**Title**
Heterogeneity in Learning



**Author**
Williams, Cole Randall

**Publication Date**
2018

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

# Heterogeneity in Learning

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Economics

by

Cole Randall Williams

Dissertation Committee:
Professor Jean-Paul Carvalho, Chair
Professor Stergios Skaperdas
Professor John Duffy

2018

# TABLE OF CONTENTS

# LIST OF FIGURES

# ACKNOWLEDGMENTS

I would like to thank my mentors Jean-Paul Carvalho and Stergios Skaperdas for guiding me into the world of economic research. My interest in microeconomic theory was inspired by them: Stergios engaging the fundamental questions of economic epistemology and Jean-Paul beautifully describing the “spooky action at a distance” of strategies lying off the equilibrium path. They have spent a considerable amount of time and effort developing me as a researcher. I have learned from John Duffy the virtue of developing rigorous theory with a pragmatic eye towards its empirical application. Michael McBride has provided many helpful insights: particularly useful to me were comments during my advancement exam which led me to expand my understanding of several important concepts.

This dissertation has also benefited from many conversations with my peers in economics, logic and the philosophy of science, and mathematical behavioral sciences.

# CURRICULUM VITAE

## Cole Randall Williams

### EDUCATION

**Doctor of Philosophy in Economics** **2018**
University of California, Irvine *Irvine, CA*

**Bachelor of Arts in Economics** **2013**
University of California, Riverside *Riverside, CA*

### TEACHING EXPERIENCE

**Teaching Assistant** **2013–2017**
University of California, Irvine *Irvine, CA*

### AWARDS

Kassouf Fellowship for excellence in coursework and research **2016–2017**
Economics Merit Fellowship **2018**
UCI Summer Research Fellowship **2016-2017**

### SERVICE

**Organizer** **2016**
Theory, History, and Political Economy Workshop *Irvine, CA*

# ABSTRACT OF THE DISSERTATION

Heterogeneity in Learning

By

Cole Randall Williams

Doctor of Philosophy in Economics

University of California, Irvine, 2018

Professor Jean-Paul Carvalho, Chair

This dissertation contributes to the understanding of heterogeneity in rational learning.

The first chapter proceeds from the observation that, in an environment of social learning under unobserved heterogeneity, those with opinions closest to one's own may in fact be the most informative for oneself. For example, a similar opinion of a restaurant may suggest similar tastes and similar political views may suggest a similarity in values or interpretations of evidence. In this environment, individuals will display rational forms of confirmation bias and other ostensibly anomalous patterns of behavior.

The second chapter is a research collaboration with Aydin Mohseni studying learning by agents in social networks. Each agent has a preference for both *accuracy* (choosing the "correct" action) and *conformity* (selecting the action taken by the majority of her neighbors) with the relative weight placed on these two concerns being heterogeneous between agents. Related literature finds the star network to possess optimality properties. In contrast, our analysis finds that agents in highly centralized networks, such as star networks, take longer to settle on the optimal action than agents in other standard networks.

To combat the replication crisis in science, a group of prominent scholars has proposed *redefining statistical significance* by reducing the $p$-value significance threshold from 0.05 to

0.005. The third chapter shows that, if researchers can exercise their "degrees of freedom" to obtain significance and if they are heterogeneous, then this proposal may exacerbate the problems with reproducibility. I provide an example demonstrating that even a small amount of researcher bias can produce this effect and give a general characterization of the conditions when it will occur.

# Chapter 1

# Echo Chambers: Social Learning under Unobserved Heterogeneity

> *The diversity of our opinions does not proceed from some men being more rational than others but solely from the fact that our thoughts pass through diverse channels and the same objects are not considered by all.*
>
> Descartes, *Discourse on Method.*

Belief formation occurs not merely through introspection but by observing the opinions of others. Social learning governs behavior in a number of domains, including consumption decisions, occupational choice, political preferences, and scientific beliefs. Yet our understanding of how we learn from others is still incomplete.

The problem hinges on precisely how one should respond to another person's opinion. The rational response is context-dependent. It depends on the underlying structure of the population, including the information available about its structure. When agents are *homogeneous*, differing only in their private information, individuals observing each others' opinions learn to agree [8, 36]. When agents are *heterogeneous*, but each agent's type is known, individuals may fail to agree. In this paper, we show that when agents are heterogeneous and individual types are *unobservable*, individuals may not only fail to agree but also display rational forms of confirmation bias and other anomalous patterns of behavior.

Social learning in a context of unobserved heterogeneity becomes a process of *dual learning*: by observing the opinions of others an agent learns both about his parameter of interest and the structure of the heterogeneity in the population. Encountering an agent with a divergent opinion could now mean that this agent is importantly different from himself. As such, learning is *local*: individuals place greater weight on opinions closer to their own and rationally discount highly divergent views. This unlocks a number of new results, including non-monotonicities in belief formation, belief polarization, and social identification through social learning.

There are many real-world examples of dual learning processes. The following presents four such examples. The first two—restaurant choice and political opinion—return several times throughout the paper to illustrate our results.

**Restaurant.** Restaurant choice is a canonical example of social learning.[1] Under unobserved heterogeneity, a negative review of a restaurant can either mean that the restaurant is of low quality or that the reviewer has different preferences to oneself. The more positive experiences one has of the restaurant, the more likely one is to discount a negative reviewer's opinion as proceeding from different preferences.

**Politics.** Individuals may form their political beliefs by sharing opinions, but may also retain distinct preferences over policies due to different normative values and/or interpretation of evidence. The larger the divergence in opinions, the more likely one is to attribute disagreement to underlying differences.

[1]Expositions employing this example include [13], [11], [53], [80], [32],[25], and [34].

**Technological Innovation.** Empirically consistent patterns in innovation diffusion are often explained as the result of heterogeneity among an innovation's potential adopters.[2] A prime example is Munshi (2004) who studies the diffusion of high yield varieties of rice and wheat during the Indian Green Revolution. Rice yields were particularly sensitive to variations in factors like soil characteristics and managerial inputs that are not easy to observe. Munshi finds evidence that growers came to place less weight on their neighbors' rice-growing decisions and outcomes than they do in the case of wheat.

**Scientific Theories.** Experts equally fluent in a scientific discipline often disagree.[3] One possible source of disagreement is the diversity in inferences drawn from evidence. Bayesian econometricians focus less on statistically significant p-values, and people may be convinced to different degrees of an instrumental variable's excludability or a theoretical model's assumptions. The different lenses through which we filter empirical observations, including scientific research, can lead to a diversity of opinion. Hence, experts may attribute disagreement to different dispositions to evidence.

The dual learning process that we study, arising from unobserved heterogeneity, unlocks a number of new results. The following is a non-exhaustive summary.

Our first result establishes the model's primary mechanism: learning is *local* in the sense that individuals place more weight on opinions closer to their own. By increasing the difference in opinions, this weight can be made arbitrarily close to zero. We then ask how one responds to changes in another's opinion. We find that the answer depends on which of the two countervailing forces of dual learning dominates. This leads to *non-monotonicity in disagreement*, whereby, encountering someone with a slight difference in opinion can have a

[2]See [48], [57], [49], and [88].

[3]Galileo battled with the Catholic church and fellow scientists alike over the heliocentric model of the solar system, the germ theory of disease was contested for centuries, and there was longstanding dissent over theories of continental drift. Contemporary science hosts disagreements over the fundamental roles of randomness and measurement in quantum mechanics [75] and the plausibility of group selection in evolutionary biology [33]. In economics there has been disagreement over topics like the efficacy of monetary policy in stimulating the real economy and the employment effects of raising the minimum wage.

larger influence on one's beliefs than if they were to hold a starkly different opinion.

In the long run, after exchanging opinions with enough other individuals, one's own opinion will converge. At this point, we observe *social identification* through social learning: learning an additional individual's opinion will serve almost entirely as a means for assessing the degree of similarity between the other individual and oneself.[4] Hence dual learning provides a basis for social identification through similarity of beliefs [5, 4].

We then consider extensions to the basic framework, beginning with introducing our model into the observational learning environment [17, 11] in which agents learn more coarsely by observing the actions performed by others. We show that our characterization of heterogeneity has *competing welfare effects in observational learning*: (1) agents will never converge with certainty to their optimal action in the limit of learning and (2) the process can avoid falling into an information cascade when it would have done so with certainty under homogeneity.

Another extension characterizes the behavior of media consumers.[5] Consumers choose to acquire information from sources that confirm their own beliefs. Some come to place enough trust in a media source to rely on its reports in place seeking out their own information. In existing models, the media's effect on public disagreement requires that the population itself not be aware of the disagreement. Otherwise, the agents will condition on the disagreement and the media's effect will cease. In our model, the public's awareness of disagreement can strengthen the disagreement.

Section 1.2 provides the background in terms of probability and decision theory for our more general approach to social learning which takes into account unobserved heterogeneity. Section 1.3 introduces the model of dual learning. Section 1.4 presents the main results

[4]For example, following many conversations about climate policy with various people, hearing an additional person's view may have a negligible effect on one's own opinion, but can be quite instructive about the similarity or differences in basic values.

[5][39] review the related literature on media bias.

of the paper, starting with the most basic setting and gradually increasing in complexity. Section 1.5 considers extensions of the model and Section 1.6 concludes.

## 1.2 Related Literature

A discussion of the rational response to someone's beliefs must begin by specifying the form that rational belief will take and how it will respond to evidence. For this, we look to seminal figures in the development of the subjectivist (or personalistic) view of probability, [71], [27], and [74], who show that, if an individual satisfies certain basic coherence requirements, then their beliefs can be characterized by probabilities and they will update according to conditioning in response to new evidence. At this baseline level of rationality, there is no imperative that individuals come to agreement upon discovering that they hold conflicting views—it depends exclusively on how the other's beliefs fit into their respective models of the world.

> "The criteria incorporated in the personalistic view do not guarantee agreement on all questions among all honest and freely communicating people, even in principle."
>
> Savage, (1954)

In the development of classical game theory, these minimal coherence requirements were insufficient to provide general tractability for games with *incomplete information*, that is, games in which some players are uncertain about the game being played. In such games, a player's optimal action will depend on an infinite hierarchy of beliefs: their first-order beliefs about the game, second-order beliefs about the other players' beliefs, third-order beliefs about the other players' beliefs about their beliefs, and so on ad infinitum. Harsanyi (1967) proposed the powerful simplifying assumption that it be common knowledge that players' beliefs are *mutually consistent*: any discrepancies between the various players' beliefs are driven solely by differences in private information.[6]

As shown by Aumann (1976), a strong implication of the mutual consistency assumption is that rational individuals cannot publicly disagree. More precisely, Aumann shows that with mutual consistency, if individuals' beliefs about an event are common knowledge, then they will agree.[7]

But of course, public disagreement is pervasive. Roughly 63% of Americans are absolutely certain of the existence of God, while 9% do not even believe in God, 48% believe that global climate change is due to human activity while 31% believe the causes to be natural, and 15% believe that the collapse of the World Trade Center resulted from controlled demolition while 75% do not.[8]

An account of the manifest public disagreement requires a weakening of the mutual consistency assumption. One weakening of the assumption that can sustain disagreement is to abandon the coherent belief paradigm altogether. For example, disagreement could be driven by confirmation bias [70, 35], motivated reasoning [56, 54, 14], bounded memory [86], or rule-of-thumb belief updating procedures [28, 29].

We could alternatively deviate from mutual consistency by allowing individuals to begin with heterogeneous prior beliefs.[9] In this case, classic results in Bayesian consistency [31] and the merging of opinions [20, 51] guarantees that agreement is almost surely reached over time.

[6]Our discussion highlights the fact that the mutual consistency assumption is stronger than the common prior assumption, though these are often treated as equivalent in the literature. Mutual consistency entails both a common prior and common knowledge of the information structure.

[7]See [36], [9], and [73] for extensions and [72] for a discussion of the *Agreeing to Disagree* results spawned by Aumann.

[8](Pew, Religious Landscape Study, 2014),(Pew, The Politics of Climate, 2016),(Angus Reid, Public Opinion, 2010)

[9]For a discussion of the rationale for using models with heterogeneous prior beliefs see **(author?)** [63] and for applications see [30], [84], [42], and [16].

As has been noted in the literature, robust disagreement can emerge when we allow for heterogeneity beyond mere differences in prior beliefs. For example, heterogeneous interpretations of public signals can explain the presence of asset trading [45, 52, 2]. In contrast, the No Trade Theorems of [61] and [82] predict that risk-averse traders with mutually consistent beliefs will not engage in trade.

Unobserved heterogeneous priors have been studied in the context of information aggregation. [76] find that if agents have heterogeneous prior beliefs that are unobservable but correlated, then the information is fully aggregated through successive declarations of beliefs. [78] show that when agents have unobservable heterogeneous prior beliefs, agents will come to favor observing the opinions of those with whom they have become most familiar.

[80] study the asymptotic beliefs and actions of a population comprised of heterogeneous types in the context of observational learning. The important difference between their model and ours is that agents in our model receive a signal of their own parameter of interest (e.g. utility from performing an action) and agents in their model receive signals of the state of the world which then determines their parameter of interest. This difference is the fundamental driver of our results and leads to distinct and interesting outcomes when applied to observational learning (see section 1.5.1).

## 1.3 Dual Learning

In our model, agents are sorted into heterogeneous and unobservable *types*. Agents of the same type seek to learn the same parameter of interest. In the examples, this translates to agents of the same type having the same tastes, values, or dispositions towards evidence. Each agent receives an informative signal of his parameter of interest and observes the opinions of the other agents from which he performs *dual learning*: he learns about his parameter as well as the likelihood that other agents are of the same type.

A way to visualize dual learning is to consider a variant of the classic ball and urn model. The following illustrates the process of dual learning and also foreshadows our result of 'non-monotonicity in disagreement.'

### 1.3.1 A Tale of Two Urns

Imagine that before you is an urn containing 100 green and red balls. You are asked to guess the number of green balls in the urn and will be paid in accordance with how close your guess $\widehat{G}$ is to the actual number of green balls $G$.[10] You are permitted to draw a ball from the urn 10 times (with replacement) and your draws come up with 8 green and 2 red balls. Suppose there is also another participant who makes 30 draws which you observe prior to making your guess. The left panel in Figure 1.1 illustrates how your guess would change by observing that the other participant, whose first 10 draws were identical to yours (8 green and 2 red), continued drawing only red balls, ending with a total of 8 green and 22 red.[11] On the $y$-axis is your guess $\widehat{G}$ (the posterior expected number of green balls) and on the $x$-axis is the additional red balls drawn by the other participant.

Now imagine there is also a second urn containing 100 green and red balls and you are uncertain which of the two urns the other participant is drawing from. Firstly, this uncertainty will lead you to place less weight on the other's draws than your own when forming your guess. Secondly, you will engage in *dual learning*—the weight you place on the other's draws will be updated based on the similarity of their draws to your own. In our example, as the other participant continues drawing red balls, it becomes increasingly likely that they are drawing from a different urn than you. Hence, you begin to place less weight on their

[10] For example, your payment could be given by the quadratic loss function $1 - (\frac{G}{100} - \frac{\widehat{G}}{100})^2$.

[11] The simulation assumes a uniform prior over the number of green balls in the urns, $\pi(G) = \frac{1}{101}$ for $G = 0, 1, 2, ..., 100$.

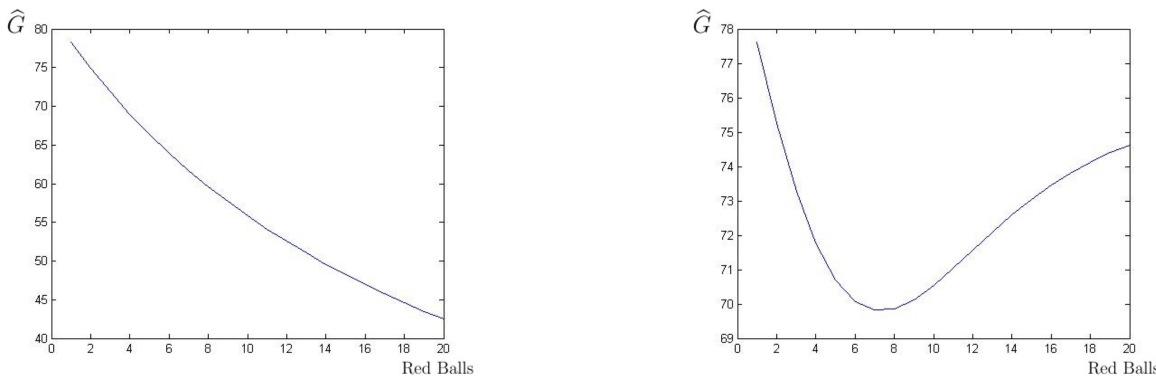


One Urn Two Urns

Figure 1.1: Both diagrams depict how your guess $\widehat{G}$ changes as the other participant draws additional red balls. In the left panel there is a single urn and in the right there are two urns. At the origin you and the other participant have each drawn 8 green and 2 red balls.

draws in forming your guess. The right panel illustrates how dual learning leads your guess to move non-monotonically as we increase the disparity in the color frequency of the other participant's draws and your own.

### 1.3.2 A Model of Dual Learning

In a standard social learning model, a population of agents $i \in N$ seek to learn a single parameter $\theta^* \in \Theta$. Our model extends this to permit unobserved heterogeneity so that agents $i$ and $j$ seek to learn possibly distinct parameters $\theta_i^*$ and $\theta_j^*$.

Nature first partitions the population $\bigcup_{t=1}^{T} N_t = N$ where agents $i$ and $j$ belonging to the same element of the partition $N_t$ are said to be of the same *type*. The partition is formed randomly, with an agent being independently assigned to $N_t$ with probability $\gamma_t$. We may assume the vector of assignment probabilities $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, ..., \gamma_T)$ to be either known or chosen by nature from a known distribution, but the realized partition is unknown.

Nature then independently assigns $\theta^t$ to each member of the partition $N_t$ according to the known probability measure $\Pi(\cdot)$ with density $\pi(\theta)$. Agents of the same type seek to learn the same parameter, that is, for $i$ and $j$ in $N_t$, $\theta_i^* = \theta_j^* = \theta^t$.

Each agent receives a signal $s_i \in S$ in accordance with a conditional distribution with density $f_{\theta_i^*}(s)$. The family of conditional densities $f_\theta(s)$ are one-to-one,[12] continuous in $\theta$ and $s$, and the accompanying measures $F_\theta$ are mutually absolutely continuous.[13] An agent then merges his signal with any of the other information he acquires (e.g. other agents' actions or opinions) to obtain his posterior probability measure $\Pi_i(\cdot|I_i)$ defined as $\Pi_i(A|I_i) \equiv Pr(\theta_i^* \in A|I_i)$ for measurable $A \subset \Theta$ with density $\pi_i(\theta_i^*|I_i)$. Observing other agents' actions or opinions can also be informative of whether they are of the same type as himself. Define $Q_{ij}(I_i) = Pr(j \text{ same type as } i|I_i)$ for $j \neq i$ to be the collection of $i$'s posterior *perceived similarity*.

Each agent then uses his information to select an action $x_i$ from the set $X$ maximizing his expected payoff $E[u(x_i; \theta_i^*)|I_i]$. The particular form of the payoff function will be specified for each application.

As in [34], our results are most crisply articulated when agents have clarity in their inferences. For this reason, our primary analysis (section 1.4) allows agents to observe each others' *opinions* $\hat{\theta}(s_i)$. When parameters are real-valued $\Theta \subset \mathbb{R}$, we follow [77, 78] and specify an agent's opinion to be the expectation of his parameter given only his private information $\hat{\theta}(s_i) = E[\theta_i^*|s_i]$. More generally, we can think of an agent's opinion as some sufficient statistic of his private information. Qualitatively, our results remain under more coarse learning.

[12]If $f_\theta(s) = f_{\theta'}(s)$ almost everywhere in $S$, then $\theta = \theta'$.

[13]Mutual absolute continuity provides that almost surely no single signal will perfectly reveal the distribution from which it is drawn.

## 1.4 Interactive Belief Formation

This section analyzes the basic workings of the model. Beginning with a characterization of disagreement in the simplest setting of two agents and two actions, we move gradually toward increasing generality. Each step toward generality provides an additional insight into the role of dual learning in belief-formation and disagreement.

### 1.4.1 Two Agents & Two Actions

Our analysis begins by considering two agents 1 and 2 who are each faced with a choice between two actions $x_i \in \{0, 1\}$. Agents are assigned one of two possible parameter values $\theta_i^* \in \{0, 1\}$ whereby the payoff to choosing the action $x_i = \theta_i^*$ exceeds the payoff to choosing $x_i \neq \theta_i^*$. Each agent assigns a prior probability of $\pi$ to $\theta_i^* = 1$ and a prior of $\tilde{\pi}$ to the other agent being of the same type. Each agent observes his signal $s_i \sim f_{\theta_i^*}$ and forms his opinion $\hat{\theta}(s_i) = E[\theta_i^*|s_i]$. In this setting, an agent's opinion coincides with what the literature calls an agent's *private belief*

$$\hat{\theta}(s_i) = p(s_i) = Pr(\theta_i^* = 1|s_i). \tag{1.1}$$

From this, we obtain a natural notion of disagreement whereby 1 and 2 *disagree* when $p(s_i) < \pi < p(s_j)$, that is, their opinions are pushed in different directions from the prior.

Each agent then observes the other's opinion. We take the perspective of agent 1 and simplify notation. Denote 1's posterior belief by $P(\mathbf{s}) = Pr(\theta_1^* = 1|s_1, p(s_2))$ and the posterior perceived similarity by $Q(\mathbf{s}) = Pr(\textit{2} \text{ same type as } \textit{1}|s_1, p(s_2))$ where $\mathbf{s} = (s_1, s_2)$.

As a benchmark for comparison, consider how agents update upon observing each other's opinions when they are certainly of the same type $\theta_1^* = \theta_2^* = \theta^*$. In this case, it is straightforward to see that both agents come to agreement on what we call the *shared opinion* $\hat{\theta}(s_1, s_2) = E[\theta^*|s_1, s_2]$ which simplifies

$$\hat{\theta}(s_1, s_2) = p(s_1, s_2) = Pr(\theta^* = 1|s_1, p(s_2)) \tag{1.2}$$

Unobserved heterogeneity adds an additional layer of complexity to the updating procedure. Fortunately, we can compute agent 1's posterior simply as the weighted average between his own opinion and the shared opinion where the weight is precisely the perceived similarity

$$P(\mathbf{s}) = p(s_1, s_2)Q(\mathbf{s}) + p(s_1)(1 - Q(\mathbf{s})). \tag{1.3}$$

By inspection of (1.3), when the perceived similarity is small, agent 1 mostly disregards 2's opinion $P(\mathbf{s}) \approx p(s_1)$. Conversely, when the perceived similarity is close to unity, 1's beliefs resemble those of the standard model $P(\mathbf{s}) \approx p(s_1, s_2)$.

The perceived similarity is itself revised upon observing opinions. It will be highest when the agents receive, in a sense, similar signals and shrink as their signals diverge. More precisely, following Bayes Rule we can write the perceived similarity

$$Q(\mathbf{s}) = \frac{f(s_1, s_2)\tilde{\pi}}{f(s_1, s_2)\tilde{\pi} + f(s_1)f(s_2)(1 - \tilde{\pi})} = \left[1 + \frac{f(s_1)}{f(s_1|s_2)} \frac{1 - \tilde{\pi}}{\tilde{\pi}}\right]^{-1} \tag{1.4}$$

where $f(\mathbf{s}) \equiv \int_\Theta f_\theta(\mathbf{s})\mathrm{d}\Pi(\theta)$ is the marginal likelihood of receiving signals $\mathbf{s}$, assuming that they were drawn from the same distribution. As seen in (1.4), 1's perceived similarity increases upon observing the signals just in case $f(s_1|s_2) > f(s_1)$, that is, the likelihood of receiving $s_1$ from a distribution from which we already obtained $s_2$ exceeds the unconditional likelihood of having drawn $s_1$. With just two parameters, this inequality simplifies to $[f_1(s_1) - f_0(s_1)][f_1(s_2) - f_0(s_2)] > 0$.

The picture of interactive belief formation under dual learning is distinct from the standard model. When two agents share their opinions, their beliefs are drawn closer, but generically, we should not expect full agreement. Letting $P_i(\mathbf{s})$ represent $i$'s posterior belief, the difference in posterior beliefs is proportional to the difference in opinions

$$P_i(\mathbf{s}) - P_j(\mathbf{s}) = \big(1 - Q(\mathbf{s})\big)\big(p(s_i) - p(s_j)\big). \tag{1.5}$$

Thus full agreement only occurs when agents hold equivalent private information $p(s_i) = p(s_j)$ or if they are certainly of the same type $Q(\mathbf{s}) = 1$.

The behavior predicted by the dual learning model can depart strongly from that predicted by the standard model. We shall see this difference in the following example.

**Restaurant Example**

Consider a new restaurant. If individual $i$ chooses to dine there $x_i = 1$, he receives either high ($u_i = 1$) or low ($u_i = 0$) satisfaction. The payoff to any given visit is random and depends on an unknown parameter $\theta_i^* \in \{0, 1\}$, with $u_i(1; \theta_i^*) = \theta_i^*$ with probability 0.75. Thus $i$'s expected payoff to dining at the new restaurant is 0.75 if $\theta_i^* = 1$ and 0.25 if $\theta_i^* = 0$. Specify $i$'s payoff to not dining at the new restaurant $x_i = 0$ to be $u_i(0) = 0.4$ with certainty. Assume agents' prior beliefs are $\pi = \tilde{\pi} = \frac{1}{2}$.

This example resembles [10] in that an agent's realized payoff operates as a signal. Before addressing the decision problem, let us first see how beliefs evolve if both agents repeatedly dine at the new restaurant and have opposed experiences.

Figure 1.2 illustrates the dynamics in agent 1's beliefs when both agents repeatedly dine at the restaurant and each time agent 1 receives high ($u_1 = 1$) satisfaction and 2 receives low ($u_2 = 0$) satisfaction. Along the horizontal axis, we increase the number of times each has dined at the restaurant. Under the standard model of learning, agents combine their

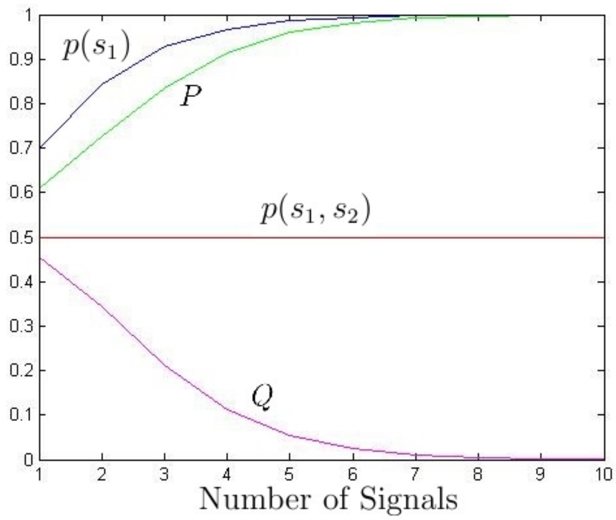


Figure 1.2: Belief Dynamics. The diagram depicts the change in 1's beliefs as we increase the number of high payoffs $s_1 = (1, 1, ...)$ received by 1 and low payoffs $s_2 = (0, 0, ...)$ received by 2.

opinions and form the shared opinion $p(s_1, s_2)$ which, due to their conflicting experiences, remains unchanged from the prior. Allowing for multiple types, the disparity in satisfaction provides increasingly strong evidence that agent 2 is of a distinct type to that of 1 and the perceived similarity vanishes $Q(\mathbf{s}) \to 0$. This observation, taken together with equation (1.3) implies that 1's posterior $P(\mathbf{s})$ quickly converges to his own opinion $p(s_1)$.

Consider how the agents' decisions are affected by their divergent experiences. For simplicity, suppose that an agent selects the dining option that maximizes his expected payoff, that is, each will dine at the new restaurant whenever it yields an expected payoff of at least 0.4.[14] The *ex ante* expected payoff to dining at the new restaurant is 0.5, and thus both agents choose this option. If the agents were homogeneous, then their persistent conflicting experiences would lead both to continue dining at the new restaurant.

After one visit to the new restaurant, before observing agent 1's opinion, agent 2's low satisfaction experience would reduce his expected payoff to dining at the new restaurant to

[14]Given our objective of illustrating opinion formation, we set aside the questions of optimal or strategic experimentation as studied in Bolton (1999) and section 1.5.4 of this paper.

0.375. With no other information, agent 2 would not choose to dine there again. Upon learning that agent 1 had received high satisfaction from the new restaurant, agent 2 would revise his expected payoff to about 0.43 and would thus be willing to give the new restaurant another chance. After a second visit to the new restaurant brings agent 2 low satisfaction, he will not choose to dine there again, regardless of agent 1's satisfaction.

We can formalize our observation from the restaurant example that disparity between opinions leads one to place less weight on another's opinion in the following proposition.

**Proposition 1.1.** *Say that agent* 1 *and* 2 *agree if* $\pi < p(s_i) \leq p(s_j)$ *or* $p(s_i) \leq p(s_j) < \pi$ *and disagree if* $p(s_i) < \pi < p(s_j)$.

*(a) If agent* 1 *and* 2 *agree, then the perceived similarity is strictly increasing as we increase the certainty of either of their opinions.*

*(b) If agent* 1 *and* 2 *disagree, an increase in the difference of their opinions reduces the perceived similarity.*

*(c) Under maximal disagreement the perceived similarity vanishes and* 1*'s posterior belief converges to his opinion:* $p(s_i) \to 0$ *and* $p(s_j) \to 1$ *imply* $Q(\mathbf{s}) \to 0$ *and thus* $|P(\mathbf{s}) - p(s_1)| \to 0$.

