 1's information. The key term, m , denotes how the signal ζ is precise relative to the prior ω_0 . When the prior becomes more diffused (σ_ω increases), therefore, the bias can decrease even when Self 2 remains just as uncertain about Self 1's action. Similarly, even when the prior becomes more precise and thus Self 2 knows Self 1's

information better, the bias can become exacerbated. This reflects the key finding of Weinstein and Yildiz (2007) that it is not vanishing noise but details of perturbation that determines the equilibrium. In this way, this effect differs substantively from removing imperfect recall by an increase in r_a .

There are three ways to interpret this result. The primary interpretation is how reflection practices, such as journaling, can affect the students' study decisions (increase in r_ε). When disadvantaged students perform well, they no longer attribute the success to luck but remember the encouraging pieces of information they used to have and thus attribute it to their past efforts; they will keep working hard for the final because they think they can do well if they exert efforts. When students with high self-imposed expectations fail, they will no longer think of themselves as a failure by assuming they have always exerted efforts. Instead, they may recall the difficulties they were feeling for the mid-term and the resulting lack of effort, and they stay resilient because they may still be able to perform if they work hard.

The other two ways are to either observe the present condition more (decrease in σ_ε) or to rely less on the autobiographical memory (increase in σ_ω). The last case corresponds to the idea of letting go in meditation, as discussed in the Introduction. Through meditation, Self 1 thinks that the state of the world is just it is, and follows her first-order belief to make the decision; since there is no "history" to conform to, each Self makes the decision that one believes is optimal. Since $\lim_{m \rightarrow 0} \mathbb{E}[\mu(x_2, \zeta)] = \frac{1}{2}$, there is no self-imposed expectation over behaviors. When the self-dialogue is based more on observations and less on assumptions, Self 1 can adapt flexibly to the current information instead of conforming to the self-imposed expectations from the prior belief.

3.3.3 Coordination Problems between Two Selves

The above results have highlighted the importance of memory and attention on various outcomes, such as belief and behaviors, and welfare and psychological outcomes. Here, we clarify the mechanism by focusing on the role of strategic interdependence by assuming perfect recall of information ($r_\varepsilon = 1$). While this common knowledge environment leads to multiple equilibria and thus loses the uniqueness of predictions in Section 3, it illuminates the basis of the strategic interdependence.

The following proposition shows that the experimentation decisions exhibit strategic complementarities; and with slight imperfection of recall, there will be multiple equilibria in the settings with common knowledge. We have defined $\mu_z \equiv \mathbb{P}(s_1 = 1|z)$

as the probability the Self 2 assigns to Self 1 for experimenting ($a_1 = 1$) prior to observing z (the key assumption is that $\mu_0 \neq \mu_1$ is possible). Let us denote $s_2 \in \{0, 1\}$ as Self 2's strategy to respond to the outcomes. Henceforth, let us denote $\nu_z \equiv \mathbb{P}(s_2 = 1|z)$ as the probability the Self 1 assigns to Self 2 for her strategy.

Proposition 3.1 Multiple Equilibria under Imperfect Recall of Actions and Perfect Recall of Information ($r_a < 1, r_\varepsilon = 1$). *There exists $\underline{\Pi}_1$ such that, for $\bar{p}_1(x_1) \in [\underline{\Pi}_1, \bar{\Pi}_2(0, y)]$, there are two (strict) equilibria:*

(i) **Equilibrium with experimentation.**

$$a_1^* = 1 \text{ and } a_2^* = \begin{cases} 1 & \text{if } y_1 = 1 \text{ and } \bar{p}_2(x_2) \in (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(1, 0)) \\ 0 & \text{if } y_1 = 0 \text{ and } \bar{p}_2(x_2) \in (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(1, 0)) \end{cases}$$

and $a_2^* = 1$ if $\bar{p}_2(x_2) \geq \bar{\Pi}_2(1, 0)$ and $a_2^* = 0$ if $\bar{p}_2(x_2) \leq \bar{\Pi}_2(1, 1)$.

(ii) **Equilibrium without experimentation.**

$$a_1^* = 0 \text{ and } a_2^* = \begin{cases} 1 & \text{if } \bar{p}_2(x_2) \in (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(0, y)) \\ 0 & \text{if } \bar{p}_2(x_2) \in (\bar{\Pi}_2(0, y), \bar{\Pi}_2(1, 0)) \end{cases}$$

and $a_2^* = 1$ if $\bar{p}_2(x_2) \geq \bar{\Pi}_2(1, 0)$ and $a_2^* = 0$ if $\bar{p}_2(x_2) \leq \bar{\Pi}_2(1, 1)$.

Sketch of proof: By verifying the equilibrium conditions of Definition 1. Suppose Self 2 does not act responsively to any outcome y_1 as she simply assumes that Self 1 has not experimented ($a_1 = 0$). Then, Self 1 will not experiment even when she believes the return to learning about θ is high since Self 2 will not utilize the information generated. Suppose, in contrast, Self 1 acts responsively to the outcome, y_1 , by assuming that $a_1 = 1$ to draw her inference. Then it will be valuable, even when the return may be low, for Self 1 to experiment since Self 2 will use the information. In both cases, the Self 2's original assumption about Self 1's actions will be consistent in equilibrium. Appendix A2.2 contains a complete proof. \square

This result of multiple equilibria suggests that, even with slight recall imperfection⁹, the equilibrium experimentation decision may not be adaptive. In some cases,

⁹Note that the equilibrium multiplicity arises even when the recall error probability is arbitrarily small (i.e. $r \simeq 1$), following the logic of Bagwell (1995). To see why, let us investigate the Bayes' rule to consider the Self 2's belief about Self 1's action when her prior belief is $\mathbb{P}(a_1^* = 1) = 0$.

$$\mathbb{P}(a_1^* = 1|z = 1) = \frac{r\mathbb{P}(a_1^* = 1)}{r\mathbb{P}(a_1^* = 1) + (1-r)\mathbb{P}(a_1^* = 0)} = \frac{r \times 0}{r \times 0 + (1-r) \times 1} = 0$$

the decisionmaker may “underexperiment”: even if the Self 1 receives high x_1 so that she knows it is worthwhile to give a try to generate information that improves future actions, so long as the Self 2 does not harness that information, she would not experiment. In other cases, the decisionmaker may “overexperiment”: even if the Self 1 receives low x_1 so that she knows trying is not worthwhile since it is too costly, so long as the Self 2 expects Self 1 to have tried, then it would be worthwhile to conform to such expectations. This illustrates how miscoordination results in welfare loss. Moreover, being in a sub-optimal equilibrium will be frustrating since each Self can improve her welfare by jointly shifting the decisions.

Without focusing on the equilibrium outcomes without asymmetric information, we can ask what conditions on strategies lead to the efficient outcomes. The following proposition shows that, due to the common interest assumption, it is necessary and sufficient to have some symmetries in strategic uncertainties to ensure adaptation. Here, we call the strategies of Self 1 as s_1 instead of a_1 . Let us call $\{s_1, s_2\}$ to be *adaptive*, if and only if $\{s_1, s_2\} \in \arg \max_{\tilde{s}_1, \tilde{s}_2} v_{\tilde{s}_1 \tilde{s}_2}$.

Proposition 3.2. Symmetry Conditions for Adaptation under Imperfect Recall. *Under imperfect recall over action ($r_a < 1$), the set of strategies, $\{s_1, s_2\}$, is adaptive for any payoff parameters if and only if $\mu_z + \nu_z = 1$.*

Proof. Let us derive the indifference conditions for each Self, and rearrange them to show the equivalence.

Given expectations $\{\nu_0, \nu_1\}$, Self 1 chooses $s_1 = 1$ if and only if $\mathbb{E}[v|s_1 = 1] > \mathbb{E}[v|s_1 = 0]$, where

$$\begin{aligned} \mathbb{E}[v|s_1 = 1] &= r[\nu_1 v_{11} + (1 - \nu_1)v_{10}] + (1 - r)[\nu_0 v_{11} + (1 - \nu_0)v_{10}] \\ \mathbb{E}[v|s_1 = 0] &= r[\nu_0 v_{01} + (1 - \nu_0)v_{00}] + (1 - r)[\nu_1 v_{01} + (1 - \nu_1)v_{00}] \end{aligned}$$

Given expectations $\{\mu_0, \mu_1\}$, Self 2 uses the Bayes’ rule to update his belief of the Self 1’s action:

$$\mathbb{P}(s_1 = 1|z = 1) = \frac{\mu_1 r}{\mu_1 r + (1 - \mu_1)(1 - r)}, \quad \mathbb{P}(s_1 = 1|z = 0) = \frac{\mu_0(1 - r)}{\mu_0(1 - r) + (1 - \mu_0)r}$$

As Self 2 chooses $s_2 = 1$ if and only if $\mathbb{E}[v|s_2 = 1] > \mathbb{E}[v|s_2 = 0]$, by applying the

for any $r < 1$. In words, since the Self 2 already knows the Self 1’s action in pure-strategy equilibrium, there is no information to gain from the signal z even when it is arbitrarily accurate. The analyses in Section 3 had introduced uncertainty in $\mathbb{P}(a_1^* = 1)$ to smooth this discontinuity, and

Bayes' rule,

$$\begin{aligned} s_2(z=1) = 1 &\Leftrightarrow \mu_1 r (v_{11} - v_{10}) \geq (1 - \mu_1) (1 - r) (v_{00} - v_{01}) \\ s_2(z=0) = 1 &\Leftrightarrow \mu_0 (1 - r) (v_{11} - v_{10}) \geq (1 - \mu_0) r (v_{00} - v_{01}) \end{aligned} \quad (3.6)$$

By defining $\Delta_z \equiv \mu_z + \nu_z - 1$, we can combine the indifference condition to obtain

$$\begin{aligned} s_1 = 1 &\Leftrightarrow v_{11} - v_{00} \geq \Delta_1 [r (v_{11} - v_{10}) + (1 - r) (v_{00} - v_{01})] \\ &\quad + \Delta_0 [(1 - r) (v_{11} - v_{10}) + r (v_{00} - v_{01})] \end{aligned}$$

Note that, for $\bar{p}(x_1) \in \tilde{\Pi}_\tau \subset \Pi_\tau$, the choice $s_1 = 1$ is adaptive if and only if $v_{11} \geq v_{00}$ since the base strategies $\{s_1, s_2\} \in \{\{1, 1\}, \{0, 0\}\}$ are the two strict equilibria. For this condition to hold regardless of the values of payoff and recall accuracy, it is necessary and sufficient to have $\Delta_1 = \Delta_0 = 0$. That is, $\mu_z + \nu_z = 1$ for both z . \square

This result shows that, when there is an imperfect recall, the degree of strategic uncertainties necessary to ensure adaptive choice stand in contrast from the benchmark case under perfect recall. In the benchmark, it was through backward induction that involved no strategic uncertainties that ensured an adaptive decision. Under perfect recall, however, there were multiple equilibria that generated inefficiencies because there were no strategic uncertainties. As Proposition 4.2 shows, it is an appropriate degree of strategic uncertainty that ensures the adaptation. This condition includes the risk-dominance criterion that assumes the maximum strategic uncertainty of $\{\frac{1}{2}, \frac{1}{2}\}$. Such maximum strategic uncertainty can satisfy the condition of symmetry, and thus, lead to adaptive decisions. This is because the risk dominance is equivalent to payoff dominance in common interest games. Even though the set-up is modified to be sequential, this equivalence remains.

