

Naive Social Learning, Mislearning, and Unlearning

Tristan Gagnon-Bartsch
Harvard University

Matthew Rabin*
Harvard University

This version: February 26, 2022

First circulated: November, 2015

Abstract

We study social learning in several natural, yet under-explored, environments among people who naively think each predecessor's action reflects solely that person's private information. Such naivete leads to striking forms of mislearning, yielding states of the world that agents always come to disbelieve even when true. In such states, even when an early generation learns the truth, later generations will “unlearn” and develop false beliefs. We demonstrate manifestations of this result in a variety of settings. When the qualities of alternatives are independent, naive inference polarizes perceptions: people overestimate the quality of one option and underestimate the quality of all the others. We show how these “extreme” perceptions lead to under-diversified investments and excessive herding among consumers with diverse tastes.

JEL Classification: D82; D83; D91.

Keywords: Observational learning; Naive inference; Herding.

*Department of Economics and Business School, Harvard University. E-mails: gagnonbartsch@fas.harvard.edu and matthewrabin@fas.harvard.edu. We are grateful to Krishna Dasaratha, Erik Eyster, Gilat Levy, Ronny Razin, Josh Schwartzstein, Neil Thakral, He Yang, and seminar participants at Harvard-MIT, London School of Economics, UC Berkeley, and the CSEF-IGIER Symposium (Capri) for helpful feedback. We thank the Centre for Experimental Research on Fairness, Inequality and Rationality (FAIR) at the Norwegian School of Economics (NHH), for financial support and hospitality while conducting research on this paper.

1 Introduction

People often rely on the behavior of others to inform their decisions. Consumers select restaurants and films based on the crowds they draw, investors adopt financial strategies based on their friends’ portfolios, and both doctors and patients consider the popularity of a drug when choosing prescriptions. Understanding how knowledge spreads through observational learning is the premise of a large literature, starting with Banerjee (1992) and Bikhchandani et al. (1992). They emphasize how rationality might lead people to imitate others. Eyster and Rabin (2014) show, however, that in virtually all settings, rationality demands limited imitation (and often *anti*-imitation), since imitation by earlier people renders the information in group behavior massively redundant. The rationality needed to fully discount this redundancy when we see all our neighbors select the same restaurant, stock, or drug seems unrealistic. As such, a growing literature (e.g., DeMarzo et al., 2003; Eyster and Rabin, 2010, 2014; Dasaratha and He, 2020) studies forms of “naive” social learning.¹ In these models, people neglect the redundancies in the information gleaned from others’ behavior, which causes excessive imitation and overconfidence. This paper explores novel implications of naive social learning that emerge in an array of environments beyond those previously studied.

To first develop an intuition for naive inference, consider privately informed agents who sequentially choose from a set of options with unknown payoffs. Each agent observes the choices of all those before acting her, but she neglects that her predecessors are themselves learning from *their* predecessors. Thus, she treats each observed action as if it reflects solely that predecessor’s private information. For instance, a naive consumer thinks *each* customer buying a popular new product has an independent positive signal about the product, and thus forms overly strong beliefs about its quality. Eyster and Rabin (2010) (henceforth ER) formalize this logic and show that in information-rich environments where rational players surely learn the truth, naive players likewise grow fully confident—but with positive probability in the wrong thing.

By examining a different, yet natural, range of environments than earlier papers, we reveal further principles of naive learning, and we use these principles to draw out welfare implications and comparative statics across several applications. Our primary insight is that naivete sharply limits the states of the world in which society may come to believe. Specifically, there can exist “abandoned states” that people *always* come to disbelieve even when they are true. In fact, even when an early generation knows for certain that such a state is true, later generations “unlearn”: they become convinced of something false. The nature of these abandoned states generates various economic im-

¹In addition to DeMarzo et al. (2003) and Dasaratha and He (2020), several other papers study heuristic rules of thumb that embed redundancy neglect in network settings. Such papers, which typically build from the canonical model by DeGroot (1974), include Golub and Jackson (2010); Levy and Razin (2018); Molavi et al. (2018); Chandrasekhar et al. (2020) and Mueller-Frank and Neri (2021). There are also papers studying other biases in social learning that are distinct from the notion of redundancy neglect; see, e.g., Guarino and Jehiel (2013); Gagnon-Bartsch (2016); Bohren and Hauser (2021) and Frick et al. (2020). We further discuss related literature in Section 2.3.

plications. For instance, when inferring the payoff difference between two independent alternatives, naive learners maximally exaggerate this difference. Not only will they grow too confident that a popular medicine (e.g., the name brand) is better than an unpopular one (e.g., a generic), they infer it is *much* better—even when in reality they are nearly identical. We show that such exaggeration can lead to both severe under-diversification in investment settings and to costly herding in markets with diverse tastes. People who naively exaggerate quality differences, for instance, may engage in wasteful queuing or spending to such a degree that they are worse off observing others than if they made decisions in isolation.²

Previous models of naivete obscure these results by focusing on environments with just two states of the world and common preferences. While rational models assume this for analytical ease, the tractability of the naive model allows us to study a broader range of settings, including those that better reflect features of many economic examples of interest. In doing so, we reveal new implications of naivete. One such result is deterministic mislearning: while ER show that people mislearn in the canonical environment only when early signals are misleading, the existence of “abandoned states” in richer settings guarantees mislearning irrespective of early signal realizations. In fact, to articulate the different nature of our mislearning, we consider environments where a large generation of players take actions each round. This implies that the collection of early signals is (essentially) never misleading. Even when we neutralize the primacy effects that drive bounded learning in rational models and drive mislearning in ER, naive inference still leads society astray through agents’ misunderstanding of the informational content of past actions. Put differently, the type of mislearning that we elucidate in this paper can be thought of as *structural mislearning*: whether society mislearns is preordained by the environment itself and does not depend on the happenstance of early signals. Moreover, although our focus on large generations helps make this point stark, our main qualitative results do not require this assumption: we also show that when players act in single file, as in the canonical model, they will still deterministically come to believe something false in any of the states we classify as “abandoned.”

Section 2 introduces our general model. In every period, a new generation of players choose from an identical set of actions and receive payoffs that depend on their own action and an unknown state of the world. Each player updates her beliefs based on predecessors’ behavior and her conditionally independent private signal. Players naively infer as if others’ actions reflect solely their private information and the common prior. While most of our results hold more generally, we focus on a setting with two key features: each generation (i) is large, and (ii) observes only the preceding generation.

²While we identify new ways in which naive learning harms welfare, earlier models already demonstrate that naive observers can be made worse off in expectation by observing others’ choices. This never happens with full rationality. In herd models such as Bikhchandani et al. (1992), information cascades prevent fully-rational players from learning the state. However, rational players are not *harmed* by observing others—the informational externality simply prevents society from reaching the first best.

Together, these assumptions simplify analysis. And the former makes long-run mislearning more surprising: because the first generation consists of a large number of people acting independently, Generation 2 becomes nearly certain of the true state.³

Although Generation 2 learns correctly, naive inference can lead Generation 3 astray. Generation 3 wrongly assumes each person in Generation 2 acts “autarkically”—relying only on her own signal—and fails to realize that each person in Generation 2 in fact has nearly perfect information. As such, Generation 3 comes to believe in the state most likely to generate those signals necessary for autarkic behavior to resemble the behavior of Generation 2.⁴ If this is the true state, then Generation 3 and all subsequent generations learn correctly. Otherwise, Generation 3 “unlearns,” and long-run beliefs never settle on the truth. Whenever public beliefs do converge, they do so to a fixed point of this process: when interpreted as autarkic, the behavior of those nearly certain in state ω is best predicted by ω .

To give an example, consider investors learning about the return on two projects, A and B . Suppose that investors’ priors are such that the payoff of B , ω^B , is a fifty-fifty draw from $\{0, m\}$ and ω^A is an independent fifty-fifty draw from $\{l, h\}$, where $h > m > l > 0$; that is, A ’s high payoff is better than B ’s, and A ’s low payoff is also better than B ’s. Suppose the signal distributions conditional on (ω^A, ω^B) have the following intuitive properties: (i) the percent of investors who would choose A based private information alone is strictly increasing in $\omega^A - \omega^B$, and (ii) this percent exceeds 50% if and only if $\omega^A > \omega^B$. If in truth $(\omega^A, \omega^B) = (h, m)$, then a majority choose A in period 1. From this, investors in period 2 correctly infer that A outperforms B , so *all* choose A . In the third round, investors’ best explanation for such a consensus is that A has the largest possible payoff advantage over B , as this would maximize the likelihood that an investor selects A when using private information alone. Following this logic, Generation 3 comes to believe the state is $(h, 0)$ whenever they see others unanimously pick A —that is, whenever $(\omega^A, \omega^B) \in \{(l, 0), (h, 0), (h, m)\}$. Thus, in states $(l, 0)$ and (h, m) , investors inevitably mislearn by way of exaggerating the payoff difference between the projects.⁵ When investors act based on the extremity of their beliefs—for instance, they can contribute more resources toward A —this exaggeration will generate costly mistakes.

Section 3 discusses general implications of naive inference for long-run beliefs. First, based on the “fixed-point” logic above, we characterize the set of states on which public beliefs can settle. We draw out several implications of this characterization, revealing the extent to which naive inference limits the conclusions society might draw. For instance, the set of stable long-run beliefs may be a

³We discuss the robustness of our results to alternative observation structures at relevant points throughout the paper.

⁴Generally, the confident behavior of Generation 2 will not match the action distribution predicted by autarkic play in any state. Because naive observers attribute such discrepancies to sampling variation, Generation 3 grows certain of the state that best predicts autarkic behavior. With naivete, this state minimizes the cross entropy between the realized action distribution and the one predicted by autarkic play.

⁵In state (l, m) investors correctly learn that B is optimal, as (l, m) is the only state in which the majority selects B .

singleton—in which case society draws the same conclusion no matter what is true—or it may be empty—in which case beliefs continually cycle.

Section 4 presents a simple yet stark implication of naive inference in settings where, as in the investment example above, individuals have common preferences and options have independent payoffs. Namely, perceptions of payoffs grow “polarized”: people come to believe the best option is as good as possible, and all lesser options are as bad as possible. Once a herd starts on option A , people think that each of their many predecessors who took A received an independent signal indicating that A is better than the alternatives. Under natural monotonicity assumptions on the signal structure, this misconception leads to these extreme, polarized beliefs.⁶

Section 5 examines the consequences of extreme beliefs when people have diverse preferences over a set of alternatives (e.g., medical treatments) relative to an outside option (e.g., no treatment). In such settings, naivete can cause excessive and costly herding where all people choose an option that is beneficial for only a minority. Since naivete leads people to exaggerate the quality of popular options, consumers with low valuations are enticed to follow the herd when they should in fact pursue their outside option. Interestingly, we show that such “over-adoption” can be triggered by large choice sets. For example, suppose patients with an illness can either experiment with unproven treatments or abstain, and most patients abstain unless they are fairly confident a treatment works. However, a minority in dire straits will experiment no matter what. If the number of available treatments is sufficiently large, then society comes to believe one is universally beneficial even when none are. In this case, Generation 2 learns that none are fully effective: those who are not desperate will abstain, and those in dire need will herd on the “least bad” treatment. Naive observers, who mistake the consensus among the desperate as independent decisions, will then grow convinced that this treatment works. Namely, with so many options to pick from, they find it unlikely that predecessors’ choices would coincide unless this treatment were truly effective. Hence, all patients—including those for which treatment is suboptimal—end up adopting a false cure.

In Section 6, we apply the model to a portfolio-choice problem and demonstrate how naivete distorts investors’ allocation of wealth across risky assets. Perceptions of an asset’s value continually grow more extreme over time, resulting in severe under-diversification. To illustrate, suppose investors know the expected return of asset A , but learn about B ’s average return, ω^B , from predecessors’ allocations.⁷ In settings where first-period allocations resolve this uncertainty among rational investors, naive investors overreact to later allocations as if they continue to reflect new information.

⁶Because ER considers a binary-state model where the payoff difference between any two actions has only two possible values, their setting obscures our result that naive players maximally exaggerate the payoff difference between the herd action and its alternatives.

⁷An important feature of our model is that people observe their predecessors’ behavior but not the outcomes of these decisions. For this reason, we consider assets that pay off long after the initial investment decision, such as real-estate or education.

The path of naive investment depends on how ω^B compares to expectations. If it beats them, then investors in period 2 correctly infer this and allocate more to B . However, because later investors neglect that this increase stems from social learning, they wrongly attribute it to new, more optimistic information. Investment in B increases yet again. As this process plays out, investors eventually allocate all their wealth to B . Moreover, when ω^B is sufficiently far from priors, investors will eventually allocate everything to the inferior asset. Our prediction of over-reaction to private information and momentum accords with Glaeser and Nathanson (2017) analysis of housing-price data, which suggests that medium-run momentum derives, in part, from naive inference based on past market prices.⁸

We conclude in Section 7. There, we note how our results differ from the predictions of some other forms of naive inference considered in the literature. (See also Section 2.3 for similar discussions.) We additionally propose a few natural extensions that are somewhat beyond the scope of our particular solution concept. These include asset pricing and investment settings where agents endogenously choose when to act.

2 Model

This section presents our baseline model. We first describe the social-learning environment (Section 2.1) and then define naive inference (Section 2.2). We compare our approach to related models in Section 2.3.

2.1 Social-Learning Environment

We consider a sequential-decision environment similar to the canonical models of observational learning developed by Banerjee (1992) and Bikhchandani et al. (1992). In those models, a new player enters in each period and makes a once-and-for-all decision. We expand on those models by allowing for heterogeneous preferences (as in, e.g., Smith and Sørensen, 2000), and we consider multiple players acting per period instead of just one.

In every period $t = 1, 2, \dots$, a new generation of $N \geq 1$ players enters, and each player simultaneously chooses an action from the set $\mathcal{A} \equiv \{A_0, \dots, A_{M-1}\}$ with $M \geq 2$. Each player is labeled by (n, t) , where t is the period in which she acts and $n \in \{1, \dots, N\}$ is her index within Generation t . Let $x_{(n,t)} \in \mathcal{A}$ denote Player (n, t) 's action and let $a_t(m)$ denote the fraction of players in period t who choose A_m . Vector $\mathbf{a}_t = (a_t(0), \dots, a_t(M-1))$ is the distribution of actions chosen in period t .⁹

⁸In a more recent paper, Bastianello and Fontanier (2021) build on this idea to derive more detailed implications of naive inference in asset markets.

⁹We extend the model to a continuous action space in one of our applications.

States and Preferences. Players aim to learn a payoff-relevant state of nature $\omega \in \{\omega_1, \omega_2, \dots, \omega_K\} \equiv \Omega$, where $K \geq 2$. Players share a common prior $\pi_1 \in \Delta(\Omega)$, where $\pi_1(k) > 0$ denotes the probability of state ω_k . Each Player (n, t) 's payoff from action A_m depends on ω and her “preference type” $\theta_{(n,t)} \in \{\theta_1, \theta_2, \dots, \theta_J\} \equiv \Theta$; payoffs are denoted by $u(A_m | \omega, \theta)$. Each player’s preference type is privately known, and we assume these types are i.i.d. across players according to a commonly known distribution $\lambda \in \Delta(\Theta)$. Given that we will focus on large generations (i.e., $N \rightarrow \infty$), we make the inconsequential simplifying assumption that the realization of preference types is fixed over time: in Generation $t = 1$, $\theta_{(n,1)}$ is drawn according to λ for each $n = 1, \dots, N$, and $\theta_{(n,t)} = \theta_{(n,1)}$ for all $t > 1$.¹⁰

Private Information. Players learn about ω from two channels: private signals and observing others. Each Player (n, t) receives a random private signal $s_{(n,t)} \in \mathbb{R}^d$, $d \geq 1$, about the state of nature. Conditional on state ω , private signals are i.i.d. across players with c.d.f. $F(\cdot | \omega)$ and density (or mass) function $f(\cdot | \omega)$. We assume these distributions have identical support $\mathcal{S} \subset \mathbb{R}^d$ for each $\omega \in \Omega$. We will place some mild restrictions on the signal structure below.

Public Information. We focus on an observation structure that differs from canonical models in two ways. First, we assume players observe the behavior of only the previous generation. Each Player (n, t) 's information set $I_{(n,t)} \equiv \{s_{(n,t)}, \mathbf{a}_{t-1}\}$ consists of her private signal and the distribution of actions in $t - 1$. Second, we focus on the limit in which each generation grows large (i.e., $N \rightarrow \infty$), which—as we show below—ensures that each generation reaches a nearly confident consensus on the state.¹¹

This setting can be interpreted as a series of large overlapping generations, where each generation is present for two periods: in the first, individuals observe the actions of the preceding generation; in the second, they take actions based on inferences from this observation. While we highlight the implications of our setup in detail below (end of Section 2.2), it is worth emphasizing the rationale for this setup. First, it will streamline the analysis by generating (nearly) deterministic belief and action dynamics. Second, it stacks the deck in favor of correct learning: if naive agents fail to learn from a large population whose behavior reveals the state to rational observers, then they certainly won’t learn in canonical observation structures that generate weaker public signals. Below, we explicitly show that our central result restricting the states that naive agents can learn indeed extends to other

¹⁰With large generations, the realized distribution of preference types within any generation is nearly identical across periods even when types are redrawn each round; this follows trivially from the law of large numbers (LLN). This assumption therefore simplifies the exposition by reducing the need to continually invoke the LLN. We drop this assumption when we consider robustness to settings without large generations.

¹¹In canonical social-learning environments (e.g., Banerjee, 1992; Bikhchandani et al., 1992; Smith and Sørensen, 2000), players observe the complete history of actions and take actions in “single file” (i.e., $N = 1$). Our observation structure with “large generations” is not typically studied because it is uninteresting under rationality: rational agents immediately learn the state after a single round of actions. We emphasize that with naive agents, this structure will lead each generation to grow confident in some state, but not necessarily the correct one.

observation structures, including the canonical one in which players act in single file and observe *all* previous actions in order (e.g., Bikhchandani et al., 1992; Smith and Sørensen, 2000; Eyster and Rabin, 2010). Hence, we focus on our “overlapping-generations” structure solely for tractability and to starkly highlight how the dynamics of rational and naive beliefs differ.

We define the *public belief* as the belief a player would form solely from observed actions (i.e., net of her private signal). Upon observing \mathbf{a}_{t-1} , Generation t forms public belief $\pi_t \in \Delta(\Omega)$ using Bayes’ rule:

$$\pi_t(j) = \frac{\Pr(\mathbf{a}_{t-1}|\omega_j)\pi_1(j)}{\sum_{k=1}^K \Pr(\mathbf{a}_{t-1}|\omega_k)\pi_1(k)}. \quad (1)$$

Our model of naivete will posit that players have erroneous beliefs about $\Pr(\mathbf{a}_{t-1}|\omega)$ —the distribution of predecessors’ behavior in a given state. Aside from this mistake, each Player (n, t) is rational: she uses Bayes’ Rule to combine the public belief with her signal to form posterior $p_{(n,t)} \in \Delta(\Omega)$, and then takes the action that maximizes her expected utility. Hence, $x_{(n,t)} = \arg \max_{A \in \mathcal{A}} \sum_{k=1}^K p_{(n,t)}(k)u(A|\omega_k, \theta_{(n,t)})$.