Proofs of this and all further propositions can be found in the appendix.

### 1.4.2 Two Agents & Continuum of Actions

We now expand the action and parameter spaces $X = \Theta = \mathbb{R}$. In doing so, we show that increasing the disagreement between opinions can lead to more interesting, non-monotonic changes in actions.

As before, an agent observes his signal $s_i$ and forms his opinion $\hat{\theta}(s_i) \equiv E[\theta_i^*|s_i]$. Assume $\theta_i^*$ to be normally distributed $\theta_i^* \sim \mathcal{N}(\theta_0, \sigma_0^2)$ and signals also normal distributed $s_i \sim \mathcal{N}(\theta_i^*, \sigma^2)$ so that an agent's opinion is a sufficient statistic of his information.

Each agent then observes the other's opinion, and updates his beliefs. If agents were homogeneous with certainty $\theta_1^* = \theta_2^* = \theta^*$, then they would come to agreement on the *shared opinion* $\hat{\theta}(\mathbf{s}) = E[\theta^*|\mathbf{s}]$, $\mathbf{s} = (s_1, s_2)$. For this section, payoffs are assumed to take the form

$$u(x_i; \theta_i^*) = -(x_i - \theta_i^*)^2. \tag{1.6}$$

Agent $i$'s optimal action coincides with the posterior expectation of his parameter

$$x_i^* = E[\theta_i^*|s_i, \hat{\theta}(s_2)] = \hat{\theta}(\mathbf{s})Q(\mathbf{s}) + \hat{\theta}(s_i)(1 - Q(\mathbf{s})). \tag{1.7}$$

How does 1's action respond to a change in 2's opinion? The answer to this will depend on which of the two countervailing forces of dual learning dominates. Observe that $s_2$ enters (1.7) first through the shared opinion $\hat{\theta}(s_1, s_2)$ and second in the perceived similarity $Q(\mathbf{s})$. We can think of the movement in the shared opinion as the *direct effect* of shifting $s_2$. This effect captures the change in 1's beliefs if he takes 2's opinion at face value and does not consider the possible differences between them. Similarly, we can think of the adjustment of the perceived similarity as the *indirect effect* of shifting $s_2$. When 2's opinion $\hat{\theta}(s_2)$ is made increasingly dissimilar to 1's opinion $\hat{\theta}(s_1)$ the indirect effect counteracts the direct effect and the net result will depend on which of these two effects dominates. In the following, we continue the political opinion example from section 1.1 to show that the net change in 1's beliefs is not a foregone conclusion, but can vary on the domain of $s_2$.

**Politics Example**

Suppose the space of political policies can be described by the real line. Let us now think of $s_i$ as $i$'s interpretation of a piece of evidence based on his epistemic and normative values and $\theta_i^*$ as $i$'s most preferred policy if he were to observe all the possible evidence. Specify $i$'s interpretation of a piece of evidence as a noisy signal $s_i = \theta_i^* + \epsilon_i$ with $\epsilon_i$ distributed standard normal. Individuals who share the relevant underlying values will be receiving signals about the same preferred policy. Assume 1's prior over his preferred policy $\theta_1^*$ is standard normal and he receives evidence that suggests a policy of $0.7 = s_1$.

Figure 1.3 shows how 1's choice of policy $x_1^* = E[\theta_1^*|\mathbf{s}]$ changes as we alter 2's opinion. The direct effect of $s_2$ on the shared opinion $\hat{\theta}(s_1, s_2)$ is represented by the positive dashed line in the top sub-figure. The indirect effect of $s_2$ on the perceived similarity is given in the bottom sub-figure. Notice that the perceived similarity peaks near the point where 2's opinions are identical to 1's ($s_2 = 0.7$) and declines as 2's opinion moves in either direction. When 1 and 2's opinions are closest, the direct effect dominates and 1's action moves in a positive and roughly linear fashion. However, as 2's opinion moves further away from 1's in either direction, the indirect effect dominates and 1's action moves negatively with $s_2$. We call this pattern *non-monotonicity in disagreement*: when in close agreement 1 responds to changes in 2's opinion in a qualitatively similar way as the standard model, pushing 2's opinion too far leads 1 to respond in precisely the opposite manner of the standard model.

**Moderating Extremists & Radicalizing Moderates**

An important implication of non-monotonicity in disagreement relates to the processes of moderating extremist or radicalizing moderate behavior. In the previous example, imagine that higher actions are deemed more extreme and socially-undesirable. When confronted by a far less extreme agent 2 ($s_2 \leq -2$) figure 1.3 shows that 1's behavior will be almost

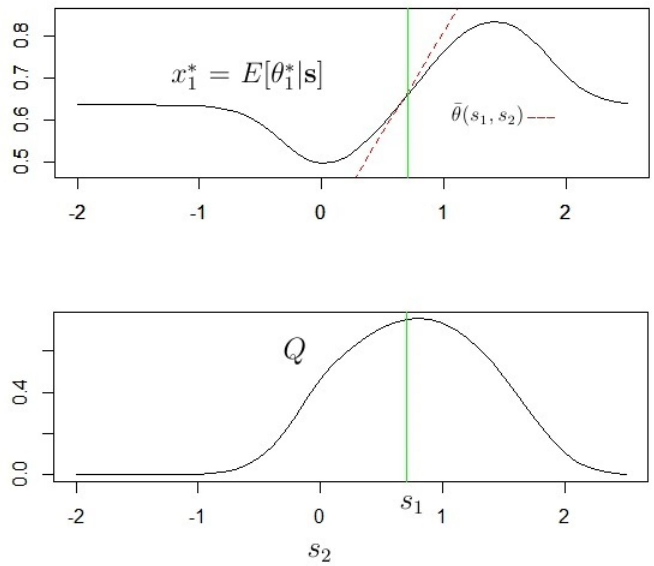


Figure 1.3: Non-Monotonicity in Disagreement. The diagrams illustrate the dependence of 1's beliefs and action on 2's signal.

entirely unchanged. If agent 2 were in fact more extreme ($s_2 \approx 0$), which may mean that 2 advocates performing some degree of socially-undesirable behavior, then 1 would reduce the extremity of his actions.

The converse observation is made by supposing instead that agent 1 is already moderate in his behavior. There is more danger in him encountering a marginally more extreme individual ($s_2 \approx 1.5$) than someone who is vastly more extreme ($s_2 \geq 2.5$).

We now formally characterize the portions of the domain on which 1's action moves either positively or negatively to changes in 2's opinion. Consider as an analogue the identity $Revenue = Price \times Quantity$ from first principles. Revenue's response to a shift in the price can be positive or negative depending on the price elasticity of demand. Similarly, the response of 1's action to changes in 2's opinion will depend on the relative elasticity of the perceived similarity. To simplify notation, let $\Delta(\mathbf{s}) \equiv \hat{\theta}(s_1, s_2) - \hat{\theta}(s_1)$.

**Definition 1.1.** *Define* $\varepsilon \equiv -\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}\Big/\frac{\Delta(\mathbf{s}')-\Delta(\mathbf{s})}{\Delta(\mathbf{s}')+\Delta(\mathbf{s})}$ *to be the elasticity of* 1*'s perceived similarity* $Q(\mathbf{s})$ *which is said to be relatively elastic if* $\varepsilon > 1$ *and relatively inelastic if* $\varepsilon < 1$.

The following proposition makes the regularity assumption that $\hat{\theta}(s_j) < \hat{\theta}(s'_j)$ implies $\hat{\theta}(s_i, s_j) < \hat{\theta}(s_i, s'_j)$. A similar (but less intuitive) statement could be made without use of this assumption.

**Proposition 1.2** (Non-Monotonicity in Disagreement)**.** *Agent 1's action* $x_1^* = E[\theta_1^*|\mathbf{s}]$ *moves positively (negatively) with a change in 2's opinion* $\hat{\theta}(s_2)$ *if the perceived similarity* $Q(\mathbf{s})$ *is relatively inelastic (elastic).*

### 1.4.3 Larger Finite Population ($n > 2$)

In this section, we will see that dual learning in a larger population produces different behavior than standard learning. In particular, we find *persuasion in numbers*: the opinions of the many outweigh the opinions of the few or one, even if both sets of opinions are equivalent in terms of information.

For example, consider customer product reviews. When many customers write reviews for a product, there is a good chance that some proportion of them will share the same preferences of the reader of these reviews. In contrast, when a single customer writes a review, the reader of this review cannot be sure if the customer's preferences match his own. Therefore, if many customers report satisfaction from single uses of a product it can more strongly influence the reader's purchasing decision than if a single customer were to report satisfaction from many uses of the product. In contrast, if it were known *ex ante* that everyone shared the same preferences, then both sets of reviews would influence the reader's beliefs identically. To illustrate this, we continue the restaurant example from section 1.4.1.

**Restaurant Example Cont.**

Suppose that agent 1 has not yet dined at the new restaurant and must solicit the opinions of his fellow agents prior to making his dining decision. We proceed by comparing three cases (1) agent 1 observes agent 2's repeated positive reviews (2) agent 1 observes the negative reviews from agents $3, 4, ..., n$, and (3) agent 1 observes both sets of reviews.

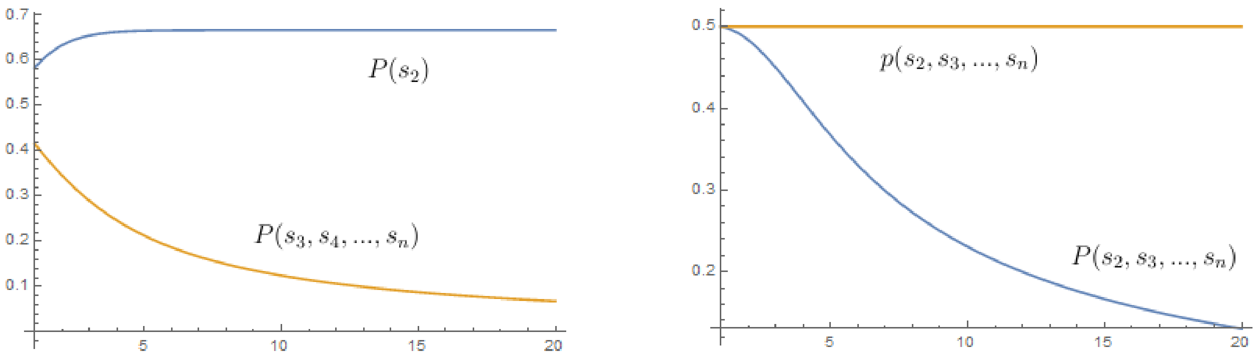


(a) (b)

Figure 1.4: Belief Dynamics. The diagrams depict the change in agent 1's beliefs as we increase the number of high payoffs $s_2 = (1, 1, ...)$ received by agent 2 and low payoffs $(s_3, s_4, ...) = (0, 0, ...)$ received by agents 3, 4, ... .

Suppose first that agent 1 observes that agent 2 repeatedly receives high satisfaction from dining at the restaurant. The plot of $P(s_2)$ in Figure 1.4.a shows how agent 1's beliefs update after each of 2's visits. Notice that, while 2's positive experiences increase 1's beliefs of his own self receiving a high expected payoff from the new restaurant, the effect tapers off. It becomes increasingly clear that 2's expected payoff from the restaurant is high, but there is no guarantee that he shares 1's tastes.

Second, suppose that agent 1 observes each agent $3, 4, ...n$ receive low satisfaction from dining at the new restaurant. The plot of $P(s_3, s_4, ..., s_n)$ in Figure 1.4.a reveals that these observations drive 1 to certainty that he will obtain a low expected payoff from the new restaurant. The effect is stronger in this case because there is a good chance that some fraction of these agents share the same tastes as agent 1.

Finally, we turn to Figure 1.4.b to see the effect of agent 1 observing agent 2 receiving ever more satisfying dining experiences and agents $3, 4, ..., n$ sequentially receiving low satisfaction experiences. The plot of the shared opinion $p(s_2, s_3, ..., s_n)$ demonstrates that, if the agents were homogeneous, then the conflicting experiences would lead 1's beliefs to remain unchanged from the prior. In contrast, the decline of $P(s_2, s_3, ..., s_n)$ reveals that the negative experiences of the many dominates the positive experiences of the one. Observing enough of these reviews will induce agent 1 to forgo dining at the new restaurant.

The following proposition describes *persuasion in numbers*. The phenomenon is most clearly identified in an environment in which $X$ and $\Theta$ are finite, the prior over the assignment probabilities $\boldsymbol{\gamma}$ takes full support in the $T$-dimensional simplex, payoffs are finite, and $x(\theta) \equiv \arg\max_{x \in X} u(x; \theta)$ varies in $\theta$.

We say that agent $j$ is *certain* $x(\theta_j^*) = x$ if $Pr\big(x(\theta_j) = x | s_j\big) = 1$ and agent $k$ is *boundedly certain* $x(\theta_k^*) = y$ if $\delta < Pr\big(x(\theta_k^*) = y | s_k\big) - Pr\big(x(\theta_k^*) = x | s_k\big) < 1$ for all $x \neq y$ and some $\delta > 0$.

**Proposition 1.3** (Persuasion in Numbers)**.** *Suppose agent $i$'s choice $x_i^*$ is informed by the opinions of $n < +\infty$ other agents, whereby $n_x \geq 1$ of these agents are certain $x(\theta_j^*) = x$ and the remaining $n_y$ agents are boundedly certain that $x(\theta_k^*) = y \neq x$.*

*(a) If the population is homogeneous, then $x_i^* = x$ for all $n_y < +\infty$.*

*(b) With a positive ex ante chance of heterogeneity and $n_y$ sufficiently large, $x_i^* = y$.*

### 1.4.4 Countably Infinite Population

What is the expected behavior of an agent's beliefs in an arbitrarily large population? Imagine that agent 1 observes his signal and the other agents' opinions in sequence. In a standard model, well-known results guarantee almost sure consistency of 1's posterior for the true $\theta^*$.

The heterogeneity in our framework precludes an immediate application of these results. To make progress on this question, we must introduce some further notation.

Let $\Omega = \Theta^T \times \Delta^T$ be the set containing the vectors of type parameters and assignment probabilities $\omega = (\theta^1, \theta^2, ..., \theta^T, \gamma_1, \gamma_2, ..., \gamma_T)$, where $\Theta^T$ is the $T$-fold product space $(\Theta \times \Theta \times ... \times \Theta)$ and $\Delta^T$ represents the $T$-dimensional simplex. The population signal density belongs the finite mixture family $g_\omega(s) \equiv \sum_{t=1}^{T} \gamma_t f_{\theta^t}(s)$, $\omega \in \Omega$ and is said to be *identified* just in case $g_\omega(s) = g_{\omega'}(s)$ a.e. implies that both $\omega$ and $\omega'$ assign the same proportion of the population $\theta_i^* = \theta$ for all $\theta \in \Theta$ [81].[15] An agent's opinion $\hat{\theta}(s_i)$ is a random quantity such that $s \mapsto \hat{\theta}(s)$ is one-to-one with realized opinions belonging to a complete separable metric space.

Define $\pi^*$ to be the true distribution of thetas throughout the population as assigned by nature, i.e. $\pi^*(\theta)$ gives the proportion of the population with $\theta_i^* = \theta$. Let $\Theta^* \equiv \text{supp}(\pi^*)$ be the finite support of $\pi^*$. Under the assumption of identifiability, as agent $i$ continues observing the opinions of the other agents, his posterior almost surely converges to a function of only his own signal and the true distribution of parameters. Denote the vector containing the signals of the first $n$ members of the population by $\mathbf{s}^n = (s_1, s_2, ..., s_n)$ and let "$\Rightarrow$" correspond to weak convergence.

**Proposition 1.4** (Belief Convergence.). *Suppose that $\Theta$ and $S$ are complete separable metric spaces endowed with their respective Borel sigma algebras with $g_\omega(s)$, $\omega \in \Theta^T \times \Delta^T$ comprising an identified finite mixture family. Then for almost all $\omega^*$, as $n \to +\infty$*

$$\Pi_i(\cdot|\mathbf{s}^n) \Rightarrow \Pi_i(\cdot|s_i, \pi^*) \;\; a.s. \tag{1.8}$$

[15]Formally, $\sum_{t=1}^{T} \gamma_t \mathbf{1}(\theta^t = \theta) = \sum_{t=1}^{T} \gamma_t' \mathbf{1}(\theta^{t'} = \theta)$ for all $\theta \in \Theta$ where $\omega = (\theta^1, \theta^2, ..., \theta^T, \gamma_1, \gamma_2, ..., \gamma_T)$ and $\omega' = (\theta^{1'}, \theta^{2'}, ..., \theta^{T'}, \gamma_1', \gamma_2', ..., \gamma_T')$. See [87] and [55] for further discussion of identified finite mixture models.

For continuity sets[16] $A \subset \Theta$, we can write $i$'s asymptotic posterior explicitly as

$$\Pi_i(A|s_i, \pi^*) = \frac{\sum_{\theta \in \Theta^*} f_\theta(s_i)\pi^*(\theta)\delta_\theta(A)}{\sum_{\theta' \in \Theta^*} f_{\theta'}(s_i)\pi^*(\theta')} \tag{1.9}$$

where $\delta_\theta(A) = \mathbf{1}_\theta(A)$ is the Dirac measure assigning point mass at $\theta$. If $\Theta$ is finite, we can write $i$'s posterior probability mass function even more simply as

$$\pi_i(\theta|s_i, \pi^*) = \frac{f_\theta(s_i)\pi^*(\theta)}{\sum_{\theta' \in \Theta^*} f_{\theta'}(s_i)\pi^*(\theta')}. \tag{1.10}$$

These expressions have a nice interpretation. If agent $i$ momentarily sets his own signal to the side and observes an infinite sequence of the other agents' opinions, he will learn the distribution of parameters throughout the population $\pi^*$. This distribution essentially becomes his new prior distribution over $\theta_i^*$ which he then updates by reintroducing his private signal.

It is worth noting that the belief convergence obtained in Proposition 1.4 does not guarantee that agents will likewise converge to their optimal action in the limit of learning. Section 1.5.1 discusses this in some detail.

In the limit of exchanging opinions, agents' beliefs converge. At this point, observing an additional agent's opinion serves almost entirely as an indicator of their underlying similarity. To examine the asymptotic behavior of $i$'s perceived similarity, let us revisit the political opinion example where we left off in 1.4.2. We state this observation formally in the proposition that follows.

[16]Sets $A$ for which the boundary has an asymptotically measure zero boundary $\Pi_i(\partial A|s_i, \pi^*) = 0$. In other words, none of the $\theta_i^*$ lie on the boundary of $A$.

**Politics Example Cont.**

Imagine now that agent 1 continues conversing with other agents, learning their opinions about the optimal policy. After enough conversations, agent 1 will learn the distribution of opinions throughout the population. With that, the effect that each additional conversation has on his own view declines to zero. However, each conversation continues to be instructive for 1 to assess the similarity between the other agents and himself. In the long run, learning an individual's opinion functions almost entirely for social identification: serving as an indicator of the similarity in values between this individual and himself.

Let $\rho$ be the *Prokhorov metric* defined over the space of measures over $\Theta$. For our purposes, it is sufficient to know that weak convergence corresponds to convergence in $\rho$. Details can be found in section 6 of Billingsley (2009).

**Proposition 1.5** (Social Identification via Social Learning)**.** *Suppose agent $i$ observes the other agents' opinions in sequence. As $n$ goes to infinity, observing $n$'s opinion has a vanishing effect on $i$'s beliefs but a non-vanishing effect on $i$'s perceived similarity of $n$,*

$$\rho\big(\Pi_i(\cdot|\mathbf{s}^n), \Pi_i(\cdot|\mathbf{s}^{n-1})\big) \to 0 \tag{1.11}$$

$$d\big(Q_{i,n}(\mathbf{s}^n), Q_{i,n}(\mathbf{s}^{n-1})\big) \to w(s_i, s_n) \tag{1.12}$$

*where $w(s_i, s_n)$ is not almost surely zero.*

## 1.5 Extensions

Now that we have studied the basic workings of the model and identified the relationship between dual learning and disagreement, let us expand the discussion and ask the model what it has to say when other modifications and features are introduced. To summarize the findings: (1) dual learning can enhance learning while disallowing optimal action convergence in the observational learning environment, (2) agents can be more persuasive if they agree with each other on auxiliary topics, (3) over-representing those with extreme views can result in polarization when there would otherwise be none, and (4) we give a basic characterization of news media's consumer behavior.

### 1.5.1 Learning From Actions

Dual Learning delivers new insights to the observational learning literature. In the observational learning framework introduced by [17] and [11], an individual chooses from a set of actions based on a private signal and information obtained through observing the actions taken by those who have chosen before. The principal finding has been the presence of *information cascades* whereby it becomes optimal for one (and hence all succeeding individuals) to follow the behavior of the individual who has chosen before oneself without regard for one's private information.

Let $X$ be a finite set of actions, $\Theta$ a finite set of parameters, and $x(\theta_i^*)$ the optimal action for an agent with parameter $\theta_i^*$. Agents choose their actions sequentially. When agent $i$ selects an action, he does so using the information contained in his private signal $s_i$ and the actions of those who have selected before him. Further technical details are reserved for the appendix. Our discussion requires the following definitions:

1. Learning is *complete* if agents asymptotically assign probability one to the true distri-

bution of parameters throughout the population $\pi^*(\theta)$.

2. A process exhibits *optimal action convergence*[17] if

$$\lim_{n\to\infty} Pr\big(x_n = x(\theta_n^*)\big) = 1. \tag{1.13}$$

3. An *information cascade* occurs whenever some agent's choice of action does not depend on his private signal.

4. A *confounding outcome* occurs when the population's limiting beliefs do not converge to certainty on $\pi^*(\theta)$ nor to a belief at which an information cascade would occur.

It will also be useful to define the following properties of the private beliefs that were introduced in 1.4.1. Letting $p_\theta(s_i) = Pr(\theta_i^* = \theta | s_i)$ be $i$'s *private belief* in $\theta$, we say that the signal structure has *unbounded private beliefs* if the support of $p_\theta$ contains 1 for all $\theta$ and *bounded private beliefs* if the support of $p_\theta$ does not contain 1 for any $\theta$.

When is the process guaranteed to produce complete learning and when is there optimal action convergence? [80] study the case in which the population seeks to learn a single parameter $\theta^*$. They show that if private beliefs are unbounded then either the process results in a confounding outcome or there is complete learning and actions converge to optimality. In contrast, there is never complete learning and thus no optimal action convergence when private beliefs are bounded. [1] extend the analysis to the case when agents only observe a subset of the history of actions. They find that with unbounded private beliefs and if agents' observations are sufficiently rich, then actions will converge in probability to the optimal action.

The introduction of unobserved heterogeneity has both a positive and a negative effect on asymptotic outcomes. Firstly, *unobserved heterogeneity completely disallows optimal action*

[17] The literature often refers to almost sure convergence or convergence in probability to the optimal action as *asymptotic learning*. See [1], [3], and [64].

*convergence.* To see why, observe that even if the population were to asymptotically learn the true distribution of thetas amongst the population $\pi^*(\theta)$, each individual $i$'s private signal is insufficient for him to deduce his own parameter of interest $\theta_i^*$. For example, suppose that after observing many agents' purchasing decisions, the population learns that half of the agents in the population enjoy consuming some product while the remaining half do not. When the next agent is tasked with deciding whether or not to consume the product, he cannot be certain as to which half of the population he belongs.

Secondly, *unobserved heterogeneity can facilitate learning when homogeneity would preclude it.* In particular, when agents in the standard model have bounded private beliefs they cannot asymptotically assign certainty to the true $\theta^*$. The reason is that as the population grows increasingly certain of the true $\theta^*$, it will eventually be the case that an agent's action carries no information about his signal. In contrast with unobserved heterogeneity, when the population grows increasingly certain of the true distribution of parameters $\pi^*(\theta)$, there is always information about an agent's signal in his action. This finding is demonstrated by the example that follows in 1.5.1 adapted from [17].

Unobserved heterogeneity also facilitates learning when private beliefs are unbounded. When agents in the standard model have unbounded beliefs, either complete learning or a confounding outcome will obtain. In our model, each agent's action provides sufficiently rich information to prohibit the possibility of a confounding outcome and complete learning will always obtain.

#### Heterogeneity Facilitating Learning

Consider first the case which we know will result in an information cascade. A countable population of agents sequentially decide whether to *adopt* or *reject* some behavior. There is either a low or high value to adopting the behavior $\theta^* \in \{L, H\}$ and the value to rejecting it is 0. For simplicity, assume $L = -1$ and $H = 1$.

Each agent receives a privately observable signal $s_i \in \{L, H\}$, with $Pr(s_i = \theta^*|\theta^*) = r > 1/2$. When agent $n+1$ is asked to choose between adopting or rejecting the behavior, he computes his expected payoff using the information contained in his private signal as well as the actions of those who have chosen before $\mathbf{x}^n = (x_1, x_2, ..., x_n)$, $x_j \in \{\text{adopt}, \text{reject}\}$. The expected payoff to adopting is $2Pr(\theta^* = H|s_{n+1}, \mathbf{x}^n) - 1$. Thus $n+1$'s best response is to adopt just in case $Pr(\theta^* = H|s_{n+1}, \mathbf{x}^n) \geq \frac{1}{2}$. Using Bayes theorem, $n+1$ will select adopt if

$$\begin{cases} r + Pr(\theta^* = H|\mathbf{x}^n) > 1 & s_{n+1} = H \\ \\ Pr(\theta^* = H|\mathbf{x}^n) > r & s_{n+1} = L \end{cases} \tag{1.14}$$

The literature refers to $Pr(\theta^* = H|\mathbf{x}^n)$ as the *public belief*—the likelihood each agent $j > n$ assigns to the parameter being $H$ after observing the first $n$ actions of the other agents, but not his own private signal. Notice that the process will enter an information cascade if the public belief exceeds $r$ as $n+1$ will choose 'adopt' regardless of his private signal. Similarly, a cascade ensues if the public belief falls below $1-r$ as $n+1$ will always choose 'reject'.

Must the process eventually enter into a cascade? There is a simple argument for why in fact it will eventually enter into a cascade with probability one. For a contradiction suppose that with positive probability the process does not at any point enter a cascade. For this to be true, the public belief could never have entered $[0, 1-r) \cup (r, 1]$. Being that the process never enters a cascade, we can infer the precise signal of each actor. [31] shows that observing an infinite sequence of IID draws will lead the public belief to almost surely converge to certainty on the true state and will thus converge to 0 or 1. This implies a contradiction as the almost sure convergence of the public belief will require it to have entered into $[0, 1-r) \cup (r, 1]$ and thus a cascade with probability one.

Now suppose that we introduce heterogeneous types into the model and define $z(\mathbf{x}^n) = Pr(\theta_j^* = H|\mathbf{x}^n)$ for $j > n$ to be the *public belief*. As before, extreme public beliefs $z(\mathbf{x}^n) \in [0, 1-r) \cup (r, 1]$ induce an information cascade. However, the process is no longer guaranteed to end up in a cascade! Notice now that if the process remains out of a cascade, the public belief will not converge to 0 or 1, but rather to the true proportion with $\theta_i^* = H$. In the appendix, we show that whenever this proportion lies within $(1-r, r)$, there is positive probability that the process never enters a cascade and complete learning will occur. The following proposition formalizes the discussion.

**Proposition 1.6.** *In a population with unobserved heterogeneity:*

*(a) Optimal action convergence does not occur.*

*(b) There is generically complete learning with unbounded beliefs.*

*(c) Complete learning outcomes robustly exist with bounded beliefs.*

### 1.5.2 Multiple Learning Problems

Up until now, we have maintained the assumption that agents form and share their opinions about a single topic. Realistically, there are many related topics we wish to learn about. The main lesson of this section is that the degree to which one can influence another's opinion is larger when they agree on auxiliary topics. Conversely, substantial disagreement on auxiliary topics can mitigate one's influence over another's opinion. Suppose that agent 1 has dined at many restaurants with agent 2 and they have shared largely the same quality of experiences each time. When 1 receives word that 2 holds a differing opinion about some new restaurant, he will be less swift in dismissing 2 as having distinct tastes. However, if 2 had a history of holding differing opinions about restaurants, he would hardly have placed any weight on 2's opinion of the new restaurant even if it had agreed with his own.

We reconsider the continuous action and parameter space of section 1.4.2. Agent $i$'s payoff to an action $x_i^* \in \mathbb{R}$ takes of the form of quadratic loss from his parameter $\theta_i^*$ as in (1.6). He has previously engaged in $L$ auxiliary learning problems, receiving private signals for parameters $\theta_i^\ell \in \Theta_\ell$ for $\ell = 1, 2, ..., L$. We assume that the type partition is constant between different learning problems $\theta_i^* = \theta_j^*$ and $\theta_j^\ell = \theta_i^\ell$ for all $\ell$, though this could be weakened to positive correlation. We want to show how the similarity in beliefs over the $L$ auxiliary issues $\theta_i^\ell$ affects the susceptibility of $i$ to be influenced in his action $x_i^* = E[\theta_i^*|\mathbf{s}]$. Assuming two agents, we can write 1's perceived similarity as

$$Q(\mathbf{s}) = \left(1 + \frac{1-\tilde{\pi}}{\tilde{\pi}} \cdot \frac{f(s_1)}{f(s_1|s_2)} \cdot \bar{R}^L\right)^{-1} \tag{1.15}$$

where $\bar{R} = \left(\prod_{\ell=1}^L \frac{f(s_{1\ell})}{f(s_{1\ell}|s_{2\ell})}\right)^{\frac{1}{L}}$ is the geometric mean of the likelihood ratios $\frac{f(s_{1\ell})}{f(s_{1\ell}|s_{2\ell})}$. Recall from section 1.4.1 that $f(s_{1\ell}|s_{2\ell}) > f(s_{1\ell})$ implies that the likelihood of observing $s_{1\ell}$ from a distribution from which we already obtained $s_{2\ell}$ exceeds the unconditional likelihood of having drawn $s_{1\ell}$. We can think of this geometric mean $\bar{R}$ as capturing the degree to which 1 and 2 agree on the $L$ auxiliary problems. When $\bar{R} < 1$ the agents tend to agree and when $\bar{R} > 1$ the agents tend to disagree. Let "sufficient auxiliary agreement (disagreement)" denote "$\bar{R}^L$ sufficiently small (large)".