3.4 Conclusion

Meditation is an ancient contemplative process that has become increasingly popular in the modern era as a habit to improve psychological well-being. While meditation is at its essence a thought process, scholars suggest their value does not arise from the type of rational thought process well known in economics involving the comparison of alternatives as formalized by expected utility theories. Daisetsu Suzuki, a scholar on Zen Buddhism, writes *“When we start to feel anxious or depressed, instead of*

asking, what do I need to get to be happy?, the question becomes, what am I doing to disturb the inner peace that I already have?." He suggests that, unlike other economics goods, some elements of psychological health are not things one can obtain through exchanges. The model in this paper suggests, consistent with writing on meditation, that letting go of unhelpful thoughts may be an important for becoming calm and content.

References

- Aumann, Robert J., Sergiu Hart, and Motty Perry. 1997. "The Absent-Minded Driver." *Games and Economic Behavior* 20, no. 1 (July 1, 1997): 102–16.
- Bagwell, Kyle. 1995. "Commitment and Observability in Games." *Games and Economic Behavior* 8 (2): 271–80.
- Beck, Aaron T. 1967. "Depression: Clinical, Experimental, and Theoretical Aspects." University of Pennsylvania Press.
- Bernheim, B. Douglas, and Raphael Thomadsen. 2005. "Memory and Anticipation." *The Economic Journal* 115 (503): 271–304. <https://doi.org/10.1111/j.1468-0297.2005.00989.x>.
- Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics* 24, no. 2 (1953): 265–72.
- Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67 (2): 349–74. <https://doi.org/10.1111/1468-0262.00022>.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2019. "Memory, Attention, and Choice" Harvard mimeo
- Caplin, Andrew, and John Leahy. 2001. "Psychological Expected Utility Theory and Anticipatory Feelings." *The Quarterly Journal of Economics* 116 (1): 55–79. <https://doi.org/10.1162/003355301556347>.
- Carlsson, Hans, and Eric van Damme. 1993. "Global Games and Equilibrium Selection." *Econometrica* 61 (5): 989–1018. <https://doi.org/10.2307/2951491>.
- Cho, In-Koo, and David M. Kreps. 1987. "Signaling games and stable equilibria." *Quarterly Journal of Economics* 102:179–221.
- Damme, Eric van, and Sjaak Hurkens. 1997. "Games with Imperfectly Observable Commitment." *Games and Economic Behavior* 21 (1–2): 282–308. <https://EconPapers.repec.org/RePEc:2:p:282-308>.
- da Silveira, Rava Azeredo and Michael Woodford. 2019. "Noisy Memory and Over-Reaction to News" AEA Papers and Proceedings
- de Quidt, Jonathan, Johannes Haushofer. 2017. "Depression through the Lens of Economics: A Research Agenda" in *The Economics of Poverty Traps*, NBER
- Dessein, Wouter, and Tano Santos. 2006. "Adaptive Organizations." *Journal of Political Economy* 114 (5): 956–95. <https://doi.org/10.1086/508031>.
- Frankel, David, Stephen Morris, Ady Pauzner 2003. "Equilibrium selection in global games with strategic complementarities", *Journal of Economic Theory*
- Gittins, John. 1979 "Bandit Processes and Dynamic Allocation Indices" *Journal of the Royal Statistical Society, Series B*, Vol. 41, No. 2.

- Harsanyi, John and Reinhard Selten 1988. "A General Theory of Equilibrium Selection in Games" Cambridge, MA: The MIT Press
- Heidhues, Paul, Botond Kőszegi, and Philipp Strack. 2018. "Unrealistic Expectations and Misguided Learning." *Econometrica* 86 (4): 1159–1214. <https://doi.org/10.3982/ECTA14084>.
- Horner, Johannes, and Andrzej Skrzypacz, 2016 "Learning, Experimentation and Information Design." Mimeo
- Kabat-Zinn, Jon 2009. "Wherever You Go, There You Are: Mindfulness Meditation In Everyday Life" Hachette Book Group
- Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68. <https://doi.org/10.1111/j.1468-0262.2005.00564.x>.
- Khaw, Mel Win, Ziang Li, Michael Woodford 2017. "Risk Aversion as a Perceptual Bias" NBER Working Paper No. 23294
- Koszegi, Botond. 2006. "Ego Utility, Overconfidence, and Task Choice." *Journal of the European Economic Association* 4 (4): 673–707. <https://www.jstor.org/stable/40004814>.
- Kőszegi, Botond and Loewenstein, George F. and Murooka, Takeshi, 2019. "Fragile Self-Esteem"
- Kohlberg, Elon, and Jean-François Mertens. 1986. "On the Strategic Stability of Equilibria". *Econometrica*. 54 (5):
- Loewenstein, George. 2000. "Emotions in Economic Theory and Economic Behavior." *American Economic Review* 90 (2): 426–32. <https://doi.org/10.1257/aer.90.2.426>.
- Maggi, Giovanni. 1999. "The Value of Commitment with Imperfect Observability and Private Information." *The RAND Journal of Economics* 30 (4): 555–74. <https://doi.org/10.2307/2556065>.
- Manski, Charles F. 2004. "Measuring Expectations." *Econometrica* 72 (5): 1329–76. <https://doi.org/10.1111/j.1468-0262.2004.00537.x>.
- Marschak, John and R. Radner 1972. "Economic Theory of Teams." Yale University Press
- Morris, Stephen, and Hyun Song Shin 2001. "Global Games: Theory and Applications" Cowles Foundation Discussion Paper
- Mullainathan, Sendhil. 2002. "A Memory-Based Model of Bounded Rationality." *The Quarterly Journal of Economics* 117 (3)
- Piccione, Michele, and Ariel Rubinstein. 1997. "On the Interpretation of Decision Problems with Imperfect Recall." *Games and Economic Behavior* 20, no. 1 (July 1, 1997): 3–24.
- Rabin, Matthew, and Joel L. Schrag. 1999. "First Impressions Matter: A Model of

- Confirmatory Bias.” *The Quarterly Journal of Economics* 114 (1): 37–82. <https://econpapers.repec.org/82..htm>.
- Rothschild, Michael 1974 “A Two-Armed Bandit Theory of Market Pricing” *Journal of Economic Theory*
- Rubinstein, Ariel. 1989. “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge.’” *The American Economic Review* 79 (3): 385–91.
- Ryan, Tim. 2012. “A Mindful Nation: How a Simple Practice Can Help Us Reduce Stress, Improve Performance, and Recapture the American Spirit”, Hay House
- Sobel, Joel. 2000. “Economists’ Models of Learning.” *Journal of Economic Theory* 94 (2): 241–61. <https://doi.org/10.1006/jeth.2000.2701>.
- Suzuki, Daisetsu. 1933 “Essays in Zen Buddhism”
- Tainya C. Clarke et al., 2018. “Use of Yoga, Meditation, and Chiropractors Among U.S. Adults Aged 18 and Over” NCHS Data Brief No. 325
- Weinstein, Jonathan, and Muhamet Yildiz. 2007. “A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements.” *Econometrica* 75 (2): 365–400.
- Wilson, Andrea. 2014. “Bounded Memory and Biases in Information Processing.” *Econometrica* 82 (6): 2257–94.
- Yanagida, Seizan. 1974. “History and Words of Zen,” Okawa et al. eds. *Great Books of the World*, Chuo Koron-sha
- Zimmermann, Florian. 2019. “The Dynamics of Motivated Beliefs.” crctr224_2019_073. CRC TR 224 Discussion Paper Series. University of Bonn and University of Mannheim, Germany.

Appendix Proofs

This Appendix shows the proof of the Propositions and Lemmas in the main text. It also provides additional characterizations of the equilibrium. A1 examines the distributional properties; A2 analyzes the model with partially imperfect recall; and A3 finally analyzes the model with fully imperfect recall.

A1. Preliminaries

This sub-Section characterizes the conditional distributions and derives the monotonicity of interim beliefs that facilitate the main analyses.

Lemma A1.1. Conditional Distributions. *Given the signals of the model, the following conditional distributions hold:*

$$\omega|x_1 \sim \mathcal{N}(\alpha x_1 + (1 - \alpha)\bar{\omega}_0, \alpha\sigma_\varepsilon^2), \text{ where } \alpha \equiv \frac{\sigma_\omega^2}{\sigma_\varepsilon^2 + \sigma_\omega^2} \quad (3.7)$$

$$\omega|x_2 \sim \mathcal{N}(\beta x_2 + (1 - \beta)\bar{\omega}_0, \beta\sigma_\varepsilon^2), \text{ where } \beta \equiv \frac{2\sigma_\omega^2}{\sigma_\varepsilon^2 + 2\sigma_\omega^2} \quad (3.8)$$

$$x_2|x_1 \sim \mathcal{N}\left(\frac{(1 + \alpha)x_1 + (1 - \alpha)\bar{\omega}_0}{2}, \frac{\sigma_\varepsilon^2}{2}\right) \quad (3.9)$$

$$x_1|x_2, \zeta \sim \mathcal{N}(\gamma x_2 + (1 - \gamma)\zeta, \gamma\sigma_\varepsilon^2), \text{ where } \gamma \equiv \frac{(1 - r_\varepsilon)^2}{(1 - r_\varepsilon)^2 + 1} \quad (3.10)$$

$$\zeta|x_1, x_2 \sim \mathcal{N}(x_1, (1 - r_\varepsilon)^2 \sigma_\varepsilon^2) \quad (3.11)$$

Proof. By the Bayes' rule. To understand (3.10), note that $x_1|x_2 \sim \mathcal{N}(x_2, \sigma_\varepsilon^2)$ because $\mathbb{E}x_1 = \mathbb{E}\xi$ and $\mathbb{E}x_1 + \mathbb{E}\xi = 2x_2$ by definition. Then, combine this belief with the signal ζ . \square

Lemma A.1.2. Monotonicity of Interim Belief. *Let $\bar{\bar{p}}_t(\bar{\omega}_t)$ denote the interim belief over θ conditional on, $\bar{\omega}_t$, the mean of the interim belief over ω . Let $\Psi(\bar{p}_2|\bar{p}_1)$ denote the distribution of \bar{p}_2 conditional on \bar{p}_1 .*

(i) $\bar{\bar{p}}_t(\bar{\omega}_t)$ is strictly increasing, and $\lim_{\bar{\omega}_t \rightarrow \infty} \bar{\bar{p}}_t(\bar{\omega}_t) = 1$ and $\lim_{\bar{\omega}_t \rightarrow -\infty} \bar{\bar{p}}_t(\bar{\omega}_t) = 0$.