Finally, to avoid trivial (and non-generic) complications arising from indifference, we assume each preference type $\theta \in \Theta$ has a unique preferred action when the state is known.

Assumption 1. *For each state $\omega \in \Omega$ and each type $\theta \in \Theta$, the set of optimal actions, $\arg \max_{A \in \mathcal{A}} u(A|\omega, \theta)$, is single-valued.*

Additionally, whenever a player is indifferent between actions due to uncertainty over ω , we assume she follows a commonly known tie breaking rule.¹²

2.2 Naive Social Inference

Following Eyster and Rabin (2010), we assume individuals naively think that any predecessor’s action depends solely on that player’s private information. This implies that a naive agent infers from past actions *as if* all her predecessors ignored the history of play and hence learned nothing from others’ actions. That is, she infers as if all her predecessors acted in “autarky”. Hence, a naive player has an erroneous model of $\Pr(\mathbf{a}_{t-1}|\omega)$ —the likelihood of predecessors’ actions conditional on the state—and consequently form biased public beliefs (Equation 1). In reality, this likelihood depends on the beliefs of Generation $t - 1$, which in turn depend on what Generation $t - 1$ observed, and so on. Naive players, however, neglect the social inference conducted by preceding generations and infer from \mathbf{a}_{t-1} as if it reflects solely the private signals of those players acting in Generation $t - 1$.¹³

¹²The specific rule is irrelevant for our results, but in specific applications we assume players break indifference by selecting the option with the lower index.

¹³Naive players fail to realize that past behavior (in $t \geq 2$) already incorporates all useful private information. In simple single-file settings, this generates over-weighting of early signals.

To formally define and analyze naive learning, we will work directly with “action distributions.” These distributions over the action space, \mathcal{A} , specify the frequency that each action A_m is chosen as a function of (i) the state, (ii) the distributions of signals and preferences, and (iii) prior beliefs. The action distributions will represent a convenient way to summarize primitives of the model (signals, preferences, and prior beliefs), and will in fact be sufficient for much of our analysis.

Two types of action distributions are of particular importance. The first is what we call an *autarkic action distribution*: the theoretical distribution of actions in state ω assuming players act solely on private signals and the prior, π_1 .

Definition 1. *The autarkic distribution $\mathbb{P}_\omega \in \Delta^M$ is the distribution of actions generated by autarkic play in state ω : $\mathbb{P}_\omega(m)$ is the probability that a player chooses option A_m when her beliefs are based solely on her signal and the prior.¹⁴*

In state ω , the actions taken by the first generation will follow the autarkic distribution \mathbb{P}_ω regardless of whether players are rational or naive. A naive player, however, expects behavior to be autarkic in every period.¹⁵

Definition 2. *A naive player infers from each predecessor’s action as if, conditional on ω , it were an independent draw from \mathbb{P}_ω .*

Although naive players expect to see behavior that follows an autarkic distribution in each round, they actually see behavior of predecessors who have themselves formed confident beliefs about the state by observing those before them. Thus, the second type of action distribution central to our analysis is what we call an *aggregated-signal distribution*: the theoretical distribution of actions taken by players who are certain (either rightly or wrongly) that the state is ω . The aggregated-signal distribution is like an autarkic distribution, but the prior belief π_1 is replaced with a degenerate belief on ω .

Definition 3. *The aggregated-signal distribution $\mathbb{T}_\omega \in \Delta^M$ is the distribution of actions generated by players who put probability 1 on ω : $\mathbb{T}_\omega(m)$ is the probability that a player chooses option A_m when certain the state is ω .¹⁶*

¹⁴More formally,

$$\mathbb{P}_\omega(m) = \int_{\mathcal{S}} \sum_{\theta \in \Theta} \lambda(\theta) \zeta(m|\theta, s, \pi_1) dF(s|\omega),$$

where $\zeta(m|\theta, s, \pi_1)$ is the probability that type θ chooses action A_m when relying solely on her private signal s and the prior π_1 . Given a fixed tie-breaking rule, $\zeta(m|\theta, s, \pi_1)$, and hence \mathbb{P}_ω , are well defined.

¹⁵The definition of naivete originally proposed by Eyster and Rabin (2008, 2010)—which they call “Best-Response Trailing Naive Inference” (BRTNI)—posits that naive players best respond to the belief that all others are fully cursed in the sense of Eyster and Rabin’s (2005) “cursed equilibrium”.

¹⁶The “aggregated-signal distribution” reflects the behavior of a player who has observed an arbitrarily large collection of independent signals drawn from $F(\cdot|\omega)$, assuming the state is identifiable from signals. More formally, letting δ_ω denote a degenerate belief on state ω , $\mathbb{T}_\omega(m) = \sum_{\theta \in \Theta} \lambda(\theta) \zeta(m|\theta, \delta_\omega)$ where $\zeta(m|\theta, \delta_\omega)$ is the probability that type θ chooses action A_m when certain the state is ω .

The families of both the autarkic distributions, $\{\mathbb{P}_\omega\}_{\omega \in \Omega}$, and the aggregated-signal distributions, $\{\mathbb{T}_\omega\}_{\omega \in \Omega}$, depend solely on the primitives of the model. A main result below shows that the relationship between these two families dictates which states a naive society can come to learn.

Finally, we make three mild assumptions on the information structure.

Assumption 2. *The collection of signal distributions $\{F(\cdot|\omega)\}_{\omega \in \Omega}$ is such that:*

1. (Full Support.) *For all $\omega \in \Omega$, \mathbb{P}_ω has full support over \mathcal{A} .*
2. (Identifiability.) *For all $\omega, \omega' \in \Omega$, $\mathbb{P}_\omega \neq \mathbb{P}_{\omega'}$ whenever $\omega \neq \omega'$.*
3. (Bounded Signals.) *There exists a finite $\beta \in \mathbb{R}_+$ such that for any signal $s \in \mathcal{S}$ and any two states ω and ω' , the log-likelihood ratio is such that*

$$-\beta \leq \log \left(\frac{\Pr(\omega|s)}{\Pr(\omega'|s)} \right) \leq \beta.$$

Assumption 2.1 implies that, for each action $A_m \in \mathcal{A}$, there exists a possible signal realization that would lead some type to choose A_m based on that signal alone. Thus, naive players think each action is taken with positive probability in autarky. We impose this assumption simply to ensure that naive players never observe actions they thought were impossible.

Assumption 2.2 implies that autarkic distributions are distinct across states. With large generations, this ensures that naive players think that \mathbf{a}_{t-1} identifies the state as N grows large.

Assumption 2.3 implies that signals have bounded informativeness. Consequently, if the public belief in any Generation t puts sufficiently high probability on state $\hat{\omega}$, then the posteriors of all agents in Generation t will continue to put high probability on $\hat{\omega}$ no matter their private signals. As we formalize in Lemma 1 below, this induces (with high likelihood) a deterministic path of public beliefs and actions across generations. In particular, if the observed behavior from Generation 1 causes Generation 2 to put sufficient weight on the true state (which happens with arbitrarily high probability as $N \rightarrow \infty$ under Assumption 2.2), then private signals will have no bearing on decisions after Generation 1. Thus, the path of public beliefs and actions across generations is entirely determined by the initial condition (i.e., the state the second generation infers from the first).¹⁷

The purpose of these assumptions is two fold. First, the deterministic dynamics they induce are very simple to analyze. Second, they illuminate how naive inference can generate learning failures

¹⁷Our bounded-signal assumption is mild given that the constant β in Assumption 2.3 is arbitrary. We therefore accommodate signals of any strength so long as there exists some upper bound to their informativeness. If we allowed for unbounded signals, our analysis would need to account for the vanishingly unlikely event that the public belief in some Generation t puts high probability on ω yet the behavior in Generation t is far from \mathbb{T}_ω (in the sense of cross entropy, defined below in Section 3). The probability of this event is non-zero with unbounded signals since it is possible for a fraction of players within a single generation to receive arbitrarily strong yet misleading signals.

in even the most favorable information structures. Our identifiability assumption (2.2) with large generations implies that learning is immediate for rational agents. Despite this, we show that naive agents may continually disbelieve the true state. Thus, the forms of misinference we identify do not stem from actions revealing insufficient data (as in, e.g., Bikhchandani et al., 1992 or Smith and Sørensen, 2000). Rather, they emerge in data-rich settings where we might expect learning to occur. Given that naive agents fail to learn in our data-rich settings, one might expect a similar failure in the canonical setting where just a single agent acts per round; we indeed verify this below in Proposition 3.

To preview the analysis in the remainder of the paper, we sketch how beliefs evolve in our model when the state is ω . Since Generation 1 acts solely on private information, their distribution of actions, \mathbf{a}_1 , converges to \mathbb{P}_ω as $N \rightarrow \infty$. Next, our identifiability assumption implies that Generation 2, whether rational or naive, grows arbitrarily certain of ω upon observing \mathbf{a}_1 . Each player in Generation 2 then takes the action that is optimal for her type in state ω , and thus \mathbf{a}_2 converges to \mathbb{T}_ω as $N \rightarrow \infty$. Rational followers understand that \mathbf{a}_2 reflects the behavior of agents who essentially know the state. Hence, so long as \mathbb{T}_ω is unique, rational public beliefs and behavior immediately converge: if they are rational, Generations 3 and beyond will continue to believe the state is ω .

While rational learning is trivial in this environment, a naive Generation 3 may grow certain of a state different from ω . Because they think others act solely on private signals, a naive Generation 3 infers from \mathbf{a}_2 as if it resulted from autarkic play. Hence, they ask themselves what distribution of signals Generation 2 must have received in order to take actions \mathbf{a}_2 in autarky. Generation 3 then becomes nearly certain of the state most likely to generate those signals. This inference is flawed since Generation 2 did not act in autarky—they learned from their predecessors. A similar mistake will play out among each successive generation: Generation $t \geq 3$ observes the behavior \mathbf{a}_{t-1} of a prior generation who is in fact (nearly) certain of some state, but wrongly interprets \mathbf{a}_{t-1} as if it were based on weak private signals. In the language of our formal setup, naive agents treat \mathbf{a}_{t-1} as a multinomial random variable governed by some autarkic distribution \mathbb{P}_ω when in fact it is governed by some aggregated-signal distribution, \mathbb{T}_ω . As such, the relationship between these classes of distributions will determine what naive agents come to believe. We formalize this relationship in Section 3.

2.3 Related Models

Before turning to long-run dynamics in Section 3, we briefly review how our analysis compares to both ER and other papers on naive social learning. Given that we adopt ER’s model of naivete, our approach differs from theirs primarily in the type of environments we consider. Other papers, however, examine quite different forms of naive learning.

ER explores a binary-state model with a continuum of actions, common preferences, and one player acting per round. Specifically, $\Omega = \{0, 1\}$, $\mathcal{A} = [0, 1]$, and $u(x|\omega) = -(x - \omega)^2$. With these preferences, a player optimally chooses $x \in [0, 1]$ equal to her belief that $\omega = 1$. This implies that actions perfectly reveal an agent’s posterior belief. ER’s main result is that, with positive probability, a naive society grows confident in the wrong state. A naive player in period t treats the announced posterior of the player in $t - 1$ as that player’s independent signal, despite the fact that it also incorporates the signals of players in all prior periods. As such, players vastly over-count early signals. If early signals are sufficiently misleading—which happens with positive probability—then players grow confident in the wrong state.¹⁸

Naivete leads society astray in ER’s environment only when early signals are misleading. However, with our “large generations” assumption, early signals are (essentially) never misleading. As such, when large populations act each round in ER’s two-state setting, a naive society always converges to the truth. In contrast, we emphasize that in other natural environments, naive observers may still converge to false beliefs even when early generations perfectly reveal the state.

Additionally, ER’s two-state framework implies that if people learn the payoff of action $x = 1$, then they implicitly learn the payoff of all other actions. In contrast, we consider settings where payoffs are independent across actions: knowledge that x is superior to x' does not reveal *by how much* x is preferred to x' . Such a distinction matters, for instance, when deciding how much to pay, or how long to wait, to obtain x over x' . As we will show, naive agents systematically overestimate the payoff of the herd action relative to those not chosen. In this sense, naive inference restricts which constellation of payoffs agents may come to believe.

Our model of naive inference—adopted from ER—is related to several other approaches that are similarly motivated by the intuition that people neglect informational redundancies when learning from others. DeMarzo et al. (2003) propose a model of “persuasion bias” in which neighbors in a network communicate posterior beliefs. Building on DeGroot’s (1974) model of consensus formation, they assume players form posteriors by taking the average of neighbors’ beliefs as if they reflect independent signals with known precision. Since players neglect that stated beliefs already incorporate signals previously shared, they over-count early signals.¹⁹ Our model is also related to Level- k thinking (e.g., Crawford and Iriberri, 2007). Naive agents act like Level-2 thinkers, as they best respond to the belief that others use only private information (Level-1).²⁰ Additionally, Bohren

¹⁸Eyster and Rabin (2014) show that this intuition holds even when players over-count predecessors’ signals by any arbitrarily small amount. These results stand in sharp contrast with the rational model, in which wrong herds are likely to occur only in settings where players remain relatively uncertain of the state. Rational models never lead to confident beliefs in a false state; they are not compelling models of society thinking it knows things it does not.

¹⁹For more recent papers that study generalizations of the DeGroot learning rule, see Golub and Jackson (2010); Levy and Razin (2018); Molavi et al. (2018); Chandrasekhar et al. (2020) and Mueller-Frank and Neri (2021).

²⁰The equivalence between naive inference as defined by ER and Level-2 thinking is not general, but it happens to hold in many models of social learning where players are concerned with others’ irrationality only to the extent that it

(2016) studies a variant of the canonical two-state model in which only a fraction $\alpha \in [0, 1]$ of players can observe past actions and players have wrong beliefs about α . This model corresponds with ours when $\alpha = 1$ and all players think $\hat{\alpha} = 0$. Similar to our argument about ER, agents in Bohren’s (2016) model converge to the truth as generations grow large.

Empirically, few experiments are designed to specifically test ER’s model of naivete. ER describe how findings from earlier experimental work, like Kübler and Weizsäcker (2004), suggest that people neglect informational redundancies in social learning.²¹ Eyster et al. (2018) find more direct evidence. They first tell each subject the difference in the number of heads and tails from 100 independent flips of a coin. Then, moving in sequence, each subject estimates the total difference in heads and tails across all predecessors—including herself—and announces this estimate. A Bayesian Nash equilibrium strategy is to add one’s own 100-trial sample to her immediate predecessors’ estimate.²² However, they find a weak tendency towards redundancy neglect: subjects fail to fully understand that the most recent predecessor’s behavior incorporates the information of earlier predecessors. They also show severe redundancy neglect in a treatment where four independent players in each round are asked to derive these sums.

3 Belief Dynamics, Long-Run Beliefs, and Abandoned States

This section characterizes belief dynamics under naive inference and presents some general implications for long-run learning. In particular, we characterize the possible limit beliefs that society may come to hold, revealing that there may exist states that people always come to disbelieve, even when they are true. When such an “abandoned state” occurs, society *unlearns*: Generation 2 is very likely to learn the true state, yet later generations become convinced of something false. Furthermore, we show that our classification of “abandoned states” is robust in the sense that these states will not be assigned high probability in the long run even in the canonical environment where only a single agent takes an action in each period.

3.1 Characterization of Naive Belief Dynamics

We first formalize the intuition that, if generations grow large, then with arbitrarily high probability the public belief will transition from one generation to the next in a deterministic manner. Fix an arbitrarily small $\epsilon > 0$. We then define a deterministic *belief-transition function* $\phi : \Omega \rightarrow \Omega$ as

interferes with inference.

²¹Enke and Zimmermann (2017) also find evidence that people fail to account for redundancies in information.

²²In fact, determining the optimal action in this setting does not even require proper use of Bayes’ Rule: the Bayesian Nash equilibrium corresponds to the unique Iterated-Weak-Dominance (IWD) outcome, and in this IWD strategy profile no players need to apply Bayes’ Rule.

follows: if all agents in Generation t believe the probability of state ω is at least $1 - \epsilon$, then all agents in Generation $t + 1$ believe the probability of state $\phi(\omega)$ is at least $1 - \epsilon$. In our first result below, we (i) show that there exists such a transition function that describes the evolution of naive beliefs when ϵ is sufficiently small; and (ii) derive its functional form.²³

Under naive inference, the transition function ϕ is characterized by the solution to a particular distance-minimization problem between the autarkic and aggregate-signal distributions. Namely, if Generation t is confident that the state is $\hat{\omega}_t$, then Generation $t + 1$ grows confident in the state $\hat{\omega}_{t+1}$ whose autarkic distribution, $\mathbb{P}_{\hat{\omega}_{t+1}}$, is closest to the observed distribution, $\mathbb{T}_{\hat{\omega}_t}$, in terms of cross entropy.

Definition 4. *The cross-entropy distance between the observed distribution $\mathbb{T} \in \Delta^M$ and the predicted distribution $\mathbb{P} \in \Delta^M$ is defined as*

$$H(\mathbb{T}, \mathbb{P}) = - \sum_{m=0}^{M-1} \mathbb{T}(m) \log \mathbb{P}(m).^{24} \quad (2)$$

To avoid trivial complications, we assume that no aggregated-signal distribution lies perfectly “between” two distinct autarkic distributions.

Assumption 3. *For any $\omega \in \Omega$, $H(\mathbb{T}_\omega, \mathbb{P}_{\omega'}) = H(\mathbb{T}_\omega, \mathbb{P}_{\omega''})$ if and only if $\omega' = \omega''$.*

Assumption 3 ensures that our transition function ϕ is single valued when ϵ is sufficiently small, and, given that it fails only for knife-edge cases, it is relatively innocuous. We can now characterize ϕ as follows.

Lemma 1. *For $\epsilon > 0$ sufficiently small and N sufficiently large, beliefs evolve according to a well-defined belief-transition function $\phi : \Omega \rightarrow \Omega$ such that if all agents in Generation t place probability exceeding $1 - \epsilon$ on some state $\hat{\omega}_t \in \Omega$, then all agents in Generation $t + 1$ place probability exceeding $1 - \epsilon$ on state $\phi(\hat{\omega}_t) \in \Omega$, where ϕ is defined by*

$$\phi(\hat{\omega}_t) = \underset{\omega \in \Omega}{\operatorname{argmin}} H(\mathbb{T}_{\hat{\omega}_t}, \mathbb{P}_\omega).$$

To provide intuition, suppose all players in Generation t think $\hat{\omega}_t \in \Omega$ is true with probability exceeding $1 - \epsilon$. For ϵ sufficiently small, these players choose actions as if $\hat{\omega}_t$ were true. Thus, \mathbf{a}_t converges

²³The function ϕ depends on ϵ . However, since all of our results will take ϵ as fixed, we will not explicitly denote this dependence for the sake of notational ease.