**Politics Example (Multiple Policies)**

The example illustrating non-monotonicity in disagreement in section 1.4.2 demonstrated that there was a necessary limit on the extent to which 2 could influence 1's view on a particular political policy. Now suppose the two agents continue their conversation and 1 discovers that he shares much common ground with 2 on a large array of other political

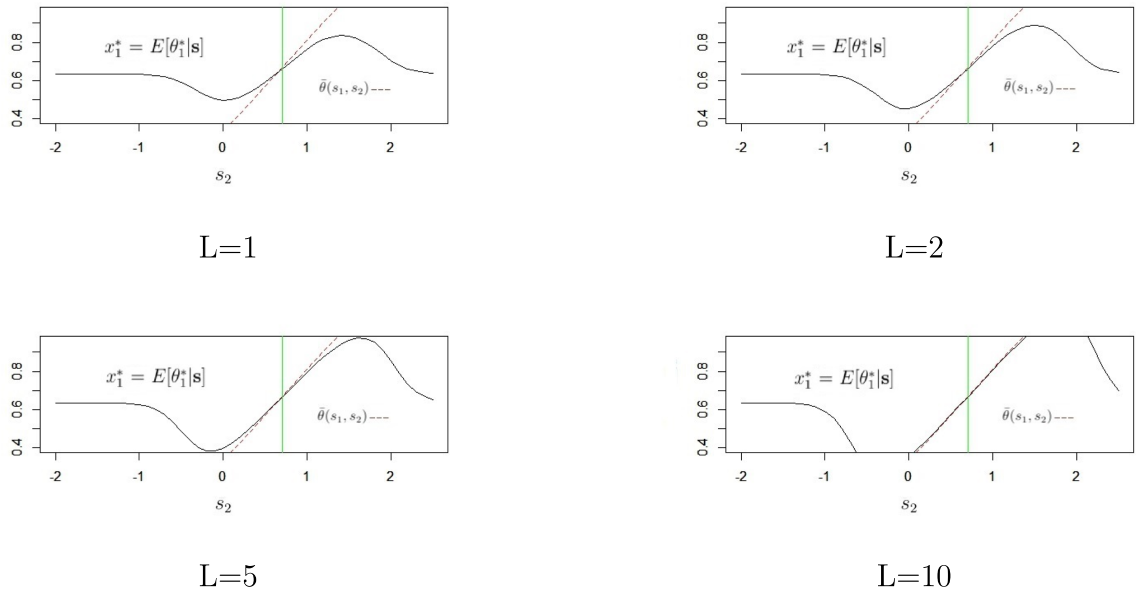


Figure 1.5: Multiple Learning Problems

policies. This discovery will open up 1's belief about the original policy to being more susceptible to influence by 2. Figure 1.5 illustrates the persuasive power of agreeing on auxiliary issues. In each diagram, we fix $s_1 = 0.7$ and vary $s_2$ just as in figure 1.3. Between the diagrams we vary the number of other issues $L$ on which the agents agree, where we specify the agreement as $\bar{R} = 0.75$.

There are a couple of different ways to express the idea that the degree to which 2 can influence 1's action is larger when they agree on auxiliary topics. Firstly, notice in figure 1.5 that by increasing $L$ we expand the domain on which 1's action moves positively with 2's opinion. More generally, we can show that for every compact subset of the signal space $S' \subset S$, there is sufficient auxiliary agreement such that 1's action will move positively with 2's opinion for all $s_1$ and $s_2$ in $S'$.

Secondly, observe in figure 1.5 that increasing the auxiliary agreement raises the peaks and lowers the troughs of the $E[\theta_1^*|\mathbf{s}]$ curve. Generally, if the shared opinion $\hat{\theta}(s_1, s_2)$ is an unbounded function of $s_2$, then $\max_{s_2} E[\theta_1^*|\mathbf{s}]$ can be made arbitrarily large and $\min_{s_2} E[\theta_1^*|\mathbf{s}]$ arbitrarily small. Unbounded shared opinions can be found in the above example with normally distributed signals and a setting in which 2's signals consist of all vectors of finitely many draws from $f_{\theta_2^*}(s)$.

Consider a modification to the above example so that instead the agents disagree on the auxiliary topics $\bar{R} > 1$. This change would result in figure 1.5 showing the opposite qualitative effect of increasing $L$. In particular, an increase in $L$ lowers the variation of $x_1^*$ in 2's opinion and thus the influence that 2 can have on 1's beliefs.

**Proposition 1.7.**

(a) For every compact subset $S' \subset S$, there is sufficient auxiliary agreement such that 1's action $x_1^*$ moves positively with a change in 2's opinion $\hat{\theta}(s_2)$ for all signals in $S'$.

(b) If the shared opinion is unbounded in $s_2$, then 1's action $x_1^*$ can be made arbitrarily large or small given sufficient auxiliary agreement.

(c) The distance between $x_1^*$ and his own opinion $\hat{\theta}(s_1)$ will be arbitrarily small under sufficient auxiliary disagreement.

### 1.5.3 (Perceived) Polarization

What happens if the media and social media skew their coverage in a way that over-represents those with more extreme views and this distortion is not accounted for by the population? We are going to look at an example of how dual learning can serve as a channel through which this type of distortion can lead to polarization where there would otherwise be none.

First, let us see why this distortion cannot drive polarization in the homogeneous case. Suppose all agents seek to learn $\theta^* \in \{L, M, R\}$ (left, moderate, right) and are faced with a

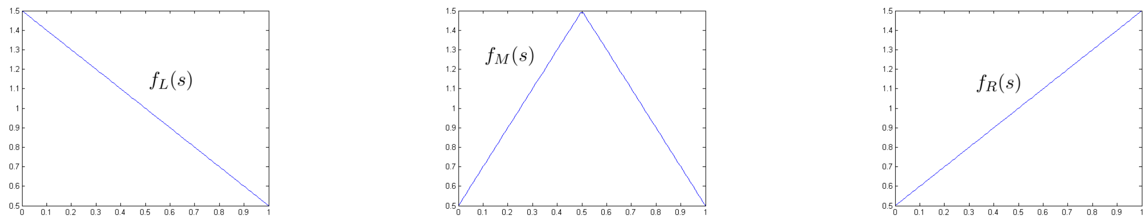


Figure 1.6: Conditional Densities

set of actions $x_i \in \{L, M, R\}$ that yield a payoff of $u_i(x_i) = \mathbf{1}(x_i = \theta^*)$. Each agent receives $s_i \sim f_{\theta^*}$. The conditional densities are represented in figure 1.6, with $f_L(s) = \frac{3}{2} - s$ skewing signals left, $f_M(s) = \frac{1}{2} + 2s$ for $s < \frac{1}{2}$ and $\frac{5}{2} - 2s$ for $s \geq \frac{1}{2}$ giving moderate signals, and $f_R(s) = \frac{1}{2} + s$ skewing signals right.[18]

Each agent observes his own signal $s_i$ and an infinite sequence of other agents' opinions. Consider the effect of a systematic distortion of the publicly observable opinions that over-represents extreme opinions (opinions assigning near certainty to $L$ and $R$). In this case, the publicly observable opinions will outweigh each individual's private signal and the population will fully agree. Of course, the particular belief that the population settles on could be affected by the distortion, but there would nonetheless be agreement.

Now introduce the ex ante possibility of unobserved heterogeneity. Suppose that in actual fact $\theta_i^* = M$ for each and every agent $i$. Then observing infinitely many undistorted opinions will reveal this to be the case and the limiting behavior will be all agents selecting $x_i = M$ regardless of their signal.

As before, suppose there is a systematic distortion that over-represents extreme opinions. Then in the limit, the population might come to believe that there is true polarization. Furthermore, this very belief will drive polarized behavior.

Let $\hat{\pi} = (\hat{\pi}_L, \hat{\pi}_M, \hat{\pi}_R)$ be the limiting estimated distribution of thetas amongst the population.

[18]These densities are used to simplify the exposition. That dual learning can serve as a conduit for generating polarization does not depend on the form of densities.

The expected payoff from action $x \in \{L, M, R\}$ for an agent with signal $s_i$ is

$$U(x; s_i, \hat{\pi}) = Pr(\theta_i^* = x | s_i, \hat{\pi}) = \frac{f_x(s_i)\hat{\pi}_x}{f_L(s_i)\hat{\pi}_L + f_M(s_i)\hat{\pi}_M + f_R(s_i)\hat{\pi}_R}. \quad (1.16)$$

By inspection of (1.16), an agent's optimal action is the one maximizing the product $f_x(s_i)\hat{\pi}_x$. Figure 1.7 demonstrates that leading the population to believe that there are in fact fewer moderates (lowering $\hat{\pi}_M$) and more of the extremes (raising $\hat{\pi}_L$ and $\hat{\pi}_R$) will lead the population to increasingly segment between agents choosing $x_i = L$ and $x_i = R$. The top left sub-figure shows the undistorted case in which all agents optimally select $x_i = M$. With the small amount of distortion in the top right sub-figure, those on fringe with the most extreme signals are induced into choosing $L$ or $R$. The bottom left shows the increase in distortion increases the proportion of the population choosing $L$ or $R$. Finally, in the bottom right sub-figure, the distortion is sufficient to induce agents with any signal to choose $L$ or $R$.

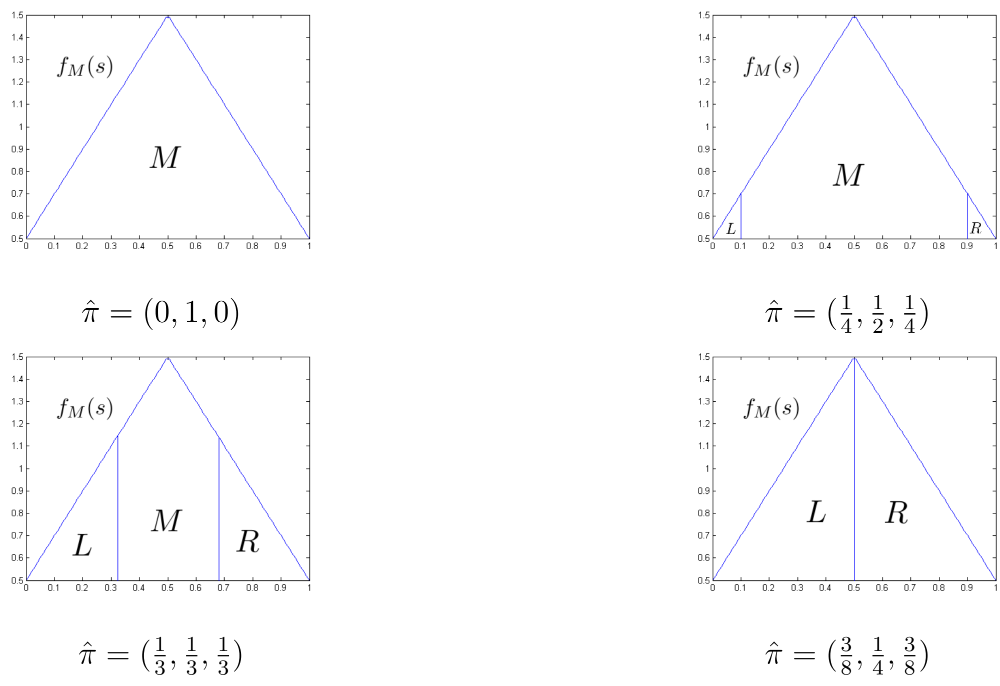


$\hat{\pi} = (0, 1, 0)$ $\hat{\pi} = (\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$

$\hat{\pi} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ $\hat{\pi} = (\frac{3}{8}, \frac{1}{4}, \frac{3}{8})$

Figure 1.7: (Perceived) Polarization. The diagrams show that an increase in the perceived polarization $\hat{\pi}$ can drive polarized actions.

.

### 1.5.4 Application: News Media

How do individuals choose between news sources from which to acquire information? [39] review the literature related to this topic. [65] model consumers with a preference for reading news that confirms their own biases. We show that, with unobserved heterogeneity, consumers will rationally choose to acquire news from sources that tend to confirm their own views—without an explicit confirmation bias. We also fill a gap in this literature by demonstrating that the media can facilitate public disagreement even when the public is aware of the disagreement.

#### Media Consumer Choice

Agents are engaged in a sequence of learning problems. In period $\ell = 1, 2, ...$ agent $i$ selects action $x_i^\ell \in \mathbb{R}$ and receives payoff

$$u_i(x_i^\ell; \theta_i^\ell) = -\alpha_i (x_i^\ell - \theta_i^\ell)^2 \tag{1.17}$$

where $\theta_i^\ell \in \mathbb{R}$ and $\alpha_i$ measures $i$'s idiosyncratic preference for holding accurate beliefs. Before selecting an action, $i$ can choose to acquire a signal $s_i^\ell \sim f_{\theta_i^\ell}$ at a cost $c_i > 0$ and choose whether to observe the opinion $\hat{\theta}(s_M^\ell)$ of media firm $M$ at opportunity cost $d$. We are interested in the case when acquiring direct information about an issue requires more effort than the time it takes to observe the media's report and hence we assume $c_i > d$. We also assume $f_\theta$ to be a normal density with mean $\theta$ and precision $\tau$.

For the moment, suppose that there is a single media firm and consider the decisions faced by $i$ in a given period $\ell$. After acquiring the information that he wishes to obtain, his optimal action will be to select $x_i^\ell(\cdot) = E[\theta_i^\ell|\cdot]$ yielding an expected payoff of $-\alpha_i \text{Var}[\theta_i^\ell|\cdot]$.

When choosing the information to acquire, $i$ must take into account both the benefits to the current period as well as the potential benefits to future periods. Assume $i$ discounts the future at the rate $0 < \beta_i < 1$.

**Proposition 1.8** (Media Consumer Choice)**.** *A consumer's optimal choice in period $\ell$ depends on $\alpha_i$ and $Q_{iM}^{\ell-1}$ as characterized by figure 1.8. Asymptotically, each $\alpha$-type will select from either their left-most or right-most column.*

Figure 1.8 charts out the optimal choice for $i$ with cost $c_i$ and discount rate $\beta_i$ for different sensitivities to accuracy $\alpha_i$ and perceived similarity $Q_{iM}^{\ell-1}$ (written more clearly as $Q$). The symbol $\emptyset$ represents obtaining no signals, $s_i$ obtaining $i$'s own signal, $s_M$ viewing the media's opinion, and $s_i, s_M$ obtaining $i$'s own and also view the media's opinion. In the "Experiment" region, a consumer will be willing to view both $s_i^\ell$ and $\hat{\theta}(s_M^\ell)$ for a period payoff that is lower than observing either only $s_i$ or no signals $\emptyset$. The term "experiment" is drawn from the literature studying bandit processes. In the language of [41], the consumers' decision problem is a bandit *superprocess*.

Consider first $\alpha_i$ in a neighborhood of zero (row ① of figure 1.8). Such an agent would never deem the value high enough to purchase any information regardless of the perceived similarity. In row ②, which vanishes if the cost $c_i$ is too small, the consumer is never willing to obtain his own signal and will only observe the media's opinion with a high enough perceived similarity.

Next, consider the other extreme of an agent with a very large $\alpha_i$ at level ⑤ who is quite intent on forming accurate beliefs. This agent would always find a benefit in obtaining $s_i^\ell$ and, so long as $Q$ is not close to zero, will also view the media's opinion.

The behavior of agents with intermediate sensitivities to accuracy $\alpha_i$ is interesting. At ③, so long as the cost of observing the media's opinion $d$ is not too large, there is a range of $Q$ for which $i$ will obtain a signal and observe that of the media for the sole purpose of experimentation. If the perceived similarity falls too low, the experimentation stops

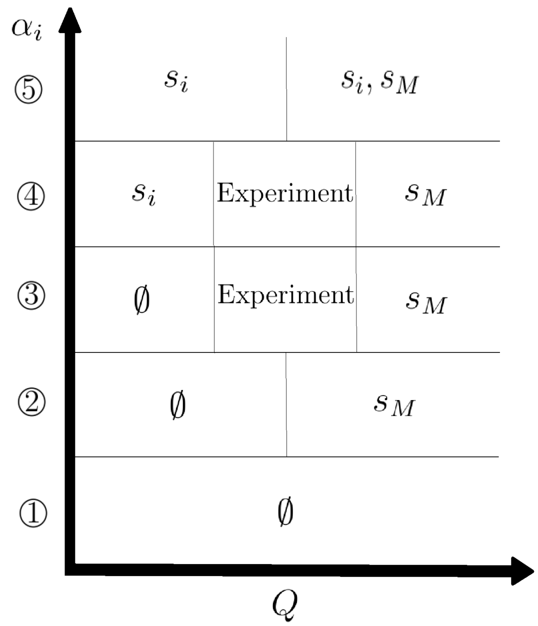


Figure 1.8: Optimal Choice. The diagram gives $i$'s optimal information acquisition for various sensitivities to accuracy $\alpha_i$ and perceived similarity $Q$. Row ② vanishes if $c_i$ is too small and there is no experimentation if $d$ is too large.

altogether and no information observed. Once $Q$ is sufficiently high, the experimentation stops and the agent depends on the media's report.

At ③, middle values of the perceived similarity also involve experimentation if. If the perceived similarity falls too low, $i$ will give up on the media and only trust his own information. As with ②, if $Q$ becomes high, $i$ will stop obtaining his own signal and depend only on that of the media.

Asymptotically, all $\alpha_i$ will at some point choose from either their respective rightmost or leftmost columns in figure 1.8. All types could learn to permanently ignore the media firm at some point when it has a track record of disagreeing too much with them. For $\alpha_i$ types ② through ④, they could find themselves in a situation where they choose to forever trust the information of this media firm.

**Further Observations**

We could conduct the same consumer analysis for the case with multiple media firms. This would show that some agents consume no media and others attend to media sources that do not have too low a perceived similarity. For those intermediate $\alpha_i$ types who rely on media without obtaining their own signals, there is a higher expected payoff when the consumed media sources tend to agree with each other than if they are discordant.

There are a couple more observations to make. For clarity, assume that there are two media firms $A$ and $B$ and the cost for a consumer to observe either of the firm's opinions is zero.

The first observation is that *a small amount of information can result in vast disagreement among the population*. An extreme example of this is found by supposing that the population has grown divided over the course of many periods with half the population assigning a perceived similarity of nearly one to firm $A$ and zero to firm $B$ and the remaining half assigning the reverse beliefs. Suppose further that no agent finds it optimal to pay $c_i$ to purchase his own signal. Then the population will be sharply divided whenever the media firm's reports are distant from each other. This would be like Fox News and MSNBC presenting distinct opinions about some issue and viewers adopting the opinion of the news source that they have agreed with most in the past.

This example also illustrates the second observation that *public disagreement facilitated by the media does not dissipate with public awareness of the disagreement*. In existing models where diverging opinions are driven by media bias, it is important that the agents in the model are themselves unaware of this divergence. Otherwise, agents will simply condition on the opinion divergence and the media's effect vanishes. In our model, agents can grow to trust certain media sources more than others. In fact, learning that other media sources collide with one's trusted source is evidence against believing these other sources. Public disagreement is not crushed by observing the disagreement. Rather, observing public disagreement today lays the groundwork for even stronger disagreement tomorrow.

## 1.6 Conclusion

One of the unexpectedly useful insights gained from Harsanyi's mutual consistency assumption and Aumann's theorem is that we must look to differences between individuals beyond mere differences in their private information to understand the disagreement we observe in the world. This paper studies the emergence and patterns in disagreement when people take into account the unobservable differences between themselves when forming their beliefs. This more complex form of social learning, what we refer to as *dual learning*, captures many phenomena that do not fit existing models.

Disagreement is an important issue to understand. The opinions of a populace culminate in voting behavior that drives political decisions. The views of those tasked with determining research funding influence the very trajectory of science. It is our hope that this discussion and analysis will provide guidance for future empirical and theoretical work in uncovering the underlying differences that drive public disagreement.

## 1.7 Mathematical Appendix

*Proof of Proposition 1.* Let $R_k \equiv \frac{f_L}{f_H}(s_k)$ be the likelihood ratio for drawing signal $s_k$. The assumption $p(s_i) < \pi_H < p(s_j)$ implies $R_i > 1 > R_j$. Expanding (1.4)

$$Q(\mathbf{s}) = \left[1 + \frac{\pi_H^2 + R_j\pi_H\pi_L + R_i\pi_H\pi_L + R_iR_j\pi_L^2}{\pi_H + R_iR_j\pi_L} \cdot \frac{1-\tilde{\pi}}{\tilde{\pi}}\right]^{-1}$$

we find $Q$ to be differentiable in $R_i$ and $R_j$. Upon differentiating, we find $\frac{dQ}{dR_i} < 0$ whenever $R_j < 1$ and $\frac{dQ}{dR_j} > 0$ for $R_i > 1$. Hence, $Q$ is reduced by any increase in $R_i$ or decrease in $R_j$, which corresponds to a decrease in $p(s_i)$ or increase in $p(s_i)$. Letting $R_i^{-1}$ and $R_j$ go to 0 sends $Q$ to zero. From (1.3), $Q(\mathbf{s}) \to 0$ implies $P(\mathbf{s}) - p(s_1) = [p(s_1, s_2) - p(s_1)]Q(\mathbf{s}) \to 0$. $\square$

The mechanism of proposition 1.1 can be found in more general environments. The private beliefs $p(s_i)$ of the two parameter environment provided an intuitive conceptualization of agreement and disagreement between 1 and 2. As we generalize, we now associate the agreement between 1 and 2 directly with the likelihood ratio $\frac{f(s_1)}{f(s_1|s_2)}$: lower values imply a higher degree of agreement. For measurable $A \subset \Theta$ we can write

$$\Pi_1(A|\mathbf{s}) = \Pi'(A|\mathbf{s})Q(\mathbf{s}) + \Pi'(A|s_1)\big(1 - Q(\mathbf{s})\big) \tag{1.18}$$

where $\Pi'(\cdot|\mathbf{s})$ is a probability measure that assumes all signals in the vector $\mathbf{s}$ come from agents of the same type. Let $\|\cdot\|$ denote the total variation metric between two probability measures, $\|\mu - \lambda\| \equiv 2\sup_B |\mu(B) - \lambda(B)|$ where the supremum is taken over measurable $B$.

**Proposition 1.9.** *Consider a change in $s_1$ and/or $s_2$.*

*(a) If $\frac{f(s_1)}{f(s_1|s_2)} \to +\infty$ then $\|\Pi_i(\cdot|\mathbf{s}) - \Pi'(\cdot|s_1)\| \to 0$.*

*(b) If $\frac{f(s_1)}{f(s_1|s_2)} \to 0$ then $\|\Pi_i(\cdot|\mathbf{s}) - \Pi'(\cdot|\mathbf{s})\| \to 0$.*

*Proof.* (a) For any $A \in \mathcal{B}(\Theta)$

$$|\Pi_1(A|\mathbf{s}) - \Pi'(A|s_i)| = |\Pi'(A|\mathbf{s}) - \Pi'(A|s_i)| \cdot Q(\mathbf{s}) \leq Q(\mathbf{s}) \tag{1.19}$$

and hence $2\sup_{A\in\mathcal{B}(\Theta)} |\Pi_1(A|\mathbf{s}) - \Pi'(A|s_i)| \leq 2 \cdot Q(\mathbf{s})$. It follows that $\frac{f(s_1)}{f(s_1|s_2)} \to +\infty$ implies $Q(\mathbf{s}) \to 0$ and thus $\|\Pi_i(\cdot|\mathbf{s}) - \Pi'(\cdot|s_1)\| \to 0$. Part (b) can be proved in a similar fashion. $\square$

*Proof of Proposition 1.2.* Consider a shift in $s_2$ to $s_2'$ and define $\mathbf{s} = (s_1, s_2)$ and $\mathbf{s}' = (s_1, s_2')$. By definition 1.1, $1-\epsilon \equiv 1+\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}\Big/\frac{\Delta(\mathbf{s}')-\Delta(\mathbf{s})}{\Delta(\mathbf{s}')+\Delta(\mathbf{s})}$ which rearranges to $(1-\epsilon)\cdot\big(\Delta(\mathbf{s}')-\Delta(\mathbf{s})\big)\cdot\big(Q(\mathbf{s}')+Q(\mathbf{s})\big) = 2\big(E[\theta_i^*|\mathbf{s}'] - E[\theta_i^*|\mathbf{s}]\big)$. By the regularity assumption $\text{sgn}\big(\Delta(\mathbf{s}') - \Delta(\mathbf{s})\big) = \text{sgn}\big(\hat{\theta}(s_2') - \hat{\theta}(s_2)\big)$ and thus

$$\text{sgn}(1-\epsilon)\cdot\text{sgn}\big(\hat{\theta}(s_2') - \hat{\theta}(s_2)\big) = \text{sgn}\big(E[\theta_i^*|\mathbf{s}'] - E[\theta_i^*|\mathbf{s}]\big) \tag{1.20}$$

where $\text{sgn}(\cdot)$ is the well-known signum function defined for any real number as $x = \text{sgn}(x)\cdot|x|$. □

*Proof of Proposition 1.3.* (a) Consider first the homogeneous case. If some agent $j$ is certain that $Pr(x(\theta^*) = x|s_j) = 1$ then $i$ too obtains this certainty and no opinions to the contrary will reduce it: $Pr(x(\theta^*) = x|s_j, \mathbf{s}) = 1$ for any $\mathbf{s}$ that is not perfectly revealing. (b) As payoffs are finite, if $i$ assigns a high enough probability to the event $x(\theta_i^*) = y$, he will indeed choose $x_i^* = y$. Let $\rho$ denote the probability with which nature assigns an agent to a type with parameter $\theta^y$ which is itself defined as the parameter at which $x(\theta^y) = y$. Let the probability with which nature assigns an agent to a type with parameter $\theta \neq \theta^y$ be denoted by $\eta_\theta(1-\rho)$ and $\boldsymbol{\eta}$ be the vector of length $|\Theta|-1$ containing all such $\eta_\theta$. Agent $i$'s posterior can be written

$$Pr(x(\theta_i^*) = y|\mathbf{s}) = \int_{\mathcal{D}} Pr(x(\theta_i^*) = y|s_i, \rho, \boldsymbol{\eta})\mathrm{d}\mu(\rho, \boldsymbol{\eta}|\mathbf{s}_{-i}) \tag{1.21}$$

where $\mu(\cdot|\mathbf{s}_{-i})$ is the posterior probability measure over $(\rho, \boldsymbol{\eta})$ given the signals of the agents other than $i$ and $\mathcal{D}$ is the $|\Theta|$-dimensional simplex. Expanding the integrand of (1.21)

$$Pr(x(\theta_i^*) = y|s_i, \rho, \boldsymbol{\eta}) = \frac{f_{\theta^y}(s_i)\rho}{f_{\theta^y}(s_i)\rho + (1-\rho)\sum_{\theta\neq\theta^y} f_\theta(s_i)\eta_\theta} \tag{1.22}$$

we see that there exists a cutoff such that, if $Pr(\rho > k^*|\mathbf{s}_{-i}) > 1-\epsilon$, then $i$ will select $x_i^* = y$.

By the Lebesgue Decomposition Theorem, we can decompose $\mu = \mu_1 + \mu_2$ such that $\mu_1$ is absolutely continuous with respect to the ($|\Theta|$-dimensional) Lebesgue measure $\mu_1 << \lambda$ and $\mu_2$ is singular with respect to the Lebesgue measure $\mu_2 \perp \lambda$. Let $v \equiv \frac{d\mu_1}{d\lambda}$ be the Radon-Nikodym Derivative of $\mu_1$ with respect to $\lambda$. Let $\rho < \rho' < 1$, $\mathbf{s}_j$ the signals of the $n_x$ agents certain that $x(\theta_j^*) = x$, and $\mathbf{s}_k$ the remaining $n_y$ signals, and write

$$\frac{v(\rho, \boldsymbol{\eta}|\mathbf{s}_j, \mathbf{s}_k)}{v(\rho', \boldsymbol{\eta}|\mathbf{s}_j, \mathbf{s}_k)} = \frac{f(\mathbf{s}_k|\rho, \boldsymbol{\eta})}{f(\mathbf{s}_k|\rho', \boldsymbol{\eta})} \cdot \frac{v(\rho, \boldsymbol{\eta}|\mathbf{s}_j)}{v(\rho', \boldsymbol{\eta}|\mathbf{s}_j)} \tag{1.23}$$

The density $v(\rho, \boldsymbol{\eta}|\mathbf{s}_j)$ is almost surely positive. Let us now write the ratio of likelihoods

$$\frac{f(\mathbf{s}_k|\rho, \boldsymbol{\eta})}{f(\mathbf{s}_k|\rho', \boldsymbol{\eta})} = \frac{\prod_{s \in \mathbf{s}_k} \left(f_{\theta^y}(s_k)\rho + (1-\rho)\sum_{\theta \neq \theta^y} f_\theta(s_k)\eta_\theta\right)}{\prod_{s \in \mathbf{s}_k} \left(f_{\theta^y}(s_k)\rho' + (1-\rho')\sum_{\theta \neq \theta^y} f_\theta(s_k)\eta_\theta\right)} \tag{1.24}$$

By assumption, $\frac{f_{\theta^y}(s_k)}{f_\theta(s_k)} > b > 1$ for all $k$ some such $b$. With some algebra it can be shown that (1.24) is less than

$$\left(\frac{\rho b + 1 - \rho}{\rho' b + 1 - \rho'}\right)^{n_y} \tag{1.25}$$

which goes to zero as $n_y \to +\infty$. It follows that $v(\rho, \boldsymbol{\eta}|\mathbf{s}_{-i})$ goes to zero as $n_y \to +\infty$. Similarly, we could show that the posterior probability on any atoms goes to zero as $n_y \to +\infty$. We can thus write

$$Pr(\rho < k^*) = \int_{\mathcal{D}|\rho < k^*} v(\rho, \boldsymbol{\eta}|\mathbf{s}_{-i}) \mathrm{d}\lambda(\rho, \boldsymbol{\eta}). \tag{1.26}$$

As the integrand of (1.26) is bounded and converges pointwise to 0, the Bounded Convergence Theorem provides that $Pr(\rho < k^*)$ likewise converges to 0 as $n_y \to \infty$. □

**Remark.** If the population were homogeneous $T = 1$, then Doob's Consistency Theorem [31] tells us that if $i$ observes an infinite sequence of signals, with prior probability one the posterior will be consistent at the true parameter. One condition of Doob's Theorem is that the family of distributions is one-to-one. This assumption is not satisfied for the case of $T$ component mixture models with component weights $\gamma_t$ and parameterized distributions $F_{\theta^t}$, $t = 1, 2, ..., T$. This is indeed the setting in which we are working.