(ii) whenever $\bar{p}'_1 > \bar{p}''_1$, $\Psi(\bar{p}_2|\bar{p}'_1) < \Psi(\bar{p}_2|\bar{p}''_1)$.

Proof. (i) For all $t = 0, 1, 2$, by the First-Order Stochastic Dominance (FOSD) of normal distributions, if $\bar{\omega}'_t > \bar{\omega}''_t$

$$\Phi(\omega|\bar{\omega}'_t) < \Phi(\omega|\bar{\omega}''_t).$$

Note that $p = C_n(\omega)$ is an increasing function. Let us write $\bar{\Phi}(p|\bar{\omega}_t)$ as the conditional distribution. Then, if $\bar{\omega}'_t > \bar{\omega}''_t$,

$$\bar{\Phi}(p|\bar{\omega}'_t) < \bar{\Phi}(p|\bar{\omega}''_t).$$

Thus, its average, $\bar{p}_t(x_t) \equiv \int_0^1 p d\bar{\Phi}(p|\bar{\omega}_t)$, also increases: $\bar{p}_t(\bar{\omega}'_t) > \bar{p}_t(\bar{\omega}''_t)$.

Note that

$$\lim_{\bar{\omega}_t \rightarrow \infty} \bar{\Phi}(p|\bar{\omega}_t) = \begin{cases} 1 & \text{if } p = 1 \\ 0 & \text{if } p < 1 \end{cases}, \quad \lim_{\bar{\omega}_t \rightarrow -\infty} \bar{\Phi}(p|\bar{\omega}_t) = \begin{cases} 1 & \text{if } p = 0 \\ 0 & \text{if } p > 0 \end{cases}.$$

Therefore, $\lim_{\bar{\omega}_t \rightarrow \infty} \bar{p}_t(\bar{\omega}_t) = 1$ and $\lim_{\bar{\omega}_t \rightarrow -\infty} \bar{p}_t(\bar{\omega}_t) = 0$.

(ii) by Lemma A1.1. (3.9) and the FOSD of normal distributions, if $\bar{\omega}'_1 > \bar{\omega}''_1$

$$\Phi(\bar{\omega}_2|\bar{\omega}'_1) < \Phi(\bar{\omega}_2|\bar{\omega}''_1). \quad (3.12)$$

Let $\varphi_t(\bar{p}_t)$ denote the inverse of $\bar{p}_t(\bar{\omega}_t)$, which exists by strict monotonicity in (i). Thus, (3.12) can be written as, $\bar{p}'_1 < \bar{p}''_1$ with $\bar{p}'_1 \equiv \bar{p}_1(\bar{\omega}'_1)$ and $\bar{p}''_1 \equiv \bar{p}_1(\bar{\omega}''_1)$,

$$\Phi(\varphi_2(\bar{p}_2) | \varphi_1(\bar{p}'_1)) < \Phi(\varphi_2(\bar{p}_2) | \varphi_1(\bar{p}''_1)). \quad (3.13)$$

Therefore, $\Psi(\bar{p}_2|\bar{p}'_1) < \Psi(\bar{p}_2|\bar{p}''_1)$. □

Note that by Lemma A1.1. and Lemma A1.2.(i), $\bar{p}_t(x_t)$ is increasing for both $t = 1, 2$.

A2. Partially Imperfect Recalls

First, we characterize the unique equilibrium under perfect recall over action and imperfect recall over information in Proposition 1; then, we characterize the multiple equilibria under the imperfect recall over action and perfect recall over information. Henceforth, based on the monotonicity derived in Lemma A1.2., we will abbreviate the notations of beliefs. In particular, $\pi_1 = p\pi_H + (1-p)\pi_L$ and q_{ay} depend on the interim belief \bar{p}_t ; in turn, \bar{p}_t depends on the prior $\bar{\omega}_0$ and signal x_t .

A2.1 Equilibrium under Imperfect Recall over Information ($r_a = 1, r_\varepsilon \leq 1$).

Proposition 1 states that, under perfect recall over action ($r_a = 1$), the equilibrium will take a threshold form with respect to the interim beliefs, and provides the ordering

for the thresholds.

Proof of Proposition 1. we prove the Proposition 1 in 2 parts by the backward induction: the Part 1 characterizes Self 2's strategy, and the Part 2 characterizes Self 1's strategy given the Self 2's strategy.

Part 1. Self 2's threshold. given any Self 1's strategy, Self 2 chooses

$$a_2^* = \begin{cases} 1 & \text{if } \bar{p}_2(x_2) \geq \bar{\Pi}_2(z, y_1) \\ 0 & \text{if o.w.} \end{cases},$$

where

$$0 < \bar{\Pi}_2(1, 1) < \bar{\Pi}_2(0, y) < \bar{\Pi}_2(1, 0) < 1. \quad (3.14)$$

The proof consists of three steps: the first step shows the threshold-form strategy holds; then, the second step show that the ordering among $\Pi(a, y)$ is the reverse of the ordering among posterior q_{ay} due to the indifference condition; and then, we derive the ordering in terms of posterior belief; finally, we combine these three steps.

Step 1 Monotonicity of action in posterior belief:

Sub-Lemma A2.1.1(i) given any posterior belief q ,

$$a_2^* = 1 \Leftrightarrow q \geq \bar{q}, \text{ where } \bar{q} \equiv \frac{c - (\pi_L - \pi_0)}{\pi_H - \pi_L} \in (0, 1) \quad (3.15)$$

Proof. By the comparison of expected welfare given each action. Let $v_{\tilde{a}}(p)$ denote the per-period welfare conditional on action $a_t = \tilde{a}$. In period 2, it is optimal to choose $a_2^* = 1$ whenever $v_1(q) \geq v_0(q)$. Reorganizing this condition, we obtain $q \geq \bar{q}$ in (.). Note that $\bar{q} \in (0, 1)$ due to the assumption that $\pi_H > \pi_L$ and $c > \pi_L - \pi_0$. ■

Step 2 Ordering of prior belief at the indifference condition:

Sub-Lemma A2.1.1(ii) given any histories $\{a, y\}$ and $\{\tilde{a}, \tilde{y}\}$,

$$q_{ay} \geq q_{\tilde{a}\tilde{y}} \Rightarrow \Pi(a, y) \leq \Pi(\tilde{a}, \tilde{y}). \quad (3.16)$$

Proof. By the definition of $\Pi(a, y)$ as the indifference prior level. Let us write $q_{ay}(p)$ as the posterior belief resulting from prior p and histories a, y . Note that by the Bayes' rule, the posterior $q_{ay}(p)$ must be strictly increasing in prior p for any histories a, y .

Let us consider some histories $\{a, y\}$ and $\{\tilde{a}, \tilde{y}\}$ such that $q_{ay}(p) \geq q_{\tilde{a}\tilde{y}}(p)$ for any fixed p . Suppose we choose p, \tilde{p} such that $q_{ay}(p) = q_{\tilde{a}\tilde{y}}(\tilde{p}) = \bar{q}$ at the indifference level as in Sub-Lemma A2.1.(i). Then, since $q_{ay}(p)$ is strictly increasing in p , $p \leq \tilde{p}$ must

hold. Since by definition, $p = \Pi(a, y)$ and $\tilde{p} = \Pi(\tilde{a}, \tilde{y})$, $q_{ay} \geq q_{\tilde{a}\tilde{y}} \Rightarrow \Pi(a, y) \leq \Pi(\tilde{a}, \tilde{y})$ for any $q_{ay}, q_{\tilde{a}\tilde{y}} \in (0, 1)$. \blacksquare

Step 3 Ordering of posterior beliefs: using assumption on π , we can order the posterior beliefs:

Sub-Lemma A2.1.1(iii) for any p ,

$$0 \underbrace{< q_{10}}_{(i)} < \underbrace{p}_{(ii)} < \underbrace{q_{11}}_{(iii)} < \underbrace{1}_{(iv)}. \quad (3.17)$$

and (v) $p = q_{00} = q_{01}$

Proof. By the Bayes' rule. The posterior beliefs conditional on $a_1 = 1$ are

$$q_{10}(p) = \frac{1}{1 + \frac{1-p}{p} \frac{1-\pi_L}{1-\pi_H}}; \quad q_{11}(p) = \frac{1}{1 + \frac{1-p}{p} \frac{\pi_L}{\pi_H}} \quad (3.18)$$

Since $\pi_L < \pi_H$, $q_{10} < p < q_{11}$ as in (ii) and (iv). Moreover, by $1 - \pi_H > 0$, (i) $q_{10} > 0$; by $\pi_L > 0$, (iv) $q_{11} < 1$. Finally, (v) holds because y_1 is uninformative about θ conditional on $a_1 = 0$. \blacksquare

By applying Sub-Lemma A2.1.(iii) to the ordering (3.17) in Sub-Lemma A2.1.(ii), we can derive the ordering (3.14). The results $0 < \Pi(1, 1)$ and $\Pi(1, 0) < 1$ hold by the property of degenerate priors. The following argument shows this for $0 < \Pi(1, 1)$, and an analogous argument applies to $\Pi(1, 0) < 1$. First, note that for any histories a and y , $q_{ay}(0) = 0$. Second, since the threshold $\bar{q} > 0$ by Assumption A2, $\bar{q} = q_{ay}(p) > q_{ay}(0) = 0$. Third, since $q_{ay}(p)$ is strictly increasing in p , $0 < \Pi(a, y)$. Thus, $0 < \Pi(1, 1)$. \square

Henceforth, let us denote $\hat{\Pi}_2^1 \equiv (\bar{\Pi}_2(1, 1), \bar{\Pi}_2(0, y))$ and $\hat{\Pi}_2^0 \equiv (\bar{\Pi}_2(0, y), \bar{\Pi}_2(1, 0))$ for notational ease.

Part 2. Self 1's threshold. given any Self 2's strategy, Self 1 chooses

$$a_1^* = \begin{cases} 1 & \text{if } \bar{p}_1(x_1) \geq \bar{\Pi}_1 \\ 0 & \text{if o.w.} \end{cases},$$

where $\bar{\Pi}_1 < \bar{\Pi}_2(0, y)$.

The proof consists of four steps: we begin by setting up the indifference condition, then we show its monotonicity. We apply the Intermediate Value Theorem to show the existence and uniqueness of the threshold. Finally, we also note that because of the informational value, the threshold of effort is lower than those in the second period without additional information.