²⁴We refer to the cross-entropy colloquially as a distance even though it is not symmetric and thus not a proper metric. The well-known Kullback-Leibler divergence (or “relative entropy”) of \mathbb{P} from \mathbb{T} in terms of cross entropy is simply $H(\mathbb{T}, \mathbb{P}) - H(\mathbb{T}, \mathbb{T})$. These measures are common in both information theory (e.g., Cover and Thomas, 2006) and the literature on learning with misspecified models (e.g., Berk, 1966; Schwartzstein, 2014; Esponda and Pouzo, 2016; Bohren and Hauser, 2021).

to $\mathbb{T}_{\hat{\omega}_t}$ as $N \rightarrow \infty$. Since the informativeness of \mathbf{a}_t swamps that of any individual private signal as N grows large (Assumption 2.3), players in Generation $t + 1$ will therefore put probability exceeding $1 - \epsilon$ on the state most likely to yield $\mathbf{a}_t \approx \mathbb{T}_{\hat{\omega}_t}$ through autarkic behavior.

Lemma 1 implies that beliefs evolve deterministically according to ϕ from Generation 2 onward so long as the required initial condition is met; that is, so long as all agents in Generation 2 assign probability exceeding $1 - \epsilon$ to some state. We can further show that this initial condition will be met with arbitrarily high probability: by the Strong Law of Large Numbers, $\mathbf{a}_1 \rightarrow \mathbb{P}_\omega$ a.s. as $N \rightarrow \infty$, and, since \mathbb{P}_ω is unique (Assumption 2.2), observing $\mathbf{a}_1 \approx \mathbb{P}_\omega$ essentially reveals the state to Generation 2. An argument analogous to the proof of Lemma 1 then implies that all agents in Generation 2 believe the probability of ω is at least $1 - \epsilon$ when N is sufficiently large. Thus, in light of Lemma 1, there is an arbitrarily high probability that public beliefs evolve deterministically from Generation 2 onward when N is sufficiently large.

For the remainder of the paper, we implicitly fix an arbitrarily small ϵ and consider a sufficiently large N such that the transition function in Lemma 1 applies. We therefore focus on the high-probability event in which public beliefs evolve according to ϕ . As such, let $\hat{\omega}_t$ denote the state that Generation t assigns high probability (i.e., above $1 - \epsilon$). We will call $\hat{\omega}_t$ the “public belief” in Generation t . It is worth noting that this terminology is slightly imprecise: the public belief in Generation t is not degenerate on $\hat{\omega}_t$, but it does put arbitrarily high weight on $\hat{\omega}_t$. This is what we mean when we say Generation t “grows confident” in $\hat{\omega}_t$.²⁵

To summarize, we study the process of public beliefs, $\langle \hat{\omega}_t \rangle$, that is (i) defined by $\hat{\omega}_{t+1} = \phi(\hat{\omega}_t)$, and (ii) starts from the initial condition in which Generation 2 puts arbitrarily high probability on the true state; that is, $\hat{\omega}_2 = \omega$.

Naivete can lead society astray beginning in Generation 3: since they neglect that Generation 2 learned from Generation 1, Generation 3 treats \mathbf{a}_2 as if it reflects autarkic play when in fact \mathbf{a}_2 reflects actions taken with near perfect knowledge of ω (i.e., $\mathbf{a}_2 \approx \mathbb{T}_\omega$). As previewed in Section 2, Generation 3 will then (i) infer what distribution of signals Generation 2 must have received in order to take actions \mathbf{a}_2 under autarkic play, and (ii) come to believe in the state $\hat{\omega}_3 \in \Omega$ most likely to generate those signals.²⁶

It is noteworthy, albeit immediate, that naivete gives rise to “unlearning” across generations. Al-

²⁵Since π_t will assign an arbitrarily high weight to $\hat{\omega}_t$ as $N \rightarrow \infty$, $\hat{\omega}_t$ is sufficient for π_t in terms of predicting the behavior of Generation t .

²⁶Agents in our model may observe a distribution of predecessors’ actions that they think is extremely unlikely (e.g., the observed behavior \mathbf{a}_2 does not resemble the autarkic distribution in any state). In general, there need not exist $\hat{\omega}$ such that $\mathbb{P}_{\hat{\omega}} = \mathbb{T}_\omega$. The updating process described by ϕ implicitly handles any such discrepancy between observations and predictions by attributing it to sampling variation. Furthermore, we focus on large yet finite generations rather than a continuum (e.g., Banerjee and Fudenberg, 2004) to ensure that Bayesian updating is well defined. Under our finite approach, the probability of any observation within a naive agent’s model may go to zero in N but remains positive short of the limit.

though Generation 2 effectively learns the state, Generation 3 maintains this knowledge if and only if $\phi(\omega) = \omega$. If $\phi(\omega) \neq \omega$, then society “unlearns” ω between Generations 2 and 3 and grows confident in something false. Furthermore, when unlearning occurs, public beliefs will never settle on the truth: even if public beliefs put high probability on ω in some later period, they will again move away. As detailed in the next section, we will classify such a state as “abandoned”.

3.2 Long-Run Beliefs

The limiting behavior of public beliefs will take one of two forms. We say that long-run beliefs are *stationary* if $\langle \hat{\omega}_t \rangle$ converges to a fixed state. In this case, successive generations eventually believe in the same (but potentially incorrect) state. Otherwise, we say that long-run beliefs are *cyclic*. In this case, $\langle \hat{\omega}_t \rangle$ becomes periodic, continually cycling over a fixed subset of states. While we primarily focus on environments with stationary long-run beliefs, we briefly discuss the case of cyclic beliefs below (Section 3.4 and Appendix A).

The fact that society may fail to settle on the truth raises two questions focal to this paper. First, as a function of the true state, what does a naive society in fact come to believe? Second, are there “abandoned states” that society continually disbelieves, regardless of what is true?

To address these questions, we now examine which beliefs are reliably passed from one generation to the next. Consider a state $\hat{\omega}$ that society will continue to believe whenever the preceding generation assigns it high probability. Let $\Omega^* \subseteq \Omega$ denote the set of all such states. More precisely, $\hat{\omega} \in \Omega^*$ if and only if there exists $\omega \in \Omega$ such that $\langle \hat{\omega}_t \rangle$ converges to $\hat{\omega}$ given the initial condition $\hat{\omega}_2 = \omega$. Thus, the only states in which society can remain certain are the fixed points of ϕ . Although this result follows directly from Lemma 1, it plays a central role in our analysis and we therefore state it here for reference.

Proposition 1. *If $\langle \hat{\omega}_t \rangle$ converges to a stationary limit belief $\hat{\omega} \in \Omega$, then $\hat{\omega} \in \{\omega \in \Omega \mid \phi(\omega) = \omega\}$. Thus, $\Omega^* = \{\omega \in \Omega \mid \phi(\omega) = \omega\}$.*

Proposition 1 highlights how naivete restricts the set of states on which public beliefs may settle. The set Ω^* consists solely of those states ω such that the predicted behavior of privately informed agents in state ω most closely—out of all possible states—resembles the behavior of agents certain of ω . If $\omega \notin \Omega^*$, we call ω *abandoned*: public beliefs will never settle on such a state. Whenever $\Omega^* \subsetneq \Omega$, there will exist abandoned states that people continually disbelieve even when they are true. Our applications below reveal several natural scenarios where abandoned states emerge.

Abandoned states do not emerge, however, in the canonical two-state environment with common preferences studied by ER. To illustrate, suppose $\Omega = \{0, 1\}$, $\mathcal{A} = \{A_0, A_1\}$, and suppose the utility function for all players is $u(A_m|\omega) = \mathbb{1}\{m = \omega\}$ —they earn a payoff equal to one when their action matches the state and zero otherwise. If the state is $\omega = 0$, then all players in Generation 2

effectively learn this and take action A_0 . What will a naive Generation 3 infer? So long as signals are informative, then the likelihood of action A_0 in autarky is higher in state 0 than it is in state 1; that is, $\mathbb{P}_0(0) > \mathbb{P}_1(0)$. Hence, \mathbb{P}_0 better fits the observed behavior of Generation 2 than \mathbb{P}_1 does, and Lemma 1 implies $\hat{\omega}_3 = 0$. Thus, $\phi(0) = 0$. An analogous argument yields $\phi(1) = 1$. Therefore, in the canonical two-state environment, Ω^* is identical to Ω . Accordingly, naive beliefs can settle on either of the two states. ER show that a naive society may settle on the incorrect state when earlier signals are misleading. However, the likelihood of such mislearning vanishes in our setting with large generations: the likelihood that the set of initial signals is *collectively* misleading goes to zero in N . In contrast, we will show that there are many environments where some states are never supported by naive beliefs, even when N grows large. Put differently, the type of mislearning that arises in ER's setting happens only when early signals are misleading, yet in alternative environments there can exist states that lead to mislearning with probability arbitrarily close to one regardless of whether early signals are misleading or not. We refer to this latter form of mislearning as *structural mislearning*: whether society mislearns is preordained by the environment itself and does not depend on the happenstance of early signals.

A simple demonstration of this result arises when we extend the example above so that there are more states than actions. In this case, there necessarily exist predetermined states that society will adamantly disbelieve even when they are true.

Proposition 2. *Suppose players have common preferences. If $|\Omega| > |\mathcal{A}|$, then there exists at least one state $\omega \in \Omega$ such that when ω is true, $\hat{\omega}_t \neq \omega$ for all $t > 2$.*

Intuitively, the behavior of autarkic agents in Generation 1 reveals the commonly-preferred action, and Generation 2 then herds on this action. However, because there is a unique state that best predicts such a herd in autarky, there are at most $M \equiv |\mathcal{A}|$ states that society can come to believe—one for each of the M possible herd actions. Thus, if $|\Omega| > |\mathcal{A}|$, then there necessarily exist states that a naive society will deterministically deem false no matter what is true.

It is worth noting that canonical models of social learning do not typically examine settings with more states than actions. For the questions these models address, it is typically sufficient to define coarse states that comprise all events in which a particular action is optimal. There are, however, economically important settings with more states than actions; e.g., settings where society aims to learn the quality difference between two goods. If queuing costs were introduced or if prices were to change, then this quality difference would become crucial for deciding whether the option with superior quality is still worthwhile. In our setting with large generations, Generation 2 may perfectly infer the magnitude of this quality differential. Despite this, Proposition 2 implies that Generation 3 will necessarily assign negligible probability to some potential values of the quality differential even

when they are indeed true.²⁷

3.3 Robustness of Abandoned States

Proposition 1 is straightforward in our large-generation setting with deterministic belief dynamics, but it is particularly important because the result extends to more familiar settings where it is less obvious. Notably, a naive society will never grow confident in any of the states deemed “abandoned”—i.e., those outside of Ω^* —irrespective of the observation structure so long as the number of predecessors that each agent observes grows large in t . Although our large-generation setting makes the logic of abandoned states tractable and transparent, the characterization of these states described by Lemma 1 and Proposition 1 is indeed more general. This robustness is critical for porting the insights derived in our simplified framework to more familiar settings.

More explicitly, we now show that Proposition 1 extends to the canonical single-file setting commonly assumed in the literature on social learning (e.g., Bikhchandani et al., 1992; Smith and Sørensen, 2000; Eyster and Rabin, 2010; Bohren and Hauser, 2021). Define the *canonical single-file environment* as one similar to ours, but with the following differences. First, a single player acts in each period $t = 1, 2, \dots$; let x_t and θ_t define the choice and type of this player, respectively. Second, an agent’s type, θ_t , is i.i.d. across periods. Third, each agent t observes the complete history of actions. The “public belief” $\pi_t \in \Delta(\Omega)$ is thus a player’s belief after observing (x_1, \dots, x_{t-1}) . Finally, we drop Assumption 2.3 that bounds the informativeness of signals.

The next proposition shows that the sequence of public beliefs never converges to certainty on any state $\omega \notin \Omega^*$ no matter what is true, where Ω^* is defined by Lemma 1 and Proposition 1.

Proposition 3. *Consider the canonical single-file environment, and consider any realized state $\omega \in \Omega$. If $\omega_k \notin \Omega^*$, then $\Pr(\pi_t \rightarrow \delta(\omega_k) | \omega) = 0$.*

Intuitively, if beliefs eventually concentrate in the neighborhood of $\hat{\omega}$, then a Player t late in the sequence will observe a distribution of behavior among her predecessors that resembles $\mathbb{T}_{\hat{\omega}}$. Because a naive observer thinks each predecessor acts independently, the observed order of actions does not influence her inference—only the *aggregate* distribution of behavior matters. Following the same logic as Lemma 1, the state that maximizes the “autarkic” likelihood of this distribution is $\hat{\omega} = \arg \min_{\omega \in \Omega} H(\mathbb{T}_{\hat{\omega}}, \mathbb{P}_{\omega}) = \phi(\hat{\omega})$. Thus, in order for Player t to remain confident in $\hat{\omega}$, it must be that $\hat{\omega} = \phi(\hat{\omega})$. This implies that the only states on which public beliefs may settle in this alternative structure are exactly those in Ω^* described in Proposition 1. As such, our results characterizing when a naive society fails to learn the truth (in the long-run) are robust to alternative structures.²⁸

²⁷Even in settings where rational agents cannot perfectly discern the payoff difference between two options, they would not assign negligible probability to plausible values—they would instead remain uncertain. In contrast, naive agents systematically put (nearly) zero probability on a predetermined set of states.

²⁸The intuition above continues to hold under more general assumptions on the set of predecessors that each player

Although our result that beliefs never settle on states outside Ω^* is robust to relaxing our large-generation assumption, this assumption is crucial for pin pointing *which* state in Ω^* society will come to believe. In the canonical single-file environment, limit beliefs depend on the sample path of signals. Hence, the single-file setting only broadens the scope for mislearning. Namely, it introduces the chance that society mislearns when the true state lies in Ω^* ; this almost never happens in our large-generation setting.²⁹

3.4 Cyclic Beliefs

Before turning to applications that exhibit abandoned states, we briefly discuss the possibility of cyclic beliefs for the sake of completeness.

In settings where public beliefs converge to a fixed state, Ω^* is sufficient to identify the states that society will necessarily dismiss: all states outside Ω^* will receive negligible weight in the long-run. However, in settings with cyclic beliefs, Ω^* no longer provides a sufficient answer to this question. The process $\langle \hat{\omega}_t \rangle$ may become periodic and continually cycle over a subset of states. Although these states lie outside Ω^* , society will frequently grows confident in them. Hence, to identify the states that society assigns negligible weight in the long run, we must additionally distinguish those states that society may assign high probability infinitely often; we denote the set of such states by Ω^{**} .³⁰ Analogous to Proposition 1, public beliefs may periodically put high probability on a state only if it is a fixed point of a composition of ϕ ; that is, $\Omega^{**} = \{\omega \in \Omega : \exists L \geq 1 \text{ such that } \phi^L(\omega) = \omega\}$. Any state ω outside Ω^{**} satisfies a strengthened notion of “abandoned”: not only do beliefs fail to converge on ω , but *all* generations eventually assign it negligible probability even when ω is the true state. Appendix A provides more details and an example where Ω^* is empty yet $\Omega^{**} = \Omega$ —public beliefs continually cycle over all possible states regardless of what is true.

Despite the caveat above, the applications we consider beyond this section all deliver stationary long-run beliefs. Thus, focusing on Ω^* will indeed be sufficient to determine which states are permanently disbelieved.

Before turning to those applications, it is worth considering whether there is a limit to how often naivete can lead to permanently false beliefs. Namely, can there exist environments where, for every realization of the state, each generation in the long run grows confident in something false? The answer is no: in every environment, there exists some state that will induce correct learning among

observes. For instance, suppose N players act in each round, and let $O_{(n,t)}$ denote the set of all the predecessors that Player (n, t) observes. Proposition 3 extends if, for all $n = 1, \dots, N$, $|O_{n,t}|$ is increasing in t and $\lim_t |O_{(n,t)}| \rightarrow \infty$.

²⁹In the single-file environment, it is possible that society can grow certain of any $\omega \in \Omega^*$ regardless of the true state. The logic follows directly from Eyster and Rabin’s (2010) Proposition 4: if $\omega \in \Omega^*$, then there exists a sample path of signals realized with positive probability such that beliefs settle on ω .

³⁰Formally, $\hat{\omega} \in \Omega^{**}$ if and only if there exists $\omega \in \Omega$ such that, given the initial condition $\hat{\omega}_2 = \omega$, there exists no $\bar{t} \in \mathbb{N}$ such that $\hat{\omega}_t \neq \hat{\omega}$ for all $t > \bar{t}$.

some generations in the long run.

Proposition 4. *There exists a state $\omega \in \Omega$ such that when ω occurs, society places arbitrarily high probability on the truth (i.e., $\hat{\omega}_t = \omega$) infinitely often.*

If Ω^* is non-empty, then all generations correctly learn ω whenever $\omega \in \Omega^*$. But even when Ω^* is empty, there still must exist some ω such that, when ω occurs, long-generations will periodically come to believe ω . Although there is some limit to structural mislearning, our applications illuminate natural environments in which a *majority* of states will lead a naive society to permanently false conclusions.

3.5 Example: Long-Run Beliefs that are Independent of the State

This section provides an initial example of abandoned states. The application here highlights that naivete can cause public beliefs to converge on a single preordained state no matter what is true. That is, naive inference funnels beliefs toward Ω^* , which happens to be a singleton.

Consider two types of agents with opposing preferences. To fix ideas, imagine farmers deciding whether use a well-known seed type (A_0) or adopt a new hybrid type (A_1). Option A_0 yields a payoff normalized to zero for all farmers. The payoff from A_1 , however, is sensitive to a farmer’s soil type and is positive only if it matches well with one’s plot.³¹ Suppose there are two types of soil, high salinity ($\theta = H$) and low ($\theta = L$). Each farmer knows both her own soil type and that a fraction $\lambda > \frac{1}{2}$ of farmers are high types. Initially, it is unknown whether seed A_1 is compatible with high salinity ($\omega = H$) or low salinity ($\omega = L$). A farmer with compatible soil earns $v > 0$ by planting A_1 , but one with incompatible soil earns $-v$. Thus, option A_1 is optimal for Farmer (n, t) if and only if $\theta_{(n,t)} = \omega$. Farmers in village t learn about ω by observing choices in the adjacent village $t - 1$.

We consider a signal structure with the natural property that that, no matter the state, adoption of the new seed increases as social learning takes place. In particular, suppose a known fraction $\psi \in (0, 1)$ of farmers receive informative private signals about ω (e.g., from past experience with seed A_1). The remaining farmers have uninformative signals. Each informed farmer receives an i.i.d. signal $s \in \{H, L\}$ with mass function $f(s = \omega | \omega) = \rho \in (\frac{1}{2}, 1)$; her signal matches the true state with chance $\rho \in (\frac{1}{2}, 1)$.