In the following, we carefully define a function $h : \Omega \to \Omega$ that generates an equivalence class of $\omega$'s, in that $g_\omega(s) = g_{\tilde{\omega}}(s)$ a.e. implies that $h(\omega) = h(\tilde{\omega})$. The family $g_{\omega'}$ defined on the image $\omega' \in \Omega' \equiv h(\Omega)$ is one-to-one and hence we can apply Doob's Theorem.

By the Borel Isormorphism Theorem, there exists a Borel isormphism $z$ between $\Theta$ and a subset of the interval $[0, 1]$ with the same cardinality as $\Theta$. Without loss of generality assume $z^{-1}(0) \in \Theta$. Define the linear order on $\Theta$ to satisfy $\theta \leq \theta'$ iff $z(\theta) \leq z(\theta')$.

**Definition.** Let $h : \Omega \to \Omega$ with $h(\omega) = \omega' = (\theta^{1'}, \theta^{2'}, ..., \theta^{T'}, \gamma_{1'}, \gamma_{2'}, ..., \gamma_{T'})$ be defined by the following:

1. Combine duplicate $\theta$'s. Starting from left to right in $\omega$, replace any of $\theta^{t+k} = \theta^t$ for some $k > 0$ with $z^{-1}(0)$ and add $\gamma_{t+k}$ to $\gamma_t$ while also replacing $\gamma_{t+k}$ with 0.
2. Permute the indices so that the $\theta$'s are in ascending order. If there is $\theta^t = z^{-1}(0)$ with $\gamma_t > 0$ for some $t$, place it to the right (a higher index) to all the $\theta^\tau = z^{-1}(0)$ with $\gamma_\tau = 0$.

**Lemma 1.1.**

1. $h(\omega)$ is Borel Measurable.
2. $\Omega' \equiv h(\Omega)$ is a Borel subset of the complete separable metric space $\Omega$.

*Proof.* (1) First decompose the domain $\Omega = \bigcup B_m$ whereby on each of the $B_m$ the order of the $\theta^t$ does not change and if $\theta^t = \theta^{t'}$ or $\gamma_t = 0$ for some $\omega \in B_m$, then it does so for all $\omega' \in B_m$. Define the function $y(\omega) \equiv \big(z(\theta^1), z(\theta^2), ..., z(\theta^T), ..., \gamma_1, \gamma_2, ...\gamma_T\big)$. It can be shown that each $y(B_m)$ is a Borel subset of $[0,1]^{2T}$ and as $y$ is Borel measurable $B_m \in \mathcal{B}_\Omega$.

By design, $h(\omega)$ is continuous on each $B_m$ and the image of these subsets $h_m \equiv h(B_m)$ can be shown to be Borel. Take any $A \subset h(\Omega)$ such that $A \in \mathcal{B}_\Omega$ and write $A = \bigcup A_m$ where $A_m \equiv A \cap h_m$. We can write,

$$h^{-1}(A) = \bigcup (B_m \cap h^{-1}(A)) = \bigcup (B_m \cap h^{-1}(A_m)). \tag{1.27}$$

As $h$ is continuous on each $B_m$, we know $h^{-1}(A_m)$ is contained in each sub-sigma algebra $\mathcal{B}_{B_m}$ and is thus also contained in $\mathcal{B}_\Omega$. It follows that $h^{-1}(A) \in \mathcal{B}_\Omega$.

(2) Follows immediately from $h(\Omega) = \bigcup h_m$ and the fact that each $h_m$ is Borel. $\square$

Using Lemma 1.1, we can extend Doob's consistency theorem to the case of finite mixture models. The statement of the theorem writes "consistency*" to emphasize the use of a qualified notion of consistency. *Consistency* of the posterior at a point $\omega_0$ entails that it will almost surely asymptotically assign probability 1 to every neighborhood of that point. For mixture models, *consistency** of the posterior at a point $\omega_0$ entails that the posterior will almost surely assign probability 1 to every neighborhood of the *set of points equivalent to* $\omega_0$. Here $\omega$ and $\omega'$ are *equivalent* if they assign the same weight to each $\theta$, $\sum_{t=1}^T \gamma_t \mathbf{1}(\theta^t = \theta) = \sum_{t=1}^T \gamma'_t \mathbf{1}(\theta'^t = \theta)$.

**Theorem 1.1** (Doob's Theorem for Finite Mixture Models.)**.** *Suppose that* $\Theta$ *and* $S$ *are complete separable metric spaces endowed with their respective Borel sigma algebras with* $g_\omega(s)$, $\omega \in \Theta^T \times \Delta^T$ *comprising an identified finite mixture family. Let* $\Pi$ *be a prior and*

$\{\Pi(\cdot|\mathbf{s}^n)\}$ *a posterior. Then there exists* $\Omega_0 \subset \Omega$ *with* $\Pi(\Omega_0) = 1$ *such that* $\{\Pi(\cdot|\mathbf{s}^n)\}_{n\geq 1}$ *is consistent* at every* $\omega \in \Omega_0$.

Theorem 1.1 is proved en route to the proof of proposition 1.4. Without loss of generality, the proof proceeds with $i$ conditioning directly on the signals of agent $-i$. We could replace the $s_j$ with $j$'s opinion $\hat{\theta}_j = \hat{\theta}(s_j)$ and the proof would otherwise be unchanged.

*Proof of Proposition 1.4.* By lemma (1.1.1) $h(\omega)$ is Borel measurable so we can induce a measure $\lambda$ on $\Omega'$ defined as $\lambda(A|\mathbf{s}^n) \equiv \widetilde{\Pi}(h^{-1}(A)|\mathbf{s}^n)$ where is $\widetilde{\Pi}(\cdot|\mathbf{s}^n)$ is the public belief over $\omega \in \Omega$ conditional on the vector $\mathbf{s}^n$, $\widetilde{\Pi}(B|\mathbf{s}^n) \equiv Pr(\omega \in B|\mathbf{s}^n)$ for $B \in \mathcal{B}_\Omega$.

First write

$$\Pi_i(A|\mathbf{s}^n) = \int_{\Omega'} Pr(\theta_i^* \in A|\mathbf{s}^n, \omega')\mathrm{d}\lambda(\omega'|\mathbf{s}^n) = \int_{\Omega'} \sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma_t'} \mathbf{1}(\theta^t \in A)\mathrm{d}\lambda(\omega'|\mathbf{s}^n). \tag{1.28}$$

*Claim 1:* $\lambda(\cdot|\mathbf{s}^n) \Rightarrow \delta_{h(\omega^*)}(\cdot)$

The family $g_{\omega'}(s)$ is one-to-one on $\Omega'$ which by lemma (1.1.2) is a Borel subset of a complete separable metric space. Hence by Doob's Theorem [31] as stated in [40], $\lambda(\cdot|\mathbf{s}^n)$ is almost surely consistent and by the Portmanteau Theorem paired with the fact that $\Omega'$ is a separable metric space, $\lambda(\cdot|\mathbf{s}^n)$ converges weakly to $\delta_{h(\omega^*)}(\cdot)$. Theorem 1.1 follows immediately by noting that if $\lambda$ assigns probability 1 to every neighborhood of $h(\omega^*)$, then $\widetilde{\Pi}$ assigns probability 1 to every open set containing $\omega$ such that $h(\omega) = h(\omega^*)$.

The Portmanteu Theorem, equation (1.28), and claim 1 imply that whenever $\sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma_t'} \mathbf{1}(\theta^t \in$

$A$) is almost surely continuous with respect to $\delta_{h(\omega^*)}(\cdot)$,

$$\Pi_i(A|\mathbf{s}^n) \to \int_{\Omega'} \sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t} \mathbf{1}(\theta^t \in A) \mathrm{d}\delta_{h(\omega^*)}(\omega). \tag{1.29}$$

*Claim 2:* For all $A$ with $\Pi_i(\partial A|s_i, \pi^*) = 0$, $\sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t}\mathbf{1}(\theta^t \in A)$ is almost surely continuous with respect to $\delta_{h(\omega^*)}(\cdot)$.

The function $\sum_t \frac{f_{\theta^t}(s_i)\gamma_t}{\sum_{t'} f_{\theta^{t'}}(s_i)\gamma'_t}\mathbf{1}(\theta^t \in A)$ is only discontinuous when some $\theta^t$ crosses the boundary of $A$, and is thus almost surely continuous with respect to $\delta_{h(\omega^*)}(\cdot)$ just in case $\delta_{h(\omega^*)}\big(D\big) = 0$ where $D \subset \Omega'$ is defined as the subset on which $\theta^t \in \partial A$ for some $\theta^t \in \omega \in D$. A set $A$ satisfies this condition if and only if $\Pi_i(\partial A|s_i, \pi^*) = 0$ implying that such an $A$ is a continuity set with respect to $\Pi_i(\cdot|s_i, \pi^*)$. As (1.29) holds for all continuity sets $A$, it follows by a final application of the Portmanteu Theorem that $\Pi_i(\cdot|\mathbf{s}^n)$ weakly converges to $\Pi_i(\cdot|s_i, \pi^*)$. □

*Proof of Corollary 5.* Section 6 in [18] shows that weak convergence corresponds to convergence in $\rho$, hence (1.11) follows from proposition 1.4 and the triangle inequality. As $Q_{in}(\mathbf{s}^n) - Q_{in}(\mathbf{s}^{n-1}) =$

$$\int_{\Omega'} \big(Q_{in}(\mathbf{s}^n|\omega') - Q_{in}(\mathbf{s}^{n-1}|\omega')\big)\mathrm{d}\lambda(\omega'|\mathbf{s}^n) \tag{1.30}$$

has a continuous integrand in $\omega'$, the difference converges to $w(s_i, s_n) \equiv Q_{in}(\mathbf{s}^n|h(\omega^*)) - Q_{in}(\mathbf{s}^{n-1}|h(\omega^*))$. The second term in this difference is constant in $s_n$ and $Q_{in}(\mathbf{s}^n|h(\omega^*))$ is not almost surely constant in $s_n$. □

Our discussion is made commensurate with the observational learning literature by assuming $\Theta$ to be finite, distinct $\theta$ and $\theta'$ prescribe different optimal choices from a finite set of actions

$X$, and the vector of assignment probabilities $\boldsymbol{\gamma}$ to be known. Assume that each action is played for some open set of beliefs.

*Proof of Proposition 1.6.* (a) Under heterogeneity, there is $\theta_i^*, \theta_j^* \in supp(\pi^*)$ with $x(\theta_i^*) \neq x(\theta_j^*)$. Optimal action convergence implies that with probability one such an agent $i$ receives signal $s_i$ such that they choose $x_i = x(\theta_i^*)$. Mutual absolute continuity of the signaling distributions would also imply $j$ receives $s_j$ inducing $x_j = x(\theta_i^*)$ with probability one. Hence, there is no optimal action convergence.

(b) This proof draws strongly from [80] (S&S). In this environment, nature chooses between only finitely many parameter vectors $\boldsymbol{\theta} = (\theta^1, \theta^2, ..., \theta^T)$ and hence only finitely many population distributions of thetas. Denote by $\tilde{\pi}(\theta) = \sum_{t=1}^{T} \gamma_t \mathbf{1}(\theta^t = \theta)$ a generic distribution of thetas and $\pi^*$ the true distribution as chosen by nature. The likelihood ratios $\ell_{\tilde{\pi}}(\mathbf{x}^n) = \frac{Pr(\tilde{\pi}|\mathbf{x}^n)}{Pr(\pi^*|\mathbf{x}^n)}$ for $\tilde{\pi} \neq \pi^*$ and $\mathbf{x}^n = (x_1, x_2, ..., x_n)$ the first $n$ actions chosen form a Martingale conditional on $\pi^*$. Define $\psi(x|\tilde{\pi}, \boldsymbol{\ell})$ to be the ex ante probability of an agent performing action $x$ conditional on $\tilde{\pi}$ being the true distribution of parameters and $\boldsymbol{\ell}$ being their prior vector of likelihood ratios $\ell_{\tilde{\pi}}$.

By the Martingale Convergence Theorem, there exists a real, nonnegative stochastic variable $\ell_{\tilde{\pi}}^{\infty}$ such that $\ell_{\tilde{\pi}}(\mathbf{x}^n) \to \ell_{\tilde{\pi}}^{\infty}$ almost surely. This implies that asymptotically for all $\tilde{\pi}$ and all actions $x$ played with positive probability

$$\ell_{\tilde{\pi}}^{\infty} = \frac{\psi(x|\tilde{\pi}, \boldsymbol{\ell}^{\infty})}{\psi(x|\pi^*, \boldsymbol{\ell}^{\infty})} \ell_{\tilde{\pi}}^{\infty} \tag{1.31}$$

If only one action is taken with positive probability at $\boldsymbol{\ell}^{\infty}$, then because private beliefs are unbounded, the public belief must assign certainty to the population being homogeneous. As the likelihood ratios almost surely will not converge to certainty on the false $\tilde{\pi}$, then the population truly is homogeneous $\pi^*$ and complete learning has occurred.

Consider the case where two actions $x$ and $x'$ are active in the limit. This would imply that at $\boldsymbol{\ell}^\infty$ agents assign positive prior probability to exactly two parameters $\theta$ and $\theta'$. If only one parameter $\theta$ were assigned positive probability only the action $x(\theta)$ would be active. If more than two parameters were assigned positive probability, then unbounded private beliefs would entail that more than two actions would be active. For equation (1.31) to be satisfied, either $\ell_{\tilde{\pi}}^\infty = 0$ or $\psi(x|\tilde{\pi}, \boldsymbol{\ell}^\infty) = \psi(x|\pi^*, \boldsymbol{\ell}^\infty)$ and we shall proceed to show that the latter equality cannot hold for $\tilde{\pi} \neq \pi^*$.

For a given $\boldsymbol{\ell}$, $i$'s best response will be $x$ just in case her private belief $p_\theta(s_i)$ exceeds some threshold $K(\boldsymbol{\ell})$. Define $\widetilde{F}_\theta(p_\theta)$ and $\widetilde{F}_{\theta'}(p_\theta)$ to be the conditional distributions of the perceived similarity $p_\theta$. The ex ante probability that $i$ chooses $x$ for a given $\tilde{\pi}$ is

$$\psi(x|\tilde{\pi}, \boldsymbol{\ell}^\infty) = \tilde{\pi}(\theta)\big(1 - \widetilde{F}_\theta(K(\boldsymbol{\ell}))\big) + \tilde{\pi}(\theta')\big(1 - \widetilde{F}_{\theta'}(K(\boldsymbol{\ell}))\big). \tag{1.32}$$

From lemma A.1 in S&S $F_\theta(p_\theta) > F_{\theta'}(p_\theta)$ whenever both are not zero or one. Thus for distinct $\tilde{\pi}$ and $\pi^*$, almost surely $\psi(x|\tilde{\pi}, \boldsymbol{\ell}^\infty) \neq \psi(x|\pi^*, \boldsymbol{\ell}^\infty)$. It follows that $\ell_{\tilde{\pi}}^\infty = 0$ and complete learning has occurred.

Generically, more than two actions cannot be active at $\boldsymbol{\ell}^\infty$. To see this, let $J$ be the number of $\tilde{\pi} \neq \pi^*$ with $\ell_{\tilde{\pi}}^\infty > 0$ and $M$ the number of actions active at $\boldsymbol{\ell}^\infty$. Because of the identity $\sum_{x \in X} \psi(x|\tilde{\pi}, \boldsymbol{\ell}^\infty) = 1$, if the equality $\psi(x|\tilde{\pi}, \boldsymbol{\ell}^\infty) = \psi(x|\pi^*, \boldsymbol{\ell}^\infty)$ holds for $M-1$ of the active actions, it must also hold for the remaining active action. Hence satisfying equation (1.31) for all $\tilde{\pi}$ and active actions generates a system of $J(M-1)$ equations in $J$ unknowns $\ell_{\tilde{\pi}}^\infty$. nown 1. As the equations generically differ, they can only be solved when $M = 2$.

(c) Assume $\pi^*(H) \in (1-r, r)$. The idea is to first isolate a positive measure of trajectories that after $M_\epsilon$ steps will never leave a small radius around $\pi^*(H)$. Then we show that, of these trajectories, a positive proportion will have never left $(1-r, r)$ in the first $M_\epsilon$ steps.

Let $z(\mathbf{s}^n) = Pr(\theta_j^* = H|\mathbf{s}^n)$ for $j > n$ be a variant of the public belief $z(\mathbf{x}^n)$ defined in 1.5.1. Prior to a cascade $z(\mathbf{s}^n) = z(\mathbf{x}^n)$. As per theorem 1.1, $z(\mathbf{s}^n) \to \pi^*(H)$ almost surely as $n \to +\infty$. Let $Z$ be the set of the trajectories of the public belief with Borel sigma algebra $\mathcal{B}_Z$ and probability measure $\mu$ induced from the signaling distribution $G_{\omega^*}$. Define $Z' \subset Z$ to be the subset of trajectories such that $z_m \in (1-r, r)$ and at all points in the sequence. We want to show that $\mu(Z') > 0$.

Let $\delta \equiv \frac{1}{2}\min(|r - \pi^*(H)|, |\pi^*(H) + 1 - r|)$. For almost all trajectories, there exists an $M < +\infty$ such that, for all $m > M$, $|z_m - \pi^*(H)| < \delta$. Let $M_\epsilon$ be the smallest integer such that $\mu(Z_\epsilon) > \epsilon$ where $Z_\epsilon = \{z \in Z : \forall m \geq M_\epsilon, |z_m - \pi^*(H)| < \delta\}$.

As each trajectory has only received finitely many signals there are only finitely many unique of signal frequencies observed in the first $M_\epsilon$ steps of each trajectory. Partition $Z_\epsilon = \bigcup Z_\epsilon^l$ where each $Z_\epsilon^l$ comprises of trajectories which received the same signal frequencies in the first $M_\epsilon$ steps. The positive measure for $Z_\epsilon$ requires that at least one member of its partition has positive measure and thus suppose $\mu(Z_\epsilon^l) = \epsilon_1 > 0$. Denote by $k$ the number of $H$ signals in the first $M_\epsilon$ steps for trajectories in $Z_\epsilon^l$.

Every permutation of $k$ "$H$" signals and $M_\epsilon - k$ "$L$" signals has the same positive probability denoted by $\epsilon_2 > 0$. Thus we can find a subset of trajectories $\hat{Z}_\epsilon^l \subset Z_\epsilon^l$ that never leave $(1-r, r)$ defined by the permutation in which the first $2 * \min\{k, M_\epsilon - k\}$ steps form an oscillating sequence of $L, H, L, H, ...$ and then including whatever signals remain at the end. At no point during the oscillation does $z \in \hat{Z}_\epsilon^l$ leave $(1-r, r)$ and if it leaves after then necessarily it's $M_\epsilon$th entry $z_{M_\epsilon}$ would too, contradicting $\hat{Z}_\epsilon^l \subset Z_\epsilon$. It follows that $\mu(\hat{Z}_\epsilon^l) = \epsilon_1\epsilon_2 > 0$. Notice that $\hat{Z}_\epsilon^l \subset Z'$ as all $z \in \hat{Z}_\epsilon^l$ never leave $(1-r, r)$.

Thus we know $\mu(Z') \geq \mu(\hat{Z}_\epsilon^l) \geq \epsilon_1\epsilon_2 > 0$ □

In the low probability event of avoiding an information cascade each agent follows the action dictated by his private signal. In the limit, the proportion choosing the correct action is $r$.

*Proof of Proposition 1.7.* (a) By proposition 1.2, it will be enough to show that for every compact $S' \subset S$, $\bar{R}^L$ can be made sufficiently small so that 1's perceived similarity is relatively inelastic $\epsilon < 1$ for all signals in $S'$. Because $S'$ is compact and the conditional densities $f_\theta(s)$ continuous in $s$, both $\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ and $\max_{s_1,s_2,s_2' \in S'} \frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ are well defined on $S'$. It can be shown that for any given signals, $\frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ goes to zero as $\bar{R}^L$ goes to zero. It follows that $\max_{s_1,s_2,s_2' \in S'} \frac{Q(\mathbf{s}')-Q(\mathbf{s})}{Q(\mathbf{s}')+Q(\mathbf{s})}$ also goes to zero as $\bar{R}^L$ goes to zero. From definition 1.1, it follows that for $\bar{R}^L$ sufficiently small, $\epsilon < 1$ on $S'$.

(b) We will show that for sufficient auxiliary agreement $E[\theta_1^*|\mathbf{s}]$ can be made larger than any $M < +\infty$. Having assumed the shared estimate to be an unbounded function of $s_2$, we can find an $s_2$ such that $\hat{\theta}(s_1, s_2) > M$ for any $M$ and $s_1$. From

$$E[\theta_1^*|\mathbf{s}] - \hat{\theta}(s_1, s_2) = \big(1 - Q(\mathbf{s})\big)\big(\hat{\theta}(s_1) - \hat{\theta}(\mathbf{s})\big) \tag{1.33}$$

and equation 1.15, $E[\theta_1^*|\mathbf{s}] - \hat{\theta}(s_1, s_2) \to 0$ as $\bar{R}^L \to 0$, completing the proof. The proof that $E[\theta_1^*|\mathbf{s}]$ can be made smaller than any $m > -\infty$ and for part (c) follow by similar arguments. □

David Blackwell has shown that with a fixed and finite set of choices in each period, there is a deterministic stationary Markov policy for which, for any initial state, the total expected reward is the supremum of the total expected rewards for the class of all policies [19]. The optimal policy satisfies the functional equation

$$V(Q) = \max_{\hat{s} \in \hat{S}} U(\hat{s}; Q) + \beta E\big[V(Q')|\hat{s}, Q\big] \tag{1.34}$$

with $\hat{S} = \{\emptyset, \hat{s}_i, \hat{s}_M, (\hat{s}_i, \hat{s}_M)\}$ the set of signal combinations the consumer can choose to observe. Define $A(Q)$, $B(Q)$, $C$, and $D$ to be the expected period payoffs to observing $(s_i, s_M)$, $s_M$, $s_i$, and $\emptyset$ (no signals) respectively for $Q^\ell_{im} = Q$ and a given $\alpha_i$. Let $\bar{A} \equiv A(1)$ and $\bar{B} \equiv B(1)$.

**Lemma 1.2.**

1. The expected period payoff from observing $s^\ell_M$ is increasing in $Q$.

2. The value function $V(Q)$ is non-decreasing in $Q$. If at $Q$, there is positive probability of ever observing $s^\ell_M$, then $V(Q)$ is increasing in $Q$.

3. There are diminishing expected period returns to information at $Q = 1$.

4. $A(Q)$ and $B(Q)$ are continuous in $Q$.

5. For a given $c_i$ and $\beta_i$, we can partition the domain for $\alpha_i \in \mathbb{R}_+ = [0, b_0] \cup (a_1, b_1] \cup (a_2, b_2] \cup (a_3, +\infty)$ where $a_m < a_{m'}$ and $b_m < b_{m'}$ whenever $m < m'$. The following inequalities hold for the various $\alpha_i$-types.

$$\begin{cases} \bar{B} - d \leq D,\ C - c_i < D,\ \bar{A} - c_i - d < D, & \alpha_i \in [0, b_0] \\ \bar{B} - d > D,\ C - c_i \leq D,\ \bar{A} - c_i - d < D, & \alpha_i \in (a_1, b_1] \\ \bar{B} - d > D,\ C - c_i > D,\ \bar{A} - c_i - d \leq D, & \alpha_i \in (a_2, b_2] \\ \bar{B} - d > D,\ C - c_i > D,\ \bar{A} - c_i - d > D, & \alpha_i \in (a_3, +\infty) \end{cases}$$

*Proof of Lemma 1.2.1.* The proof follows the same form as the proof for 1.2.2. □

*Proof of Lemma 1.2.2.* Expand the value function

$$V(Q) = Q\, V^1(Q) + (1 - Q)\, V^2(Q) \tag{1.35}$$

with $V^1(Q)$ to be value at $Q$ if in fact $i$ and $M$ are of the same type and $V^2(Q)$ if $i$ and $M$ are of different types. Let $\widetilde{V}(Q)$ to be the value function if we remove the option of the agent observing $s_M^\ell$. We want to show

$$V^1(Q) \geq \widetilde{V}(Q) \geq V^2(Q) \tag{1.36}$$

with the last inequality holding strictly whenever the optimal policy assigns positive probability to the consumer ever observing $s_M^\ell$.

The second inequality in (1.36) comes from the fact that $\widetilde{V}(Q)$ follows the policy that maximizes the flow of utility when observing $s_M^\ell$ is not an option. If $i$ and $M$ are not of the same type, then they can do not better than by following the policy of $\widetilde{V}(Q)$, hence, $\widetilde{V}(Q) \geq V^2(Q)$. If the consumer ever observes $s_M^\ell$, then they both pay the cost $d$ and also receive information that will almost surely lead them to choose a suboptimal action. In this case, $\widetilde{V}(Q) > V^2(Q)$.

To demonstrate that the first inequality in (1.36) holds, assume for a contradiction that it does not hold $V^1(Q) < \widetilde{V}(Q)$. From what was previously shown, this would imply $V(Q) < \widetilde{V}(Q)$. This implies a contradiction as the policy for $V$ could be modified to never acquire $s_M^\ell$ guaranteeing $V(Q) = \widetilde{V}(Q)$.

Finally, it follows from the inequalities in (1.36) that if $Q' > Q$

$$V(Q') \geq Q'V^1(Q) + (1 - Q')V^2(Q) \geq V(Q) \tag{1.37}$$

with the last inequality holding strictly if there is a positive probability of the consumer ever observing $s_M^\ell$. $\square$

*Proof of Lemma 1.2.3.* This is where the assumed normal-normal conjugate environment comes into play. Recall that a consumer's expected payoff to $\hat{\mathbf{s}}$ (ignoring the costs) is $-\alpha_i Var(\theta_i^*|\hat{\mathbf{s}})$. Without loss of generality, set $\alpha_i = 1$.

We want to show $\bar{A} - \bar{B} < C - D$. The payoff to $\hat{s}_i$ with precision $\tau$ is $-(\tau + \tau_0)^{-1}$ and the payoff to not observing $s_i$ is $-\tau_0^{-1}$. The effect of observing $s_j$ is to increase $\tau_0$. The desired inequality is obtained by differentiating

$$\frac{d}{d\tau_0}\big[-(\tau+\tau_0)^{-1} + \tau_0^{-1}\big] = (\tau+\tau_0)^{-2} - \tau_0^{-2} \tag{1.38}$$

which is less than zero whenever $\tau > 0$. □

*Proof of Lemma 1.2.4.* As $A(Q)$ and $B(Q)$ give the maximized values for objective functions that are differentiable in $x$ and continuous in $Q \in [0, 1]$, the *Theorem of the Maximum* implies that $A(Q)$ and $B(Q)$ are continuous in $Q$. □

*Proof of Lemma 1.2.5.* $\alpha_i \in [0, b_0]$. All inequalities hold strictly at $\alpha_i = 0$. Both sides of each inequality decrease linearly in $\alpha_i$, with the right side decreasing at a faster rate. Hence, there exists a unique $\alpha_i$ at which point the each inequality becomes an equality. That $\bar{B} - d = D$ implies $C - c_i < D$ follows from $\bar{B} = C$ and $d < c_i$. That $\bar{B} - d = D$ implies $\bar{A} - c_i - d < D$ follows from part 3 of the lemma and some algebra.

$\alpha_i \in (a_1, b_1]$. That $C - c_i = D$ implies $\bar{A} - c_i - d < D$ follows again from part 3 of the lemma. □

*Proof of Proposition 1.8.* For $\alpha_i \in [0, b_0]$, any policy that acquires a signal yields an expected payoff that is less than (strictly so if $Q^0 < 1$) the one that does not select any in each period.

For $\alpha_i \in (a_3, +\infty)$, the inequalities in Lemma 1.2 entail that no policy that ever selects $\emptyset$ or $\hat{s}_m$ is optimal. Hence, we can consider only plans that select $\hat{s}_i$ or $(\hat{s}_i, \hat{s}_M)$ in each period.