Step 1 : Self 1's indifference condition:

Sub-Lemma A2.1.2(i) *given Self 2's strategy in Proposition 1, Self 1's indifference condition is given by $d\bar{w}_a(\bar{p}_1) = 0$, where $d\bar{w}_a(\bar{p}_1) = \bar{w}_1(\bar{p}_1) - \bar{w}_0(\bar{p}_1)$ satisfies*

$$\begin{aligned} d\bar{w}_a(\bar{p}_1) &= v_1(\bar{p}_1) - v_0 \\ &+ \int_{\hat{\Pi}_2^1} \pi_1 [v_1(q_{11}) - v_0] \\ &+ \int_{\hat{\Pi}_2^0} (1 - \pi_1) [v_0 - v_1(\bar{p}_2)], \end{aligned} \tag{3.19}$$

where the integral is over $\Psi(\bar{p}_2|\bar{p}_1)$ but abbreviated for notational ease.

Proof. Given Self 2's strategy, the welfare from choosing $a_1 = 0$ is

$$\begin{aligned} \bar{w}_0(\bar{p}_1) &= v_0 \\ &+ \int_0^{\bar{\Pi}_2(0,y)} v_0 + \int_{\bar{\Pi}_2(0,y)}^1 v_1(\bar{p}_2), \end{aligned} \tag{3.20}$$

since the outcome y_1 will not affect a_2 ; the welfare from choosing $a_1 = 1$ is

$$\begin{aligned} \bar{w}_1(\bar{p}_1) &= v_1(\bar{p}_1) \\ &+ \int_0^{\bar{\Pi}_2(1,1)} v_0 + \int_{\bar{\Pi}_2(1,1)}^{\bar{\Pi}_2(1,0)} \{\pi_1 v_1(q_{11}) + (1 - \pi_1) v_0\} + \int_{\bar{\Pi}_2(1,0)}^1 v_1(\bar{p}_2) \end{aligned} \tag{3.21}$$

since the outcome y_1 affects a_2 as it provides information about the productivity type. By taking the difference, we obtain the indifference condition (3.19). \blacksquare

Step 2 : Monotonicity of Self 1's indifference condition:

Sub-Lemma A2.1.2(ii) Self 1's indifference condition, $d\bar{w}_a(\bar{p}_1)$, is strictly increasing in \bar{p}_1 .

Proof. To show monotonicity, let us consider the payoff difference in period t , $d\bar{v}^t(p)$, conditional on the confidence level $p \in (0, 1)$:

$$d\bar{v}(p) = d\bar{v}^1(p) + d\bar{v}^2(p),$$

where

$$\begin{aligned} d\bar{v}^1(p) &= v_1(p) - v_0 \\ d\bar{v}^2(p) &= \pi_1 v_{a^*}(q_{11}) + (1 - \pi_1) v_{a^*}(q_{10}) - v_{a^*}(p) \end{aligned}$$

where a^* is the optimal action conditional on the posterior beliefs. Note that by the Law of Iterated Expectations, $d\bar{w}_a(\bar{p}_1) = \int_0^1 d\bar{v}(p) d\Psi(p|\bar{p}_1)$.

Thus, the derivatives of $d\bar{v}^t(p)$ with respect to p satisfies, denoted as $d\bar{v}^{t'}(p)$,

$$d\bar{v}^{1'}(p) = \pi_H - \pi_L$$

$$d\bar{v}^{2'}(p) = \begin{cases} 0 & \text{for } p \leq \bar{\Pi}_2(1, 1) \\ (\pi_H + \pi_L - c - \pi_0)(\pi_H - \pi_L) & \text{for } p \in \hat{\Pi}_2^1 \\ -[1 - (\pi_H + \pi_L - c - \pi_0)](\pi_H - \pi_L) & \text{for } p \in \hat{\Pi}_2^0 \\ 0 & \text{for } p > \bar{\Pi}_2(1, 0) \end{cases}$$

Note that the $d\bar{v}^2(p)$ is increasing for $\hat{\Pi}_2^1$ because $\pi_H - c > 0, \pi_L - \pi_0 > 0$ but decreasing for $\hat{\Pi}_2^0$ because $1 - \pi_H > 0, c - \pi_L > 0, \pi_0 > 0$. This inverse V-shape arises because the value of experimentation is the highest when there is large uncertainty, that is, when the value of p takes an intermediate value. By combining the expressions of $d\bar{v}^{1'}(\cdot)$ and $d\bar{v}^{2'}(\cdot)$ above,

$$d\bar{v}'(p) = \begin{cases} \pi_H - \pi_L & \text{for } p \leq \bar{\Pi}_2(1, 1) \\ (1 + \pi_H + \pi_L - c - \pi_0)(\pi_H - \pi_L) & \text{for } p \in \hat{\Pi}_2^1 \\ (\pi_H + \pi_L - c - \pi_0)(\pi_H - \pi_L) & \text{for } p \in \hat{\Pi}_2^0 \\ \pi_H - \pi_L & \text{for } p > \bar{\Pi}_2(1, 0). \end{cases}$$

Thus, $d\bar{v}(p)$ is strictly increasing everywhere in $p \in (0, 1)$. That is, the increase in $t = 1$ dominates the decrease in $t = 2$ in $p \in \hat{\Pi}_2^0$. By Lemma A1.2.(ii), if $\bar{p}'_1 > \bar{p}''_1, \Psi(p_2|\bar{p}'_1) < \Psi(p_2|\bar{p}''_1)$, and thus, $d\bar{w}_a(\bar{p}'_1) > d\bar{w}_a(\bar{p}''_1)$. That is, $d\bar{w}_a(\bar{p}_1)$ is strictly increasing in \bar{p}_1 by the FOSD. \blacksquare

Step 3: Existence and uniqueness of Self 1's threshold:

First, let us consider the extreme values, where there is no uncertainty over \bar{p}_2 : when $p = 0$ so that $\theta = \theta_L, d\bar{w}_a(0) < 0$ because $\pi_0 > \pi_L - c$, and when $p = 1$ so that $\theta = \theta_H, d\bar{w}_a(1) > 0$ because $\pi_0 < \pi_H - c$.

Since $d\bar{w}_a(\bar{p}_1)$ is continuous and strictly increasing by Sub-Lemma A2.1.2(ii), and

$\lim_{\bar{p}_1 \rightarrow 0} d\bar{w}_a(\bar{p}_1) < 0$ and $\lim_{\bar{p}_1 \rightarrow 1} d\bar{w}_a(\bar{p}_1) > 0$, there exists a unique threshold $\bar{\Pi}_1$ such that $d\bar{w}_a(x_1) = 0$ by the Intermediate Value Theorem.

Step 4: Bounds on Self 1's threshold:

Note that $d\bar{w}_a(\bar{\Pi}_2(0, y)) > 0$ since $v_1(\bar{\Pi}_2(0, y)) = v_0$ by definition, and there is an informational gain in the second period. Thus, $\bar{\Pi}_1 < \bar{\Pi}_2(0, y)$. □

A2.2 Equilibria under Imperfect Recall over Action ($r_a < 1, r_\varepsilon = 1$).

Proposition 3 suggests that, for $\bar{p}_1(x_1) \in [\underline{\Pi}_1, \bar{\Pi}_2(0, y)]$ for some $\underline{\Pi}_1$, there are two strict equilibria. Let $a'_2 \in \{0, 1\}$ denote the Self 2's "actions" conditional on outcome y_1 and belief $\bar{p}_2(x_2)$. Let $a'_2 = 1$ denote the conditional action in (i) equilibrium with experimentation, and $a'_2 = 0$ denote the conditional action in (ii) equilibrium without experimentation so that $\{a_1, a'_2\} = \{1, 1\}, \{0, 0\}$ are the equilibria that correspond with (i) and (ii) in Proposition 3 respectively.

Proof. The proof compares the welfare under the two equilibria against the welfare under other sets of actions, $\{a_1, a'_2\} = \{1, 0\}, \{0, 1\}$. Henceforth, let $\hat{w}_{a_1 a'_2}(\bar{p}_1)$ denote the welfare attained with actions $\{a_1, a'_2\}$ given belief \bar{p}_1 .

Step 1 welfare under sets of actions: the welfares under each set of actions are

$$\begin{aligned} \hat{w}_{11}(\bar{p}_1) &= v_1(\bar{p}_1) \\ &\quad + \int_0^{\bar{\Pi}_2(1,1)} v_0 + \int_{\bar{\Pi}_2(1,1)}^{\bar{\Pi}_2(1,0)} \{\pi_1 v_1(q_{11}) + (1 - \pi_1) v_0\} + \int_{\bar{\Pi}_2(1,0)}^1 v_1(\bar{p}_2) \\ \hat{w}_{10}(\bar{p}_1) &= v_1(\bar{p}_1) \\ &\quad + \int_0^{\bar{\Pi}_2(0,y)} v_0 + \int_{\bar{\Pi}_2(0,y)}^1 v_1(\bar{p}_2) \\ \hat{w}_{01}(\bar{p}_1) &= v_0 \\ &\quad + \int_0^{\bar{\Pi}_2(1,1)} v_0 + \int_{\bar{\Pi}_2(1,1)}^{\bar{\Pi}_2(1,0)} \{\pi_0 v_1(\bar{p}_2) + (1 - \pi_0) v_0\} + \int_{\bar{\Pi}_2(1,0)}^1 v_1(\bar{p}_2) \\ \hat{w}_{00}(\bar{p}_1) &= v_0 \\ &\quad + \int_0^{\bar{\Pi}_2(0,y)} v_0 + \int_{\bar{\Pi}_2(0,y)}^1 v_1(\bar{p}_2) \end{aligned}$$

Step 2 comparison of welfare: we use the expressions in Step 1 to show the strict equilibria attain higher welfare than when either Self deviates to another action:

- for $\{a_1, a'_2\} = \{1, 1\}$ to be a strict equilibrium: $\hat{w}_{11} > \hat{w}_{10}$ and $\hat{w}_{11} > \hat{w}_{01}$ must

hold at \bar{p}_1 .

$$\hat{w}_{11}(\bar{p}_1) - \hat{w}_{10}(\bar{p}_1) = \int_{\hat{\Pi}_2^1} \pi_1 [v_1(q_{11}) - v_0] + \int_{\hat{\Pi}_2^0} (1 - \pi_1) [v_0 - v_1(\bar{p}_2)]$$

$$\begin{aligned} \hat{w}_{11}(\bar{p}_1) - \hat{w}_{01}(\bar{p}_1) &= v_1(\bar{p}_1) - v_0 \\ &\quad + \int_{\hat{\Pi}_2^1} \{\pi_1 [v_1(q_{11}) - v_0] - \pi_0 [v_1(\bar{p}_2) - v_0]\} \end{aligned}$$

– $\hat{w}_{11}(\bar{p}_1) > \hat{w}_{10}(\bar{p}_1)$ by the optimality of $a'_2 = 1$ in each range $\hat{\Pi}_2^1$ and $\hat{\Pi}_2^0$.