To determine what naive farmers come to believe, we must compare \mathbb{P}_ω and \mathbb{T}_ω across states. From the setup above, the autarkic distributions are $\mathbb{P}_L(1) = \psi[(1 - \lambda)\rho + \lambda(1 - \rho)]$ and $\mathbb{P}_H(1) = \psi[\lambda\rho + (1 - \lambda)(1 - \rho)]$. Among a generation certain of ω , all those with type $\theta_{(n,t)} = \omega$ choose A_1 ; this implies $\mathbb{T}_L(1) = 1 - \lambda$ and $\mathbb{T}_H(1) = \lambda$. Hence, if sufficiently few farmers have private information—

³¹This example is inspired by Munshi (2004) and Foster and Rosenzweig (1995), who study social learning among Indian farmers trying to deduce the optimal inputs for new “high-yield” strains of rice and wheat. Munshi (2004) notes that rice is quite sensitive to soil characteristics, but wheat is not.

that is, if ψ is not too large—then social learning always leads to higher adoption rates relative to autarky: those who are initially uninformed might adopt based on information gleaned from their neighbors.³² Naive observers, however, misattribute this learning-based increase in adoption. When a large share adopts, they conclude that the new seed must be optimal for the majority type. Essentially, people may mistake “confident” behavior in the low state for autarkic behavior in the high state.

To illustrate, suppose $\lambda = 0.7$, $\rho = 0.8$, $\psi = 0.5$, and $\omega = L$, so only the less-common low types should adopt. In the first period, the adoption rate is $\mathbb{P}_L(1) = 19\%$, which reveals $\omega = L$ to Generation 2. Thus in period 2, all of those with type $\theta_{(n,t)} = L$ choose A_1 , yielding an adoption rate of $\mathbb{T}_L(1) = 30\%$. A naive Generation 3, however, expects to see either the autarkic rate $\mathbb{P}_L(1) = 19\%$ (if the state L) or $\mathbb{P}_H(1) = 31\%$ (if the state H). Hence, they come to believe in whichever state is most likely to yield 30% as a result of sampling variation. Since 30% is “closer” (in terms of cross entropy) to 31% than it is to 19%, Generation 3 wrongly becomes convinced that $\omega = H$.³³

In fact, $\hat{\omega}_3 = H$ whenever ψ is not too large relative to λ and ρ . In such cases, all high types—the majority of farmers—adopt A_1 in round 3. Since this rate is higher than predicted in either state, all subsequent farmers will continue to believe $\omega = H$. Mislarning results from people intuitively (but wrongly) reasoning that high adoption rates indicate that the new technology is best for the majority type. The following proposition summarizes these results.

Proposition 5. *Consider the technology-adoption example above where two types of agents have opposing preferences conditional on the state.*

1. *Suppose adoption is optimal only for the majority type (i.e., $\omega = H$). Then public beliefs settle on the truth: $\hat{\omega}_t = H$ for all $t \geq 2$.*
2. *Suppose adoption is optimal only for the minority type (i.e., $\omega = L$). There exists a threshold value $\bar{\psi} \in (0, 1)$ such that if the fraction of informed agents falls below $\bar{\psi}$, then public beliefs wrongly settle on the false belief that adoption is optimal only for the majority type. That is, $\psi < \bar{\psi}$ implies $\hat{\omega}_t = H$ for all $t > 2$. Additionally, $\bar{\psi}$ is continuously decreasing in both λ and ρ .*

In summary, if $\psi < \bar{\psi}$, then Ω^ is a singleton: $\Omega^* = \{H\}$.*

4 Extreme Beliefs

This section presents a simple yet stark implication of naive learning in settings where individuals with common preferences choose among options with independent payoffs. Namely, people grow

³²The relevant condition on ψ is $\psi < (1 - \lambda)/[(1 - \lambda)\rho + \lambda(1 - \rho)]$. This threshold monotonically increases to 1 as the precision of signals, ρ , increases from .5 to 1.

³³This “distance” calculation follows from Equation 2: $H(\mathbb{T}_L, \mathbb{P}_L) = -0.3 \log(0.19) - 0.7 \log(0.81) \approx 0.6457 > 0.6111 \approx H(\mathbb{T}_L, \mathbb{P}_H) = -0.3 \log(0.31) - 0.7 \log(0.69)$.

certain that one option yields its highest possible payoff while all others yield the lowest possible. Under naive inference, people think a herd on option A_m results from N independent signals indicating A_m is better than all of its alternatives. We show that, under a natural signal structure in which an autarkic agent is more likely to try an option when its relative payoff is higher, this misinterpretation of the herd convinces onlookers that the payoff difference between A_m and each of the other options takes on its most extreme value.

Actions and States. Suppose players have common preferences and each option $A_m \in \{A_0, \dots, A_{M-1}\}$ yields a payoff $u(A_m|\omega) = q^m$. We refer to q^m as the “quality” of option m . The payoff-relevant state is the “quality vector” $\omega = (q^0, \dots, q^{M-1})$. Each q^m is drawn from a support Q^m with $|Q^m| \geq 2$, and is drawn independently from the quality of any other option, q^j . For instance, imagine consumers learning about the relative quality of two products from independent firms, or investors learning about the returns to assets in unrelated sectors. Thus, any quality profile $\omega \in \times_{m=0}^{M-1} Q^m$ is a feasible state.

Signals. We consider signal structures that satisfy the familiar Monotone Likelihood Ratio Property (MLRP). Building on our setup from Section 2, we further assume that $s = (s^0, \dots, s^{M-1}) \in \mathbb{R}^M$, where each s^m depends on q^m alone and is independent of all other q^j ; that is, news about q^m provides no information about q^j . Conditional on q^m , let $F^m(\cdot|q^m)$ and $f^m(\cdot|q^m)$ denote the marginal c.d.f. and associated density (or mass) function of s^m , respectively, and suppose they obey the following property:

Definition 5. F^m satisfies the Monotone Likelihood Ratio Property (MLRP) if for every $q > q'$, $f^m(s|q)/f^m(s|q')$ is strictly increasing in s .

MLRP means that higher signals unambiguously indicate higher expected quality. In this setting, MLRP naturally implies that the share of players who choose A_m in autarky strictly increases as q^m increases. This suggests that naive social learners will conflate high demand with high quality. Indeed, our next proposition shows that they will come to believe in an “extreme state”:

Definition 6. Extreme state m , denoted by ω_e^m , is the state in which $q^m = \max Q^m$ and $q^j = \min Q^j$ for all $j \neq m$.

Proposition 6. Suppose $\arg \max_j (q^1, \dots, q^M)$ is unique and equal to m . If signals about each option (i) satisfy MLRP and (ii) are independent of signals about any other option, then $\hat{\omega}_t = \omega_e^m$ for all $t > 2$. That is, public beliefs settle on the extreme state in which option m yields its highest possible payoff and all other options yield their lowest possible payoff. Hence, Ω^* is to equal the set of extreme states and all other states are abandoned.

To illustrate the logic, suppose $q^0 = \max(q^0, \dots, q^{M-1})$. Because Generation 2 infers that A_0 is optimal, they unanimously choose A_0 . Generation 3 comes to believe in the state most likely to induce a herd on A_0 under autarkic play. Intuitively, this state must maximize the chance of good

news about A_0 , but minimize the chance of good news about any other option. Under MLRP, this happens in state ω_e^0 . As such, Generation 3 infers $\hat{\omega}_3 = \omega_e^0$ and again herds on A_0 , which implies that ω_e^0 is a steady state: $\hat{\omega}_t = \omega_e^0$ for all $t > 2$. In essence, naive observers mistake herds caused by social learning as an overly optimistic signal about the chosen action’s relative quality.

Proposition 6 clearly demonstrates how naivete restricts the hypotheses society may come to believe. No matter which of the $\prod_{m=0}^{M-1} |Q^m|$ possible states is realized, beliefs eventually settle on one of only M extreme states.³⁴ This setting also exemplifies the prediction of Proposition 2: since there are more states than actions, some states are surely abandoned.

Proposition 6 implies that naive learners tend to exaggerate the quality difference between options. Imagine that firm A ’s product has only slightly higher quality than its competitor, B . When droves of consumers choose A over B due to social learning, the naive onlooker infers not only that A is better than B , but that it is *much* better. Within a naive agent’s model, A ’s predicted market share is strictly increasing in its quality advantage over B , and thus they attribute a large market share to a large advantage. However, this naive logic neglects the fact that, due to initial social learning, predecessors would choose A no matter how small of a quality advantage it has.³⁵ Such beliefs may be costly. For instance, if queuing costs arise, then extreme beliefs may generate inefficiently high congestion: since naive observers exaggerate the quality difference between the best option and all others, they are less willing—relative to rational observers—to switch to the next-best option. Consumers would be willing to wait in long queues or pay relatively high prices for firm A even when B being nearly as good. The following sections explicitly show two additional ways that extreme beliefs lead to sub-optimal actions.

5 Large Choice Sets can Trigger Excessive Adoption

This section shows how the formation of “extreme” beliefs can induce society to over-adopt a new technology. We consider a setting where people with heterogeneous preferences can choose from several unproven technologies, and we show that naivete can drive *all* people to adopt one of them even in states where only a minority should. In particular, we emphasize how the number of available options determines whether such over-adoption will occur: with more options to evaluate, it becomes more likely that society will wrongly conclude that one of them is worthwhile for all types.

³⁴The canonical setting assumed in ER precludes exaggerated perceptions of quality differences because the only states they consider are “extreme” to begin with. Only when we first allow non-extreme states to arise with positive probability do we see how naive inference blocks society from learning these states.

³⁵In reality, we of course do not expect all consumers to patronize the higher-quality firm—pricing, heterogeneous tastes, and various frictions would prevent this. The point of this application is to demonstrate a more general, directional prediction. Namely, social learning will typically cause the better firm to capture *more* of the market than if consumers were to make decisions in autarky. Yet, since naive learners neglect the effect of social learning on market shares, they wrongly attribute these larger-than-expected shares to the firm’s quality.

We first develop intuition for this result and then formalize it below. Imagine a community learning about the efficacy of two costly medical practices, A and B . Patients can either choose one of these remedies or abstain from treatment altogether. The payoff from abstaining varies across individuals—some have more severe cases than others. In truth, suppose that both remedies have low efficacy, and A is only marginally better than B . As society discovers this, those with mild cases abstain from treatment while those with dire cases adopt A . This “split” pattern of behavior may lead naive followers astray, as it sends two conflicting signals. The first is what we will call the “abstention signal”: the fact that many abstain reflects negatively on both treatments since naive observers expect many to abstain only when both are ineffective. At the same time, the observed behavior also reflects positively on treatment A . We will call this force the “consensus signal”: the fact that those who do choose a treatment unanimously select option A over B reflects positively on A since naive observers expect the demand for A to dwarf that for B only when A is significantly better. Society will wrongly conclude that A is highly effective whenever the consensus signal sends a stronger message than the abstention signal.

This happens, for instance, when the number of alternative treatments is large. With more options available, it becomes increasingly surprising that autarkic actors would identically choose A . In turn, this consensus behavior is increasingly convincing to naive onlookers that their predecessors received strongly optimistic signals about A , which indicates that A is truly beneficial.

Actions and States. For the sake of clarity, we frame our general setup in terms of the medical example above. We build on that example by considering an arbitrary number of potential treatments in order to highlight how this number influences learning. Each patient can either try one of M unproven treatments, $\{A_1, \dots, A_M\}$, or choose the outside option, A_0 , which represents no treatment. Each treatment m yields a payoff $q^m \in \{\underline{q}^m, 0\}$ for all patients. We call treatment m “effective” if $q^m = 0$; otherwise it yields $\underline{q}^m < 0$, meaning it is only partially effective or harmful. To avoid trivialities arising from indifference (Assumption 3), let $\underline{q}^m \neq \underline{q}^j$ for all $m \neq j$, and, without loss of generality, $\underline{q}^1 = \max_m \{\underline{q}^m\}_{m=1}^M$. That is, A_1 is the best option whenever none are fully effective.

The payoff of the outside option, A_0 , depends on a patient’s type, θ , and is denoted by q_θ^0 . Suppose there are two types: type $\theta = H$ represents a “dire” patient who highly values treatment, and type $\theta = L$ represents a patient with a mild case and hence a low demand for treatment. We assume that the outside option of a high-valuation type ($\theta = H$) is bleak enough that she always picks some treatment over abstaining; i.e., $q_H^0 < \underline{q}^1$. In contrast, a low-valuation type ($\theta = L$) selects a treatment only when she is sufficiently confident it will work and otherwise abstains; i.e., $q_L^0 \in (\underline{q}^1, 0)$. Let $\lambda \in (0, 1)$ denote the fraction of high-valuation (i.e., “dire”) patients.³⁶

Signals. The signal structure over the uncertain options matches Section 4: patients receive condi-

³⁶Our results in this section hold regardless of whether a player’s type is publicly observable or not, and it easily extends with additional types.

tionally independent signals $s^m \sim F^m(\cdot|q^m)$ about each treatment, where each F^m satisfies MLRP with respect to q^m (Definition 5). Additionally, suppose that, for each option $m = 0, \dots, M$, there exist signal realizations (occurring with positive probability) that lead a low-value patient to select m in autarky. Thus, a low-value patient ($\theta = L$) refuses treatment in autarky whenever she receives sufficiently pessimistic signals about all the treatments, while a high-value patient ($\theta = H$) always chooses some treatment.

We analyze public beliefs in the state where all treatments are ineffective and thus low-valuation types ($\theta = L$) should abstain. The next proposition shows that having more treatments to choose from makes it more likely that society will mislearn the state, causing people to wrongly believe some treatment is effective for all. Let ω^0 denote the state in which no treatment is effective, and let ω^m denote that in which only treatment m is effective. To ensure that increasing the number of options does not mechanically increase the likelihood of an effective treatment, we fix the prior probability of ω^0 at $\chi \in (0, 1)$, which is independent of M .

Proposition 7. *Suppose the state is ω^0 . There exists a finite value $\bar{M} \geq 2$ such that:*

1. *Public beliefs settle on the true state if and only if the number of uncertain options, M , is less than \bar{M} .*
2. *If $M \geq \bar{M}$, then $\hat{\omega}_t = \omega^1$ for all $t > 2$, and society wrongly believes option A_1 is optimal for all types.*

For intuition, consider inference and behavior among Generations 2 and 3. Since Generation 2 correctly infers the state, *all* patients in Generation 2 choose treatment m whenever it is truly effective (i.e., in state ω^m). This herd clearly indicates to future generations that A_m is effective and society correctly learns. However, when no treatment works (i.e., in state ω^0), Generation 2 sends a more opaque message to followers: low-value types abstain while high-value types select the least bad treatment, A_1 . Naive followers who observe this split behavior confront the conflicting “abstention” and “consensus” signals introduced in the opening example. The consensus signal suggests that option A_1 is effective, while the abstention signal tempers this optimism. Thus, when the consensus signal is relatively strong, Generation 3 and all future generations will wrongly infer that A_1 is optimal for all.³⁷

One feature of the environment that determines the strength of the consensus signal is the size of the choice set—it becomes relatively stronger as the number of options, M , grows. Increasing the

³⁷Perhaps the simplest case in which the consensus signal dominates is when the abstention signal provides no information at all. This happens whenever there exist no private signals strong enough to entice low-valuation patients to experiment with risky treatments in autarky. Thus, naive observers expect all low types to abstain irrespective of their information. Generation 3 consequently thinks that only the actions of dire types reveal information and thus update their beliefs entirely from the consensus signal.

set of options makes the consensus behavior more surprising: with more options to pick from, it becomes less likely that autarkistic choices would coincide unless A_1 were truly effective. Furthermore, Proposition 7 holds for any λ . Thus, no matter how few dire patients there are, a sufficiently large number of options will allow this small minority to greatly sway public beliefs.³⁸

More specifically, Lemma 1 implies that the conflicting behavior observed by Generation 3—dire types select A_1 while all others abstain—is more likely in state ω^1 than ω^0 if and only if

$$\underbrace{\left(\frac{\mathbb{P}_{\omega^1}(1)}{\mathbb{P}_{\omega^0}(1)}\right)^\lambda}_{\text{Consensus signal}} > \underbrace{\left(\frac{\mathbb{P}_{\omega^0}(0)}{\mathbb{P}_{\omega^1}(0)}\right)^{1-\lambda}}_{\text{Abstention signal}}. \quad (3)$$

The left-hand side of (3) represents the signal in favor of state ω^1 due to the consensus among dire patients, and the right-hand side represents the signal in favor of ω^0 due to abstention. The strength of the “consensus signal”, $\mathbb{P}_{\omega^1}(1)/\mathbb{P}_{\omega^0}(1)$, is increasing in M . To see why, first note that a patient chooses A_1 in autarky only when her private signal about that treatment, s^1 , is sufficiently high relative to her private signals about the other $M - 1$ options. That is, s^1 must satisfy $M - 1$ threshold conditions, where each depends on the realized signal about an alternative treatment, $m > 1$. Now consider how the likelihood that these conditions are met—and hence that A_1 is chosen—differs between the state where no treatment is effective, ω^0 , and the one where A_1 is uniquely effective, ω^1 . The private signal distributions for options $m = 2, \dots, M$ are identical in each of these states, but the distribution of s^1 in state ω^1 first-order stochastically dominates that in state ω^0 . Hence, each of the $M - 1$ threshold conditions on s^1 is more likely to hold in state ω^1 than in state ω^0 . In this sense, moving from ω^0 to ω^1 has an “ $M - 1$ fold” effect toward increasing the likelihood that an autarkic agent chooses A_1 over any other treatment. As a result, increasing M makes the “consensus signal” more indicative of ω^1 relative to ω^0 . On the other hand, M has no effect on the relative informativeness of the “abstention signal”. To see why, note that a low type chooses A_0 only if each of her M signals is low. Out of these M conditions, only one—the one about A_1 —is less likely satisfied in state ω^1 relative to ω^0 . Thus, the likelihood of observing no treatment in ω^1 relative to ω^0 is independent of M .

Summarizing, the informativeness of abstention behavior is independent of M , but the counter-vailing informativeness of the consensus behavior is increasing in M : a herd on one of the options becomes more surprising when there are more alternatives to choose from. Consequently, for any signal distributions, there exists an M large enough such that this consensus signal dominates inference and society wrongly concludes that A_1 is optimal for all types. Since it appears as if those

³⁸The population share of dire patients, λ , is another feature of the environment that determines the strength of the consensus signal—the larger is this population, the more surprising it becomes that they congregate on an identical treatment. The threshold \bar{M} identified in Proposition 7 is therefore decreasing in λ : fewer options are needed to trigger an unwarranted herd when dire patients are more prevalent.

adopting the new technology have precise private information in favor of A_1 , their consensus behavior lures the rest of society to imitate them. Thus, even a small minority of individuals with distinct preferences can lead society astray.^{39,40} Interestingly, this mistake never happens when there is only one uncertain option available: naive observers expect all dire types to choose the lone treatment no matter the state, so they appropriately infer nothing from their actions. It is the availability of many viable options that leads naive observers to over-infer from consensus behavior.