Stationarity of the optimal policy implies $V(Q|$ if $\hat{s}_i$ is chosen$) = \frac{C-c_i}{1+\beta_i}$ and is constant in $Q$. $V(0|$ if $\hat{s}_i$ is chosen$) > V(0|$ if $(\hat{s}_i\hat{s}_M)$ is chosen$)$, $V(1|$ if $\hat{s}_i$ is chosen$) < V(1|$ if $(\hat{s}_i\hat{s}_M)$ is chosen$)$, and by Lemma 1.2.2, $V(Q|$ if $(\hat{s}_i\hat{s}_M)$ is chosen$)$ is increasing in $Q$. Hence, there exists a $\bar{Q}$ such that $\hat{s}_i$ is the optimal control for $Q \leq \bar{Q}$ and $(\hat{s}_i, \hat{s}_M)$ the optimal control for $Q \geq \bar{Q}$.

For $\alpha_i \in (a_1, b_1)$, no policy that ever selects only $\hat{s}_i$ in a period is optimal. We can thus restrict our attention to policies that only ever select from $\{\emptyset, \hat{s}_M, (\hat{s}_i, \hat{s}_M)\}$. As before, there exists cutoff $\bar{Q}$ such that $\emptyset$ is optimal at $Q \leq \bar{Q}$ and either $\hat{s}_M$ or $(\hat{s}_i, \hat{s}_M)$ are optimal for $Q \geq \bar{Q}$

There exists another interior cutoff $\widehat{Q}$ such that if $Q > \widehat{Q}$ then $\hat{s}_M$ is optimal. This follows by noting that the value function, when $(\hat{s}_i, \hat{s}_M)$ is optimal, is bounded from above

$$V(Q| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) < A(Q) - c_i - d + \beta_i \frac{\bar{B} - d}{1 - \beta_i}. \tag{1.39}$$

As the optimal policy is stationary, if only $\hat{s}_M$ is ever chosen, it will always be chosen thereafter, yielding a value

$$\begin{aligned} V(Q| \text{ if } \hat{s}_M \text{ is chosen}) &= B(Q) - d + \beta_i V(Q| \text{ if } \hat{s}_M \text{ is chosen}) \\ \iff V(Q| \text{ if } \hat{s}_M \text{ is chosen}) &= \frac{B(Q) - d}{1 - \beta_i} \end{aligned} \tag{1.40}$$

The right side of (1.40) exceeds the right side of (1.39) when

$$B(Q) - A(Q) > -c_i + \frac{\beta_i}{1-\beta_i}(\bar{B} - B(Q)). \tag{1.41}$$

As the above inequality holds at $Q = 1$ for this domain of $\alpha_i$ and $A(Q)$ and $B(Q)$ are continuous in $Q$, the aforementioned cutoff $\widehat{Q}$ exists.

Let us show that for $c_i$ sufficiently high $\widehat{Q} < \bar{Q}$, i.e. the consumer will never experiment, either always selecting $\emptyset$ or $\hat{s}_M$ depending on the relation of $Q^0$ and $\bar{Q}$. If this is true, then the stationarity of the policy implies

$$V(\bar{Q}) = \frac{D}{1-\beta_i} = \frac{B(\bar{Q}) - d}{1-\beta_i} \tag{1.42}$$

Appealing to the upper bound on $V(\bar{Q}|$ if $(\hat{s}_i, \hat{s}_M)$ is chosen)), we can select $c_i$ sufficiently large such that the right side of (1.39) is less both terms in the equality of (1.42).

Next, we show that if $d$ is not too large, there will be an open interval under which the consumer experiments, i.e. $\bar{Q} < \hat{Q}$. Suppose for a contradiction that for all $\alpha_i \in (a_1, b_1)$ and all $d > 0$, $\hat{Q} \leq \bar{Q}$. This implies again that equation (1.42) holds.

For all $\epsilon_1 > 0$, we can choose $\alpha_i$ arbitrarily close to $b_1$, such that for

$$C - c_i < D < C - c_i + \epsilon_1. \tag{1.43}$$

Suppose we modify the policy to select $(\hat{s}_i, \hat{s}_M)$ at $\bar{Q}$. Then the value at $\bar{Q}$ is given by

$$V(\bar{Q}| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) = A(\bar{Q}) - c_i - d + \beta_i E[V(Q')|\bar{Q}, \hat{s}_i, \hat{s}_M] \tag{1.44}$$

We can write

$$E[V(Q')|\bar{Q}, \hat{s}_i, \hat{s}_M] =$$

$$\Pr(Q' > \bar{Q}) \cdot E[V(Q')|Q' > \bar{Q}] + Pr(Q' \leq \bar{Q}) \cdot E[V(Q')|Q' \leq \bar{Q}] \tag{1.45}$$

where all terms in the above equality condition on $\bar{Q}$ and $\hat{s}_i, \hat{s}_M$. As $E[V(Q')|Q' > \bar{Q}] = E[B(Q') - d|Q' > Q] > \frac{B(\bar{Q})-d}{1-\beta_i}$ and $E[V(Q')|Q' \leq \bar{Q}] = \frac{D}{1-\beta_i}$,

$$V(\bar{Q}| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) > A(\bar{Q}) - c_i - d + \beta_i \frac{D}{1-\beta_i} > A(\bar{Q}) - C + D - \epsilon_1 - d + \beta_i \frac{D}{1-\beta_i} \tag{1.46}$$

where the inequality $A(\bar{Q}) - C + C - c_i > A(\bar{Q}) - C + D - \epsilon_1$ follows from (1.43). For $d$ and $\epsilon_1$ sufficiently small, $V(\bar{Q}| \text{ if } (\hat{s}_i, \hat{s}_M) \text{ is chosen})) > V(\bar{Q}| \text{ if } \emptyset \text{ is chosen}))$. Hence, the assumed optimal policy is in fact suboptimal, implying a contradiction.

The proof for "experimentation" at $\alpha_i = b_1$ and $\alpha_i \in (a_2, b_2]$, is nearly identical to the proof above. □

# Chapter 2

# Truth and Conformity on Networks

One of the most important domains of social inquiry is that of broad public discourse. Which social policy will lead to better outcomes? Which political candidate is more qualified for office? Typically, public discussion on such questions of import is influenced by the human tendency of conformity. Individual decisions are informed and influenced by peers; the presence of conformist bias in social discourse is well-studied, and well-supported.[1]

We present a model of social inquiry where it exhibits two properties endemic to matters of public discussion: (1) individuals are subject to varying degrees to conformity bias, and (2) the influence of the pressure to conformity is expressed via social networks. We examine how the structure of social ties in tandem with conformity bias can influence the flow and reliability of information in matters of public opinion.

In our model, heterogeneous agents express public opinions where those expressions are driven by the competing priorities of accuracy and of conformity. Agents learn, by Bayesian conditionalization, from private evidence from nature, and from the public declarations of other agents.

[1]See [7], [23], and [62].

Several key findings emerge. We see that the most influential public declarations are made by agents when they go against the consensus of their neighbors, but that the most informative declarations, on average, are made by agents when their social influences are balanced. This provides a unifying explanation for our results: networks that produce configurations of social relationships that sustain a diversity of opinions empower honest communication and hence reliable acquisition of the truth.

In related literature on network epistemology [89, 90], less connected networks are shown, under the right conditions, to increase the reliability of inquiry. In those cases, greater connectivity can cause premature lock-in to consensus in epistemic communities dealing with an exploration-exploitation trade-off. We arrive at a similar moral by different means.

We show that networks are differentially conducive to informative communication depending on the degree to which a community is divided in its publicly stated opinions. When communities are most divided, more connected networks, such as complete networks, do best. Whereas, when communities are near consensus, less connected networks exhibiting low degree-centrality, such as circle networks, are optimal.

Across the networks literature, star networks have been shown to possess certain optimality properties: they emerge as the product of various processes of strategic network formation [43, 12], can lead to efficient division of cognitive labor [44, 90], and can provide optimal conditions for information dissemination [43]. In contrast, we find that, in the presence of a modicum of conformity bias, star networks produce to the worst possible conditions for social learning.

In §2.2, we explain our model. In §2.3, we present the long run success of learning in the presence of conformist bias. In §2.4, we present simulations illustrating our central results. In §2.5, we provide an analysis of the deeper patterns that unify and explain our results. In §2.6, we conclude.

## 2.2 The Model: Caesar or Pompeia?

To animate our model, let us consider an anachronistic allegorical vignette. A community of Roman citizens has come together to discuss which candidate is better qualified for office. The candidates are Caesar and Pompeia. In discussing their beliefs, the citizens are influenced, to varying degrees, by two competing motivations: the motivation to say, honestly, who they believe is the better qualified candidate, and the motivation to agree with their neighbors, or, more particularly, those with whom they share social or economic ties.

Each citizen varies with respect to the weight she places on each honesty and conformity. On one extreme, we may find Titus the Truth-Teller, who speaks his mind, come what may. Titus has come to believe—both from what is public knowledge, and from his own private information and experiences—that it is Pompeia who is more likely to be a better candidate. And so he declares as much, and he does so without any thought or worry concerning the impact of his declaration on the social regard of his peers.

On the other extreme, we find Cassius the Conformist, for whom harmony with peers is his sole concern. Making his true beliefs known does not enter the picture. Whichever candidate his peers favor, Cassius favors. Now, it happens to be Caesar.

Most of the remaining citizens, however, are not so extreme in their dispositions, but rather fall somewhere between Titus and Cassius. They care about making their true beliefs known, to some degree, and also about harmony with their peers, to some degree. Most make their declarations in a way that is contingent both on the strength of their beliefs, and on the weight of the social pressures around them.

Imagine that there are two states of the world: $\theta$ and $\neg\theta$. We can think of these as corresponding to where one social policy will lead to better outcomes, or one political candidate will be better suited to office.

Agents are interested in learning the true state of the world. This proceeds in two ways: (1) They get private evidence $\sigma \sim f_\theta(\sigma)$ from Nature; we can think of these as hearing some piece of news, or reasoning about an argument. And (2), they observe the public declarations $\mathbf{x}_{-i} \in \{\theta, \neg\theta\}^{N-1}$ of other agents across the network. Declarations indicate to others the state a declaring agent ostensibly believes to be true.

For each agent $i$, her payoffs are a convex combination of her *truth-seeking* orientation $\alpha_i$ and desire for *conformity* to her neighbors $(1-\alpha_i)$. Her payoff for a declaration $x \in \{\theta, \neg\theta\}$ then is given by

$$U_i(x) = \alpha_i P_i(x) + (1-\alpha_i) N_i(x)$$

where $P_i(x)$ is agent $i$'s expectation of the truth of $x$ given her current information, and $N_i(x)$ is the proportion of her neighbors that have also declared $x$.[2]

We can think of an agent $i$ as engaged in two games simultaneously which determine her payoffs in proportion to her type: a Bayesian learning game that contributes $\alpha_i$ of her payoff, where the data are the agent's private evidence $\sigma$ and others' public declarations $\mathbf{x}_{-i}$, and an $n_i$-player pure coordination game that constitutes the remaining $(1-\alpha_i)$ of her payoff, where $n_i$ is the count of agent $i$'s neighbors.

Our epistemic community of $N$ agents inhabits a society where their patterns of shared social influence are described by a network. Here, nodes represent agents, and neighbors are connected by edges. Standard networks include complete, circle, star, and random networks (see FIGURE 2.1).

Networks vary with respect to the patterns of social influence they capture. The complete network describes a social structure in which each agent has social ties with every other. The circle describes a social structure in which each agent shares social ties with exactly two other

[2]Note that an agent's truth-seeking payoff for a declaration is for based on her expectation that it corresponds to the true state of the world—agents do not know, and do not find out, whether their assessments are accurate.

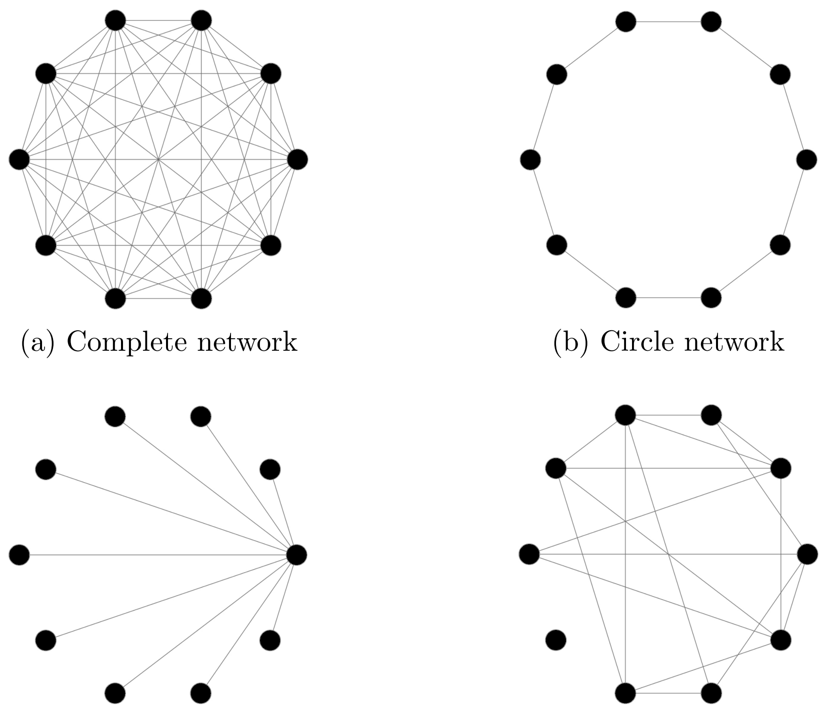

(a) Complete network (b) Circle network

(c) Star network (d) Random network

Figure 2.1: Social networks with 10 agents.

individuals. Note that complete and circle networks are special cases of regular networks,[3] where the regular network is of degree $N-1$ and degree 2, respectively. In contrast, the star network describes a centralized social structure, where one individual (a central agent) has far more connections than the rest (the peripheral agents), who are otherwise socially isolated.

Before the game, agent types (truth-seeking/conformity orientations) are drawn from a continuous distribution:

$$\alpha_1, \ldots, \alpha_N \overset{iid}{\sim} G \text{ with } supp(G) = [0, 1]$$

And Nature randomly chooses the state of the world to be $\theta$ or $\neg\theta$. Each state of the world induces a distinct distribution from which evidence $\sigma \sim f_\theta(\sigma)$ may be drawn.

The distributions $f_\theta(\sigma) = 2\sigma$ and $f_{\neg\theta}(\sigma) = 2 - 2\sigma$, depicted in FIGURE 2.2, are used in our

[3]Regular networks are those in which all nodes are of the same degree, or number of edges. Here, this will correspond to all agents having the same number of neighbors.

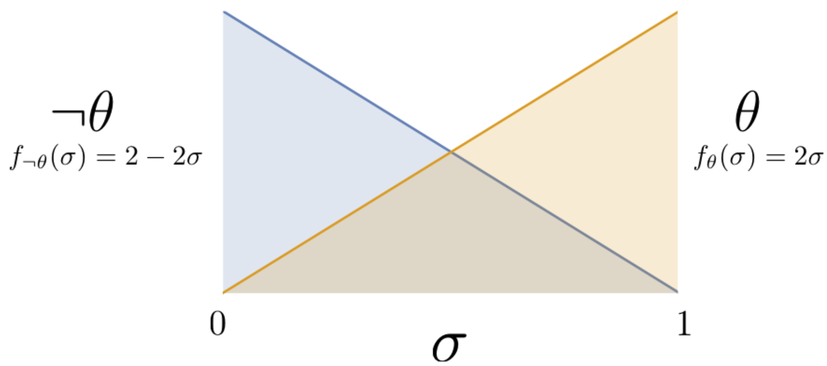


Figure 2.2: Densities $f_\theta(\sigma)$, $f_{\neg\theta}(\sigma)$, of evidence $\sigma$, for each state of the world $\theta$ and $\neg\theta$.

simulations due to their convenient functional form. More generally, however, the distributions need only satisfy: *mutual absolute continuity* and *unbounded evidence*. Mutual absolute continuity requires that both distributions agree on what subsets of possible evidence have positive probability, meaning that no single piece of evidence can falsify one or the other. And unbounded evidence give us that evidence has the potential, in principle, to make one arbitrarily (though not completely) confident of either state. We take this to be a reasonable assumption, as we want to allow that, for any degree of belief shy of absolute certainty, there can—in principle—exist some evidence, however unlikely, which is sufficiently compelling to produce that belief.

In each round, an agent is chosen at random to receive private evidence from Nature, and to make a public declaration to be observed by the network. Upon receiving her evidence, an agent updates her beliefs, via conditionalization, about the true state of the world. This is done in the normal way, using Bayes' rule

$$P(\theta|\sigma, \boldsymbol{h}^t) = \frac{P(\sigma|\theta)P(\theta|\boldsymbol{h}^t)}{P(\sigma|\theta)P(\theta|\boldsymbol{h}^t) + P(\sigma|\neg\theta)P(\neg\theta|\boldsymbol{h}^t)}$$

where $P(\sigma|\theta)$ is the likelihood of her new evidence $\sigma$ given the state $\theta$, and $P(\theta|\boldsymbol{h}^t)$ is her prior on $\theta$ given the history of declarations at that time $\boldsymbol{h}^t$. Note that $P(\theta|\boldsymbol{h}^t)$ is also the public belief at that time—the shared portion of individual beliefs about the true state of the world constituted by the history of learning from public declarations. To simplify exposition, assume the population begins with ignorance priors.[4]

Next, the agent calculates her utilities, given her truth-seeking orientation, chooses her best response as a function of her private evidence and public prior (which, together, form her posterior probability $P(\theta|\sigma, \boldsymbol{h}^t)$ over $\theta$), and the composition of her neighbors:[5]

$$BR_i(\sigma, N_i(x)) = \arg\max\{U_i(\theta), U_i(\neg\theta)\}.$$

Following this, the other agents in the network observe her declaration, and update their beliefs. To do so, they must consider the likelihood of her having made her declaration given the composition of her neighbors, her likely evidence, her possible truth-seeking orientations, and their own prior beliefs about the state of the world.

So, what precisely do agents learn from one another's declarations? Well, when agent $i$ declares $x = \theta$, others know that it was her best response to do so. It follows that

$$\begin{aligned} U_i(\theta) &> U_i(\neg\theta) \\ \alpha_i P_i(\theta) + (1-\alpha)N_i(\theta) &> \alpha_i(1 - P_i(\theta)) + (1-\alpha_i)(1 - N_i(\theta)) \\ \alpha_i(2P_i(\theta) - 1) + (1-\alpha_i)(2N_i(\theta) - 1) &> 0. \end{aligned} \tag{†}$$

We can get an intuitive grasp of this inequality (†) by considering fixed values of the proportion of the declaring agent's neighbors who are also declaring $\theta$ (depicted in FIGURE 2.3.). The shaded area captures the values of agent $i$'s truth-seeking orientation $\alpha_i$ (on the horizontal axis), and posterior belief $P_i$ (on the vertical axis), that are compatible with her having declared $\theta$. That is, the region in which (†) is satisfied.

[4] An *ignorance prior* is a probability distribution assigning equal probability to all possibilities. Our proofs will require only non-degenerate priors, and our simulations will employ a range of priors.

[5] In the case of payoff ties, the agent chooses among her best responses at random.

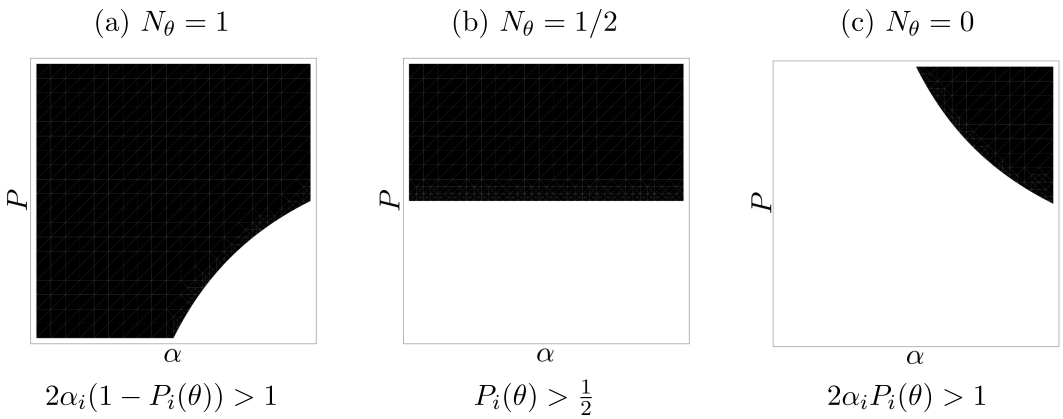


Figure 2.3: What is inferred from agent $i$'s declaration of $\theta$, as captured by condition (†), about her posterior belief $P_i(\theta)$ and truth-seeking orientation $\alpha_i$, when different proportions of her neighbors $N_i(\theta) = 1, 1/2, 0$ are making the same declaration.

Consider the $N_i(\theta) = 1$ case (Figure 2.3a). This is where all of the focal agent's neighbors are also declaring $\theta$. Here, we see that a broad range of beliefs and truth-seeking orientations are compatible with her having declared $\theta$. What can be ruled out (the area in white) is that it was *not* the case that she was both highly truth-seeking and strongly believed in the truth of $\theta$. Here, others do not learn much from observing the focal agent's declaration.

Consider the $N_i(\theta) = 1/2$ case (Figure 2.3b). This is where the focal agent's neighbors are evenly split; half declaring $\theta$ and half $\neg\theta$. Here, the other agents infer that the focal agent's social influences are balanced, and so her truth-seeking orientation $\alpha_i$ is no longer relevant. Her declaration is now determined purely by her posterior belief. If $P_i(\theta) > 1/2$, then she would make the declaration she did, if not, she would not. Here, others learn the direction of the focal agent's belief, but not much about its strength.

Next, consider the $N_i(\theta) = 0$ case (Figure 2.3c). This is where none of the focal agent's neighbors are declaring $\theta$. Here, we see that only a narrow range of beliefs and truth-seeking orientations are compatible with her declaration of $\theta$. It must have been the case that she was both highly truth-seeking, and possess a strong belief in the truth of $\theta$. Now, others learn both the direction and strength of the focal agent's belief, and through it about the

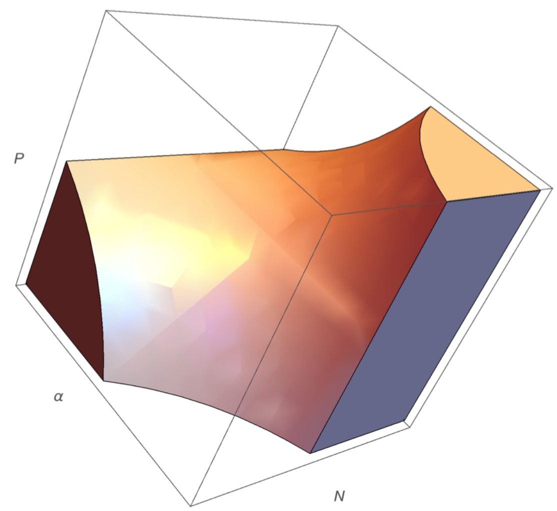


Figure 2.4: The content of an agent's declaration, visualized as a hyperboloid—corresponding to (†)—of the values of $N_\theta$, $\alpha_i$, and $P_i(\theta)$ for which she could have made her declaration.

strength of her evidence.

What happens when we put all this together? The hyperboloid in FIGURE 2.4 gives us the delimited domain of the likelihood of an agent's declaration. This captures the reasoning we just covered as to the qualitative inferences agents make about one another's beliefs from their declarations. From this, the agents in the population update their beliefs about the state of the world, in the normal way, using Bayes' rule. (See APPENDIX A for the mathematical details.)

In this way, rational agents learn from their own private evidence, the declarations of other agents in the network, and the public belief about the true state evolves through discussion and across the network.

## 2.3 Truth in the Long Run

Our primary interest lies in dynamical analysis of the short-to-medium-run behavior of social inquiry under conformist bias. Before we proceed to this analysis, however, it may help us in this to understand the long-run trajectory of social learning under conformity. What we find is that, in the long run, irrespective of social structure or conformist bias, epistemic communities like the ones we have described will converge to believing in, and publicly declaring, the true state of the world.

More precisely, given any social network, unbounded evidence, and the possibility of sufficiently truth-seeking agents $1 \in \text{supp}(\alpha)$, a community of Bayesian learners will, with probability one, converge to knowing and declaring the truth in the long run. This is captured by the following proposition and its corollary.

**Proposition 2.1.** *An epistemic community learning about the state of the world will, in the long run, converge in belief to the true state.*[6]

**Corollary 2.1.** *For such an epistemic community, converging in belief to the true state implies converging to consensus in declaring the true state.*

Convergence in beliefs follows from the fact that our agents learn via Bayesian conditioning, that the true state is contained in each agent's hypothesis set, and that agent declarations are always to some degree informative as to the state of the world. Given this, classical convergence results for Bayesian learning[7] guarantee long run acquisition of the truth.

Convergence in declarations follows from the fact that, given convergence in beliefs, the community's beliefs will inevitably pass a threshold such that a consensus on declaring the true state cannot be escaped. Moreover, with enough time following the passing of this threshold of belief, the population will almost surely traverse a positive probability path to consensus on the true state, whereupon it will never leave this consensus.

It may well be that "in the long run we are all dead," [Keynes, 1923, p. 80] but it can be helpful to confirm where we are headed. We have seen that our epistemic communities will

[6] All proofs can be found in Appendix A.

[7] For an excellent exposition of the classic results, see [**?** ].

arrive at the truth in the limit of time, so we turn to short and medium run analysis of social learning for a richer and more pressing picture of inquiry.

## 2.4 Truth and Conformity in the Short and Medium Run

What can be said about the short and medium run behavior of learning under conformity? What role does social structure play in the reliable acquisition of true beliefs? To answer these questions, we ran simulations of our model of epistemic communities engaged in social learning and discourse. We recorded and analyzed the resulting behavior over a parameter sweep of network types, population sizes, initial declarations, prior beliefs, and distributions of the individuals' truth-seeking and conformity orientations.

For the simulations, we varied the structure of social influences by placing our agents on each complete, regular (of degree $N/2$), circle, star, and random (of mean degree $N/2$) networks. We varied the number of agents $N$ in the network from 2 agents (at which all networks are essentially identical) to 20 agents. We considered when the initial declarations of the society were at a consensus on the true state, a consensus on the false state, and an even split. We varied the shared prior beliefs of the population between relative confidence in the true state ($P(\theta) = 0.75$), skepticism toward the true state ($P(\theta) = 0.25$), and ambivalence about the true state ($P(\theta) = 0.5$). Each combination of network structure, population size, initial declarations, and prior beliefs composed one parameter setting.

For each parameter setting we ran 10,000 simulations where each simulation was composed of 100 turns, and where each turn consisted of the following phases: (1) a randomly selected agent receives her private evidence from Nature; (2) the agent updates her private belief in light of this evidence; (3) the agent chooses her best response given her beliefs, her

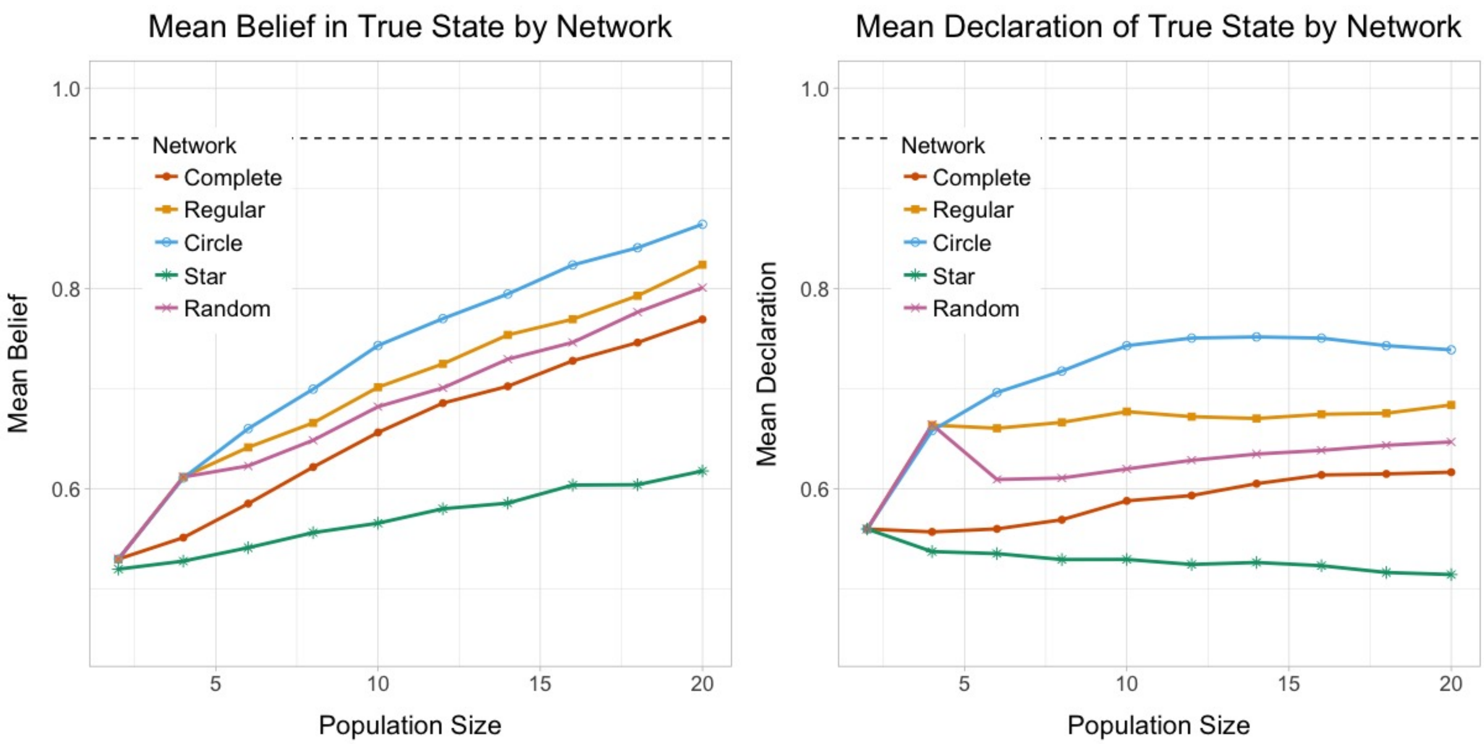


(a) (b)

Figure 2.5: Plots of the mean belief in the true state $P(\theta)$ (A), and declaration of the true state $\theta$ (B), for each network type, and for network sizes from 2 agents to 20. Note that the networks only become fully distinct at $N = 6$. The dashed line represents performance in total absence of any conformity.

neighbors' declarations, and her truth-seeking/conformity orientation; (4) the agent makes her declaration to the network; (5) the other agents in the network update their beliefs in light of her declaration.