– $\hat{w}_{11}(\bar{p}_1) > \hat{w}_{01}(\bar{p}_1)$ requires that $\bar{p}_1 \geq \underline{\Pi}_1$. To see this, note that the first-period difference, $v_1(\bar{p}_1) - v_0$, is negative. However, there is a gain from experimentation in the second period since $v_1(q_{11}) - v_0$ and $-[v_1(\bar{p}_2) - v_0]$ are positive. When $a'_2 = 1$, this gain from experimentation is also maximized.

- for $\{a_1, a'_2\} = \{0, 0\}$ to be a strict equilibrium: $\hat{w}_{00} > \hat{w}_{10}$ and $\hat{w}_{00} > \hat{w}_{01}$ must hold at \bar{p}_1 .

$$\hat{w}_{00}(\bar{p}_1) - \hat{w}_{10}(\bar{p}_1) = v_0 - v_1(\bar{p}_1)$$

$$\hat{w}_{00}(\bar{p}_1) - \hat{w}_{01}(\bar{p}_1) = \int_{\hat{\Pi}_2^1} \pi_0 [v_0 - v_1(\bar{p}_2)] + \int_{\hat{\Pi}_2^0} (1 - \pi_0) [v_1(\bar{p}_2) - v_0]$$

– $\hat{w}_{00}(\bar{p}_1) > \hat{w}_{10}(\bar{p}_1)$ by $\bar{p}_1(x_1) < \bar{\Pi}_2(0, y)$.

– $\hat{w}_{00}(\bar{p}_1) > \hat{w}_{01}(\bar{p}_1)$ because for $\bar{p}_1(x_1) < \bar{\Pi}_2(0, y)$, $v_0 > v_1(\bar{p}_2)$; for $\bar{p}_1(x_1) > \bar{\Pi}_2(0, y)$, $v_0 < v_1(\bar{p}_2)$.

□

A3. Fully Imperfect Recalls

To characterize the equilibrium under imperfect recall over both action and information, we first derive the expressions of thresholds in Lemma 1. After examining their properties, we prove the Lemma 1 and other Propositions.

A3.1 Threshold Formulas

Proposition A3.1.1 Self 1's Equilibrium Thresholds. *Given Self 2's equilibrium strategies $\{\zeta_{zy}^r(x_2)\}_{z,y}$, the Self 1's indifference condition Π_1^r satisfies*

$$d\bar{w}_a(\Pi_1^r) + dw_a^r(\Pi_1^r) = 0,$$

where $d\bar{w}_a(\bar{p}_1) \equiv \bar{w}_1(\cdot) - \bar{w}_0(\cdot)$ is defined in Proposition 1 (3.19). $dw_a^r(\bar{p}_1) \equiv w_1^r(\cdot) - w_0^r(\cdot)$ is

$$\begin{aligned} dw_a^r(\bar{p}_1) = & \int_{\hat{\Pi}_2^1} \{\pi_1(1 - \lambda_{11})[v_0 - v_1(q_{11})] - \pi_0\lambda_{01}[v_1(\bar{p}_2) - v_0]\} \\ & + \int_{\hat{\Pi}_2^0} \{(1 - \pi_1)\lambda_{10}[v_1(q_{10}) - v_0] - (1 - \pi_0)(1 - \lambda_{00})[v_0 - v_1(\bar{p}_2)]\}, \end{aligned} \quad (3.22)$$

where, for $\tilde{a} \in \{0, 1\}$,

$$\lambda_{\tilde{a}y}(x_1^r, \bar{x}_2) \equiv \mathbb{E}[\mathbb{P}(a_2 = 1 | z, y, \bar{x}_2) | x_1^r, a_1 = \tilde{a}]$$

is the expected probability that $a_2 = 1$ from Self 1's perspective at x_1^r . Here, $\{x_1^r, \bar{x}_2\}$ are

$$x_1^r = \frac{\varphi_1(\Pi_1^r) - (1 - \alpha)\bar{w}_0}{\alpha}, \quad \bar{x}_2 = \frac{\varphi_2(\bar{p}_2) - (1 - \beta)\bar{w}_0}{\beta} \quad (3.23)$$

where $\alpha \equiv \frac{\sigma_\omega^2}{\sigma_\varepsilon^2 + \sigma_\omega^2}$ and $\beta \equiv \frac{2\sigma_\omega^2}{\sigma_\varepsilon^2 + 2\sigma_\omega^2}$, by the Lemma A1.1 Conditional Distributions. $\lambda_{\tilde{a}y}(\cdot)$ satisfy

$$\lambda_{\tilde{a}0}(\cdot) = r_a \Phi\left(\frac{\zeta_{\tilde{a}0}^r - x_1^r}{(1 - r_\varepsilon)\sigma_\varepsilon}\right) + (1 - r_a) \Phi\left(\frac{\zeta_{(1-\tilde{a})0}^r - x_1^r}{(1 - r_\varepsilon)\sigma_\varepsilon}\right) \quad (3.24)$$

$$\lambda_{\tilde{a}1}(\cdot) = r_a \Phi\left(\frac{x_1^r - \zeta_{\tilde{a}1}^r}{(1 - r_\varepsilon)\sigma_\varepsilon}\right) + (1 - r_a) \Phi\left(\frac{x_1^r - \zeta_{(1-\tilde{a})1}^r}{(1 - r_\varepsilon)\sigma_\varepsilon}\right). \quad (3.25)$$

Proof. Under imperfect recall over action, there is a positive chance that Self 2's choice becomes miscoordinated, deviating from the strict equilibria in Appendix A2.2. By the additive separability of expected utility, the equilibrium welfare can be written as

$$w_a(\bar{p}_1) = \bar{w}_a(\bar{p}_1) + w_a^r(\bar{p}_1),$$

where $w_a^r(x_1) < 0$ indicates the resulting welfare loss.

- When choosing $a_1 = 0$, the welfare loss occurs when Self 2 (i) chooses $a_2 = 1$ instead of $a_2 = 0$ given $y_1 = 1$ and $\bar{p}_2 \in \hat{\Pi}_2^1$, and (ii) chooses $a_2 = 0$ instead of $a_2 = 1$ given $y_1 = 0$ and $\bar{p}_2 \in \hat{\Pi}_2^0$:

$$w_0^r(\bar{p}_1) = \int_{\hat{\Pi}_2^1} \pi_0 \lambda_{01} [v_1(\bar{p}_2) - v_0] \\ + \int_{\hat{\Pi}_2^0} (1 - \pi_0) (1 - \lambda_{00}) [v_0 - v_1(\bar{p}_2)]$$

- When choosing $a_1 = 1$, the welfare loss occurs when Self 2 (i) chooses $a_2 = 0$ instead of $a_2 = 1$ given $y_1 = 1$ and $\bar{p}_2 \in \hat{\Pi}_2^1$, and (ii) chooses $a_2 = 1$ instead of $a_2 = 0$ given $y_1 = 0$ and $\bar{p}_2 \in \hat{\Pi}_2^0$:

$$w_1^r(\bar{p}_1) = \int_{\hat{\Pi}_2^1} \pi_1 (1 - \lambda_{11}) [v_0 - v_1(q_{11})] \\ + \int_{\hat{\Pi}_2^0} (1 - \pi_1) \lambda_{10} [v_1(q_{10}) - v_0]$$

By taking the difference, we obtain the condition (3.22). Using the distribution of ζ conditional on x_1 and x_2 derived in Lemma A1.1. Conditional Distributions and using $\varphi_t(\bar{p}_t)$ as the inverse of $\bar{p}_t(\bar{\omega}_t)$ as in Lemma A1.2. Monotonicity of Interim Belief, we obtain (3.23), (3.24) and (3.25). \square

Proposition A3.1.2. Self 2's Equilibrium Thresholds. *Given Self 1's equilibrium strategy, x_1^r , and information, $\{x_2, y_1, z, \zeta\}$, Self 2's thresholds as defined in Lemma 1 is*

$$\zeta_{zy}^r(x_2) \equiv \frac{x_1^r - \gamma x_2 + \sqrt{\gamma} \sigma_\varepsilon \Phi^{-1}(M_{zy}(\bar{p}_2(x_2)))}{1 - \gamma}, \quad (3.26)$$

where

$$M_{zy}(\bar{p}_2) \equiv \frac{1}{1 - P_y \left(\frac{r_a}{1 - r_a} \right)^{2z-1} \frac{q_{1y} - \bar{q}}{\bar{p}_2 - \bar{q}}}, \quad (3.27)$$

and

$$P_1 \equiv \frac{\pi_1}{\pi_0}, P_0 \equiv \frac{1 - \pi_1}{1 - \pi_0} \quad (3.28)$$

and $\pi_1 \equiv \bar{p}_2 \pi_H + (1 - \bar{p}_2) \pi_L$.

Proof. At the threshold, the indifference condition (3.15) must be satisfied: $q(x_2, \zeta, z, y_1) = \bar{q}$, where \bar{q} is defined in Proposition 1. Defining Self 2's belief over Self 1's action it-

eratively as

$$\begin{aligned}\tilde{\mu}(x_2, \zeta) &\equiv \mathbb{P}(a_1 = 1 | x_2, \zeta) \\ \mu_{zy}(\tilde{\mu}) &\equiv \mathbb{P}(a_1 = 1 | \tilde{\mu}, z, y_1),\end{aligned}$$

the posterior belief satisfies

$$q(x_2, \mu, z, y_1) = \mu_{zy}(\tilde{\mu}) q_{1y}(\bar{p}_2) + [1 - \mu_{zy}(\tilde{\mu})] \bar{p}_2. \quad (3.29)$$

The expression of $\tilde{\mu}(x_2, \zeta)$ given distributions will be analyzed next. The expression of $\mu_{zy}(\tilde{\mu})$ satisfies, by the Bayes' rule,

$$\mu_{zy}(\tilde{\mu}) = \frac{1}{1 + \frac{1}{P_y} \left(\frac{r_a}{1-r_a} \right)^{1-2z} \frac{1-\tilde{\mu}}{\tilde{\mu}}}. \quad (3.30)$$

Substituting (3.29) and (3.30) into the indifference conditions (3.15), we can rearrange the formula:

$$\tilde{\mu}(x_2, \zeta_{zy}^r) = M_{zy}(\bar{p}_2(x_2)). \quad (3.31)$$

Note that since focused on interior of $\hat{\Pi}_2^y$, $\bar{p}_2 - \bar{q} = 0$ or $q_{1y} - \bar{q} = 0$ is ruled out in the formula (3.27).