6 Extremism in Investment Decisions

While previous applications show how extreme beliefs can cause some types to make suboptimal decisions, this section considers a portfolio-choice problem where naive learning will necessarily lead all players to make harmful mistakes. In particular, when naive investors use predecessors' allocations to predict the returns on two investments, they will suffer from two forms of inefficiency: (i) even when it is optimal to diversify, they eventually allocate all resources to a single asset, and (ii) when the true returns are far from initial expectations, naive investors allocate all resources to the worse of the two assets. Although the logic parallels previous extremism results (Section 4), this investment setting illuminates some additional implications of naivete. First, investment dynamics reflect forms of over-extrapolation and momentum observed in “bubble” markets, where perceptions of an asset’s value grow more extreme over time. Second, we emphasize how the direction of these dynamics hinge on how the true state compares to initial expectations.

Suppose there are two assets that pay off in terms of a consumption good in some final period T .⁴¹ We call one “safe” and the other “risky”. Suppose the two assets yield constant expected payoffs across generations, but their realized payoffs in any Generation t is subject to noise. Specifically, the safe and risky asset yield payoffs $d_t^s = 1 + \eta_t^s$ and $d_t^r = 1 + \omega + \eta_t^r$, respectively, where the η_t^m terms represent i.i.d. mean-zero random shocks. Thus, the safe asset has an expected payoff of 1, and the risky asset has an expected payoff of $1 + \omega$, where the “fundamental” ω is initially unknown. We include aggregate shocks, η_t^m , to model a scenario where rational risk-averse investors diversify

³⁹This intuition extends beyond the specific medical example used for exposition. For instance, consider investors who vary in risk aversion learning about the returns on various assets. More risk-averse agents are analogous to mild patients in our medical example because they require relatively higher signals in order to invest. Thus, when investors are evaluating a sufficiently large number of options, Proposition 7 then implies that the more risk-averse agents will be unduly swayed to follow the strategies of those seeking high-risk positions.

⁴⁰Gagnon-Bartsch (2016) predicts a similar form of costly herding in a social-learning model with agents who exaggerate the extent to which others share their tastes. Namely, society will wrongly adopt a single action when in fact those with different tastes are better off choosing distinct options.

⁴¹While we focus on the limit case where T is arbitrarily large, the dynamics we describe are identical for finite T . We consider long-term investments that do not payoff immediately to ensure that the current generation does not observe the outcomes of the previous generation’s choices. The model alternatively applies to cases where each investor’s outcome is realized immediately but is privately observed.

even when ω is perfectly known, and we assume these shocks are distributed $\eta_t^m \sim N(0, \rho_\eta^{-1})$ for both $m = r, s$. Additionally, we assume private signals and the prior over the fundamental, ω , are normally distributed as well: $\omega \sim N(\bar{\omega}, \rho_\omega^{-1})$, and, conditional on ω , $s_{(n,t)} \sim N(\omega, \rho_s^{-1})$. An investor updates this prior over ω based on her private signal and her predecessors' allocations (as described below).⁴²

Each Investor (n, t) has initial wealth $W_0 \in \mathbb{R}$ and allocates a fraction $x_{(n,t)} \in [0, 1]$ to the risky asset. For tractability, we consider investors with exponential utility $u(W) = -\exp(-\alpha W)$ over wealth, where α measures absolute risk aversion. It is well-known that this generates mean-variance preferences over the optimal allocation (see, e.g., Grossman, 1976).⁴³ It is then straightforward to show that a naive Investor (n, t) with information $I_{(n,t)}$ chooses

$$x_{(n,t)} = \frac{1}{2 + \rho_\eta \widehat{\text{Var}}[\omega | I_{(n,t)}]} \left\{ 1 + \frac{\rho_\eta}{\alpha W_0} \widehat{\mathbb{E}}[\omega | I_{(n,t)}] \right\}, \quad (4)$$

where the expectation and variance terms above represent those perceived by Investor (n, t) under her naive autarkic model. Finally, we assume Investor (n, t) 's information set $I_{(n,t)}$ consists of her private signal and the aggregate amount of wealth invested in the risky asset during the preceding period, denoted by \bar{x}_{t-1} .⁴⁴

Observed behavior among the first generation, \bar{x}_1 , efficiently aggregates information about ω . As we show in the proof of our next proposition, the aggregate allocation in Generation 1 is determined by an ‘‘autarkic’’ demand function, $D_A(\omega)$, that is strictly increasing in ω and takes interior values for all $\omega \in \mathbb{R}$. Hence, D_A is invertible and thus perfectly reveals the fundamental. Since $\bar{x}_1 = D_A(\omega)$, an onlooker in Generation 2 can invert this relationship to learn ω . Demand among Generation 2 then

⁴²We assume the assets' prices are fixed in order to avoid complications that arise when extending our model of naive inference into a full equilibrium concept (see, e.g., Eyster and Rabin, 2008). While we abstract from asset pricing for ease of exposition, we discuss in the conclusion how our results would naturally extend to settings with endogenous prices. Namely, our prediction that traders form extreme perceptions of an asset's value will generate inflated prices.

⁴³Specifically, letting W denote an investor's final wealth, Investor (n, t) with information $I_{(n,t)}$ chooses an allocation x to maximize

$$\mathbb{E}[W | x, I_{(n,t)}] - \frac{1}{2} \alpha \text{Var}[W | x, I_{(n,t)}],$$

where the expectation and variance terms are with respect to the investor's subjective distribution of wealth conditional on her chosen allocation.

⁴⁴Note that \bar{x}_t is the expected value of $x_{(n,t)}$ across individuals, conditional on ω . In this setting, we are able to drop some of our running assumptions for the sake of simplification. In particular, we consider a continuum of states and a continuum of investors in each generation instead of an arbitrarily large, yet finite, group. We can do so because any observed behavior \bar{x}_{t-1} of Generation $t - 1$ is perfectly consistent with Generation t 's naive model; that is, there exists some state $\hat{\omega}_t$ inferred by Generation t that perfectly predicts \bar{x}_{t-1} in autarky. This is not typically the case in a more general setup. Furthermore, because we consider a continuum of agents, we can also accommodate unbounded normally-distributed signals. It is worth noting that our qualitative results in this section do not depend on these slightly modified assumptions; we take this approach solely to leverage the simple closed-form solution for the optimal investment under normally-distributed beliefs. Additionally, the results extend to a market where the same participants repeatedly make investment decisions and observe the full history of demand.

adjusts to reflect their new knowledge of ω .

Before examining how naive investors will eventually mislearn ω , we first consider the rational benchmark. Rational demand for the risky asset will remain fixed from Generation 2 onward. Rational investors in period $t > 2$ understand that their predecessors know ω and appropriately imitate them. Thus, they will allocate the optimal amount x^r to the risky asset in each period beyond the first. Furthermore, due to aggregate uncertainty, this rational choice will involve diversification—that is, $x^r \in (0, 1)$ —so long as ω is not too large in magnitude.

In contrast, naive investors neglect that the aggregate behavior of the previous generation already incorporates all the available information about ω and wrongly uses it to update their beliefs. Consequently, naive beliefs and allocations evolve over time, and they do so until investors allocate either all or nothing to the risky asset.

In particular, since naive investors think each generation acts in autarky, they wrongly think that the autarkic demand function from above determines behavior in each period. A naive Generation t who observes \bar{x}_{t-1} thinks $\bar{x}_{t-1} = D_A(\omega)$. Attempting to infer ω from this relationship, they wrongly grow certain that $\hat{\omega}_t = D_A^{-1}(\bar{x}_{t-1})$. From Equation 4, the aggregate demand for the risky asset in period t is then

$$\bar{x}_t = \frac{1}{2} \left(1 + \frac{\rho_\eta}{\alpha W_0} \hat{\omega}_t \right), \quad (5)$$

so long as this value lies in $[0, 1]$. These conditions recursively define the allocation process, $\langle \bar{x}_t \rangle$.

Over time, naive inference polarizes investors' perceptions of the payoff difference between the two assets: beliefs about ω diverge toward positive or negative infinity, and $\langle \bar{x}_t \rangle$ consequently converges to either 1 or 0.⁴⁵

Proposition 8. *Fix the true value of the fundamental, ω .*

1. *There exists a threshold value ω^* such that if $\omega > \omega^*$, then perceptions of the fundamental, $\hat{\omega}_t$, and demand for the risky asset, \bar{x}_t , are both increasing in t . Otherwise, if $\omega < \omega^*$, then $\hat{\omega}_t$ and \bar{x}_t are both decreasing in t .*
2. *The threshold value ω^* is increasing in $\bar{\omega}$ and is such that $\omega^* < \bar{\omega}$.*
3. *Demand for the risky asset, $\langle \bar{x}_t \rangle$, monotonically converges to 1 or 0. Hence, players eventually invest in a single asset.*

Whether beliefs about ω grow excessively optimistic or pessimistic depends on whether Generation 2—who learns the true value of ω —invests more or less in the risky asset than Generation 1. This

⁴⁵Since we have assumed a prior on ω with full support on \mathbb{R} , $\hat{\omega}_t$ converges to $\pm\infty$. Of course, this prediction is unrealistic—the presence of some rational investors in the market would moderate this result. Instead, we think the more important feature of this result is the qualitative prediction that perceptions tend to change monotonically over time. The extreme outcomes of “infinite” perceptions and zero diversification simply help make the logic clear.

initial revision in demand creates momentum that propagates through all future periods. To illustrate, suppose demand for the risky asset increases from period 1 to 2 (i.e., $\bar{x}_2 > \bar{x}_1$). Since Generation 3 treats \bar{x}_2 as if it reflects autarkic demand, Generation 3 must infer a higher value of ω than Generation 2 did. As such, the demand for the risky asset increases yet again. This logic plays out across all periods: each new generation observes a larger “autarkic demand” than the last, which leads them to allocate even more to the risky asset. Likewise, if the initial revision in demand decreases from period 1 to 2 (i.e., $\bar{x}_2 < \bar{x}_1$), then subsequent demand (and beliefs) decrease over time.

This inferential error is driven by investors continually using past demand as if it reflects new information. Investors neglect that observed demand already incorporates all information in the economy, and hence attribute any changes to new private information. When the current generation incorporates this “new” information, the allocation moves yet again in the same direction as the initial (rational) adjustment. Hence, naivete predicts momentum even when no new information is realized, offering a plausible explanation for the sort of unwarranted swings in group beliefs that appear to be a hallmark of financial markets. This qualitative prediction accords with the empirical findings of Glaeser and Nathanson (2017), who suggest that momentum in the housing market derives, in part, from naive inference based on past market prices. Indeed, an extension of our model that incorporates pricing would predict momentum in prices, which in turn leads to price bubbles.⁴⁶

The direction of this naivete-fueled momentum—and thus whether investors eventually allocate all or nothing to the risky asset—depends on how the true fundamental compares to the threshold value, ω^* , derived in Proposition 8. The value ω^* is such that the initial change in demand for the risky asset is positive if and only if $\omega > \omega^*$. As noted above, ω^* is increasing in investors’ prior expectation, $\bar{\omega}$, but lies somewhat below $\bar{\omega}$. Intuitively, demand initially increases when investors learn that ω exceeds expectations. However, due to risk aversion, there is a range of values $\omega \in (\omega^*, \bar{\omega})$ for which demand increases upon learning ω even though it falls short of expectations: in such cases, the reduction in uncertainty over ω more than compensates for the risky asset’s lower-than-expected return, causing its demand to increase.⁴⁷ Furthermore, the fact that ω^* is increasing in $\bar{\omega}$ leads to a perverse effect of high expectations: fixing the true fundamental, when initial expectations are higher, it becomes more likely that agents will under-invest in the risky asset. Even when the true fundamental is such that the risky asset is worthwhile, the initial change in demand will be negative

⁴⁶In this investment context, the implications of naive inference seemingly run opposite those of cursedness (Eyster and Rabin, 2005). As explored in Eyster et al. (2019), cursed traders in asset markets fail to infer from price. Investors in our model, however, *over*-infer from past behavior—they revise their beliefs even when demand provides no new information. For a recent analysis of how these biased beliefs influence equilibrium prices, see Bastianello and Fontanier (2021).

⁴⁷Interestingly, the speed at which perceptions diverge from the fundamental value is increasing in the level of risk aversion, α . As risk aversion increases, observers expect greater conservatism among the (supposedly) autarkic investors acting before them. Hence, observers think previous investors require stronger signals in order to allocate a majority of their resources to the risky option. Fixing $\bar{x}_t > 1/2$, Generation $t + 1$ thus infers a greater expected return the larger is α .

when expectations are sufficiently high. Thus, this downward revision will create further momentum that drives investment in the risky asset to zero.




Figure 1: The two figures describe investors' long-run beliefs and allocations as a function of the true fundamental (x -axis). The two frames consider different initial expectations: Frame (a) considers high initial expectations—and thus a high cutoff value, ω^* —while Frame (b) considers low expectations. The region highlighted in red indicates values of the fundamental that lead investors to fully invest in the asset that yields the lower expected utility. The truly optimal allocation involves diversification for all $\omega \in (-\alpha W_0/\rho_\eta, \alpha W_0/\rho_\eta)$.

Part 3 of Proposition 8 establishes that aggregate allocations indeed increase or decrease until investors devote all or no wealth to the risky asset. Hence, naive risk-averse investors are worse off relative to their rational counterparts whenever it is optimal to diversify. This is the case whenever $\omega \in (-\alpha W_0/\rho_\eta, \alpha W_0/\rho_\eta)$; diversification is optimal in a wider range of states when either risk aversion or the variance in returns (i.e., $1/\rho_\eta$) is larger.

Moreover, the asset that naive investors come to demand is sometimes the worse one—the one that yields the lower expected utility. For instance, if the optimal strategy is to invest 80% in the risky asset, naive investors may end up investing 0%.

Corollary 1. For any collection of parameters $(\bar{\omega}, \rho_\omega, \rho_s, \rho_\eta, \alpha, W_0)$, there exists an open interval $\Omega' \subset \mathbb{R}$ such that whenever $\omega \in \Omega'$, naive investors eventually allocate all resources to the asset that yields the lower expected utility.

The intuition is straightforward in light of Proposition 8. Consider the case where $\bar{\omega} > \omega^* > 0$; that is, investors expect the risky asset to outperform the safe asset. While this expectation is fulfilled whenever $\omega \in (0, \omega^*)$, the fact that ω falls sufficiently short of expectations implies that perceptions of ω decrease over time (Proposition 8, part 3). Eventually, $\hat{\omega}_t$ will fall so low that investors allocate no wealth to the risky asset. Panel (a) of Figure 1 depicts this case: when the fundamental value lies in the the red interval, investors will eventually hold none of the risky asset despite it having a higher

expected return than the safe one. Panel (b) shows the opposite case where investors may end up holding only the risky asset despite it being dominated by the safe one.

7 Conclusion

This paper explores new predictions of Eyster and Rabin’s (2010) model of naive inference that emerge in a broader array of environments than previously studied. In many environments, there exist “abandoned states” that naive agents *necessarily* disbelieve in the long run even when they are true. For instance, under natural assumptions on preferences and signals, naive agents will systematically form extreme beliefs about the quality of a new technology relative to the status quo—they will either think it has the highest or lowest possible quality while being blocked from believing it has some intermediate value. As we have shown, these extreme beliefs can lead to severe under-diversification in investment settings, or to the costly over-adoption of goods beneficial for only a minority of agents. This logic also suggests that, relative to more direct methods of information transmission, individuals may be more susceptible to overpay or inefficiently queue for products when information spreads through observational learning.

This is not the first paper to study “redundancy neglect” in environments beyond than the canonical two-state model. However, it does provide results distinct from earlier work. Notably, our predictions regarding extreme beliefs and unlearning distinctly follow from Eyster and Rabin’s (2010) formulation of naive inference, and do not readily emerge in other models of naive learning based on the DeGroot model. For instance, in DeMarzo et al. (2003), agents repeatedly share their beliefs about a normally distributed state within a network and use a naive updating rule based on DeGroot (1974): each agent treats her neighbors’ reports as independent signals, neglecting that their posteriors already incorporate information previously shared amongst each other. Consequently, agents over-count signals and grow confident in some false state whenever initial signals are misleading. The nature of this mislearning, however, differs from ours in two ways. First, agents in DeMarzo et al. (2003) do not gravitate toward extreme perceptions over time. Since the DeGroot heuristic is an averaging rule, beliefs converge to a weighted average of initial signals instead of tending to extreme values.⁴⁸ Second, when the number of observed neighbors grows large, agents in DeMarzo et al. (2003) learn correctly. The law of large numbers implies that the first round of communication sends agents directly to a confident and correct posterior. In our setting, even if players correctly learn the state after one round of observation, later generations mislearn by reinterpreting confident behavior as if it were autarkic.

Several interesting applications of naive learning remain beyond the scope of our simple frame-

⁴⁸A similar distinction can be drawn with the many other papers on naive inference that build from DeGroot’s (1974) framework (e.g., Golub and Jackson (2010)).

work. Since our model of naive inference only specifies how a player best responds to her observations, we lack a solution concept to analyze settings where a player’s payoff depends directly on others’ actions.⁴⁹ For instance, a natural extension of our extremism results might consider how firms set prices in order to exploit—or undermine—naive consumers’ tendency to exaggerate quality differences across products. Similarly, in the investment setting of Section 6, we believe an equilibrium model incorporating prices would generate a bubble: as the perceived returns to an asset continually increase, so will its price.⁵⁰

Finally, novel predictions may arise with endogenous timing and costly delay.⁵¹ Instead of moving in sequence, suppose investors can choose when to act. This creates an incentive for strategic delay in order to glean information from others’ investments. But a naive investor, who wrongly thinks others rely solely on private signals, expects those with optimistic signals to invest immediately and those with pessimistic signals to exit the market. A naive investor therefore neglects others’ incentive to delay and expects high initial investment whenever returns are high. However, because investors will in fact delay, naive observes wrongly attribute the resulting low initial investment to low returns. We conjecture that such environments systematically promote pessimism among naive agents. Analyzing the equilibria of these extensions is part of an on-going research agenda incorporating both naive inference and cursed thinking (Eyster and Rabin, 2005) into models of dynamic learning.

References

- BANERJEE, A. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- BANERJEE, A. AND D. FUDENBERG (2004): “Word-of-mouth learning,” *Games and Economic Behavior*, 46, 1–22.
- BASTIANELLO, F. AND P. FONTANIER (2021): “Partial Equilibrium Thinking in General Equilibrium,” Mimeo.
- BERK, R. H. (1966): “Limiting Behavior of Posterior Distributions when the Model is Incorrect,” *The Annals of Mathematical Statistics*, 37, 51 – 58.
- BIKHCHANDANI, S., D. HIRSHLEIFER, AND I. WELCH (1992): “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, 100, 992–1026.

⁴⁹Eyster and Rabin (2008) define an equilibrium concept called “Inferentially-Naive Information Transmission” (INIT) that incorporates the form of naivete assumed in this paper. INIT reduces to naivete in social-learning environments where players care about others’ actions solely for their informational content.

⁵⁰See Bastianello and Fontanier (2021) for a more recent analysis that verifies this conjecture.