Three regularities readily emerged from the data (see FIGURES 2.5A, 2.5B): (1) In all simulations, the star network performed worse than all other standard networks in terms of generating reliable belief in, and declaration of, the true state. (2) The circle network, on the other hand, performed better than other standard networks on all counts. (3) The other networks—complete, regular (of degree $N/2$), and random (of mean degree $N/2$)—yielded middling performances, neither as good as the circle, nor as poor as the star, with the regular network typically outperforming the random network, and the random network outperforming the complete network.[8]

To make sense of these regularities in our simulation results, analytic treatment of the model and its dynamics is needed. What should be obvious is that conformity bias muddies the waters with respect to the information content of individuals' declarations. In the absence of conformity, our epistemic communities would rapidly and reliably acquire the truth, and the underlying network structure would make no difference to this learning.

What we will find is that different networks induce social configurations more or less conducive to honest communication, and that this will also depend on the degree to which the population is divided or unified in their public declarations.

## 2.5 Influence, Information, and Social Structure

To understand why different social networks are more or less conducive to the reliable acquisition of true beliefs, we first need a measure of informativeness. For this, we introduce the concepts of influence and informativeness of declarations, and show how they are related.

We define the *influence* of a declaration $x \in \{\theta, \neg\theta\}$ as the difference between the public belief in $x$ before and after its declaration to the network, $q(x|x) - q(x)$, where $q$ is the public belief. Next, we define the *informativeness* of a declaration $x \in \{\theta, \neg\theta\}$, as the reduction in uncertainty it produces with respect to its corresponding state when starting from a maximal entropy prior, $H(q|q(x) = 1/2) - H(q|x)$, where $H$ is the Shannon entropy function.

We now derive the fact that the informativeness of a declaration is monotonically increasing in its influence on the public belief (see Lemma 2.1 in Appendix A). This gives us that a

[8]In our simulation plots (FIGURE 2.5), we mark the performance of learning in the absence of any conformity bias—that is, of unimpeded Bayesian learning—with a dashed line. We will continue to compare our results to this control case, denoting the case of learning in the absence of conformity bias in further plots (FIGURE 2.6, 2.7, 2.8) each time with a dashed line.

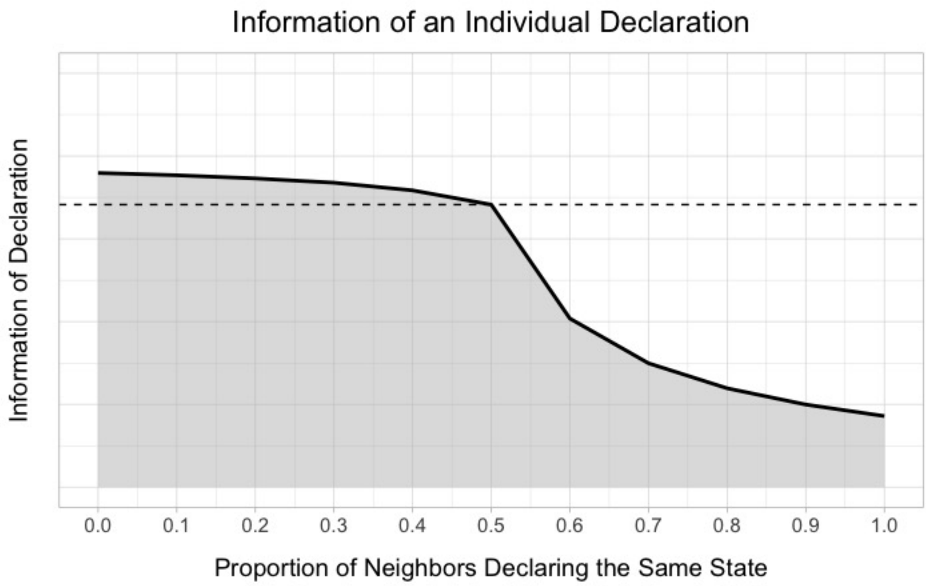


Figure 2.6: The influence and informativeness of an agent's declaration, as a function of the proportion of her neighbor's who are declaring the same state.

declaration will be (minimally) maximally influential just in case it is (minimally) maximally informative. We will use this fact repeatedly to infer the relative informativeness of declarations from their influence.

## Optimal Information From Going Against the Grain

Given our measures of influence and informativeness, our first insight follows straightforwardly from our model of agents learning via Bayesian conditioning under uncertainty about one another's evidence and truth-seeking orientations. It is that the most informative declarations—those that have the most significant effect on the public belief—are those that "go against the grain." That is, those made by agents exactly when they deviate from the consensus of their neighbors.

This insight is captured by the following proposition:

**Proposition 2.2.** *The informativeness of an agent's declaration is monotonically increasing in proportion of her neighbors who are declaring the opposing state.*

And since the minimum proportion of an agent's neighbors who may declare in favor of any state is zero, we have the following as an immediate corollary:

**Corollary 2.2.** *The most informative declaration in favor of a state is one made by an agent when she goes against the consensus of her neighbors.*

This corresponds to the case in FIGURE 2.3, where $N_i(\theta) = 0$, and is visualized by the plot of information of declarations in FIGURE 2.6 where we see the change in belief by the population in response to an agent's declaration as a function of the proportion of that agent's neighbors who are declaring the same state.

When an agent deviates from the consensus of her immediate peers, it is inferred by the broader network that she is both likely to be more truth-seeking and that she has received sufficiently strong evidence to justify the loss in social payoffs she incurred. No other declaration is more influential on the public belief.

## Optimal Expected Information From Conflicted Neighbors

We have seen what that the most informative declarations occur when an agent goes against the consensus of her peers. But such declarations are rare, as it takes highly truth-seeking agents with good evidence to be willing to make them. We should ask then: under what conditions, *on average*, do we expect to find the most informative declarations?

These turns out to be the obverse of where we find the most influential declaration. The most informative declarations, on average, must come from individuals whose neighbors are perfectly divided in terms of their declarations.

This is captured by the following observation:

**Observation 2.1** (Observation 1)**.** *The most influential and informative declaration, in*

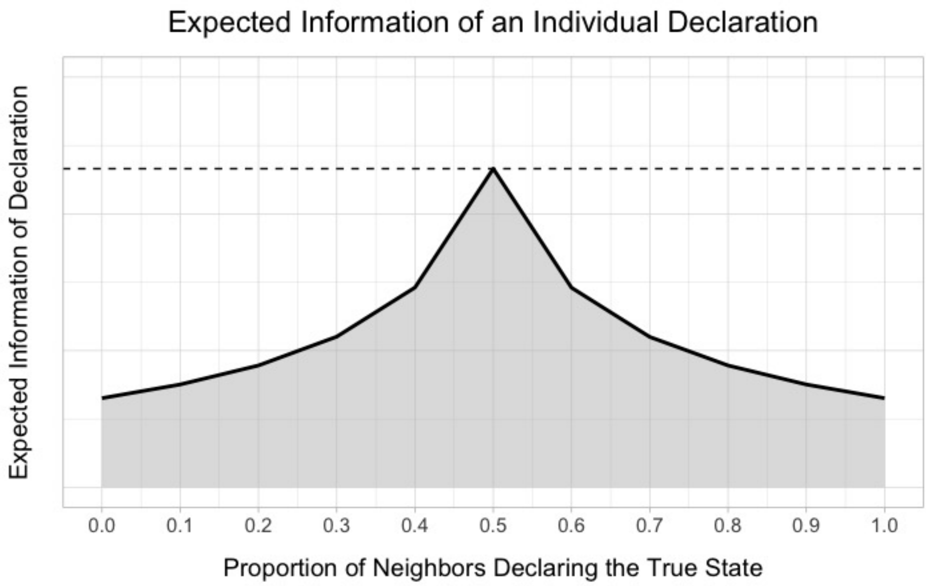


Figure 2.7: The *expected* influence and informativeness of an agent's declaration, as a function of the proportion of her neighbor's who are declaring the same state.

*expectation, is that made by an agent when her neighbors are evenly divided in their declarations.*[9]

**Observation 2.2** (Observation 2)**.** *The expected information of declarations is convex and increasing for* $N_i(\theta) \in (0, 1/2)$ *and convex and decreasing for* $N_i(\theta) \in (1/2, 1)$*.*

This corresponds to the case in Figure 2.3, where $N_i(\theta) = 1/2$, and is visualized by the plot of expected information of declarations in Figure 2.7. In Figure 2.7 we see the expected change in belief of the population in response to an agent's declaration, as a function of the proportion of that agent's neighbors who are declaring the true state. Our propositions make use of these observations.

It is when an agent's social influences equally represent each viable position that she is most free to declare her honest belief, and in such a case others infer that she is most likely doing

[9] Our observations are computationally verified for the following distributions of types and evidence: the distribution of truth-seeking orientations in the population was varied from Beta(1,5) (corresponding to high conformism), to uniform, and Beta(5,1) (corresponding to high truth-seeking). And the distributions of evidence induced by each state of the world were varied between the linear case described before, and Gaussian distributions with means of 1 and -1, and variances of 1, 10, and 100.

so.

### 2.5.1 Informativeness of Networks

Which networks then are most conducive to the social configurations that yield honest communication? Using the insights developed so far, we extend the concept of expected informativeness to the level of social networks.

Assume that $\theta$ is the true state of the world, then *expected influence* of declarations $X = \{\theta, \neg\theta\}$ for an $N$-agent network $\mathcal{G}$ is given by

$$E_X[q(\theta|\mathcal{G}) - q(\theta)] \propto \sum_{k=0}^{N} \sum_{j=1}^{\binom{N}{k}} \sum_{i=1}^{N} E_X[q(\theta|x_i) - q(\theta)]$$

where the first sum is over the number of the agents in the network declaring the true state, the second sum is over the possible configurations of declarations in the network given the number of agents declaring the true state, and the third sum is over the individuals in the network.[10] In this way, we infer the informativeness of a network in aggregate as well as for fixed proportions of the community declaring the true state.

With a generalized measure of expected informativeness, we compute the expected informativeness of 10-agent networks for different proportions of the population declaring the true state (see FIGURE 2.8).

From this, several observations emerge. Denote the proportion of the community declaring $\theta$ by $N_\theta$. For all networks, then, the least informative state is that of consensus, $N_\theta = 0$ or 1, and the most informative state is when there is an even split in declarations $N_\theta = 1/2$. Given

[10]Note that we suppress the normalizing term from the definition of the influence of a declaration. The reason for this is that both terms are constants, and are therefore irrelevant for determining maxima or minima.

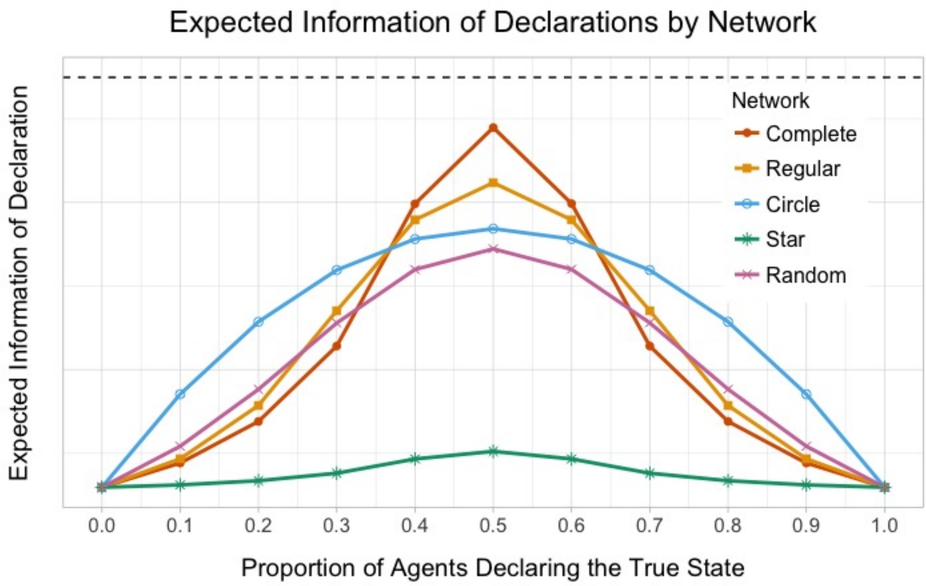


Figure 2.8: The expected informativeness of the next declaration for 10-agent networks as a function of the proportion of the population which is declaring of the true state. The dashed line denotes the expected information in the absence of any conformity.

Observation 1, it should be clear why this is so. Declarations are expected to be informative in measure to the presence of balanced dissent.

Next, we observe that, when the population is nearly split, the complete network produces the most informative declarations among the networks considered, while the circle network produces the most informative declarations when the population is near consensus. Finally, the star network provides the least informative social configuration no matter the proportion of the population making either declaration.

We may understand these results in terms of our previous insights, and sharpen them by considering large networks. On a star network, when the population is large, practically every individual has merely one neighbor. Hence, for any proportion of declarations in the population, the star network will be in the minimally informative state. That is, $\boldsymbol{I}(\mathcal{G}_{star}|N_\theta) = N_\theta \boldsymbol{I}(1) + (1 - N_\theta)\boldsymbol{I}(0) = \boldsymbol{I}(0)$.[11]

[11] Given the assumption of symmetry of expected informativeness across $N_\theta = 1/2$, we have that $\boldsymbol{I}(0) = \boldsymbol{I}(1)$, and, more generally, that $\boldsymbol{I}(1/2 - c) = \boldsymbol{I}(1/2 + c)$ for all $c \in [0, 1/2]$.

**Proposition 2.3.** *For large networks, the star network is minimally informative in any state.*

On a complete network, when the population is evenly divided, $N_\theta = 1/2$, each individual is in the optimal position to make informative declarations. When all individuals are neighbors and the population is sufficiently large, the expected informativeness of the network as a whole recapitulates the expected informativeness of individual declarations given in FIGURE 2.7. That is, $\boldsymbol{I}(\mathcal{G}_{complete}|N_\theta) = \boldsymbol{I}(N_\theta)$. Given Observation 1, we show that no network can be more informative in such a state.

**Proposition 2.4.** *For large networks, when the population is evenly split in declarations, the complete network is maximally informative.*

On a circle network, when the population is near consensus, a single dissenting individual can make it possible for both her their neighbors to declare their honest beliefs. That is, given that each individual has two neighbors, their neighbors' declarations are binomially distributed with the success parameter given by the population proportion of declarations, $\boldsymbol{I}(\mathcal{G}_{circle}|N_\theta) = N_\theta^2\boldsymbol{I}(0)+2N_\theta(1-N_\theta)\boldsymbol{I}(1/2)+(1-N_\theta)^2\boldsymbol{I}(1)$. Contrast this with the complete network, where near consensus, every individual faces strong incentives to conform.

**Proposition 2.5.** *For large connected networks, for a range of states near consensus in declarations, the circle network is the maximally informative network.*

More generally, we can express the expected informativeness of the declaration of any individual with $d$ connections and proportion $N_\theta$ of her neighbors declaring the true state as

$$E_{N_\theta}[\boldsymbol{I}_d] = \sum_{k=0}^{d} \binom{d}{k} N_\theta^k (1 - N_\theta)^{d-k} \boldsymbol{I}\left(\frac{k}{d}\right). \tag{$*$}$$

From this, we can derive the informativeness of a any network, when we conceive of networks as admixtures of proportions of individuals with different numbers of neighbors.

Given any large network, it can be represented as a distribution $\mu = \langle \mu_d \rangle$ over the degree $d$ of individuals within the network, Thus, the expected informativeness of the network will be $\boldsymbol{I}(\mathcal{G}_\mu | N_\theta) = \sum_d \mu_d \cdot E_{N_\theta}[\boldsymbol{I}_d]$. Using this, we provide bounds for the informativeness of epistemic networks near consensus.

**Proposition 2.6.** *For large networks, near consensus, any network (including any regular or random network) of minimum degree at least two will be intermediate in informativeness between the circle and complete network.*

## 2.6 Conclusion

When social learning proceeds under the influence of conformity bias, the structure of social relationships underpinning the epistemic community becomes crucial to reliable acquisition of truth. That disagreement and diversity in publicly held opinions can be optimal for honest communication gives us our key insight into understanding the effects of different social networks. The question as to which social networks lead to reliable beliefs becomes a question as to which social networks produce and sustain optimal patterns of disagreement throughout the process of learning.

In sum, we find that in the presence of even a modicum of conformity bias the star network always provides the worst conditions for informative communication, the complete network provides optimal conditions exactly when the population is evenly divided, the circle network provides optimal conditions near consensus, and that, in such a state, all sufficiently connected networks will be intermediate in informativeness between the circle and complete networks.

This has implications for real-world social networks, which tend to exhibit low average degree and high degree-centrality [85]. We may conjecture that, when we suspect conformity bias at play in social discourse and decision-making, interventions which reduce the density of connections of a social network while still keeping it connected, and interventions which decrease its centralization by reducing the relative influence of central individuals, may lead to more informative communication—and so to more reliable beliefs—for the epistemic community as a whole.

## 2.7 Mathematical Appendix

### Learning from others' declarations

When agent $i$ declares $x = \theta$, we know that it was her best response. As previously mentioned, this implies that the following condition holds:

$$\alpha_i(2P_i(\theta) - 1) + (1 - \alpha_i)(2N_i(\theta) - 1) > 0. \tag{$\dagger$}$$

We plug agent $i$'s (publicly unknown) posterior probability $P(\theta|\sigma)$ into (†) to get the elaborated condition

$$\alpha_i \left( \frac{2}{1 + \frac{1-\bar{P}}{\bar{P}} \frac{1-\sigma}{\sigma}} - 1 \right) + (1 - \alpha_i)(2N_i(\theta) - 1) > 0 \tag{$\ddagger$}$$

where $\bar{P}$ denotes the (publicly known) prior $P(\theta|\boldsymbol{h}^t)$. We then compute the likelihood of agent $i$'s declaration $\theta$, given our public prior, as follows.

Let $\phi$ denote the left-hand term of our elaborated condition (‡), under which our agent would have declared $\theta$, so that $\mathbb{I}[\phi > 0]$ is its indicator function. We then get the likelihood of the declaration given each possible state of the world,

$$P(x=\theta|\theta,\bar{P}) = \int_A \int_\Sigma \mathbb{I}[\phi>0]\mathrm{d}F_\theta(\sigma)\mathrm{d}G(\alpha),$$
$$P(x=\theta|\neg\theta,\bar{P}) = \int_A \int_\Sigma \mathbb{I}[\phi>0]\mathrm{d}F_{\neg\theta}(\sigma)\mathrm{d}G(\alpha).$$

From these, we obtain the posterior—the belief of the other agents in the network in light of agent $i$'s declaration of $\theta$—using Bayes' rule as follows

$$P(\theta|x=\theta,\bar{P}) = \left(1+\frac{\int_A\int_\Sigma \mathbb{I}[\phi>0]\mathrm{d}F_{\neg\theta}(\sigma)\mathrm{d}G(\alpha)}{\int_A\int_\Sigma \mathbb{I}[\phi>0]\mathrm{d}F_\theta(\sigma)\mathrm{d}G(\alpha)}\frac{1-\bar{P}}{\bar{P}}\right)^{-1}$$

which yields the new public belief.

*Proof of Proposition 2.1.* There are two states of the world $\theta$ and $\neg\theta$. Without loss of generality, suppose $\theta$ to be the true state of the world. Let $q(\boldsymbol{h}^t) = P(\theta|\boldsymbol{h}^t)$ be the *public belief* and $\boldsymbol{h}^t$ the history of declarations up to time $t$. As is well-known, the likelihood ratio

$$\ell(\boldsymbol{h}^t) \equiv \frac{1-q(\boldsymbol{h}^t)}{q(\boldsymbol{h}^t)}$$

is a martingale conditional on $\theta$. Let $X$ be the finite set of declarations. For any given declaration $x \in X$,

$$\ell(\boldsymbol{h}^t, x) = \ell(\boldsymbol{h}^t)\frac{P(x|\boldsymbol{h}^t,\neg\theta)}{P(x|\boldsymbol{h}^t,\theta)}$$

and thus the martingale property follows:

$$E\big[\ell(\boldsymbol{h}^{t+1})|\theta\big] = \sum_{x\in X}\ell(\boldsymbol{h}^t,x)P(x|\boldsymbol{h}^t,\theta) = \sum_{x\in X}\ell(\boldsymbol{h}^t)P(x|\boldsymbol{h}^t,\neg\theta) = \ell(\boldsymbol{h}^t).$$

By Theorem 3(b) of [**?** ], when evidence is unbounded, individuals almost surely converge in belief to the true state. □

We show that convergence in beliefs implies a convergence in declarations. In particular, we show that convergence in beliefs implies that the community's belief in the true state will be bounded from bellow over time. We then observe, using simple probabilistic arguments, that given sufficient time the community will almost surely arrive at a consensus state where all individuals are declaring the true state. Finally, we show that, having arrived at such a consensus with individual beliefs in the true state appropriately bounded from bellow, the community must remain at this consensus forever.

*Proof of Corollary 2.1.* Let $q$ and $q'$ denote the public belief before and after hearing a declaration, respectively. Consider a focal agent $i$ having received her evidence from Nature on a given turn. Let $P_i$ denote the focal agent's posterior belief $P(\theta|\sigma, \boldsymbol{h}^t)$, and suppose that this agent declared $x = \neg\theta$. It is straightforward to show that if the population could observe the focal agent's posterior, the public belief would be precisely equal to her posterior

$$q'(\neg\theta, q, N_i(\theta), P_i) = P_i. \tag{*}$$

Let $\Pi(\cdot|\neg\theta, q, N_\theta)$ be the distribution over the focal agent's posterior belief given her declaration of $\neg\theta$, $q$ the public belief when she selected her action, and $N_i(\theta)$ the proportion of her neighbors declaring $\theta$. By (*) we can write

$$q'(\neg\theta, q, N_i(\theta)) = \int_0^1 P_i \, \mathrm{d}\Pi(P_i|\neg\theta, q, N_i(\theta)).$$

We can thus interpret the public belief as the public's expectations of the focal agent's posterior. As the public belief almost surely converges to certainty on the truth, for almost all trajectories of the public belief $\{q_t\}_{t=0}^{+\infty}$, for all $\epsilon > 0$ there exists a time $T_\epsilon$ such that, if $t > T_\epsilon$ then $q_t > 1 - \epsilon$. That is, there is a time after which the public belief in $\theta$ will always be at least $1 - \epsilon$. Then choose $\epsilon = 1/2$.

With probability 1 at some point along the trajectory after $T_\epsilon$ all agents will be declaring $\theta$. To see this, let $\lambda$ be the probability all $N$ agents choose declarations in sequence, each has an $\alpha$ sufficiently high such that they declare the state they believe to be more likely regardless of their neighbors' declarations, and they receive evidence such that their posterior assigns higher probability on $\theta$. However small the probability $\lambda$ might be, it exceeds 0. Hence, the probability that this event does not occur goes to zero as $t \to +\infty$.

Assume, for the sake of contradiction, that at some point after $T_\epsilon$ an agent goes against the consensus and declares $\neg\theta$, then her posterior must satisfy

$$P_i \leq -\frac{1-\alpha_i}{2\alpha_i} + \frac{1}{2}.$$

But then we get that $E[P_i|\neg\theta, \cdot] \leq 1/2$. That is, her belief in $\theta$ was less than 1/2, which contradicts the fact that her belief was bounded from bellow. Hence, no agent can deviate from the consensus after time $T_\epsilon$, and convergence in belief implies convergence in declaration.

□

**Lemma 2.1** (Monotonicity of Informativeness in Influence)**.** *The informativeness of a declaration about a state is monotonically increasing in its influence on the public belief.*

*Proof.* Without loss of generality, let the focal agent declare $x = \theta$. We show that the informativeness of her declaration, $H(q|q(\theta) = 1/2) - H(q|x = \theta))$, is monotonically increasing in its influence, $q(\theta|x = \theta) - q(\theta)$.

First, we unpack the definition of informativeness, temporarily omitting the assumption of the maximal entropy prior $q(\theta) = 1/2$, to get

$$\begin{aligned} H(q) - H(q(\theta|x=\theta) &= \mathrm{E}[-\ln(q(\theta|x=\theta))] - \mathrm{E}[-\ln(q(\theta))] \\ &= \mathrm{E}[\ln(q(\theta)) - \ln(q(\theta|x=\theta))] \\ &= \mathrm{E}\left[\ln\left(\frac{q(\theta)}{q(\theta|x=\theta)}\right)\right] \\ &= q(\theta)\cdot\ln\left(\frac{q(\theta)}{q(\theta|x=\theta)}\right) + q(\neg\theta)\cdot\ln\left(\frac{q(\neg\theta)}{q(\neg\theta|x=\theta)}\right) \end{aligned}$$

Now, let $A \equiv q(\theta)$ and $B \equiv q(\theta|x=\theta)$, so that $C \equiv B - A$ denotes the influence of the declaration $x = \theta$. Then we can re-write the preceding expression as

$$A\cdot\ln\left(\frac{A}{A+C}\right) + (1-A)\cdot\ln\left(\frac{1-A}{1-(A+C)}\right)$$

Taking the partial derivative with respect to influence $C$, and solving for when it is positive—i.e., for when informativeness is increasing—yields

$$A + C - 1 > 0 \quad \text{or} \quad B > 1/2.$$

And when $q(\theta) = 1/2$, we have that $B = q(\theta|x=\theta) \geq 1/2$, and so informativeness is monotonically increasing in influence, as desired. □

We will show that $q'(\theta, N_i(\theta)') < q'(\theta, N_i(\theta))$ whenever $N_i(\theta)' > N_i(\theta)$. From this it follows straightforwardly that, given $N_i(\theta) \in [0,1]$, the most influential declaration occurs just when $N_i(\theta) = 0$.

To do so, consider a given focal agent $i$ having received evidence $\sigma \sim f_\theta(\sigma)$ from Nature. Let $r = r(\sigma) \equiv P_i(\neg\theta|\sigma)$ be one minus her private belief, $G_{\neg\theta}(r)$ and $G_\theta(r)$ the conditional cdf's for $r$, and $g(r) \equiv \frac{dG_{\neg\theta}}{dG_\theta}(r)$ the Radon-Nikodym derivative of $G_{\neg\theta}$ with respect to $G_\theta$.

**Lemma 2.2.** $g(r) = \frac{r}{1-r}$ *almost surely.*

*Proof.* If an agent updates her belief after observing $r$, it will remain unchanged. Thus from Bayes' theorem $r = P_i(\neg\theta|r) = \frac{g(r)}{g(r)+1}$. □

**Lemma 2.3.** *The ratio $\frac{G_{\neg\theta}}{G_\theta}(r)$ is strictly increasing for $r$ in the common support of $G_\theta$ and $G_{\neg\theta}$.*

*Proof.* Let $r' > r$. From Lemma 2.2 we have that $g(r)$ is strictly increasing, hence,

$$G_{\neg\theta}(r) = \int_0^r g(x)dG_\theta(x) < g(r)G_\theta(r)$$

And thus

$$\begin{aligned} G_{\neg\theta}(r') - G_{\neg\theta}(r) &= \int_r^{r'} g(x)dG_\theta(x). \\ &> [G_\theta(r') - G_{\neg\theta}(r)]g(r) \\ &> [G_\theta(r') - G_{\neg\theta}(r)]\frac{G_{\neg\theta}(r)}{G_\theta(r)}. \end{aligned}$$

It follows that $\frac{G_{\neg\theta}(r')}{G_\theta(r')} > \frac{G_{\neg\theta}(r)}{G_\theta(r)}$. □

*Proof of Proposition 2.2.* Now, we proceed to show that $q'(\theta, N_i(\theta)') < q'(\theta, N_i(\theta))$ whenever $N_i(\theta)' > N_i(\theta)$. Define $q'$ to be the posterior public belief, $q$ the prior public belief, $N_i(\theta)$ the proportion of the focal agent's neighbors declaring $\theta$, and $\Pi(\cdot|x_i, q, N_i(\theta))$ the posterior belief over the declaring agent's truth-seeking orientation $\alpha_i \in [0,1]$. Then

$$q'(\theta, N_i(\theta)) = \int_0^1 q'(\theta, N_i(\theta), \alpha_i)\mathrm{d}\Pi(\alpha_i|\theta, N_i(\theta), q).$$

For a given $\alpha_i$ in the support of $\Pi(\cdot|\theta, N_i(\theta), q)$, there exists a threshold $\bar{r} = \bar{r}(\alpha_i, q, N_i(\theta))$

such that the agent only selects $x_i = \theta$ if $r \leq \bar{r}$. From Bayes' theorem,

$$q'(\theta, N_i(\theta), \alpha_i) = \left(1 + \frac{1-q}{q}\frac{G_{\neg\theta}(\bar{r})}{G_\theta(\bar{r})}\right)^{-1}.$$

If $\bar{r}(\alpha_i, N_i(\theta)', q) \geq \bar{r}(\alpha_i, N_i(\theta), q)$ holds, and further holds strictly for a subset of $\alpha_i$ with positive posterior probability, then, by Lemma 2.3, $q'(\theta, N_i(\theta)') < q'(\theta, N_i(\theta))$.