By the Self 1's strategy given in Lemma 1, and by Lemma A1.1. Conditional Distributions, $\tilde{\mu}(x_2, \zeta)$ is given by:

$$\begin{aligned}\tilde{\mu}(x_2, \zeta) &= \mathbb{P}(x_1 \geq x_1^r | x_2, \zeta) \\ &= 1 - \Phi \left(\frac{x_1^r - [\gamma x_2 + (1 - \gamma) \zeta]}{\sqrt{\gamma} \sigma_\varepsilon} \right)\end{aligned}$$

Substituting this into (3.31) and using the symmetry of normal distribution's cumulative distribution around 0, we obtain (3.26). \square

A3.2 Properties of Thresholds

Given the expressions of the thresholds, we first characterize their properties that will be key inputs into the proof of the Lemma 1 and other Propositions in the main text.

Lemma A3.2.1. Monotonicity with Another Self's Threshold.

- (i) Self 1's Threshold: *when Self 2's thresholds, $\{\zeta_{zy}^r(x_2)\}$, are raised uniformly, Self 1's threshold, x_1^r , also increases.*

(ii) Self 2's Threshold: when Self 1's thresholds, x_1^r , is raised, every Self 2's thresholds, $\{\zeta_{zy}^r(x_2)\}$, increases uniformly.

Proof. By taking the derivatives.

(i) Self 1's Threshold: let us consider a uniform increase of Self 2's thresholds by

$$\frac{\partial [\zeta_{zy}^r(x_2) + \Delta]}{\partial \Delta}.$$

We take a uniform change because individual changes will be measure zero for Self 1. Taking the derivatives from the Self 1's thresholds as in (3.24) and (3.25),

$$\begin{aligned} \frac{\partial \lambda_{\bar{a}0}(x_1^r)}{\partial \Delta} &= \frac{1}{(1-r_\varepsilon)\sigma_\varepsilon} \left[r_a \phi \left(\frac{\zeta_{\bar{a}0}^r - x_1^r}{(1-r_\varepsilon)\sigma_\varepsilon} \right) + (1-r_a) \phi \left(\frac{\zeta_{(1-\bar{a})0}^r - x_1^r}{(1-r_\varepsilon)\sigma_\varepsilon} \right) \right] > 0 \quad (3.32) \\ \frac{\partial \lambda_{\bar{a}1}(x_1^r)}{\partial \Delta} &= \frac{-1}{(1-r_\varepsilon)\sigma_\varepsilon} \left[r_a \phi \left(\frac{x_1^r - \zeta_{\bar{a}1}^r}{(1-r_\varepsilon)\sigma_\varepsilon} \right) + (1-r_a) \phi \left(\frac{x_1^r - \zeta_{(1-\bar{a})1}^r}{(1-r_\varepsilon)\sigma_\varepsilon} \right) \right] < 0. \end{aligned} \quad (3.33)$$

Thus,

$$\begin{aligned} \frac{\partial [dw_a^r(\Pi_1^r)]}{\partial \Delta} \Big|_{\Delta=0} &= \int_{\hat{\Pi}_2^1} \left\{ -\pi_1 \frac{\partial \lambda_{11}}{\partial \Delta} [v_0 - v_1(q_{11})] - \pi_0 \frac{\partial \lambda_{01}}{\partial \Delta} [v_1(\bar{p}_2) - v_0] \right\} \quad (3.34) \\ &+ \int_{\hat{\Pi}_2^0} \left\{ (1-\pi_1) \frac{\partial \lambda_{10}}{\partial \Delta} [v_1(q_{10}) - v_0] + (1-\pi_0) \frac{\partial \lambda_{11}}{\partial \Delta} [v_0 - v_1(\bar{p}_2)] \right\} \end{aligned}$$

Since the utility differences, $(v_0 - v_1(\cdot))$, are negative in the domains of $\hat{\Pi}_2^1$ and $\hat{\Pi}_2^0$, combined with (3.32) and (3.33),

$$\frac{\partial [dw_a^r(\Pi_1^r)]}{\partial \Delta} < 0.$$

Since the threshold is defined by $d\bar{w}_a(\Pi_1^r) + dw_a^r(\Pi_1^r) = 0$, and $d\bar{w}_a(\cdot)$ is strictly increasing, this implies $\frac{\partial x_1^r}{\partial \Delta} > 0$. \square

(ii) Self 2's Threshold: given Self 2's threshold (3.26)

$$\frac{\partial \zeta_{zy}^r(x_2)}{\partial x_1^r} = \frac{1}{1-\gamma} = (1 + (1-r_\varepsilon)^2) > 0$$

for every Self 2. \square

Lemma A3.2.2. Monotonicity in Prior Belief. combining Self 2's threshold

(3.26) in Proposition A3.1.2 into Self 1's indifference (3.22) in Proposition A3.1.1,

$$\frac{\partial [dw_a^r(x_1^r)]}{\partial \bar{\omega}_0} > 0.$$

and $\lim_{\bar{\omega}_0 \rightarrow \infty} dw_a^r(x_1^r) > 0$ and $\lim_{\bar{\omega}_0 \rightarrow -\infty} dw_a^r(x_1^r) < 0$.

Proof. The proof takes three steps: first, we substitute the indifference conditions to express the condition in terms of the prior belief; then, we take its derivative and consider its limit.

Step 1 substitution. when substituting Self 2's threshold to Self 1's indifference,

$$\frac{\zeta_{zy}^r - x_1^r}{(1 - r_\varepsilon) \sigma_\varepsilon} = \frac{\gamma (x_1^r - x_2) + \sqrt{\gamma} \sigma_\varepsilon \Phi^{-1}(M_{zy}(\bar{p}_2(x_2)))}{(1 - \gamma)(1 - r_\varepsilon) \sigma_\varepsilon}$$

Note that,

$$x_1^r - x_2 = \frac{\varphi_1(\Pi_1^r)}{\alpha} - \frac{\varphi_2(\bar{p}_2)}{\beta} - \frac{\sigma_\varepsilon^2}{2\sigma_\omega^2} \bar{\omega}_0 \quad (3.35)$$

because $\frac{1}{\alpha} - \frac{1}{\beta} = \frac{\sigma_\varepsilon^2}{2\sigma_\omega^2}$. Therefore,

$$\frac{\zeta_{zy}^r - x_1^r}{(1 - r_\varepsilon) \sigma_\varepsilon} = -\frac{(1 - r_\varepsilon) \sigma_\varepsilon}{2\sigma_\omega^2} \bar{\omega}_0 + \frac{1 - r_\varepsilon}{\sigma_\varepsilon} \left(\frac{\varphi_1(\Pi_1^r)}{\alpha} - \frac{\varphi_2(\bar{p}_2)}{\beta} \right) + \frac{\Phi^{-1}(M_{zy}(x_2))}{1 - \gamma} \quad (3.36)$$

Step 2 derivative in threshold formula. For the strategic implications, we can use (3.36) and

$$\frac{\partial \lambda_{\bar{a}0}(\cdot | \bar{\omega}_0)}{\partial \bar{\omega}_0} = \frac{(1 - r_\varepsilon) \sigma_\varepsilon}{2\sigma_\omega^2} \left[r_a \phi \left(\frac{\zeta_{\bar{a}0}^r - x_1^r}{(1 - r_\varepsilon) \sigma_\varepsilon} \right) + (1 - r_a) \phi \left(\frac{\zeta_{(1-\bar{a})0}^r - x_1^r}{(1 - r_\varepsilon) \sigma_\varepsilon} \right) \right] < 0 \quad (3.37)$$

$$\frac{\partial \lambda_{\bar{a}1}(\cdot | \bar{\omega}_0)}{\partial \bar{\omega}_0} = -\frac{(1 - r_\varepsilon) \sigma_\varepsilon}{2\sigma_\omega^2} \left[r_a \phi \left(\frac{x_1^r - \zeta_{\bar{a}1}^r}{(1 - r_\varepsilon) \sigma_\varepsilon} \right) + (1 - r_a) \phi \left(\frac{x_1^r - \zeta_{(1-\bar{a})1}^r}{(1 - r_\varepsilon) \sigma_\varepsilon} \right) \right] > 0. \quad (3.38)$$

Thus,

$$\begin{aligned} \frac{\partial [dw_a^r(\Pi_1^r)]}{\partial \bar{\omega}_0} &= \int_{\hat{\Pi}_2^1} \left\{ -\pi_1 \frac{\partial \lambda_{11}}{\partial \bar{\omega}_0} [v_0 - v_1(q_{11})] - \pi_0 \frac{\partial \lambda_{01}}{\partial \bar{\omega}_0} [v_1(\bar{p}_2) - v_0] \right\} \\ &+ \int_{\hat{\Pi}_2^0} \left\{ (1 - \pi_1) \frac{\partial \lambda_{10}}{\partial \bar{\omega}_0} [v_1(q_{10}) - v_0] + (1 - \pi_0) \frac{\partial \lambda_{11}}{\partial \bar{\omega}_0} [v_0 - v_1(\bar{p}_2)] \right\} \end{aligned} \quad (3.39)$$

Since the utility differences are negative in the domain, combined with (3.37) and (3.38),

$$\frac{\partial [dw_a^r(x_1^r)]}{\partial \bar{\omega}_0} > 0.$$

so long as $\frac{(1-r_\varepsilon)\sigma_\varepsilon}{2\sigma_\omega^2} \neq 0$. Note that the distributions remain unchanged because the indifference condition keeps \bar{p}_1 and \bar{p}_2 constant.

Step 3 Limit. Since $\frac{\zeta_{zy}^r - x_1^r}{(1-r_\varepsilon)\sigma_\varepsilon}$ is decreasing linearly in $\bar{\omega}_0$, $\lim_{\bar{\omega}_0 \rightarrow \infty} \frac{\zeta_{zy}^r - x_1^r}{(1-r_\varepsilon)\sigma_\varepsilon} = -\infty$ and $\lim_{\bar{\omega}_0 \rightarrow -\infty} \frac{\zeta_{zy}^r - x_1^r}{(1-r_\varepsilon)\sigma_\varepsilon} = \infty$. By the indifference condition, $\lim_{\bar{\omega}_0 \rightarrow \infty} \lambda_{\hat{a}0}(\cdot | \bar{\omega}_0) = 1$ and $\lim_{\bar{\omega}_0 \rightarrow \infty} \lambda_{\hat{a}1}(\cdot | \omega_0) = 0$. Thus, $\lim_{\bar{\omega}_0 \rightarrow \infty} dw_a^r(\Pi_1^r) > 0$. Analogously, $\lim_{\bar{\omega}_0 \rightarrow -\infty} dw_a^r(\Pi_1^r) < 0$. \square

Lemma A3.2.3. Limit of Almost Perfect Recall over Information ($r_\varepsilon \rightarrow 1$). *For all parameters,*

$$\lim_{r_\varepsilon \rightarrow 1} dw_a^r(\bar{p}_1) = 0$$

Proof. By combining Self 1 and Self 2's indifference conditions. By (3.22), (3.24) and (3.25), this condition is equivalent to the following two conditions satisfied simultaneously:

$$\begin{aligned} \frac{\pi_0}{\pi_1} \frac{\lambda_{01}}{1 - \lambda_{11}} &= -\frac{v_1(q_{11}) - v_0}{v_1(\bar{p}_2) - v_0} \text{ for } \bar{p}_2 \in \hat{\Pi}_2^1 \\ \frac{1 - \pi_1}{1 - \pi_0} \frac{\lambda_{10}}{1 - \lambda_{00}} &= -\frac{v_1(\bar{p}_2) - v_0}{v_1(q_{10}) - v_0} \text{ for } \bar{p}_2 \in \hat{\Pi}_2^0 \end{aligned}$$

Now we show that the left-hand sides (LHSs) and the right-hand sides (RHSs) of these two formula will be equal.