⁵¹Such settings have been analyzed with rational agents by Chamley and Gale (1994) and Gul and Lundholm (1995).

- BOHREN, J. A. (2016): “Informational herding with model misspecification,” *Journal of Economic Theory*, 163, 222–247.
- BOHREN, J. A. AND D. HAUSER (2021): “Learning with Heterogeneous Misspecified Models: Characterization and Robustness,” *Econometrica*, 89, 3025–3077.
- CHAMLEY, C. AND D. GALE (1994): “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, 62, 1065–1085.
- CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2020): “Testing Models of Social Learning on Networks: Evidence From Two Experiments,” *Econometrica*, 88, 1–32.
- COVER, T. M. AND J. A. THOMAS (2006): *Elements of Information Theory*, Wiley-Interscience.
- CRAWFORD, V. P. AND N. IRIBERRI (2007): “Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner’s Curse and Overbidding in Private-Value Auctions?” *Econometrica*, 75, 1721–1770.
- DASARATHA, K. AND K. HE (2020): “Network Structure and Naive Sequential Learning,” *Theoretical Economics*, 15, 415–444.
- DEGROOT, M. H. (1974): “Reaching a Consensus,” *Journal of the American Statistical Association*, 69, 118–121.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *Quarterly Journal of Economics*, 118, 909–968.
- ENKE, B. AND F. ZIMMERMANN (2017): “Correlation Neglect in Belief Formation,” *Review of Economic Studies*, 86, 313–332.
- ESPONDA, I. AND D. POUZO (2016): “Berk-Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models,” *Econometrica*, 84, 1093–1130.
- EYSTER, E. AND M. RABIN (2005): “Cursed Equilibrium,” *Econometrica*, 73, 1623–1672.
- (2008): “Naive Herding,” Mimeo.
- (2010): “Naive Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–243.
- (2014): “Extensive Imitation is Irrational and Harmful,” *Quarterly Journal of Economics*, 129, 1861–1898.
- EYSTER, E., M. RABIN, AND D. VAYANOS (2019): “Financial Markets Where Traders Neglect the Informational Content of Prices,” *The Journal of Finance*, 74, 371–399.
- EYSTER, E., M. RABIN, AND G. WEIZSÄCKER (2018): “An Experiment on Social Mislearning,” Mimeo.
- FOSTER, A. D. AND M. R. ROSENZWEIG (1995): “Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture,” *Journal of Political Economy*, 1176–1209.

- FRICK, M., R. IJIMA, AND Y. ISHII (2020): “Misinterpreting Others and the Fragility of Social Learning,” *Econometrica*, 88, 2281–2328.
- GAGNON-BARTSCH, T. (2016): “Taste Projection in Models of Social Learning,” Mimeo.
- GLAESER, E. L. AND C. G. NATHANSON (2017): “An Extrapolative Model of House Price Dynamics,” *Journal of Financial Economics*, 126, 147–170.
- GOLUB, B. AND M. O. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–49.
- GROSSMAN, S. J. (1976): “On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information,” *Journal of Finance*, 31, 573–585.
- GUARINO, A. AND P. JEHIEL (2013): “Social Learning with Coarse Inference,” *American Economic Journal: Microeconomics*, 5, 147–74.
- GUL, F. AND R. LUNDHOLM (1995): “Endogenous Timing and the Clustering of Agents’ Decisions,” *Journal of Political Economy*, 103, 1039–66.
- KÜBLER, D. AND G. WEIZSÄCKER (2004): “Limited Depth of Reasoning and Failure of Cascade Formation in the Laboratory,” *The Review of Economic Studies*, 71, 425–441.
- LEVY, G. AND R. RAZIN (2018): “Information diffusion in networks with the Bayesian Peer Influence heuristic,” *Games and Economic Behavior*, 109, 262–270.
- MILGROM, P. (1981): “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 12, 380–391.
- MOLAVI, P., A. TAHBAZ-SALEHI, AND A. JADBABAIE (2018): “A Theory of Non-Bayesian Social Learning,” *Econometrica*, 86, 445–490.
- MUELLER-FRANK, M. AND C. NERI (2021): “A General Analysis of Boundedly Rational Learning in Social Networks,” *Theoretical Economics*, 16, 317–357.
- MUNSHI, K. (2004): “Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution,” *Journal of Development Economics*, 73, 185–213.
- SCHWARTZSTEIN, J. (2014): “Selective Attention and Learning,” *Journal of the European Economic Association*, 12, 1423–1452.
- SMITH, L. AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371–398.

Appendix

A Cyclic Beliefs

This appendix provides details on cyclic beliefs and provides an example. As noted in Section 3.4, if the process $\langle \hat{\omega}_t \rangle$ does not converge to a fixed state, then will eventually become periodic. That is, there exists some period $t^* \geq 1$ and a subset of states $\{\hat{\omega}^1, \dots, \hat{\omega}^L\}$ with $L \geq 2$ elements such that (i) for each $l = 1, \dots, L$, we have $\hat{\omega}_{t^*+l} = \hat{\omega}^l$, and (ii) for all $t > t^*$, we have $\hat{\omega}_t = \hat{\omega}_{t+L}$. The process therefore converges to an *absorbing set*, $\{\hat{\omega}^1, \dots, \hat{\omega}^L\}$, comprising fixed points of the L -fold composition of ϕ . That is, for each $\hat{\omega} \in \{\hat{\omega}^1, \dots, \hat{\omega}^L\}$ $\hat{\omega} = \phi^L(\hat{\omega})$.

As discussed in Section 3.4, in an environment that generate cyclic beliefs, the set Ω^* is insufficient for identifying states that are permanently disbelieved. Doing so requires us to instead distinguish all states that are part of an absorbing set. We denote the collection of all such states by $\Omega^{**} \equiv \{\omega \in \Omega \mid \exists L \geq 1 \text{ such that } \phi^L(\omega) = \omega\}$. For any state in $\hat{\omega} \in \Omega^{**}$, it is possible—given the appropriate realized state $\omega \in \Omega$ and initial condition $\hat{\omega}_2 = \omega$ —that society thinks $\hat{\omega}$ is true infinitely often. Thus, any state ω outside Ω^{**} is *strongly abandoned*: in the long-run, *every* generation puts arbitrarily small probability on ω even when it is true.

We now provide an example demonstrating that Ω^* can be empty. As a result, public opinion never settles on a fixed belief and instead continually cycles. While the example is admittedly contrived, it adequately illustrates how the model may give rise to non-convergence. The example concludes by noting how this non-convergence is robust to alternative social-learning structures (e.g., the canonical single-file environment).

Suppose that on a fixed day each week, new inventory arrives at a market. Shoppers must choose a day to visit the market, and thus seek to learn the delivery day $\omega \in \{1, 2, \dots, 7\}$. Conditional on ω , shoppers earn a payoff $u(A|\omega)$ from going to the market on day A . Assume $u(\omega|\omega) = 1$, $u(\omega + 1|\omega) = \frac{9}{10}$, and $u(x|\omega) = 0$ for all $A \notin \{\omega, \omega + 1\}$. That is, it is best to shop on the delivery day; the next day is slightly worse since inventory may be depleted; and all remaining days have no inventory. Additionally, conditional on state ω , each shopper receives a signal $s \in \{\omega - 1, \omega\}$ with mass function $f(s = \omega|\omega) = \frac{2}{3}$. Given this setup, an uncertain autarkic shopper prefers to risk arriving a day late rather than arriving a day early. More precisely, upon observing signal s , a shopper prefers to go the day *after* her signal suggests: $\mathbb{E}[u(s|\omega)|s] = \frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 = \frac{2}{3}$ and $\mathbb{E}[u(s+1|\omega)|s] = \frac{2}{3} \cdot \frac{9}{10} + \frac{1}{3} \cdot 1 = \frac{14}{15}$. This structure generates autarkic distributions with $\mathbb{P}_\omega(\omega) = \frac{1}{3}$, $\mathbb{P}_\omega(\omega + 1) = \frac{2}{3}$, and $\mathbb{P}_\omega(x) = 0$ otherwise.⁵² Hence, in autarky, shoppers expect to see crowds the day *after* delivery—not the day of.

To see how naive social learning evolves, suppose the true delivery day is Thursday. Initially, most shoppers arrive on Friday, and observers correctly deduce the delivery day. The following week, the crowd arrives on Thursday: since there is no longer uncertainty, it is optimal to go on the precise day of delivery. However, when people interpret the Thursday crowd as if it were based solely on private information, they think the delivery must have arrived on Wednesday. Hence, in the third week, the

⁵²For the sake of simplicity, this environment violates Assumption 2.1, which assumes that autarkic distributions have full support over the set of actions. That assumption guarantees that naive agents do not observe individual actions they thought were impossible. This possibility is still ruled out in this particular example despite the violation of Assumption 2. Furthermore, the example would yield identical results if we instead assumed $\mathbb{P}_\omega(x) = \nu$ for $x \notin \{\omega, \omega + 1\}$ for some ν arbitrarily small.

crowd shows up on Wednesday. Rolling forward, it is clear that for all $t \geq 2$, $\hat{\omega}_{t+1}$ is the day before $\hat{\omega}_t$. Beliefs about the delivery day continually cycle through the days of the week. Furthermore, like the example in Section 3.5, the long-run distribution of beliefs is the same no matter the true state.

Although contrived, the lack of convergence demonstrated by this example does not depend on our assumption that each generation is large and observes only the previous generation. In the canonical single-file environment (defined in Section 3.3), public beliefs will still fail to converge to a stationary point belief. Intuitively, a herd on one action always suggests a state in which it is optimal to take an action *different* from the herd. As a result, a naive society cannot maintain any fixed confident belief.

B Proofs

Proof of Lemma 1. Suppose that all agents in Generation t believe the probability of state $\hat{\omega}_t$ is at least $1 - \epsilon$. For ϵ small enough, each agent of type $\theta \in \Theta$ from Generation t chooses the unique action $\arg \max_{A \in \mathcal{A}} u(A | \hat{\omega}_t, \theta)$ (uniqueness stems from Assumption 1). Thus, Generation t 's action distribution \mathbf{a}_t is such that $a_t(m) \rightarrow \mathbb{T}_{\hat{\omega}_t}(m)$ a.s. in N .

Naive observers in period $t + 1$ think actions in t conditional on ω are independent draws from \mathbb{P}_ω : they think $N\mathbf{a}_t \sim \text{Multinomial}(N, \mathbb{P}_\omega)$ in state ω , implying

$$\Pr(\mathbf{a}_t | \omega) = C(N, \mathbf{a}_t) \prod_{m=1}^M \mathbb{P}_\omega(m)^{Na_t(m)},$$

where $C(N, \mathbf{a}) \equiv N! / \prod_{m=1}^M Na(m)!$ is a normalization constant independent of ω . Thus

$$\frac{\pi_{t+1}(j)}{\pi_{t+1}(k)} = \frac{\Pr(\mathbf{a}_t | \omega_j) \pi_1(j)}{\Pr(\mathbf{a}_t | \omega_k) \pi_1(k)} = \left(\frac{\prod_{m=1}^M \mathbb{P}_{\omega_j}(m)^{a_t(m)}}{\prod_{m=1}^M \mathbb{P}_{\omega_k}(m)^{a_t(m)}} \right)^N \frac{\pi_1(j)}{\pi_1(k)}.$$

Since $a_t(m) \rightarrow \mathbb{T}_{\hat{\omega}_t}(m)$ a.s. in N , we have that $\frac{\pi_{t+1}(j)}{\pi_{t+1}(k)}$ converges to 0 in N for all $\omega_j \neq \omega_k \Leftrightarrow \omega_k = \arg \max_{\omega \in \Omega} \prod_{m=1}^M \mathbb{P}_\omega(m)^{\mathbb{T}_{\hat{\omega}_t}(m)}$. By Assumption 3, this state ω_k is unique. So, for any $\beta' \in \mathbb{R}$, we can choose N large enough so that for all $\omega_j \in \Omega \setminus \{\omega_k\}$,

$$\log \left(\frac{\pi_{t+1}(k)}{\pi_{t+1}(j)} \right) \geq \beta'. \quad (\text{B.1})$$

We now show that each agent in Generation $t + 1$ places probability exceeding $1 - \epsilon$ on ω_k . An arbitrary agent $(n, t + 1)$ forms a posterior likelihood ratio $p_{(n,t+1)}(k) / p_{(n,t+1)}(j)$ after observing her private signal, $s_{(n,t+1)}$, and previous actions, \mathbf{a}_t , which is given by

$$\frac{p_{(n,t+1)}(k)}{p_{(n,t+1)}(j)} = \frac{\Pr(\omega_k | s_{(n,t+1)}) \pi_{t+1}(k)}{\Pr(\omega_j | s_{(n,t+1)}) \pi_{t+1}(j)}. \quad (\text{B.2})$$

By Assumption 2 (Part 3),

$$\log \left(\frac{\Pr(\omega_k | s_{(n,t+1)})}{\Pr(\omega_j | s_{(n,t+1)})} \right) \geq -\beta$$

for all signals $s_{(n,t+1)} \in \mathcal{S}$. From Condition (B.1), the posterior likelihood ratio $p_{(n,t+1)}(k)/p_{(n,t+1)}(j)$ must satisfy

$$\log \left(\frac{p_{(n,t+1)}(k)}{p_{(n,t+1)}(j)} \right) = \log \left(\frac{\Pr(\omega_k | s_{(n,t+1)}) \pi_{t+1}(k)}{\Pr(\omega_j | s_{(n,t+1)}) \pi_{t+1}(j)} \right) \geq \beta' - \beta$$

for all j . Taking β' large enough, all agents in Generation $t + 1$ believe the state is ω_k with probability at least $1 - \epsilon$. So fixing $\hat{\omega}_t \in \Omega$, $\pi_{t+1}(\hat{\omega}_{t+1}) > 1 - \epsilon$, where

$$\hat{\omega}_{t+1} = \operatorname{argmax}_{\omega \in \Omega} \prod_{m=1}^M \mathbb{P}_\omega(m)^{\mathbb{T}_{\hat{\omega}_t}(m)} = \operatorname{argmin}_{\omega \in \Omega} \left(- \sum_{m=1}^M \mathbb{T}_{\hat{\omega}_t}(m) \log \mathbb{P}_\omega(m) \right) = \operatorname{argmin}_{\omega \in \Omega} H(\mathbb{T}_{\hat{\omega}_t}, \mathbb{P}_\omega).$$

The proof is completed by defining $\phi(\hat{\omega}_t) \equiv \operatorname{argmin}_{\omega \in \Omega} H(\mathbb{T}_{\hat{\omega}_t}, \mathbb{P}_\omega)$. ■

Proof of Proposition 2. Since each Generation $t \geq 2$ acts purely on the public belief under the large-generation assumption (i.e., private signals do not influence actions), common preferences implies that for all $t \geq 2$, the action distribution \mathbf{a}_t is degenerate (a “herd”). Denote by \mathbf{a}^m the action distribution degenerate on A_m , and let $\mathcal{A}^h = \{\mathbf{a}^m\}_{m=1}^M$, where $M \equiv |\mathcal{A}|$. For each $\mathbf{a}^m \in \mathcal{A}^h$, let $\hat{\omega}^m = \operatorname{argmin}_{\omega \in \Omega} H(\mathbb{P}_\omega, \mathbf{a}^m)$ denote the updated public belief upon observing \mathbf{a}^m , and let Ω^h be the set of distinct values of $\hat{\omega}^m$ across $m = 1, \dots, M$. Since $\hat{\omega}^m$ is unique fixing \mathbf{a}^m , $|\Omega^h| \leq M$. Because $\mathbf{a}_t \in \mathcal{A}^h$ for all $t \geq 2$, $\hat{\omega}_t \in \Omega^h$ for all $t \geq 3$. Since $|\Omega^h| \leq M < |\Omega|$, there must exist at some $\omega' \in \Omega \setminus \Omega^h$ and thus $\hat{\omega}_t \neq \omega'$ for all $t \geq 3$. The fact that $\hat{\omega}_t \in \Omega^h$ for all $t \geq 2$ also implies $|\Omega^*| \leq M$ and thus $\Omega^* \subsetneq \Omega$. ■

Proof of Proposition 3. We will show this result by working with likelihood ratios rather than probabilities, so we first introduce some notation. Note that the public belief entering period t can be written recursively as

$$\pi_t(k) = \frac{\widehat{\Pr}(x_{t-1} | \omega_k, \pi_{t-1}) \pi_{t-1}(k)}{\sum_{\omega_j \in \Omega} \widehat{\Pr}(x_{t-1} | \omega_j, \pi_{t-1}) \pi_{t-1}(j)}, \quad (\text{B.3})$$

where $\widehat{\Pr}(x_{t-1} | \omega_k, \pi_{t-1})$ is a naive observer’s assessment of the likelihood that Player $t - 1$ takes action x_{t-1} in state ω_k given the prior public belief π_{t-1} . Fixing a reference state ω_k , we let $\ell_t^k(j)$ denote the public likelihood ratio of state $\omega_j \in \Omega$ relative to ω_k at the start of round t ; that is, $\ell_t^k(j) \equiv \pi_t(j)/\pi_t(k)$. The vector containing the likelihood ratios of all possible states $\omega_j \neq \omega_k$ relative to ω_k is sufficient for π_t . We denote this vector in \mathbb{R}^{K-1} by $\ell_t^k = (\ell_t^k(1), \dots, \ell_t^k(k-1), \ell_t^k(k+1), \dots, \ell_t^k(K))$. From (B.3), each element of this likelihood-ratio vector updates in a recursive fashion:

$$\ell_t^k(j) = \frac{\widehat{\Pr}(x_{t-1} | \omega_k, \ell_{t-1}^k)}{\widehat{\Pr}(x_{t-1} | \omega_k, \ell_{t-1}^k)} \ell_{t-1}^k(j), \quad (\text{B.4})$$

where $\widehat{\Pr}(x_{t-1} | \omega_k, \ell_{t-1}^k) = \widehat{\Pr}(x_{t-1} | \omega_k, \pi_{t-1})$ is simply the naive perceived likelihood that Player $t - 1$ takes action x_{t-1} in state ω_k given the likelihood-ratio vector rather than π_{t-1} .

Suppose the realized state is ω . To reach a contradiction, suppose that (i) $\Pr(\pi_t \rightarrow \delta(\omega_k) | \omega) > 0$ and that (ii) $\omega_k \notin \Omega^*$. Furthermore, from Lemma 1 and Proposition 1, it follows that $\omega_k \notin \Omega^*$ iff

there exists some $\omega_j \neq \omega_k$ such that:

$$\begin{aligned}
\phi(\omega_k) = \omega_j &\Leftrightarrow H(\mathbb{T}_{\omega_k}, \mathbb{P}_{\omega_k}) > H(\mathbb{T}_{\omega_k}, \mathbb{P}_{\omega_j}) \\
&\Leftrightarrow -\sum_{m=1}^M \mathbb{T}_{\omega_k}(m) \log \mathbb{P}_{\omega_k}(m) > -\sum_{m=1}^M \mathbb{T}_{\omega_k}(m) \log \mathbb{P}_{\omega_j}(m) \\
&\Leftrightarrow \sum_{m=1}^M \mathbb{T}_{\omega_k}(m) [\log \mathbb{P}_{\omega_j}(m) - \log \mathbb{P}_{\omega_k}(m)] > 0 \\
&\Leftrightarrow \sum_{m=1}^M \mathbb{T}_{\omega_k}(m) \log \left(\frac{\mathbb{P}_{\omega_j}(m)}{\mathbb{P}_{\omega_k}(m)} \right) > 0. \tag{B.5}
\end{aligned}$$

The assumption that $\omega_k \notin \Omega^*$ therefore implies that (B.5) holds for some state(s), and let $\omega_j \neq \omega_k$ denote a specific state for which (B.5) holds.