It can be shown that the threshold $\bar{r}(\alpha_i, N_i(\theta), q)$ is strictly increasing in $N_i(\theta)$. This gives us that $q'(\theta, N_i(\theta)', \alpha_i) \leq q'(\theta, N_i(\theta), \alpha_i)$. Furthermore, having assumed that $\alpha_i$ and $r$ take full support in $[0, 1]$, we can find a neighborhood of $\alpha_i = 1$ with positive probability such that $\bar{r}(\alpha_i, N_i(\theta), q) > 0$ for all $\alpha_i$ in this neighborhood. Hence, in this neighborhood $q'(\theta, N_i(\theta)', \alpha_i) < q'(\theta, N_i(\theta), \alpha_i)$. □

*Proof of Corollary 2.2.* We have, from proposition 2.2, that $q'(\theta, N_i(\theta)') > q'(\theta, N_i(\theta))$ whenever $N_i(\theta)' < N_i(\theta)$. It follows directly that

$$\arg \max_{N_i(\theta) \in [0,1]} q'(\theta, N_i(\theta)) = 0.$$

Thus, the most influential declaration is made just when $N_i(\theta) = 0$. And we have, from Lemma 2.1, that this is also the most informative declaration. □

*Proof of Proposition 2.3.* On a large star network, proportion 1 of individuals have a single neighbor. Thus for any proportion of the population declaring $\theta$, every individual is in the minimally informative state where either $N_\theta = 0$ or 1. Therefore, for all $N_\theta \in [0, 1]$, and symmetric $\boldsymbol{I}$, $\boldsymbol{I}(\mathcal{G}_{star}) = \boldsymbol{I}(0) \leq \boldsymbol{I}(\mathcal{G})$ for any connected network $\mathcal{G}$. □

*Proof of Proposition 2.4.* On a complete network, every individual individual is neighbors with every other. Hence, the proportions of an individuals neighbors declaring $\theta$ is the same as the proportion of the population declaring $\theta$. The expected informativeness is maximized when an individual's neighbors are equally split $N_i(\theta) = 1/2$. Thus, when exactly half the population is declaring $\theta$, the declaration of every individual in the population is at maximal expected informativeness. Hence, no other network can be more informative in this state. That is, when $N_\theta = 1/2$, $\boldsymbol{I}(\mathcal{G}_{complete}) = \boldsymbol{I}(1/2) \geq \boldsymbol{I}(\mathcal{G})$ for any connected network $\mathcal{G}$. □

To show that the circle is maximally informative near consensus, first we show that for regular networks of degree at least 2 informativeness is decreasing in degree near consensus. This implies that any regular network of degree greater than two is less informative than the circle network. We combine this with Proposition 2.3, which implies that networks of degree 1 are also less informative than the circle network, to show that the circle network is the maximally informative regular network. Next, using the fact that any network can be formulated as an admixture of individuals of various degrees we derive that the circle network is maximally informative near consensus.

**Lemma 2.4.** *For regular networks of degree at least 2, informativeness is decreasing in degree near consensus.*

*Proof.* Take the derivative of the informativeness of any regular network $\mathcal{G}_d$ of degree $d \geq 2$ with respect to the proportion of the population declaring the true state.

$$\frac{d}{dN_\theta}\left[\boldsymbol{I}(\mathcal{G}_d)\right] = \frac{d}{dN_\theta}\left[\sum_{k=0}^{d}\binom{d}{k}N_\theta^k(1-N_\theta)^{d-k}\boldsymbol{I}\left(\frac{k}{d}\right)\right].$$

Let $N_\theta$ go to 0. This makes it so only the constant terms of the derivative remain, and the expression simplifies to

$$\lim_{N_\theta \to 0^+}\frac{d}{dN_\theta}\left[\boldsymbol{I}(\mathcal{G}_d)\right] = d[\boldsymbol{I}(1/d) - \boldsymbol{I}(0)].$$

This term corresponds to the slope of the secant line connecting $\boldsymbol{I}(0)$ and $\boldsymbol{I}(1/d)$. Since $\boldsymbol{I}$ is an increasing function, this term must be decreasing in $d$. Thus, for networks of degree two and greater, informativeness is decreasing in degree near consensus. □

**Lemma 2.5.** *The circle is the maximally informative regular network near consensus.*

*Proof.* This follows from Lemma 2.4 and Proposition 2.3, which state that a regular network of degree 2 (the circle) is more informative than any network of greater degree near consensus, and that a regular network of degree 1 is less informative than any other at any state. Taken together, they imply that, near consensus, regular networks of degree two are maximally informative among regular networks. □

*Proof of Proposition 2.5.* Now, recall that any large connected network $\mathcal{G}_\mu$ can be formulated as an admixture $\mu = \langle \mu_d \rangle$ of proportions of individuals of degree $d \geq 1$, where $\sum_d \mu_d = 1$ and $\mu_d \geq 0$. The expected informativeness of any network then is a proportion-weighted sum of the expected informativeness of the individuals of each degree contained in the network. That is, $\boldsymbol{I}(\mathcal{G}_\mu | N_\theta) = \sum_d \mu_d \cdot E_{N_\theta}[\boldsymbol{I}_d]$. It follows from Lemma 2.5 that, near consensus, any network not entirely composed of individuals of degree two is strictly less informative than one which is in fact composed entirely of individuals of degree two. Thus, when $N_\theta = 0$ or 1, $\boldsymbol{I}(\mathcal{G}_{circle}) > \boldsymbol{I}(\mathcal{G}_\mu)$ for any $\mathcal{G}_\mu$ such that $\mu_0 = 0$ and $\mu_2 \neq 1$, as desired. □

*Proof of Proposition 2.6.* It follows directly from Lemma 2.4 that, near consensus, the maximally and minimally informative regular networks of degree at least two are the circle and complete network, respectively. We combine this with the fact that any large network $\mathcal{G}_\mu$ can be formulated as an admixture $\mu = \langle \mu_d \rangle$ of regular networks of degree $d$, and the linearity of expected informativeness, to adduce that the informativeness of any network is bounded above by that of the circle network and bounded bellow by the complete network. That is, when $N_\theta = 0$ or 1, $\boldsymbol{I}(\mathcal{G}_{circle}) \geq \boldsymbol{I}(\mathcal{G}_\mu) \geq \boldsymbol{I}(\mathcal{G}_{complete})$ for any $\mathcal{G}_\mu$ such that $\min\{d : \mu_d > 0\} \geq 2$, as desired. □

# Chapter 3

# How Redefining Statistical Significance will Worsen the Replication Crisis

The strength of an empirical science is in the ability for the evidence of its discoveries to be reproduced—independently and with statistically-satisfying consistency. For this reason, the low success rate of replication studies in the social [68, 24], biological [21], and medical [69] sciences has been seen as a threat to the credibility of the scientific enterprise.

In a recent, high-profile proposal aimed to ameliorate the crisis in replication, a large group of prominent scientists and statisticians have called for a reduction in the $p$-value significance threshold from its conventional level of 0.05 to 0.005 [50, 15]. They argue that the higher evidential burden would have the effect of lowering the *false positive rate*—the rate at which claimed discoveries are in fact untrue—thereby improving reproducibility.

However, this argument rests on the assumption that researchers follow sound statistical protocol. As such, it does not account for the ways in which researchers exploit their "degrees of freedom"[1] to obtain statistical significance (e.g. $p$-hacking, multiple comparisons, selective reporting).

In this article, I show that these degrees of freedom can generate unintended consequences wherein redefining statistical significance actually leads to an *increase* in the false positive rate, in turn exacerbating problems with reproducibility. The adverse effect will occur when: (1) false positive rates differ across studies and (2) lowering the significance threshold reduces the number of significant outcomes more for studies with the lowest false positive rates.

To illustrate the result, consider the following example. Imagine that some fraction of studies are sound, the remaining are unsound, and assume an unsound study always obtains significance regardless of the truth of the hypothesis being tested. In this case, redefining statistical significance will result in lowering the false positive rate of sound studies (as in [15]), but also increasing the proportion of significant outcomes that are unsound. Furthermore, once the significance threshold is made sufficiently small, the latter effect will dominate the former, resulting in an increase in the false positive rate.

The identification of this mechanism contributes another reason for growing concern with the proposal to redefine statistical significance [38, 6, 26, 60, 59, 83]. In particular, it justifies apprehension about the interplay between redefining statistical significance and researcher degrees of freedom. However, it also suggests caution for counterproposals to abandon $p$-value significance thresholds or 0-1 decision rules altogether as these expand the domain of researcher degree of freedom and will thus further threaten the reproducibility of claimed findings.

Section 3.2 begins by introducing the first version of the model in which all studies are perfectly homogeneous, as is the case in [15]. Here we find lowering the significance threshold can only reduce the false positive rate. Section 3.3 expands the model to allow studies to be

[1] See [47], [79], and [37] for extended discussions of researcher degrees of freedom.

of heterogeneous types, varying along dimensions which affect their propensity to produce false positives. Under heterogeneity, lowering the significance threshold will increase the false positive rate if and only if doing so sufficiently increases the proportion of significant outcomes of types with the highest false positive rates. Section 3.3.1 provides an example of this mechanism and section 1.6 concludes.

## 3.2 Homogeneity

Consider a unit mass of independent studies.[2] In each study, a researcher conducts a hypothesis test between a pair of null ($H_0$) and alternative ($H_1$) hypotheses, with the prior chance the alternative hypothesis is true for any particular study being $\pi \in (0,1)$. Let $f_\theta(x)$ denote the sampling density of test statistic $x \in \mathbb{R}$ under both the null ($\theta = 0$) and alternative ($\theta = 1$) hypotheses. Assume $f_\theta(x)$ satisfies the *monotone likelihood ratio property* (MLRP):

$$\frac{f_1}{f_0}(x) \leq \frac{f_1}{f_0}(x'), \quad x \leq x'. \tag{3.1}$$

The $p$-value $p(x)$ from observing $x$ is

$$p\left(x\right) \equiv Pr\left(X \geq x | \theta = 0\right) = 1 - F_0(x). \tag{3.2}$$

The MLRP provides that lower $p$-values provide stronger evidence against the null, in favor of the alternative. Let $G_\theta(p) \equiv Pr\left(p\left(x\right) \leq p | \theta\right)$ denote the CDF of the p-value under hypothesis $\theta$ and $g_\theta(p)$ the corresponding PDF. Using (3.2), we can express these functions

$$G_\theta(p) = 1 - F_\theta\left(F_0^{-1}\left(1-p\right)\right), \quad g_\theta(p) = \frac{f_\theta(F_0^{-1}(1-p))}{f_0(F_0^{-1}(1-p))}. \tag{3.3}$$

[2]Assuming a continuum allows for simple statements of our results. Alternatively, one could insert the addendum "as the number of studies goes to infinity almost surely" in each result.

It follows from (3.3) that the $p$-value is distributed uniform(0,1) under the null. The following properties of the $p$-value distributions will be used in our results. Omitted proofs can be found in the appendix.

**Lemma 3.1.**

1. $G_0$ has first order stochastic dominance over $G_1$.
2. $g_1(p)$ and $\frac{G_1}{G_0}(p)$ are non-increasing in $p$.

The researcher conducting study $i$ selects an action $a$ from the set $A$ and receives the payoff

$$U(y, a) \tag{3.4}$$

where $y = 1$ if $i$ is significant, $y = 0$ if $i$ is non-significant, and significance is preferred to non-significance $U(1, \cdot) \geq U(0, \cdot)$. The outcome for $i$ is significant if and only if the reported p-value $\hat{p}$ is less than or equal to the significance threshold $\alpha$. The reported $p$-value $\hat{p}$ is determined by both the true $p$-value $p$ and the researcher's action $a$ and is assumed to satisfy the following monotonicity condition

$$\hat{p}(p, a) \leq \hat{p}(p', a) \text{ if } p \leq p'. \tag{3.5}$$

That is, if we fix a researcher's action and change the data so that the true $p$-value decreases, so too does the reported $p$-value.

To avoid trivialities, assume the set of payoff maximizing actions to be well-defined. For simplicity, if a researcher is indifferent between actions yielding significance and others that do not, then an action yielding significance is chosen.[3]

[3] The appendix shows that the results do not depend on this assumption.

Suppose a researcher, upon obtaining a $p$-value of $p$, finds it optimal to select an action yielding significance. Then, if another researcher obtains a $p$-value of $p' \leq p$, he too must find it optimal to select an action yielding significance. By this logic, significance is obtained for all $p$-values below some cutoff $b(\alpha)$. Furthermore, if a researcher optimally obtains significance at $\alpha$, then significance must also be optimal at a less stringent $\alpha' \geq \alpha$, and thus the cutoff $b(\alpha)$ must be nondecreasing.

**Lemma 3.2.** *There exists a (weakly) increasing function $b(\alpha)$ such that significance is obtained if $p < b(\alpha)$ and non-significance obtained if $p > b(\alpha)$.*

We turn to the primary object of interest, the false positive rate $R(\alpha)$. A false positive occurs when a hypothesis test yields a significant outcome when the null is in fact true. The false positive rate is equal to the number of false positives divided by the total number of significant outcomes

$$R(\alpha) = \frac{(1-\pi)\, Pr\,(y=1|\alpha, H_0)}{(1-\pi)\, Pr\,(y=1|\alpha, H_0) + \pi Pr\,(y=1|\alpha, H_1)} \tag{3.6}$$

and by lemma 3.2

$$R(\alpha) = \left(1 + \frac{\pi}{1-\pi}\frac{G_1}{G_0}\,(b\,(\alpha))\right)^{-1}. \tag{3.7}$$

The false positive rate of sound research is given by replacing $b(\alpha)$ with $\alpha$ in (3.7). In general, $b(\alpha)$ may be less than, greater than, or equal to $\alpha$. In the special case where researchers may choose sound reporting $\big(\hat{p}\,(p, a_0) = p$ for some $a_0 \in A\big)$ and deviations from this action are costly $\big(U(1,a) \leq U(1,a_0)\big)$, the cutoff exceeds the significance threshold $b(\alpha) \geq \alpha$ and the resulting false positive rate in (3.7) exceeds the false positive rate of sound research.

Observe also in (3.7) whenever $b(\alpha)$ is non-constant, the qualitative response of $R(\alpha)$ to a change in $\alpha$ is the same as with sound research.

**Theorem 3.1.** *Under homogeneity, reducing the significance threshold $\alpha$ (weakly) decreases the false positive rate $R(\alpha)$.*

Thus, if studies are homogeneous, redefining statistical significance will, at the very least, not worsen reproducibility. As we shall now see, this is no longer true with a departure from homogeneity.

## 3.3 Heterogeneity

Realistically, studies can vary along many dimensions that influence their likelihood of producing a false positive. We capture this by allowing studies to be of distinct *types* $t \in T$. Examples include heterogeneous prior chances $\pi_t$, statistical power $G_1^t$, preferences $U_t$, or capacity to exercise degrees of freedom $A_t$ or $\hat{p}_t$.

Let types be distributed according to probability measure $\mu$ with probability space $(T, \Sigma, \mu)$ and maintain all the previous assumptions of section 3.2. For a particular type $t$, lemma 3.2 ensures the existence of a weakly increasing function $b_t(\alpha)$ such that a study of type $t$ obtains significance if and only if $p \leq b_t(\alpha)$. For the false positive rate to be well-defined, we require the mapping $t \mapsto (\pi_t, G_\theta^t(b_t(\alpha)))$ to be measurable for all $\alpha$.

To obtain the false positive rate under heterogeneity, define $R_t(\alpha)$ to be the false positive rate for studies of type $t$ and $\eta(D|\alpha)$ to be the proportion of significant outcomes that are of a type in $D \in \Sigma$ at significance threshold $\alpha$,[4] and write

$$R(\alpha) = \int R_t(\alpha) \mathrm{d}\eta(t|\alpha). \tag{3.8}$$

[4] The appendix provides a formal derivation of $\eta$.

The false positive rate is thus a weighted average of the false positive rates for the individual types, with the weights being endogenously determined by $\alpha$. Theorem 3.1 guarantees each $R_t(\alpha)$ will not increase with a reduction in $\alpha$. It follows that, if a reduction in $\alpha$ is to increase $R(\alpha)$, it must be due to a change in the type-composition of significant outcomes $\eta(\cdot|\alpha)$.

Define $B_{\delta,\alpha} = \{t \in T | R_t(\alpha) > \delta\}$ to be the set of types with false positive rates exceeding $\delta$ at significance threshold $\alpha$ so that $(B_{\delta,\alpha}, B^C_{\delta,\alpha})$ partitions types between those with the highest false positive rates $t \in B_{\delta,\alpha}$ and those with the lowest false positive rates $t \in B^C_{\delta,\alpha}$.

The following theorem states that, lowering the significance threshold *increases* the false positive rate if and only if, by some partitioning of types $(B_{\delta,\alpha}, B^C_{\delta,\alpha})$, those with the highest false positive rates increase sufficiently as a proportion of the total significant outcomes.

**Theorem 3.2.** *Reducing the significance threshold from $\alpha$ to $\alpha'$ increases the false positive rate $R(\alpha) < R(\alpha')$ if and only if there is a sufficient increase in the proportion of significant outcomes that are of types with false positive rates exceeding $\delta$, for some $\delta \in (0,1)$. In particular,*

$$\eta(B_{\delta,\alpha'}|\alpha') - \eta(B_{\delta,\alpha}|\alpha) > \psi(\delta, \alpha, \alpha') > 0. \tag{3.9}$$

*The precise form of $\psi(\delta, \alpha, \alpha')$ can be found in the appendix.*

In other words, if the studies most prone to producing false positives are the least affected by reducing $\alpha$, then redefining statistical significance will worsen reproducibility. This condition will naturally hold when the same factors that originally led a study to have a higher false positive rate also induce a lower response to changes in the significance threshold. This point is illustrated in the following example.

### 3.3.1 Example

Consider the motivating example from the introduction. Suppose there are two types of studies $T = \{t_s, t_u\}$ assigned with probabilities $\mu(t_s) = 1 - \lambda$ and $\mu(t_u) = \lambda$. The available actions and preferences for a researcher of type $t_s$ induce *sound* statistical protocol $\hat{p}_{t_s} = p$. In contrast, the available actions and preferences of a researcher of type $t_u$ induce *unsound* protocol so that significance is obtained $\hat{p}_{t_u} \leq \alpha$ at all $\alpha$. Assume all hypotheses have the same prior chances $\pi$ and $f_\theta(x)$ is differentiable.[5]

The false positive rates for each type are $R_s(\alpha) = \left(1 + \frac{\pi}{1-\pi} \cdot \frac{G_1(\alpha)}{\alpha}\right)^{-1}$ and $R_u(\alpha) = \left(1 + \frac{\pi}{1-\pi}\right)^{-1}$ so that sound research produces a lower false positive rate $R_s(\alpha) < R_u(\alpha)$ for all $\alpha < 1$. The proportion of significant outcomes that are unsound is $\eta(t_u|\alpha) = \left(1 + \frac{1-\lambda}{\lambda}(\pi G_1(\alpha) + (1-\pi)\alpha)\right)^{-1}$.

Reducing $\alpha$ results in (1) lowering the false positive rate of sound studies $R_s(\alpha)$ and (2) increasing the proportion of significant outcomes that are unsound $\eta(t_u|\alpha)$. As shown in figure 3.1, for $\alpha$ below some threshold $\alpha^* \in (0, 1)$, the former effect becomes dominated by the latter, so that reducing $\alpha$ increases $R(\alpha)$.

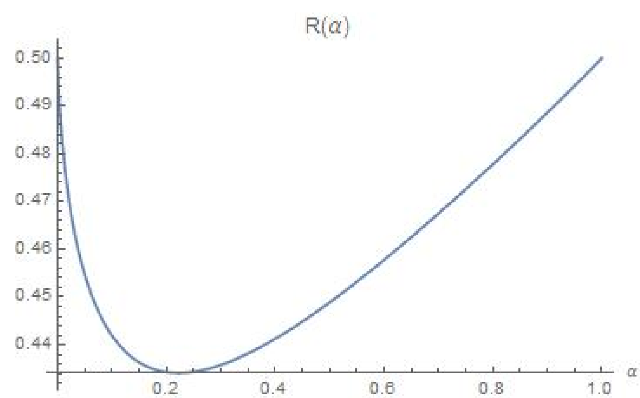


Figure 3.1: False Positive Rate. The diagram depicts the false positive rate $R(\alpha)$ as a function of the significance threshold $\alpha$ when $f_\theta \sim \mathcal{N}(\theta, 1)$ and $\lambda = \pi = \frac{1}{2}$.

[5]There is an alternative description of this example relating to the model of bias in science by [58]. Maintain the MLRP assumption and assume all studies are sound except that with chance $\lambda$ an insignificant outcome is reported as significant. The first description provides a decision-theoretic basis for this simpler one.

**Proposition 3.1.** *For any $\lambda \in (0,1)$, lowering the significance threshold $\alpha$ decreases the false positive rate $R(\alpha)$ if $\alpha > \alpha^*$ and increases the false positive rate if $\alpha < \alpha^*$ for some $\alpha^* \in (0,1)$.*

And thus if any fraction of studies are unsound, reducing the significance threshold by too much ($\alpha < \alpha^*$) will worsen reproducibility. Furthermore, even in the complementary domain ($\alpha \geq \alpha^*$) the presence of unsound studies will mitigate the positive effects of redefining statistical significance.[6]

## 3.4 Conclusion

This article has shown how redefining statistical significance will worsen reproducibility if doing so disproportionately reduces the number of significant outcomes for studies with the lowest false positive rates. Furthermore, this will occur even in the presence of arbitrarily small researcher bias. And while the analysis was motivated as a response to [15], the findings apply more generally to any 0-1 decision rule in science, so long as the appropriate monotonicity assumptions hold.

These findings emerge from mild monotonicity assumptions, but hold even more generally.[7] The online appendix analyzes extensions of the model. It shows that when researchers have the option to preregister their studies [67], redefining statistical significance produces a qualitatively similar effect as in section 3.3.1 and that multiple-hypothesis testing will only negatively impact reproducibility in the presence of some form of heterogeneity.

It is not clear what the path towards a more reliable science will entail, whether it be

[6] If $\hat{R}(\alpha)$ is the false positive rate when all studies are sound, then $\frac{\partial R}{\partial(-\alpha)} > \frac{\partial \hat{R}}{\partial(-\alpha)}$.

[7] Theorem 3.2 is still obtained if we dispense with all the assumptions in section 3.2 and simply take as primitive that $R_t(\alpha)$ is nondecreasing. The subsequent example requires only that data satisfies MLRP and that a positive fraction of studies are unsound.

more stringent statistical requirements, increased adoption of Bayesian methods, further proliferation of preregistration, or even more radical changes than these. What is clear is that the incentives and preferences underlying both the supply and demand of science cannot be ignored. The consumers of science demand crisp and clean conclusions from research and the producers of science are incentivized to meet this demand. Future proposals that fail to account for the confluence of these forces are likely to be less effective or, as we have seen in this article, exacerbate the problem.

## 3.5 Mathematical Appendix

*Proof of Lemma 3.1.* (1) By the MLRP, $F_1$ has first-order stochastic dominance over $F_0$, $F_1(x) \leq F_0(x)$ for all $x$. This further entails that $G_0$ has first-order stochastic dominance over $G_1$, $G_0(p) \leq G_1(p)$ for all $p \in [0, 1]$.

(2) That $g_1(p)$ is non-increasing follows immediately from its definition the MLRP. The derivative of $\frac{G_1}{G_0}(p)$ is proportional to $g_1(p)G_0(p) - g_0(p)G_1(p)$ which is non-positive if and only if $\frac{g_1}{g_0}(p) \leq \frac{G_1}{G_0}(p)$ which holds because $G_0$ has FOSD over $G_1$. □

*Proof of Lemma 3.2.* Choose any $p'$ such that there exists an action $\tilde{a}$ yielding significance $\hat{p}(p', \tilde{a}) \leq \alpha$ and this action is optimal $U(1, \tilde{a}) \geq U\left(y\left(p', a'\right), a'\right)$ for all $a' \in A$. If no such $p'$ exists, set $b(\alpha) = 0$.

Let $A^0_p(\alpha) \subset A$ be the actions yielding nonsignificance for $p$ at $\alpha$. Notice for $p \leq p'$, $\tilde{a}$ yields significance at $p'$ it must also at $p$, $\hat{p}(p, \tilde{a}) \leq \hat{p}(p', \tilde{a}) \leq \alpha$ and $A^0_p(\alpha) \subseteq A^0_{p'}(\alpha)$. This implies $U(1, \tilde{a}) \geq U\left(0, a'\right)$ for all $a' \in A^0_p(\alpha)$ and thus an action yielding significance must be optimal at $p$. Thus, we can set $b(\alpha)$ equal to the supremum of the set of $p'$ at which at which an action yielding significance is optimal.

To establish that $b(\alpha)$ is weakly increasing, notice that if $\alpha \leq \alpha'$, then $A_p^0(\alpha) \subset A_p^0(\alpha')$, and thus if $U(1,a) \geq U(0,a')$ for all $a' \in A_p^0(\alpha')$ then $U(1,a) \geq U(0,a')$ for all $a' \in A_p^0(\alpha)$. □

Let us justify the claim that assuming an action yielding significance is chosen when faced with indifference can be greatly weakened. Uninteresting technical complications arise in the case of indifference when: (1) unknown to the researchers themselves, the decisions of researchers with precisely the same $p$-value $p$ are correlated with the truth of the hypotheses $\theta$ (2) a researcher changes his decision after a reduction in $\alpha$ to $\alpha'$, even though he was indifferent at both significance thresholds. We wish to rule these cases out.

Let $\gamma(p,\alpha)$ be the likelihood of obtaining significance at $p$ and $\alpha$. Assume (1) $\gamma(p,\alpha)$ does not depend on $\theta$ (2) If a researcher at $p$ is indifferent at both $\alpha$ and $\alpha'$, then $\gamma(p,\alpha) = \gamma(p,\alpha')$ and (3) $\gamma(p,\alpha)$ is $G_\theta$ measurable. Define

$$\tilde{g}_\theta(p) = \frac{\gamma(p)g_\theta(p)}{\int_0^1 \gamma(s)g_\theta(s)\mathrm{d}s} \text{ and } \tilde{G}_\theta(p) = \int_0^p \tilde{g}_\theta(s)\mathrm{d}s. \tag{3.10}$$

First observe that $\tilde{g}$ satisfies the (inverted) MLRP. Consider a decrease in the significance threshold from $\alpha$ to $\alpha' < \alpha$.

$$R(\alpha') \leq R(\alpha) \iff \frac{\tilde{G}_1}{\tilde{G}_0}(b(\alpha')) \leq \frac{\tilde{G}_1}{\tilde{G}_0}(b(\alpha)). \tag{3.11}$$

By lemma 3.2 $b(\alpha') \leq b(\alpha)$ and as the (inverted) MLRP guarantees $\frac{\tilde{G}_1}{\tilde{G}_0}(p)$ is decreasing, the inequality holds. Thus Theorem 3.1 holds under these more general conditions.

*Deriving $\eta$.*

Let $M_0(\alpha,t)$ and $M_1(\alpha,t)$ be the false and true positives for type $t$ respectively. The measurability of $t \mapsto (\pi_t, G_\theta^t(b_t(\alpha)))$ guarantees $M_0(\alpha) = \int M_0(\alpha,t)\mathrm{d}\mu(t)$ and $M_1(\alpha) =$

$\int M_1(\alpha, t)\mathrm{d}\mu(t)$ are well-defined.

$$R(\alpha) = \int R_t(\alpha)\, m_t(\alpha)\, \mathrm{d}\mu(t) \tag{3.12}$$

with $m_t(\alpha) = \frac{M_0(\alpha,t)+M_1(\alpha,t)}{M(\alpha)}$ and $M(\alpha) = M_0(\alpha) + M_1(\alpha)$.

Define the probability measure $\eta(\cdot|\alpha)$

$$\eta(D|\alpha) = \int_D m_t(\alpha)\mathrm{d}\mu(t) \tag{3.13}$$

and thus

$$R(\alpha) = \int R_t(\alpha)\mathrm{d}\eta(t|\alpha). \tag{3.14}$$

*Proof of Theorem 3.2.* The measurability of $B_{\delta,\alpha}$ follows from the measurability of $t \mapsto (\pi_t, G_\theta^t(b_t(\alpha)))$. Denote the conditional expectation of $R_t(\alpha)$ for types in $D \in \Sigma$ by $\bar{R}(D|\alpha) = \mathbb{E}[R_t|D, \alpha]$. Making the condition in (3.9) explicit

$$\psi(\delta, \alpha, \alpha') = \frac{\bar{R}(B_{\delta,\alpha}|\alpha) - \bar{R}(B_{\delta,\alpha}^C|\alpha)}{\bar{R}(B_{\delta,\alpha'}|\alpha') - \bar{R}(B_{\delta,\alpha'}^C|\alpha')}\eta(B_{\delta,\alpha}|\alpha) + \frac{\bar{R}(B_{\delta,\alpha}^C|\alpha) - \bar{R}(B_{\delta,\alpha'}^C|\alpha')}{\bar{R}(B_{\delta,\alpha'}|\alpha') - \bar{R}(B_{\delta,\alpha'}^C|\alpha')}. \tag{3.15}$$

Expand the false positive rate

$$R(\alpha) = \bar{R}(B_\delta|\alpha)\eta(B_\delta|\alpha) + \bar{R}(B_\delta^C|\alpha)\eta(B_\delta^C|\alpha) \tag{3.16}$$

and note $R(\alpha') > R(\alpha)$ if and only if the first inequality in (3.9) is satisfied. All that remains to be shown is, if $R(\alpha') > R(\alpha)$, then there exists $\delta$ satisfying $\psi(\delta, \alpha, \alpha') > 0$.