LHS. Let us combine the Self 1 and Self 2's conditions: since

$$\frac{\zeta_{zy}^r - x_1^r}{(1 - r_\varepsilon)\sigma_\varepsilon} = \frac{1 - r_\varepsilon}{\sigma_\varepsilon} (x_1^r - x_2) + (1 + (1 - r_\varepsilon)^2) \Phi^{-1}(M_{zy}(\bar{p}_2(x_2)))$$

in equilibrium, for any $\{x_1^r, x_2\}$,

$$\lim_{r_\varepsilon \rightarrow 1} \frac{\zeta_{zy}^r - x_1^r}{(1 - r_\varepsilon) \sigma_\varepsilon} = \Phi^{-1}(M_{zy}(\bar{p}_2(x_2))).$$

Since $\Phi(\Phi^{-1}(M)) = M$ and by continuity of normal distributions, we can plug in the limit:

$$\begin{aligned} \lim_{r_\varepsilon \rightarrow 1} \lambda_{a0}(x_2) &= r_a M_{a0} + (1 - r_a) M_{(1-a)0} \\ \lim_{r_\varepsilon \rightarrow 1} \lambda_{a1}(x_2) &= r_a (1 - M_{a0}) + (1 - r_a) (1 - M_{(1-a)0}) \end{aligned}$$

where M_{zy} depends on $\bar{p}_2(x_2)$. Note that, by the domain restrictions $\bar{p}_2 \in \hat{\Pi}_2^y$, $\{M_{zy}\}$ exists.

Substituting the formula of M_{zy} in (3.27),

$$\begin{aligned} \lambda_{01} &= r_a (1 - r_a) \pi_1 (q_{11} - \bar{q}) D_1 \\ 1 - \lambda_{11} &= -r_a (1 - r_a) \pi_0 (\bar{p}_2 - \bar{q}) D_1 \\ \lambda_{10} &= r_a (1 - r_a) (1 - \pi_0) (\bar{p}_2 - \bar{q}) D_0 \\ 1 - \lambda_{00} &= -r_a (1 - r_a) (1 - \pi_1) (q_{10} - \bar{q}) D_0, \end{aligned}$$

where

$$\begin{aligned} D_1 &\equiv \frac{1}{(1 - r_a) \pi_0 (\bar{p}_2 - \bar{q}) - r_a \pi_1 (q_{11} - \bar{q})} + \frac{1}{r_a \pi_0 (\bar{p}_2 - \bar{q}) - (1 - r_a) \pi_1 (q_{11} - \bar{q})} \\ D_0 &\equiv \frac{1}{(1 - r_a) (1 - \pi_0) (\bar{p}_2 - \bar{q}) - r_a (1 - \pi_1) (q_{10} - \bar{q})} \\ &\quad + \frac{1}{r_a (1 - \pi_0) (\bar{p}_2 - \bar{q}) - (1 - r_a) (1 - \pi_1) (q_{10} - \bar{q})} \end{aligned}$$

Therefore, the LHS of the conditions can be written as:

$$\begin{aligned} \frac{\pi_0}{\pi_1} \frac{\lambda_{01}}{1 - \lambda_{11}} &= -\frac{q_{11} - \bar{q}}{\bar{p}_2 - \bar{q}} \text{ for } \bar{p}_2 \in \hat{\Pi}_2^1 \\ \frac{1 - \pi_1}{1 - \pi_0} \frac{\lambda_{10}}{1 - \lambda_{00}} &= -\frac{\bar{p}_2 - \bar{q}}{q_{10} - \bar{q}} \text{ for } \bar{p}_2 \in \hat{\Pi}_2^0 \end{aligned}$$

RHS. By the definition of \bar{q} in (3.15), $v_1(p) - v_0 = (\pi_H - \pi_L)(p - \bar{q})$ for any p .

Therefore, the RHS of the conditions can be written as:

$$\begin{aligned} -\frac{v_1(q_{11}) - v_0}{v_1(\bar{p}_2) - v_0} &= -\frac{q_{11} - \bar{q}}{\bar{p}_2 - \bar{q}} \text{ for } \bar{p}_2 \in \hat{\Pi}_2^1 \\ -\frac{v_1(\bar{p}_2) - v_0}{v_1(q_{10}) - v_0} &= -\frac{\bar{p}_2 - \bar{q}}{q_{10} - \bar{q}} \text{ for } \bar{p}_2 \in \hat{\Pi}_2^0 \end{aligned}$$

Thus, the LHS and RHS equal for all parameters. \square

A3.3 Proof of Lemmas and Propositions

Here, we provide the proofs of the Lemma 1 and Propositions 2.1 and 2.2 by employing the properties of indifference conditions so far derived.

Lemma 1. Equilibrium Thresholds under Imperfect Recall over Action and Information ($r_a < 1, r_\varepsilon < 1$). Lemma 1 states that, for sufficiently precise signal ζ , the unique equilibrium will take a threshold form, and the threshold depends on the recall parameters and the prior.

Proof of Lemma 1. The proof consists of two steps: first, we show that every rationalizable strategies must be bounded by highest and lowest strategies. Second, we show that these bounds must coincide so long as the signal, ζ , is sufficiently precise. This is a multi-dimensional version of the uniqueness proof in Carlsson and van Damme (1993). The argument is adopted from Morris and Shin (2001)

Sub-Lemma. A3.3(i) Bounds on Rationalizable Strategies by Thresholds. *By the Definition 1 of equilibrium, let the strategies of Self 1 and Self 2 be denoted by*

$$\begin{aligned} s_1(x_1) &= \mathbb{P}(a_1 = 1 | x_1) \\ s_2(\zeta | z, y_1, x_2) &= \mathbb{P}(a_2 = 1 | \zeta, z, y_1, x_2). \end{aligned}$$

All rationalizable strategies must satisfy, for some thresholds $\{\underline{\chi}^*, \bar{\chi}^*\}$ of Self 1 and $\{\underline{\eta}_{zy}^*(x_2), \bar{\eta}_{zy}^*(x_2)\}$ of Self 2, the following:

- Self 1's strategy:

$$s_1(x_1) = \begin{cases} 1 & \text{if } x_1 > \bar{\chi}^* \\ 0 & \text{if } x_1 < \underline{\chi}^* \end{cases} \quad (3.40)$$

- Self 2's strategy:

– $y_1 = 1$, for each z , for $\bar{p}_2(x_2) \in \hat{\Pi}_2^1$

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } \zeta > \bar{\eta}_{z1}^*(x_2) \\ 0 & \text{if } \zeta < \underline{\eta}_{z1}^*(x_2) \end{cases} \quad (3.41)$$

and for $\bar{p}_2(x_2) \notin \hat{\Pi}_2^1$

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } \bar{p}_2(x_2) \geq \bar{\Pi}_2(0, y) \\ 0 & \text{if } \bar{p}_2(x_2) \leq \bar{\Pi}_2(1, 1) \end{cases} \quad (3.42)$$

– $y_1 = 0$, for each z , for $\bar{p}_2(x_2) \in \hat{\Pi}_2^0$

$$s_2(\zeta, x_2) = \begin{cases} 0 & \text{if } \zeta > \bar{\eta}_{z0}^*(x_2) \\ 1 & \text{if } \zeta < \underline{\eta}_{z0}^*(x_2) \end{cases} \quad (3.43)$$

and for $\bar{p}_2(x_2) \notin \hat{\Pi}_2^1$

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } \bar{p}_2(x_2) \geq \bar{\Pi}_2(1, 0) \\ 0 & \text{if } \bar{p}_2(x_2) \leq \bar{\Pi}_2(0, y) \end{cases} \quad (3.44)$$

Note that $\underline{\chi}^* \leq \bar{\chi}^*$ and $\underline{\eta}_{zy}^*(x_2) \leq \bar{\eta}_{zy}^*(x_2)$ for both $y \in \{0, 1\}$.

Proof. By the iterative elimination of strictly dominated strategies with contagion from the “dominance regions” of Self 2. Henceforth, let $\{\underline{\chi}^n, \bar{\chi}^n\}$ and $\{\underline{\eta}_{zy}^n(x_2), \bar{\eta}_{zy}^n(x_2)\}$ denote the thresholds for strategies that have survived the n th stage of elimination. The proof consists of four steps: first, we consider Stage 1 elimination of Self 1 and 2 that generate infinite values of the thresholds; second, we observe that their thresholds take finite values in Stage 2; third, we observe that the sequences of thresholds are monotone, and thus, can apply the Mathematical Induction; finally, we apply the Monotone Convergence Theorem.

Step 1. elimination in Stage 1. Let us consider Self 2’s rationalizable strategies in the initial stage.

- Highest strategy: Suppose $s_1(x_1) = 0$ for all x_1 . Then,

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } \bar{p}_2(x_2) \geq \bar{\Pi}_2(0, y) \\ 0 & \text{if o.w.} \end{cases}$$

- Lowest strategy: Suppose $s_1(x_1) = 1$ for all x_1 . Then,

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } y_1 = 1 \text{ and } \bar{p}_2(x_2) \geq \bar{\Pi}_2(1, 1) \\ 0 & \text{if } y_1 = 0 \text{ and } \bar{p}_2(x_2) \geq \bar{\Pi}_2(1, 0) \end{cases}$$

Thus, the strategies other than the following are eliminated in this initial stage, for any x_2 and z :

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{when } y_1 = 1 \text{ and } \bar{p}_2(x_2) \geq \bar{\Pi}_2(0, y), \\ & \text{or when } y_1 = 0 \text{ and } \bar{p}_2(x_2) \leq \bar{\Pi}_2(0, y) \\ 0 & \text{when } y_1 = 1 \text{ and } \bar{p}_2(x_2) \leq \bar{\Pi}_2(1, 1) \\ & \text{or when } y_1 = 0 \text{ and } \bar{p}_2(x_2) \geq \bar{\Pi}_2(1, 0) \end{cases}$$

This proves (3.42) and (3.44) of the Lemma. In the notations of thresholds,

$$\underline{\chi}^1 = -\infty, \bar{\chi}^1 = \infty.$$

When $y_1 = 1$, for each z , for $\bar{p}_2(x_2) \in \hat{\Pi}_2^1$,

$$\underline{\eta}_{z1}^1(x_2) = -\infty, \bar{\eta}_{z1}^1(x_2) = \infty;$$

When $y_1 = 0$, for each z , for $\bar{p}_2(x_2) \in \hat{\Pi}_2^0$,

$$\underline{\eta}_{z0}^1(x_2) = -\infty, \bar{\eta}_{z0}^1(x_2) = \infty.$$

Note that for this Stage, Self 1's strategy does not depend on x_1 and Self 2's strategies do not depend on ζ .