Note that $\pi_t \rightarrow \delta(\omega_k)$ iff $\ell_t^k(i) \rightarrow 0$ for all $i \neq k$ (i.e., $\ell_t^k \rightarrow \mathbf{0}$). Thus, for our assumption that $\Pr(\pi_t \rightarrow \delta(\omega_k)|\omega) > 0$ to hold, we require $\Pr(\ell_t^k(i) \rightarrow 0) > 0$ for all $i \neq k$. For a contradiction, however, we now show that $\Pr(\ell_t^k(j) \rightarrow 0|\omega) = 0$ for ω_j satisfying (B.5). The proof follows along the lines of Smith and Sørensen (2000) and Bohren and Hauser (2021) who similarly consider necessary conditions for the local stability of public beliefs in social-learning settings.

Toward the desired contradiction, suppose $\Pr(\ell^k \rightarrow \mathbf{0}|\omega) > 0$, and consider a sample path where $\ell_t^k \rightarrow \mathbf{0}$. For a fixed $\epsilon > 0$, let $B_\epsilon \equiv [0, \epsilon]^{K-1}$ be the neighborhood about the limit belief $\ell^k = \mathbf{0}$. Let T be the first time such that $\ell_t^k \in B_\epsilon$ for all $t > T$. If T is not finite, then we have the desired contradiction. Thus, suppose T is finite. We now consider how the process $\langle \ell_t^k \rangle$ behaves in B_ϵ for $t > T$, and demonstrate that in fact $\langle \ell_t^k(j) \rangle$ must eventually exit $[0, \epsilon]$ and thus $\langle \ell_t^k \rangle$ exits B_ϵ . From (B.4), the log likelihood ratio $\log \ell_t^k(j)$ updates recursively according to

$$\log \ell_t^k(j) = \log \ell_{t-1}^k(j) + \log \left(\frac{\widehat{\Pr}(x_{t-1}|\omega_k, \ell_{t-1}^k)}{\widehat{\Pr}(x_{t-1}|\omega_k, \ell_{t-1}^k)} \right). \tag{B.6}$$

The sequence in (B.6) is potentially complicated by the fact that a naive observer's perceived probability that Player $t-1$ takes action x_{t-1} could potentially depend on the public beliefs held at the start of round $t-1$. Although this dependence plays a key role in the analysis of the rational belief process (as in Smith and Sørensen 2000), it is irrelevant here because a naive observer assumes that previous players infer nothing from others' actions. Hence, $\widehat{\Pr}(x_{t-1}|\omega_k, \pi_{t-1}) = \Pr(x_{t-1}|\omega_k, \pi_1)$; that is, the perceived probability of x_{t-1} is simply the probability that a player chooses x_{t-1} in autarky given the initial prior π_1 . This corresponds precisely with our definition of the autarkic distribution: for any option $A_m \in \mathcal{A}$, if $x_{t-1} = A_m$, then $\widehat{\Pr}(x_{t-1}|\omega_k, \pi_{t-1}) = \mathbb{P}_{\omega_k}(m)$, where \mathbb{P}_{ω_k} is defined in Definition 1. Thus, the naive public belief (in log-likelihood form) is characterized by

$$\log \ell_t^k(j) = \log \ell_{t-1}^k(j) + \log \left(\frac{\mathbb{P}_{\omega_j}(m)}{\mathbb{P}_{\omega_k}(m)} \right) \quad \text{if } x_{t-1} = A_m. \tag{B.7}$$

In truth, the transition probabilities of sequence (B.7)—that is, the likelihoods of actions $x_{t-1} = A_m$ —depend on Player $t-1$'s inferred beliefs from previous actions, ℓ_{t-1}^k , in addition to her private signal s_{t-1} and preference type θ_{t-1} . Hence, the transition probabilities of $\langle \ell_t^k(j) \rangle$ between any given

periods $t - 1$ and t may depend on the value of ℓ_{t-1}^k within B_ϵ . We will therefore construct a simpler unidimensional sequence $\langle \hat{\ell}_t \rangle$ which has i.i.d. transitions over time, and then compare $\langle \ell_t^k(j) \rangle$ with the simplified process $\langle \hat{\ell}_t \rangle$. First, define $\hat{\ell}_{T+1} = \ell_{T+1}^k(j)$ so that $\hat{\ell}_{T+1} < \epsilon$. Next, let

$$c \equiv \min_{m \in \{1, \dots, M\}} \log \left(\frac{\mathbb{P}_{\omega_j}(m)}{\mathbb{P}_{\omega_k}(m)} \right) \quad (\text{B.8})$$

denote the greatest possible downward change $\ell_t^k(j)$ can attain from one period to the next. Let $x^*(\theta_t, s_t, \ell^k)$ be the optimal action for Player t given type and signal (θ_t, s_t) and public belief ℓ^k , and let $x^*(\theta_t, \mathbf{0})$ denote the optimal action for type θ_t when certain of state ω_k . For $t > T + 1$, we define $\hat{\ell}_t$ by:

$$\log \hat{\ell}_{t+1} = \begin{cases} \log \hat{\ell}_t + \log \left(\frac{\mathbb{P}_{\omega_j}(m)}{\mathbb{P}_{\omega_k}(m)} \right) & \text{if } x^*(\theta_t, s_t, \ell_t^k) = x^*(\theta_t, \mathbf{0}) \text{ for all } \ell^k \in B_\epsilon \\ \log \hat{\ell}_t + c & \text{otherwise.} \end{cases} \quad (\text{B.9})$$

Essentially, $\log \hat{\ell}_t$ updates identically to $\log \ell_t^k(j)$ within B_ϵ so long as the signal s_t does not strongly contradict the public belief $\ell_t^k \in B_\epsilon$. Otherwise, $\log \hat{\ell}_t$ updates in favor of ω_k in the most favorable possible way. Denote the probability of the latter event by $r(\emptyset)$. Similarly, for each $A_m \in \mathcal{A}$, let $r(m)$ be the probability (conditional on the true state ω) of a type-signal combination (θ_t, s_t) such that $x^*(\theta_t, s_t, \ell_t^k) = x^*(\theta_t, \mathbf{0}) = A_m$ for all $\ell^k \in B_\epsilon$. Since these probabilities encompass all possible events dictating the evolution of $\hat{\ell}_t$, we have $r(\emptyset) + \sum_{m=1}^M r(m) = 1$. Furthermore, note that as $\epsilon \rightarrow 0$, the probability of a signal s_t strong enough to induce $x^*(\theta_t, s_t, \ell_t) \neq x^*(\theta_t, \mathbf{0})$ for any $\theta_t \in \Theta$ and $\ell \in B_\epsilon$ converges to zero, and thus $r(m) \rightarrow T_{\omega_k}(m)$ and $r(\emptyset) \rightarrow 0$ as $\epsilon \rightarrow 0$. As such, the assumption that (B.5) holds implies that for ϵ sufficiently small:

$$r(\emptyset)c + \sum_{m=1}^M r(m) \log \left(\frac{\mathbb{P}_{\omega_j}(m)}{\mathbb{P}_{\omega_k}(m)} \right) > 0. \quad (\text{B.10})$$

Consider an ϵ such that (B.10) holds. Since $(\hat{\ell}_{t+1} - \ell_t)_{t=T+1}^\infty$ is i.i.d. given the construction above, $(\hat{\ell}_{t+1} - \ell_t)/t \xrightarrow{a.s.} \mathbb{E}[\log \hat{\ell}_{t+1} - \ell_t] = r(\emptyset)c + \sum_{m=1}^M r(m) \log \left(\frac{\mathbb{P}_{\omega_j}(m)}{\mathbb{P}_{\omega_k}(m)} \right) > 0$. Thus $\lim_{t \rightarrow \infty} \hat{\ell}_t = \lim_{t \rightarrow \infty} [\ell_{T+1}^k(j) + \sum_{\tau=T+1}^t (\hat{\ell}_{\tau+1} - \ell_\tau)] = \infty$. By construction $\hat{\ell}_t$ updates upward weakly less than $\ell_t^k(j)$ so long as $\ell^k \in B_\epsilon$. Thus, $\hat{\ell}_t \leq \ell_t^k(j)$ for all $t \geq T + 1$ such that $\ell^k \in B_\epsilon$. Therefore, our assumption that $\ell_t \in B_\epsilon$ for all $t > T$ implies that $\lim_{t \rightarrow \infty} \ell^k(j) \geq \lim_{t \rightarrow \infty} \hat{\ell}_t = \infty$, a contradiction. Thus, $\Pr(\pi_t \rightarrow \delta(\omega_k) | \omega) = 0$. ■

Proof of Proposition 4. First consider the case where $\Omega^* \neq \emptyset$. If the true state is $\omega \in \Omega^*$, then $\hat{\omega}_2 = \omega$ and $\hat{\omega}_3 = \phi(\omega) = \omega$. Rolling forward, $\hat{\omega}_t = \omega$ for all $t \geq 2$. Next consider the case where $\Omega^* = \emptyset$. Since long-run beliefs never settle on a fixed state for any initial condition $\hat{\omega}_2 \in \Omega$, it follows that public beliefs necessarily cycle over some absorbing set that depends on the initial condition $\hat{\omega}_2 \in \Omega$. Consider any potential absorbing set $\hat{\Omega}$. Given that $\Omega^* = \emptyset$, this set cannot be a singleton, so $|\hat{\Omega}| \equiv L$ such that $2 \leq L \leq K$. By the definition of an absorbing set, the sequence of public beliefs over $\hat{\Omega}$ must be periodic, and therefore each $\hat{\omega} \in \hat{\Omega}$ is a fixed point of the L -fold

composition of ϕ . Thus if the true state is ω , then $\hat{\omega}_2 = \omega$ and $\hat{\omega}_{2+L} = \phi^L(\hat{\omega}_2) = \omega$. It follows from induction on τ that $\hat{\omega}_{2+\tau L} = \omega$ for all $\tau = 1, 2, \dots$, and thus public beliefs put arbitrarily high probability on the true state infinitely often. ■

Proof of Proposition 5. Part 1. Suppose $\omega = H$, so $\hat{\omega}_2 = H$ and $\mathbf{a}_2 = \mathbb{T}_H$, where $\mathbb{T}_H(1) = \lambda$. Using Lemma 1, $\hat{\omega}_3 = H \Leftrightarrow \mathbb{P}_H(1)^\lambda \mathbb{P}_H(0)^{1-\lambda} > \mathbb{P}_L(1)^\lambda \mathbb{P}_L(0)^{1-\lambda}$. Note that $\mathbb{P}_H(1) = \psi[\lambda\rho + (1-\lambda)(1-\rho)]$ and $\mathbb{P}_L(1) = \psi[(1-\lambda)\rho + \lambda(1-\rho)]$. Letting $\ell \equiv \lambda\rho + (1-\lambda)(1-\rho)$, the autarkic frequencies simplify to $\mathbb{P}_H(1) = \psi\ell$ and $\mathbb{P}_L(1) = \psi(1-\ell)$. Thus $\mathbb{P}_H(1)^\lambda \mathbb{P}_H(0)^{1-\lambda} > \mathbb{P}_L(1)^\lambda \mathbb{P}_L(0)^{1-\lambda}$ iff

$$k_H(\ell, \psi) \equiv \left(\frac{\ell}{1-\ell}\right)^\lambda \left(\frac{1-\psi\ell}{1+\psi\ell-\psi}\right)^{1-\lambda} > 1.$$

Since $k_H(\ell, \psi)$ is decreasing in ψ , $k_H(\ell, \psi) > 1$ for all $\psi \in (0, 1)$ if $k_H(\ell, 1) > 1$. Since $\lambda > 1/2$ and $\rho > 1/2$ imply $\ell \in (1/2, 1)$, it follows that $k_H(\ell, 1) = \left(\frac{\ell}{1-\ell}\right)^{2\lambda-1} > 1$. Hence $\hat{\omega}_3 = H$, which implies $\omega = H$ is a fixed point of ϕ , and thus $\hat{\omega}_t = H$ for all $t > 2$.

Part 2. Suppose $\omega = L$, so $\hat{\omega}_2 = L$ and $\mathbf{a}_2 = \mathbb{T}_L$, where $\mathbb{T}_L(1) = 1-\lambda$. Following the same logic as Part 1, $\hat{\omega}_3 = H$ iff

$$k_L(\ell, \psi) \equiv \left(\frac{\ell}{1-\ell}\right)^{1-\lambda} \left(\frac{1-\psi\ell}{1+\psi\ell-\psi}\right)^\lambda > 1.$$

Fixing $\lambda > 1/2$, $k_L(\ell, \psi) > 1 \Leftrightarrow \psi < \frac{1-\Lambda}{\lambda-\Lambda(1-\ell)} \equiv \bar{\psi}(\ell)$, where $\Lambda \equiv \left(\frac{1-\ell}{\ell}\right)^{\frac{1-\lambda}{\lambda}}$. Note that $\bar{\psi}(\ell)$ is decreasing in ℓ : $\frac{\partial}{\partial \ell} \bar{\psi}(\ell) < 0 \Leftrightarrow \left(\frac{2\ell-1}{\ell^2}\right) \left(\log\left(\frac{1-\lambda}{\lambda}\right) + 1\right) \Lambda < 1$. This holds for any $\ell \in (1/2, 1)$ because: (i) $\frac{2\ell-1}{\ell^2} < 1$, (ii) $\log\left(\frac{1-\lambda}{\lambda}\right) + 1 < 1$, and (iii) $\Lambda < 1$. Finally, it's straightforward to verify that $\bar{\psi}(.5) = 1$ and $\bar{\psi}(1) = 0$. Thus, $\ell \in (1/2, 1)$ implies that $\bar{\psi}(\ell) \in (0, 1)$ and $\psi < \bar{\psi}(\ell) \Leftrightarrow \hat{\omega}_3 = H$. Hence, if $\psi > \bar{\psi}$, then $\hat{\omega}_3 = L$. This implies that $\omega = L$ is a fixed point of ϕ . Thus, $\hat{\omega}_t = L$ for all $t > 1$. If $\psi < \bar{\psi}(\ell)$, then $\hat{\omega}_3 = H$. As shown in Part 1, $\omega = H$ is a fixed point of ϕ for all values of ψ , meaning $\hat{\omega}_t = H$ for all $t > 1$. ■

Proof of Proposition 6. We first prove a lemma that we use to prove both this proposition some that follow.

Lemma B.1. *If each F^m satisfies MLRP, then $\mathbb{P}_\omega(m)$ is strictly increasing in q^m and strictly decreasing in q^j for all $j \neq m$.*

Proof of Lemma B.1. Without loss of generality, we prove the result for $\mathbb{P}_\omega(1)$. We make use of well-known implications of MLRP (see Milgrom, 1981, Proposition 2):

Remark B.1. *Suppose F^m satisfies MLRP.*

1. $\mathbb{E}[q^m | s^m]$ is strictly increasing in $s^m \in S^m$.
2. $F^m(s | q^m)$ satisfies first-order stochastic dominance in s : if $q^m > \tilde{q}^m$, then for all $s \in S^m$, $F^m(s | q^m) \leq F^m(s | \tilde{q}^m)$.

In autarky (i.e., $t = 1$), Player n with signal realization $s = (s^1, \dots, s^M)$ chooses A_1 if $1 = \arg \max_{m \in \{1, \dots, M\}} \mathbb{E}[q^m | s]$. Since signals are independent across options, $\mathbb{E}[q^m | s] = \mathbb{E}[q^m | s^m]$ for each m , and MLRP implies that each $\mathbb{E}[q^m | s^m]$ is strictly increasing in s^m . (All remaining instances of “increasing” and “decreasing” within this proof are meant in the strict sense.) For each

m , let $k_m(\cdot)$ be the increasing function implicitly defined by $\mathbb{E}[q^1|s^1] > \mathbb{E}[q^m|s^m]$ iff $s^1 > k_m(s^m)$. Thus, action A_1 is chosen iff $s^1 > k_m(s^m)$ for all $m > 1$. This implies that in state $\omega = (q^1, \dots, q^M)$, the autarkic probability of choice A_1 is

$$\mathbb{P}_\omega(1) = \prod_{m=2}^M \Pr(s^1 > k_m(s^m)|\omega) = \prod_{m=2}^M \int_{S^m} 1 - F^1(k_m(s^m)|q^1) dF^m(s^m|q^m). \quad (\text{B.11})$$

We first show that $\mathbb{P}_\omega(1)$ is (strictly) increasing in q^1 . From Remark B.1, $F^1(k_m(s^m)|q^1)$ is decreasing in q^1 , which implies that each term $\int_{S^m} 1 - F^1(k_m(s^m)|q^1) dF^m(s^m|q^m)$ of the product in Equation B.11 is increasing in q^1 , and thus $\mathbb{P}_\omega(1)$ is increasing in q^1 . Next, we show $\mathbb{P}_\omega(1)$ is decreasing in q^m for all $m \geq 2$. For any arbitrary $m \geq 2$, note that each term of the product in Equation B.11 can be expressed as $\mathbb{E}[h(s^m)|q^m]$ where $h(s^m) = 1 - F^1(k_m(s^m)|q^1)$ is a decreasing function of s^m independent of q^m . It is well known that if random variable X first-order stochastically dominates X' , then $\mathbb{E}[h(X)] < \mathbb{E}[h(X')]$ for any decreasing function $h(\cdot)$ provided these expectations are finite. Since s^m conditional on q^m first-order stochastically dominates s^m conditional on \tilde{q}^m iff $q^m > \tilde{q}^m$, $\mathbb{E}[h(s^m)|\tilde{q}^m] > \mathbb{E}[h(s^m)|q^m]$ iff $q^m > \tilde{q}^m$, which implies that $\mathbb{E}[h(s^m)|q^m] = \int_{S^m} 1 - F^1(k_m(s^m)|q^1) dF^m(s^m|q^m)$ is decreasing in q^m . Thus, from Equation B.11, $\mathbb{P}_\omega(1)$ is decreasing in q^m . Since $\mathbb{P}_\omega(1)$ is increasing in q^1 and decreasing in q^m for all $m \geq 2$, it follows that ω_e^1 uniquely maximizes $\mathbb{P}_\omega(1)$. This concludes the proof of Lemma B.1.