To obtain a contradiction, assume $\psi(\delta, \alpha, \alpha') \leq 0$ for all $\delta$. Define the CDF $\nu(\hat{\delta}|\alpha) = Pr(\delta \leq$

$\hat{\delta}|\alpha) = \eta(B_{\hat{\delta}}^C|\alpha)$. By assumption, $\nu(\cdot|\alpha)$ dominates $\nu(\cdot|\alpha')$ in terms of first-order stochastic dominance. Recalling $R_t(\alpha') \leq R_t(\alpha)$,

$$R(\alpha') = \int R_t(\alpha')\mathrm{d}\eta(t|\alpha') \leq \int R_t(\alpha)\mathrm{d}\eta(t|\alpha') = \int \delta\mathrm{d}\nu(\delta|\alpha') \leq \int \delta\mathrm{d}\nu(\delta|\alpha) = R(\alpha) \quad (3.17)$$

contradicting $R(\alpha') < R(\alpha)$. □

*Proof of Proposition 3.1.*

$$R(\alpha) = \left(1 + \frac{\pi}{1-\pi} \cdot \frac{\lambda + (1-\lambda)\, G_1(\alpha)}{\lambda + (1-\lambda)\, G_0(\alpha)}\right)^{-1} \quad (3.18)$$

$R'(\alpha)$ is of the opposite sign as

$$h(\alpha) = \lambda\left(g_1(\alpha) - g_0(\alpha)\right) + (1-\lambda)\left(G_0(\alpha)\, g_1(\alpha) - G_1(\alpha)\, g_0(\alpha)\right). \quad (3.19)$$

Upon differentiating, $h'(\alpha) = (\lambda + (1-\lambda)\,\alpha)\, g_1'(\alpha) \leq 0$. Finally, noting $\lim_{\alpha\to 0^+} h(\alpha) > 0$ and $\lim_{\alpha\to 1^-} h(\alpha) < 0$ completes the proof. □

## 3.6 Appendix: Extensions

### Uncertainty & Preregistration

We expand on the example presented in the body of the paper to show that our findings do not rely on the exogeneity of unsound practices. We do this by supposing that prior to observing the data, a researcher chooses whether or not to preregister his study [67]. A preregistered study must follow sound statistical protocol. A study that is not preregistered is free to engage in unsound protocol. Preregistration offers a payoff premium but significance is still always preferred: $U(y, \text{Preregister}) = u_y$, $U(\text{Do not Preregister}) = \bar{u}$, and $u_0 < \bar{u} < u_1$.

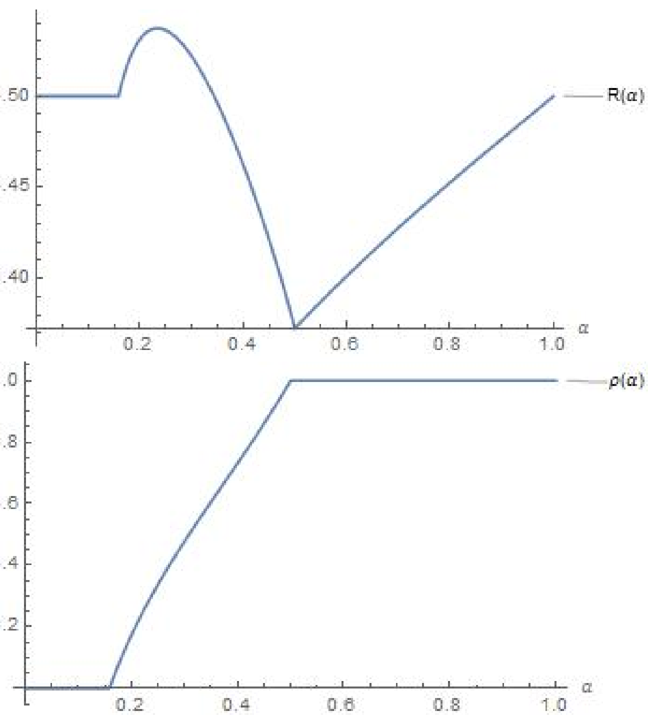


Figure 3.2: False Positive Rate under Uncertainty. The diagram depicts the false positive rate $R(\alpha)$ and preregistration rate $\rho(\alpha)$ as functions of the significance threshold $\alpha$ when $s_i|\theta \sim 2s_i + \theta(2-4s_i)$, $\pi = \frac{1}{2}$, $x \sim \mathcal{N}(\theta, 1)$, and $\bar{u} = \frac{u_1+u_0}{2}$.

Prior to choosing, each researcher receives an informative signal $s_i$ of the hypothesis being tested $H^i_\theta$. Assume the beliefs induced by the signals $\pi(s_i)$ are smoothly distributed within some interval.[8] Writing a researcher's expected payoff to preregistration

$$EU\left(\text{Preregister}|s_i\right) = Pr\left(p \leq \alpha|s_i\right)u_1 + Pr\left(p \geq \alpha|s_i\right)u_0 \tag{3.20}$$

so that a researcher's optimal action is to preregister whenever $Pr\left(p \leq \alpha|s_i\right) \geq \frac{\bar{u}-u_0}{u_1-u_0}$.

Figure 3.2 portrays both the researcher choices and the false positive rate as a function of the significance threshold when signals are drawn from the conditional density $2s_i + \theta(2-4s_i)$ with support $s_i \in [0,1]$, $x \sim \mathcal{N}(\theta, 1)$, and $\bar{u} = \frac{u_1+u_0}{2}$. For $\alpha > 1/2$, all researchers choose preregistration, conduct sound studies, thus a reduction in $\alpha$ reduces $R(\alpha)$. Once $\alpha$ is

[8] In particular, the distribution of beliefs is differentiable with a non-degenerate, convex support.

reduced below 1/2, researchers with the lowest expectations on their hypotheses $\pi(s_i)$ will abstain from preregistration and opt for the latitude to exercise degrees of freedom. Once $\alpha$ is made less than approximately 0.15, no researchers preregister and thus $R(\alpha)$ is constant. More generally, in this environment, $R(0) = R(1)$ and $R(\alpha)' > 0$ in the neighborhood of 1, yielding the following proposition.

**Proposition 3.2.** *In the preregistration example, there exists an open set on which lowering the significance threshold $\alpha$ increases the false positive rate $R(\alpha)$.*

## Multiple Hypothesis Testing

In the ideal picture of science, a scientist makes a prediction and then gathers data to test the prediction. Hypothesis testing intends to capture this ideal by requiring the null and alternative hypotheses to be specified prior to observing the data. It may thus be seen as "anti-scientific" to invert the order of operations—for a scientist to first observe the data and then to choose a hypothesis. Borrowing from [67], refer to ex post hypothesis choice as *postdiction*.

Does postdiction contribute to the false positive rate *per se*? We shall see that the answer to this question will turn on whether there is an adequate degree of heterogeneity among the hypotheses the researcher chooses from.

To illustrate the point, consider the following example of homogeneity. Suppose a researcher observes some data and obtains the $p$-values for $L$ independent and *ex ante* identical hypothesis tests. Consider the following strategies.

**Strategies.**

1. *Prediction* $R^{(1)}(\alpha)$: Test single hypothesis.

2. *Postdiction* $R^{(2)}(\alpha)$: Randomly select from significant hypotheses.

3. *Selective Postdiction* $R^{(3)}(\alpha)$: Report hypothesis with the lowest $p$-value.

Consider the posterior probability when significance is obtained in these strategies. Under prediction and postdiction, the posterior probability of the alternative hypothesis is determined by the chance of obtaining a significant $p$-value under the alternative relative to the null. Whether a hypothesis was declared before or after observing the data has no bearing on the posterior and thus the false positive rate.

The posterior under selective postdiction is determined by the chance of obtaining a $p$-value that is not only significant, but also smaller than the $p$-values of the other available hypotheses. The additional evidential burden selective postdiction places on a hypothesis functions in qualitatively the same way as lowering the significance threshold. This gives rise to the following proposition.

**Proposition 3.3.** *Under homogeneity,*

1. *Prediction yields the same false positive rate as postdiction, both exceeding the false positive rate of selective postdiction*

$$R^{(1)}(\alpha) = R^{(2)}(\alpha) \geq R^{(3)}(\alpha) \tag{3.21}$$

2. *The false positive rate for selective postdiction is decreasing in the number of hypotheses $L$.*

The proof can be found at the end of the appendix.

**Heterogeneous Hypotheses.**

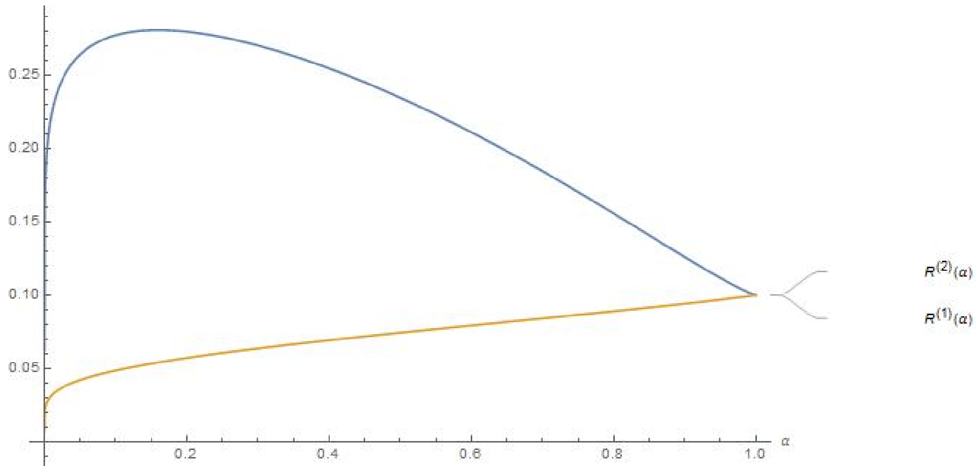


Figure 3.3: False Positive Rate under Multiple Testing. The diagram depicts the false positive rate when only the primary hypothesis is reported $R^{(1)}(\alpha)$ and under multiple testing $R^{(2)}(\alpha)$ as functions of the significance threshold $\alpha$ when $\pi_{\ell^*} = 0.9$, $\pi_{\ell'} = 0.1$, and $\sigma = 2$.

Now suppose a researcher begins with a primary hypothesis $H^*$ but may also choose from $L-1$ auxiliary hypotheses $\{H^l\}_{\ell=1}^{L-1}$. The primary hypothesis has the benefit of prior empirical or theoretical support while the auxiliary hypotheses do not. We capture this notion but specifying $\pi^* > \pi^\ell$ for $\ell = 1, 2, ..., L-1$.

Compare the two strategies: (1) test primary hypothesis (prediction) (2) report primary hypothesis if significant, otherwise randomly select from significant auxiliary hypotheses (postdiction).

Figure 3.3 shows that, naturally, if researchers always stick to their primary hypotheses, then lowering $\alpha$ reduces the false positive rate $R^{(1)}(\alpha)$. If researchers are prone to report auxiliary hypotheses with low priors when their primary hypothesis fails: (1) the false positive rate will be larger than if they stayed strictly with the primary hypothesis (2) lowering $\alpha$ may, once again, lead to an increase in the false positive rate $R^{(2)}(\alpha)$. The intuition for the last observation is that, while lowering $\alpha$ reduces the false positive rates for both the primary and auxiliary hypotheses, it also increases the propensity for researchers to select auxiliary hypotheses.

*Proof of Proposition 3.3.* Under prediction $R^{(1)}(\alpha) = \left(1 + \frac{\pi}{1-\pi}\frac{G_1(\alpha)}{\alpha}\right)^{-1}$. For postdiction, let $z_\ell$ be a random variable such that $z_\ell = 1$ if hypothesis $\ell$ is selected to be reported and $z_\ell = 0$ otherwise.

$$R^{(2)}(\alpha) = \left(1 + \frac{M_1^{(2)}(\alpha)}{M_0^{(2)}(\alpha)}\right)^{-1}. \tag{3.22}$$

$$M_j^{(2)} = \mathbb{E}_\ell\left[Pr(z_\ell = 1 \wedge H_j^\ell)\right] = \mathbb{E}_\ell\left[Pr(z_\ell = 1 | p_\ell \leq \alpha, H_j^\ell) Pr(p_\ell \leq \alpha | H_j^\ell) Pr(H_j^\ell)\right]. \tag{3.23}$$

Given that the choice of significant result to report is independent of the $p$-value

$$Pr(z_\ell = 1 | p_\ell \leq \alpha, H_j^\ell) = Pr(z_\ell = 1 | p_\ell \leq \alpha). \tag{3.24}$$

Using both this observation and $Pr(p_\ell \leq \alpha | H_j^\ell) = G_j(\alpha)$, the equality $R^{(2)}(\alpha) = R^{(1)}(\alpha)$ immediately follows.

To consider selective postdiction, let $p_{-\ell}^* = \inf_{\ell' \neq \ell} p_{\ell'}$ and also

$$R^{(3)}(\alpha) = \mathbb{E}_\ell\left[R_\ell^{(3)}(\alpha)\right] \text{ and } R_\ell^{(3)}(\alpha) = \mathbb{E}_{p_{-\ell}^*}\left[R_\ell^{(3)}\left(\alpha; p_{-\ell}^*\right)\right]. \tag{3.25}$$

For all $p_{-\ell}^*$

$$R_\ell^{(3)}(\alpha; p_{-\ell}^*) = \left(1 + \frac{\pi}{1-\pi}\frac{G_1(\min\{\alpha, p_\ell^*\})}{G_0(\min\{\alpha, p_\ell^*\})}\right)^{-1} \leq \left(1 + \frac{\pi}{1-\pi}\frac{G_1(\alpha)}{G_0(\alpha)}\right)^{-1} = R^{(1)}(\alpha) \tag{3.26}$$

and thus $R^{(3)}(\alpha) \leq R^{(1)}(\alpha) = R^{(2)}(\alpha)$.

To prove the second claim, first note that $R_\ell^{(3)}\left(\alpha; p_{-\ell}^*\right)$ is increasing in $p_{-\ell}^*$ for $p_{-\ell}^* < \alpha$ and constant otherwise. The CDF of $p_{-\ell}^*$ given the number of hypotheses $L$ is $Pr(p_{-\ell}^* \leq$

$t|L) = Pr(p_{\ell'} \leq t)^{L-1}$ and thus $Pr(p^*_{-\ell} \leq t|L)$ is first-order stochastically dominated by $Pr(p^*_{-\ell} \leq t|L')$ for $L' < L$. It follows that $R^{(3)}(\alpha)$ is decreasing in $L$. Furthermore, as $L$ tends to infinity, the distribution of $p^*_{-\ell}$ weakly converges to the dirac measure $\delta_0(\cdot)$ and thus

$$\mathbb{E}_{p^*_{-\ell}}\left[R^{(3)}_\ell\left(\alpha; p^*_{-\ell}\right)|L\right] \to R^{(3)}_\ell\left(\alpha; 0\right) \quad \text{as} \quad L \to +\infty. \tag{3.27}$$

If $g_1(p)$ is unbounded then $R^{(3)}_\ell\left(\alpha; 0\right) = 0$ and thus $R^{(3)}(\alpha|L) \to 0$ as $L \to +\infty$. □

# Bibliography


[1] D. Acemoglu, K. Bimpikis, and A. Ozdaglar. Dynamics of information exchange in endogenous social networks. *Chemistry & biodiversity*, 1(11):1829–1841, 2010.

[2] D. Acemoglu, V. Chernozhukov, and M. Yildiz. Fragility of Asymptotic Agreement Under Bayesian Learning. *Theoretical Economics*, 11:187–225, 2016.

[3] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian Learning in Social Networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.

[4] G. A. Akerlof and R. Kranton. Identity Economics. *The Economists' Voice*, 7(2):11–13, 2010.

[5] G. A. Akerlof and R. E. Kranton. Economics and Identity. *Quarterly Journal of Economics*, 115:715–753, 2000.

[6] V. Amrhein and S. Greenland. Remove, rather than redefine, statistical significance. *Nature Human Behaviour*, 2018.

[7] S. E. Asch. Opinions and Social Pressure. *Scientific American*, 193(5):31–35, 1955.

[8] R. J. Aumann. Agreeing to Disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976.

[9] M. Bacharach. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, 1985.

[10] V. Bala and S. Goyal. Learning from neighbors. *Review of Economic Studies*, 65(3):1–41, 1998.

[11] A. V. Banerjee. A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, 107(3):797–817, 1992.

[12] J. Barrett, B. Skyrms, and A. Mohseni. Self-Assembling Networks. *British Journal for the Philosophy of Science*, 2017.

[13] G. S. Becker. A Note on Restaurant Pricing and Other Examples of Social Influences on Price. *Journal of Political Economy*, 1991.

[14] R. Bénabou and J. Tiróle. Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 30(3), 2016.

[15] D. J. Benjamin, J. O. Berger, M. Johannesson, B. A. Nosek, E. J. Wagenmakers, R. Berk, K. A. Bollen, B. Brembs, L. Brown, C. Camerer, D. Cesarini, C. D. Chambers, M. Clyde, T. D. Cook, P. De Boeck, Z. Dienes, A. Dreber, K. Easwaran, C. Efferson, E. Fehr, F. Fidler, A. P. Field, M. Forster, E. I. George, R. Gonzalez, S. Goodman, E. Green, D. P. Green, A. G. Greenwald, J. D. Hadfield, L. V. Hedges, L. Held, T. Hua Ho, H. Hoijtink, D. J. Hruschka, K. Imai, G. Imbens, J. P. Ioannidis, M. Jeon, J. H. Jones, M. Kirchler, D. Laibson, J. List, R. Little, A. Lupia, E. Machery, S. E. Maxwell, M. McCarthy, D. A. Moore, S. L. Morgan, M. Munafó, S. Nakagawa, B. Nyhan, T. H. Parker, L. Pericchi, M. Perugini, J. Rouder, J. Rousseau, V. Savalei, F. D. Schönbrodt, T. Sellke, B. Sinclair, D. Tingley, T. Van Zandt, S. Vazire, D. J. Watts, C. Winship, R. L. Wolpert, Y. Xie, C. Young, J. Zinman, and V. E. Johnson. Redefine statistical significance. *Nature Human Behaviour*, 2018.

[16] J.-P. Benoit and J. Dubra. A Theory of Rational Attitude Polarization. *Working Paper*, 2014.

[17] S. Bikhchandani, D. Hirshleifer, and I. Welch. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5):992, 1992.

[18] P. Billingsley. *Convergence of Probability Measures*. 2009.

[19] D. Blackwell. Discounted Dynamic Programming. *The Annals of Mathematical Statistics*, 36(1):226–235, 1965.

[20] D. Blackwell and L. Dubins. Merging of Opinions with Increasing Information. *The Annals of Mathematical Statistics*, 1962.

[21] R. Bogdan, B. J. Salmeron, C. E. Carey, A. Agrawal, V. D. Calhoun, H. Garavan, A. R. Hariri, A. Heinz, M. N. Hill, A. Holmes, N. H. Kalin, and D. Goldman. Imaging Genetics and Genomics in Psychiatry: A Critical Review of Progress and Potential, 2017.

[22] P. Bolton and C. Harris. Strategic Experimentation. *Econometrica*, 1999.

[23] R. Bond and P. B. Smith. Culture and conformity: A meta-analysis of studies using asch's (1952b, 1956) line judgment task. *Psychological Bulletin*, 119(1):111–137, 1996.

[24] C. F. Camerer, A. Dreber, E. Forsell, T. H. Ho, J. Huber, M. Johannesson, M. Kirchler, J. Almenberg, A. Altmejd, T. Chan, E. Heikensten, F. Holzmeister, T. Imai, S. Isaksson, G. Nave, T. Pfeiffer, M. Razen, and H. Wu. Evaluating replicability of laboratory experiments in economics. *Science*, 2016.

[25] C. Chamley. *Rational herds: Economic models of social learning*. Cambridge University Press, Cambridge, U.K., 2004.

[26] H. Crane. Why redefining statistical significance will not improve reproducibility and could make the replication crisis worse. 2018.

[27] B. de Finetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut H. Poincaré*, pages 1–68, 1937.

[28] M. H. DeGroot. Reaching a Consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.

[29] P. M. DeMarzo, D. Vayanos, and J. Zwiebel. Persuasion bias, social influence and unidimensional opinions. *The Quarterly Journal of Economics*, 2003.

[30] A. K. Dixit and J. W. Weibull. Political Polarization. *Proceedings of the National Academy of Sciences of the United States of America*, 2007.

[31] J. L. Doob. Application of the Theory of Martingales. *Le calcul des probabilites et ses applications*, pages 23–27, 1949.

[32] G. Ellison and D. Fudenberg. Word-of-mouth communication and social learning. *The Quarterly Journal of Economics*, 110(1):93–125, 1995.

[33] P. et al. Abbot. Responses to The Evolution of Eusociality, Nowak et al. Nature. 466. *Nature*, 471(7339):E1–4; author reply E9–10, 2011.

[34] E. Eyster and M. Rabin. Extensive imitation is irrational and harmful. *Quarterly Journal of Economics*, 2014.

[35] R. G. J. Fryer, P. Harms, and M. O. Jackson. Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization. *Working paper*, 2015.

[36] J. D. Geanakoplos and H. M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.

[37] A. Gelman and E. Loken. The statistical Crisis in science. *American Scientist*, 102(6):460–465, 2014.

[38] A. Gelman and E. Loken. The statistical Crisis in science. *American Scientist*, 102(6):460–465, 2014.

[39] M. Gentzkow, J. M. Shapiro, and D. F. Stone. Media Bias in the Marketplace. In *Handbook of Media Economics*. 2016.

[40] J. Ghosh and R. Ramamoorthi. *Bayesian Nonparametrics*. Springer, New York, 2003.

[41] J. C. Gittins. Bandit processes and dynamic allocation indices. *Statistics*, 41(2):148–177, 1979.

[42] E. L. Glaeser and C. R. Sunstein. Why does balanced news produce unbalanced views? 2013.

[43] J. K. Goeree, A. Riedl, and A. Ule. In search of stars: Network formation among heterogeneous agents. *Games and Economic Behavior*, 67(2):445–466, 2009.

[44] S. Goyal. *Connections: An introduction to the economics of networks*. Number 1. 2007.

[45] M. Harris and A. Raviv. Differences of Opinion Make a Horse Race. *Review of Financial Studies*, 6(3):473–506, jul 1993.

[46] J. C. Harsanyi. Games with Incomplete Information Played by "Bayesian" Players, I-III Part I. The Basic Model. *Management Science*, 14(3):159–182, 1967.

[47] J. P. A. Ioannidis. Why most published research findings are false, 2005.

[48] R. Jensen. Adoption and diffusion of an innovation of uncertain profitability. *Journal of Economic Theory*, 1982.

[49] A. P. Jeuland. Incorporating heterogeneity into parsimonious models of diffusion of innovation. *University of Chicago, Center for Research in Marketing*, 1987.

[50] V. E. Johnson. Revised standards for statistical evidence. *Proceedings of the National Academy of Sciences*, 2013.

[51] E. Kalai and E. Lehrer. Weak and strong merging of opinions. *Journal of Mathematical Economics*, 23(1):73–86, 1994.

[52] E. Kandel and N. D. Pearson. Differential Interpretation of Public Signals and Trade in Speculative Markets. *Journal of Political Economy*, 103(4):831, 1995.

[53] A. Kirman. Ants, Rationality, and Recruitment. *The Quarterly Journal of Economics*, 1993.

[54] Z. Kunda. The case for motivated reasoning. *Psychological Bulletin*, 1990.

[55] B. G. Lindsay. Mixture Models: Theory, Geometry, and Applications. *NSF-CBMS Regional Conference Series in Probability and Statistics*, 1995.

[56] C. G. Lord, L. Ross, and M. R. Lepper. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 1979.

[57] V. Mahajan and R. A. Peterson. *Models for Innovation Diffusion*. Sage Publications, Beverly Hills, 1985.

[58] Z. Maniadis, F. Tufano, and J. A. List. To Replicate or Not To Replicate? Exploring Reproducibility in Economics through the Lens of a Model and a Pilot Study. *Economic Journal*, 2017.

[59] R. A. J. Matthews. Beyond 'significance': principles and practice of the Analysis of Credibility. *Royal Society Open Science*, 5(1), 2018.

[60] B. B. McShane, D. Gal, A. Gelman, C. Robert, and J. L. Tackett. Abandon Statistical Significance. 2018.

[61] P. Milgrom and N. Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

[62] T. J. H. Morganand and K. N. Laland. The biological bases of conformity, 2012.

[63] S. Morris. The Common Prior Assumption in Economic Theory. *Economics and Philosophy*, 11(02):227, 1995.

[64] E. Mossel, A. Sly, and O. Tamuz. Asymptotic learning on Bayesian social networks. *Probability Theory and Related Fields*, 2014.

[65] S. Mullainathan and A. Shleifer. The market for news. *American Economic Review*, 95(4):1031–1053, 2005.

[66] K. Munshi. Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution. *Journal of Development Economics*, 2004.

[67] B. A. Nosek, C. R. Ebersole, A. C. DeHaven, and D. T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 2018.

[68] Open Science Collaboration. The reproducibility of psychological science. *Science*, 2012.

[69] F. Prinz, T. Schlange, and K. Asadullah. Believe it or not: How much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery*, 2011.

[70] M. Rabin and J. L. Schrag. First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics*, 114(1):37–82, 1999.

[71] F. P. Ramsey. Truth and Probability. *The Foundations of Mathematics and Other Logical Essays*, 1931.

[72] A. Rubinstein and A. Wolinsky. On the logic of "agreeing to disagree" type results. *Journal of Economic Theory*, 51(1):184–193, 1990.

[73] D. Samet. Ignoring ignorance and agreeing to disagree. *Journal of Economic Theory*, 52(1):190–207, 1990.

[74] L. J. Savage. *The Foundations of Statistics*. John Wiley and Sons, 1954.

[75] M. Schlosshauer, J. Kofler, and A. Zeilinger. A snapshot of foundational attitudes toward quantum mechanics. *Studies in History and Philosophy of Science Part B - Studies in History and Philosophy of Modern Physics*, 2013.

[76] R. Sethi and M. Yildiz. Public disagreement. *American Economic Journal: Microeconomics*, 2012.

[77] R. Sethi and M. Yildiz. Perspectives, Opinions, and Information Flows. *SSRN Electronic Journal*, 2013.

[78] R. Sethi and M. Yildiz. Communication With Unknown Perspectives. *Econometrica*, 84(6), 2016.

[79] J. P. Simmons, L. D. Nelson, and U. Simonsohn. False-Positive Psychology. *Psychological Science*, 22(11):1359–1366, 2011.

[80] L. Smith and P. Sorensen. Pathological Outcomes of Observational Learning. *Econometrica*, 68(2):371–398, 2000.

[81] H. Teicher. Identifiability of Finite Mixtures. *The Annals of Mathematical Statistics*, 34(4):1265–1269, 1963.

[82] J. Tirole. On the Possibility of Speculation under Rational Expectations. *Econometrica*, 1982.

[83] D. Trafimow, V. Amrhein, C. N. Areshenkoff, C. Barrera-Causil, E. J. Beh, Y. Bilgiç, R. Bono, M. T. Bradley, W. M. Briggs, H. A. Cepeda-Freyre, S. E. Chaigneau, D. R. Ciocca, J. Carlos Correa, D. Cousineau, M. R. de Boer, S. S. Dhar, I. Dolgov, J. Gómez-Benito, M. Grendar, J. Grice, M. E. Guerrero-Gimenez, A. Gutiérrez, T. B. Huedo-Medina, K. Jaffe, A. Janyan, A. Karimnezhad, F. Korner-Nievergelt, K. Kosugi, M. Lachmair, R. Ledesma, R. Limongi, M. T. Liuzza, R. Lombardo, M. Marks, G. Meinlschmidt, L. Nalborczyk, H. T. Nguyen, R. Ospina, J. D. Perezgonzalez, R. Pfister, J. J. Rahona, D. A. Rodríguez-Medina, X. Romão, S. Ruiz-Fernández, I. Suarez, M. Tegethoff, M. Tejo, R. van de Schoot, I. Vankov, S. Velasco-Forero, T. Wang, Y. Yamada, F. C. Zoppino, and F. Marmolejo-Ramos. Manipulating the alpha level cannot cure significance testing. *PeerJ Preprints*, 6:e3411v2, apr 2018.

[84] E. Van den Steen. Overconfidence by Bayesian-Rational Agents. *Management Science*, 57(5):884–896, 2011.

[85] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

[86] A. Wilson. Bounded Memory and Biases in Information Processing. *Econometrica*, 82(6), 2014.

[87] S. Yakowitz and J. Spragins. On the Identifiability of Finite Mixtures. *Annals of Mathematical Statistics*, 39:209–214, 1968.

[88] H. P. Young. Innovation diffusion in heterogeneous populations: Contagion, social influence, social learning. *American Economic Review*, 99(5):1899–1924, 2009.

[89] K. J. S. Zollman. The Communication Structure of Epistemic Communities. *Philosophy of Science*, 74(5):574–587, 2007.

[90] K. J. S. Zollman. Network epistemology: Communication in epistemic communities. *Philosophy Compass*, 8(1):15–27, 2013.