Step 2. elimination in Stage 2. Let us consider Self 1's response to Self 2's elimination in Stage 1.

- Highest strategy: since Self 2 is never responsive to y_1 , there is no informational gain from experimentation. Thus, the resulting threshold is equivalent to the myopic threshold:

$$s_1(x_1) = \begin{cases} 1 & \text{if } \bar{p}_1(x_1) \geq \bar{\Pi}_2(0, y) \\ 0 & \text{if o.w.} \end{cases}$$

Thus, $\bar{\chi}^2 = \varphi_1(\bar{\Pi}_2(0, y))$.

- Lowest strategy: since Self 2 is maximally responsive to y_1 , and Self 1 shares the common interest with Self 2, there is a maximal informational gain from experimentation. By Proposition 3 proven in A2.2, the resulting threshold corresponds to $\underline{\Pi}_1$:

$$s_1(x_1) = \begin{cases} 1 & \text{if } \bar{p}_1(x_1) \geq \underline{\Pi}_1 \\ 0 & \text{if o.w.} \end{cases}$$

Thus, $\underline{\chi}^2 = \varphi_1(\underline{\Pi}_1)$.

Together, these show

$$\underline{\chi}^2 < \bar{\chi}^2. \quad (3.45)$$

Let us in turn consider Self 2's response to Self 1's elimination in Stage 2. The thresholds are determined by the Self 2's indifference condition (3.26).

- When $y_1 = 1$, for each z , for $\bar{p}_2(x_2) \in \hat{\Pi}_2^1$

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } x_2 \geq \eta_{z1}^1(x_2) \\ 0 & \text{if o.w.} \end{cases}$$

and the highest and lowest thresholds satisfy $\bar{\eta}_{z1}^2(x_2) > \underline{\eta}_{z1}^2(x_2)$ by (3.45).

- When $y_1 = 0$, for each z , for $\bar{p}_2(x_2) \in \hat{\Pi}_2^0$

$$s_2(\zeta, x_2) = \begin{cases} 1 & \text{if } x_2 \leq \eta_{z0}^1(x_2) \\ 0 & \text{if o.w.} \end{cases}$$

and the highest and lowest thresholds satisfy $\bar{\eta}_{z0}^2(x_2) > \underline{\eta}_{z0}^2(x_2)$ by (3.45).

Step 3. elimination in Stage 3 and Mathematical Induction.

Let us now consider Self 1's response to Self 2's elimination in Stage 2.

- Highest strategy: since Self 2 may be responsive to y_1 , there is some informational gain from experimentation. By applying the indifference condition (3.22), the resulting threshold is lower than the myopic threshold: $\bar{\chi}^3 < \bar{\chi}^2$.
- Lowest strategy: since Self 2 may be less responsive to y_1 than in Stage 1, the informational gain from experimentation is reduced. By applying the indifference condition (3.22), the resulting threshold is higher than the myopic threshold: $\underline{\chi}^3 > \underline{\chi}^2$.

By Lemma A3.2.1. Monotonicity with Another Self's Threshold, the following ordering holds for any n :

(i) if $\chi^{n-1} < \chi^n$, $\eta_{zy}^n(x_2) = \eta_{zy}^{n-1}(x_2) + \Delta$ for some Δ for all $\{z, y, x_2\}$

(ii) if $\eta_{zy}^n(x_2) = \eta_{zy}^{n-1}(x_2) + \Delta$ for all $\{z, y, x_2\}$, $\chi^n < \chi^{n+1}$.

By (i) and (ii) and the Mathematical Induction, $\bar{\chi}^n$ is a decreasing sequence since $\bar{\chi}^3 < \bar{\chi}^2$, and $\underline{\chi}^n$ is an increasing sequence since $\underline{\chi}^3 > \underline{\chi}^2$. Analogously, $\bar{\eta}_{zy}^n(x_2)$ is decreasing and $\underline{\eta}_{zy}^n(x_2)$ is increasing in n .

Step 4. Monotone Convergence Theorem. Note that $\bar{\chi}^n$ is bounded below by $\underline{\chi}^1$ and $\underline{\chi}^n$ is bounded above by $\bar{\chi}^1$. Thus, by the Monotone Convergence Theorem, the sequence of highest and lowest strategies must converge as $n \rightarrow \infty$: $\{\bar{\chi}^n, \underline{\chi}^n\} \rightarrow \{\bar{\chi}^*, \underline{\chi}^*\}$ and $\{\bar{\eta}_{zy}^n(x_2), \underline{\eta}_{zy}^n(x_2)\} \rightarrow \{\bar{\eta}_{zy}^*(x_2), \underline{\eta}_{zy}^*(x_2)\}$. By construction, both at the highest and lowest sets of strategies, $\{\bar{\chi}^*, \bar{\eta}_{zy}^*(x_2)\}$ and $\{\underline{\chi}^*, \underline{\eta}_{zy}^*(x_2)\}$, the indifference conditions of Self 1 (3.22) and Self 2 (3.26) are satisfied. ■

Lemma. A3.3.(ii) Uniqueness of Threshold-form Strategies given $r_\varepsilon \geq \bar{r}_\varepsilon$. *There exists $\bar{r}_\varepsilon < 1$ such that for $r_\varepsilon \geq \bar{r}_\varepsilon$, there is a unique value of \bar{p}_1 that satisfies the indifference conditions of Self 1 (3.22) and Self 2 (3.26).*

Proof. The indifference condition is satisfied when

$$d\bar{w}_a(\bar{p}_1) + dw_a^r(\bar{p}_1) = 0$$

Note that both $d\bar{w}_a(\cdot)$ and $dw_a^r(\cdot)$ are continuous and differentiable in \bar{p}_1 . Moreover, $d\bar{w}_a(\cdot)$ is strictly increasing, and

$$d\bar{w}_a(0) + dw_a^r(0) = d\bar{w}_a(0) < 0$$

$$d\bar{w}_a(1) + dw_a^r(1) = d\bar{w}_a(1) > 0$$

since there is no strategic uncertainty at $\bar{p}_1 \in \{0, 1\}$. Thus, by the Intermediate Value Theorem, there is a unique solution if the derivative

$$d\bar{w}'_a(\bar{p}_1) + dw_a^{r'}(\bar{p}_1) > 0$$

for all \bar{p}_1 .

By the Lemma A3.2.3, $\lim_{r_\varepsilon \rightarrow 1} dw_a^r(\bar{p}_1) = 0$. Thus, $\lim_{r_\varepsilon \rightarrow 1} dw_a^{r'}(\bar{p}_1) = 0$. By continuity with respect to r_ε , there exists $\bar{r}_\varepsilon < 1$ such that for $r_\varepsilon \geq \bar{r}_\varepsilon$, $dw_a^{r'}(\bar{p}_1) > -d\bar{w}'_a(\bar{p}_1)$ evaluated at all \bar{p}_1 . Therefore, for $r_\varepsilon \geq \bar{r}_\varepsilon$, the indifference condition is satisfied at the unique value of \bar{p}_1 . ■

Since all rationalizable strategies must be bounded by the highest and lowest thresholds, and a threshold form strategy is unique if $r_\varepsilon \geq \bar{r}_\varepsilon$ for some $\bar{r}_\varepsilon < 1$, there is a unique rationalizable set of strategies that constitute the unique equilibrium. \square

Proposition 2.1. Equilibrium with Bias to Conform to Prior Beliefs under Imperfect Recall over Action and Information ($r_a < 1, r_\varepsilon < 1$). Proposition 2.1. shows that there is a unique prior level, ω_0^* , such that the equilibrium threshold coincides with the threshold under perfect recall over action: $\Pi_1^r(\omega_0^*) = \bar{\Pi}_1$. moreover, the threshold is decreasing in the prior: $\frac{\partial \Pi_1^r(\bar{\omega}_0)}{\partial \bar{\omega}_0} < 0$. This effect arises because of the change in distribution of μ .

Proof of Proposition 2.1. By Proposition 1, given the model parameters of $\{\pi., c, \sigma.\}$, there exists a unique threshold $\bar{\Pi}_1$ such that $d\bar{w}_a(\bar{\Pi}_1) = 0$. Let us call ω_1^* a mean interim belief that satisfies $\omega_1^* = \varphi_1(\bar{\Pi}_1)$.

By Lemma A3.2.2. Monotonicity and the Intermediate Value Theorem, there exists a unique ω_0^* such that $dw_a^r(\omega_0^*) = 0$ with the implied $x_1^r = \frac{1}{\alpha}[\omega_1^* - (1 - \alpha)\bar{\omega}_0]$. Given ω_0^* , at the threshold x_1^r , both $d\bar{w}_a(\bar{p}_1) = 0$ and $dw_a^r(\bar{p}_1) = 0$. Thus, such unique level of prior exists by construction.

Since the threshold is defined by $d\bar{w}_a(\bar{p}_1(x_1^r)) + dw_a^r(\bar{p}_1(x_1^r)) = 0$, and $d\bar{w}_a(\bar{p}_1(x_1^r))$ is strictly increasing, the lemma also implies $\frac{\partial \Pi_1^r(\bar{\omega}_0)}{\partial \bar{\omega}_0} < 0$. \square

- the change in distribution of μ remains to be proven.

Proposition 2.2. Role of Reflection over Information. Proposition 2.2 states that when $r_\varepsilon \rightarrow 1$, the bias is eliminated because $\mathbb{E}\mu = \frac{1}{2}$.

Proof. a sufficient condition that guarantees the optimality for any parameters is $dw_a^r(\bar{p}_1) = 0$ for all \bar{p}_1 . By the Lemma A3.2.3, $\lim_{r_\varepsilon \rightarrow 1} dw_a^r(\bar{p}_1) = 0$ for any model parameters. Since $d\bar{w}_a(\bar{p}_1)$ is independent of r_ε , this ensures optimality in the limit at all values of \bar{p}_1 , including $\bar{\Pi}_1$: $d\bar{w}_a(\bar{\Pi}_1) = 0$. \square