We now use Lemma B.1 to prove Proposition 6. Without loss of generality, index options such that $q^1 = \arg \max_m \{q^m\}$. By Assumption 2, \mathbf{a}_1 reveals ω to Generation 2, implying $a_2(1) = 1$ and $a_2(m) = 0$ for $m \geq 2$. From Lemma 1, $\hat{\omega}_3 = \phi(\hat{\omega}_2) = \arg \max_{\tilde{\omega} \in \Omega} \prod_{m=1}^M \mathbb{P}_{\tilde{\omega}}(m)^{a_2(m)} = \arg \max_{\tilde{\omega} \in \Omega} \mathbb{P}_{\tilde{\omega}}(1)$. Hence, Lemma B.1 implies $\hat{\omega}_3 = \omega_e^1$. Since $\hat{\omega}_3 = \omega_e^1$, all players in Generation 3 choose A_1 , implying $\hat{\omega}_4 = \phi(\omega_e^1) = \arg \max_{\omega \in \Omega} \mathbb{P}_\omega(1) = \omega_e^1$. Since ω_e^1 is a fixed point of ϕ , $\hat{\omega}_t = \omega_e^1$ for all $t > 2$. \blacksquare

Proof of Proposition 7. Suppose the state is ω^0 and suppose Generation 2 puts arbitrarily high probability on ω^0 , so $\hat{\omega}_2 = \omega^0$. So long as M is finite, which implies that the prior likelihood ratio $\pi_1(\omega^0)/\pi_1(\omega^m) = \frac{\chi^{\frac{1}{M}}}{1-\chi^{\frac{1}{M}}}$ is finite, we can invoke Lemma 1 to determine $\langle \hat{\omega}_t \rangle$. Since $\hat{\omega}_2 = \omega^0$, we have $\mathbf{a}_2 \rightarrow \mathbb{T}_{\omega^0} = (1 - \lambda, \lambda, 0, \dots, 0)$ as $N \rightarrow \infty$. From Lemma 1, $\hat{\omega}_3 = \phi(\hat{\omega}_2) = \arg \max_{\omega \in \Omega} \prod_{m=0}^M \mathbb{P}_\omega^{a_2(m)} = \arg \max_{\omega \in \Omega} \mathbb{P}_\omega(0)^{1-\lambda} \mathbb{P}_\omega(1)^\lambda$. Hence, $\hat{\omega}_3$ is the unique state satisfying

$$\left(\frac{\mathbb{P}_\omega(0)}{\mathbb{P}_{\hat{\omega}_3}(0)} \right)^{1-\lambda} \left(\frac{\mathbb{P}_\omega(1)}{\mathbb{P}_{\hat{\omega}_3}(1)} \right)^\lambda < 1 \quad (\text{B.12})$$

for all $\omega \in \Omega \setminus \{\hat{\omega}_3\}$. Given that actions $m = 2, \dots, M$ are not chosen, it is immediate that $\hat{\omega}_3 \in \{\omega^0, \omega^1\}$. Hence, we simply need to consider when

$$\left(\frac{\mathbb{P}_{\omega^0}(0)}{\mathbb{P}_{\omega^1}(0)} \right)^{1-\lambda} \left(\frac{\mathbb{P}_{\omega^0}(1)}{\mathbb{P}_{\omega^1}(1)} \right)^\lambda < 1; \quad (\text{B.13})$$

if B.13 holds, then $\hat{\omega}_3 = \omega^1$, otherwise $\hat{\omega}_3 = \omega^0$. Define $\mathcal{L}(m|M) \equiv \mathbb{P}_{\omega^0}(m)/\mathbb{P}_{\omega^1}(m)$ as a function of M . The remainder of the proof characterizes values of M for which $\mathcal{L}(0|M)^{1-\lambda} \mathcal{L}(1|M)^\lambda < 1$.

We first derive $\mathcal{L}(0|M)$. Type $\theta = H$ never chooses A_0 . Type $\theta = L$ chooses $A_0 \Leftrightarrow q_L^0 \geq \mathbb{E}[q^m|s^m]$

for all m . $\mathbb{E}[q^m|s^m] = (1 - p^m)\underline{q}^m$ where p^m is the posterior belief that $q^m = 0$ conditional on s^m . Note that

$$p^m = \Pr(q^m = 0|s^m) = \left[1 + \left(\frac{\chi^{\frac{1}{M}}}{1 - \chi^{\frac{1}{M}}} \right) \frac{f(s^m|\underline{q}^m)}{f(s^m|0)} \right]^{-1}.$$

Let $\tilde{F}^m(p|q^m)$ denote the distribution of posterior beliefs p^m conditional on q^m induced by the underlying signal distribution $F^m(s|q^m)$. Since $q_L^0 \geq \mathbb{E}[q^m|s^m]$ iff p^m is less than threshold $\bar{p}_L^m \equiv \frac{|q^m - q_L^0|}{|q^m|}$, type L chooses A_0 in state ω^0 with probability

$$\mathbb{P}_{\omega^0}(0) = \prod_{m=1}^M \tilde{F}^m(\bar{p}_L^m|\underline{q}^m).$$

In state ω^1 , this probability is

$$\mathbb{P}_{\omega^1}(0) = \tilde{F}^1(\bar{p}_L^1|0) \prod_{m=2}^M \tilde{F}^m(\bar{p}_L^m|\underline{q}^m).$$

Hence, the likelihood ratio of observing A_0 in ω^0 relative to ω^1 is

$$\mathcal{L}(0|M) = \frac{\mathbb{P}_{\omega^0}(0)}{\mathbb{P}_{\omega^1}(0)} = \frac{\tilde{F}^1(\bar{p}_L^1|\underline{q}^1)}{\tilde{F}^1(\bar{p}_L^1|0)} > 1,$$

where the inequality follows from MLRP (see Remark B.1). Because $\mathcal{L}(0|M)$ is independent of M , we write it simply as $\mathcal{L}(0)$.

We now derive $\mathcal{L}(1|M)$. Type θ chooses A_1 if both $\mathbb{E}[q^1|s^1] \geq \mathbb{E}[q^m|s^m]$ for all $m > 1$ and $\mathbb{E}[q^1|s^1] > q_\theta^0$. Note that $\mathbb{E}[q^1|s^1] \geq \mathbb{E}[q^m|s^m] \Leftrightarrow p^1 \geq 1 - \left(\frac{q^m}{q^1}\right)(1 - p^m) \equiv k_m(p^m)$. This happens with probability

$$\int_0^1 1 - \tilde{F}^1(k_m(p)|q^1) d\tilde{F}^m(p|q^m),$$

which implies that in state ω^0 , type $\theta = H$ chooses A_1 with probability

$$\prod_{m=2}^M \int_0^1 1 - \tilde{F}^1(k_m(p)|\underline{q}^1) d\tilde{F}^m(p|\underline{q}^m),$$

and type $\theta = L$ chooses A_1 with probability

$$\left(1 - \tilde{F}^1(\bar{p}_L^1|\underline{q}^1)\right) \prod_{m=2}^M \int_0^1 1 - \tilde{F}^1(k_m(p)|\underline{q}^1) d\tilde{F}^m(p|\underline{q}^m).$$

Hence,

$$\mathcal{L}(1|M) = \frac{\mathbb{P}_{\omega^0}(1)}{\mathbb{P}_{\omega^1}(1)} = \left(\frac{1 - (1 - \lambda)\tilde{F}^1(\bar{p}_L^1|\underline{q}^1)}{1 - (1 - \lambda)\tilde{F}^1(\bar{p}_L^1|0)} \right) \prod_{m=2}^M \frac{\int_0^1 1 - \tilde{F}^1(k_m(p)|\underline{q}^1) d\tilde{F}^m(p|\underline{q}^m)}{\int_0^1 1 - \tilde{F}^1(k_m(p)|0) d\tilde{F}^m(p|\underline{q}^m)}. \quad (\text{B.14})$$

MLRP implies that each term in the product above (Equation B.14) is strictly less than 1. This implies that $\mathcal{L}(1|M)$ is strictly decreasing in M and the sequence $(\mathcal{L}(1|M))_{M=1}^{\infty}$ converges to 0. Let \bar{M} be the smallest integer such that $\mathcal{L}(0)^{1-\lambda}\mathcal{L}(1|\bar{M})^\lambda < 1$, and note that \bar{M} is finite since $\mathcal{L}(0)^{1-\lambda}$ is finite and constant in M .

Thus, whenever $M \geq \bar{M}$, $\hat{\omega}_3 = \omega^1$. Furthermore, if $\hat{\omega}_3 = \omega^1$, then $a_3(1) = 1$. By Lemma B.1, such autarkic behavior is most likely in ω^1 . Hence, $\hat{\omega}_4 = \omega^1$, implying ω^1 is absorbing. Thus, $M \geq \bar{M}$ implies $\hat{\omega}_t = \omega^1$ for all $t \geq 3$. If $M < \bar{M}$, then $\hat{\omega}_3 = \omega^0$. Since ω^0 is thus an absorbing state in this case, $M < \bar{M}$ implies $\hat{\omega}_t = \omega^0$ for all $t \geq 2$. ■

Proof of Proposition 8. We first derive the autarkic distribution of actions and the function determining aggregate demand in autarky. We will then use this to prove the various parts of the proposition. Consider period $t = 1$ and an arbitrary Player $(n, 1)$ with signal $s_{(n,1)}$. Since the player's prior and signal about ω are normally distributed, it follows that Player $(n, 1)$ has a normally-distributed posterior with mean $\widehat{\mathbb{E}}[\omega|I_{(n,1)}] = \frac{\rho_s}{\rho_s + \rho_\omega} s_{(n,1)} + \frac{\rho_\omega}{\rho_s + \rho_\omega} \bar{\omega}$ and variance $\widehat{\text{Var}}[\omega|I_{(n,1)}] = \frac{1}{\rho_s + \rho_\omega}$ (see, e.g., DeGroot 1970). It then follows from Equation 4 that Player $(n, 1)$'s optimal allocation is

$$z_{(n,1)} \equiv \frac{\rho_s + \rho_\omega}{2(\rho_s + \rho_\omega) + \rho_\eta} \left\{ 1 + \frac{\rho_\eta}{\alpha W_0} \left(\frac{\rho_s}{\rho_s + \rho_\omega} s_{(n,1)} + \frac{\rho_\omega}{\rho_s + \rho_\omega} \bar{\omega} \right) \right\}, \quad (\text{B.15})$$

so long as $z_{(n,1)} \in (0, 1)$. Thus, Player $(n, 1)$'s chosen allocation is

$$x_{(n,1)} = \begin{cases} 0 & \text{if } z_{(n,1)} \leq 0, \\ z_{(n,1)} & \text{if } z_{(n,1)} \in (0, 1), \\ 1 & \text{if } z_{(n,1)} \geq 1. \end{cases} \quad (\text{B.16})$$

Aggregate demand in autarky is thus $\bar{x}_1 = \mathbb{E}[x_{(n,1)}|\omega]$. Since $s_{(n,1)} \sim N(\omega, \rho_s^{-1})$, it follows from Equation B.15 that $z_{(n,1)} \sim N(\mu, \sigma^2)$ with $\mu = \nu_0 + \nu_1\omega$, where

$$\nu_0 \equiv \frac{1}{2(\rho_s + \rho_\omega) + \rho_\eta} \left\{ \rho_\omega \left(\frac{\rho_\eta}{\alpha W_0} \bar{\omega} + 1 \right) + \rho_s \right\}, \quad (\text{B.17})$$

$$\nu_1 \equiv \frac{\rho_\eta}{\alpha W_0} \left(\frac{\rho_s}{2(\rho_s + \rho_\omega) + \rho_\eta} \right), \quad (\text{B.18})$$

and $\sigma = \nu_1/\sqrt{\rho_s}$. Letting Φ and ϕ denote the standard normal CDF and PDF, respectively, we then have

$$\begin{aligned} \mathbb{E}[x_{(n,1)}|\omega] &= \left[\Phi \left(\frac{1-\mu}{\sigma} \right) - \Phi \left(\frac{-\mu}{\sigma} \right) \right] \mathbb{E}[z_{(n,1)}|z_{(n,1)} \in (0, 1), \omega] \\ &\quad + \left[1 - \Phi \left(\frac{1-\mu}{\sigma} \right) \right] \cdot 1 + \Phi \left(\frac{-\mu}{\sigma} \right) \cdot 0, \end{aligned} \quad (\text{B.19})$$

where

$$\mathbb{E}[z_{(n,1)}|z_{(n,1)} \in (0, 1), \omega] = \mu - \sigma [\phi((1-\mu)/\sigma) - \phi(-\mu/\sigma)] / [\Phi((1-\mu)/\sigma) - \Phi(-\mu/\sigma)]. \quad (\text{B.20})$$

Thus, Equation B.19 reduces to

$$\mathbb{E}[x_{(n,1)}|\omega] = 1 - \sigma[\psi((1 - \mu)/\sigma) - \psi(-\mu/\sigma)], \quad (\text{B.21})$$

where $\psi(x) \equiv x\Phi(x) + \phi(x)$. For notational convenience, let $D_A(\omega) \equiv \mathbb{E}[x_{(n,1)}|\omega]$ denote the autarkic aggregate demand for the risky asset in state ω .

Note that $D_A(\omega)$ is strictly increasing in ω and takes a value on the interior of $[0, 1]$ for all finite values of ω . This implies that D_A is invertible and \bar{x}_1 perfectly identifies ω : Generation $t = 2$ reaches a public belief $\hat{\omega}_2 = D_A^{-1}(\bar{x}_1) = \omega$. Similarly, each following naive Generation $t \geq 2$ reaches a public belief $\hat{\omega}_t$ such that $\hat{\omega}_t = D_A^{-1}(\bar{x}_{t-1})$. However, this inference will be incorrect since \bar{x}_{t-1} is not determined by D_A for $t > 2$. Instead, for $t > 2$, x_{t-1} results from a demand function, denoted by $D_F(\hat{\omega})$, that specifies the aggregate demand for the risky asset when investors are fully informed that the fundamental is equal to $\hat{\omega}$. From Equation 4, $D_F(\hat{\omega}) = \frac{1}{2} + \frac{\rho_\eta}{2\alpha W_0}\hat{\omega}$ when this quantity is in $(0, 1)$, and it is 0 or 1 otherwise. The path of public beliefs is then determined by $\hat{\omega}_t = D_A^{-1}(\bar{x}_{t-1}) = D_A^{-1}(D_F(\hat{\omega}_{t-1}))$, where the second equality follows from the fact that $\bar{x}_{t-1} = D_F(\hat{\omega}_{t-1})$. Put differently, $D_A(\hat{\omega}_t) = D_F(\hat{\omega}_{t-1})$. Define ω^* by $D_A(\omega^*) = D_F(\omega^*)$.

We now describe the evolution of beliefs. In particular, we will show that $\omega > \omega^*$ implies $\langle \hat{\omega}_t \rangle$ (and hence $\langle x_t \rangle$) is increasing in t , while $\omega < \omega^*$ implies that $\langle \hat{\omega}_t \rangle$ (and hence $\langle x_t \rangle$) is decreasing in t . Toward that end, let $H(\omega) \equiv D_F(\omega) - D_A(\omega)$ denote the change in demand that occurs when investors move from autarky to full information about ω . We will first show that H is strictly increasing on the interval of ω such that $D_F(\omega) \in (0, 1)$, and that H has a unique root of ω^* on this domain.

Note that the interval of ω for which $D_F(\omega) \in (0, 1)$ is $\Omega^I \equiv (-\alpha W_0/\rho_\eta, \alpha W_0/\rho_\eta)$, which follows from the expression for D_F , above. For $\omega \in \Omega^I$ we have

$$\begin{aligned} \frac{\partial H}{\partial \omega} &= \frac{\rho_\eta}{2\alpha W_0} + \sigma \left[\Phi\left(\frac{1-\mu}{\sigma}\right) - \Phi\left(\frac{-\mu}{\sigma}\right) \right] \left(\frac{-\partial\mu/\partial\omega}{\sigma} \right) \\ &= \frac{\rho_\eta}{2\alpha W_0} - \left[\Phi\left(\frac{1-\mu}{\sigma}\right) - \Phi\left(\frac{-\mu}{\sigma}\right) \right] (\partial\mu/\partial\omega) > 0, \end{aligned}$$

which follows from the fact that the difference of CDFs in brackets is strictly less than 1 and $\partial\mu/\partial\omega = \nu_1 < \frac{\rho_\eta}{2\alpha W_0}$ where ν_1 is defined in Equation B.18. The fact that H is strictly increasing implies that the root ω^* is unique. Moreover, it implies that the path of beliefs and allocations will be either monotonically increasing or decreasing in t , as we show next.

Part 1. Note that agents in generation $t + 1$ only form point beliefs when $\bar{x}_t \in (0, 1)$ and hence $\hat{\omega}_t \in \Omega^I$. Otherwise, $\bar{x}_t \in \{0, 1\}$ and thus the generation similarly allocates all wealth to the same asset as generation t , resulting in $\bar{x}_{t+1} = \bar{x}_t$. Thus, the system reaches a steady state once $\hat{\omega}_t$ exits Ω^I . Thus, we show that the sequences $\langle \hat{\omega}_t \rangle$ and $\langle \bar{x}_t \rangle$ are strictly increasing in t while beliefs are in Ω^I (and they are constant otherwise). Let $\Delta_\omega(t) \equiv \hat{\omega}_t - \hat{\omega}_{t-1}$ and $\Delta_x(t) \equiv \bar{x}_t - \bar{x}_{t-1}$. First note that the expression for $D_F(\hat{\omega})$ along with the fact that $\bar{x}_t = D_F(\hat{\omega}_t)$ implies that if $\Delta_\omega(t) > 0$ then $\Delta_x(t) > 0$, and if $\Delta_\omega(t) < 0$, then $\Delta_x(t) < 0$. Thus, allocations are monotonic in t whenever beliefs are. We now show that $\hat{\omega}_t$ is indeed monotonic in t . First, suppose that $\omega \in \Omega^I$ and $\omega > \omega^*$. We will show by induction that $\Delta_\omega(t) > 0$ for all $t > 2$ such that $\hat{\omega}_{t-1} \in \Omega^I$. First consider the base case. Since $\hat{\omega}_2 = \omega > \omega^*$, it follows that $H(\hat{\omega}_2) > 0$ and hence $D_F(\hat{\omega}_2) > D_A(\hat{\omega}_2)$. Note that $\hat{\omega}_3$ is such that $D_A(\hat{\omega}_3) = D_F(\hat{\omega}_2)$; thus, the preceding inequality implies $D_A(\hat{\omega}_3) > D_A(\hat{\omega}_2)$ and thus $\hat{\omega}_3 > \hat{\omega}_2$

since D_A is strictly increasing. Thus, $\Delta_\omega(3) > 0$. Now suppose $\Delta_\omega(t)$ is positive for all $t \leq t^*$ for any $t^* > 3$ and suppose $\hat{\omega}_t \in \Omega^I$. We have $\Delta_\omega(t^* + 1) > 0$ if $D_F(\hat{\omega}_{t^*+1}) - D_F(\hat{\omega}_{t^*}) > 0$. But since $D_F(\hat{\omega}_{t^*}) = D_A(\hat{\omega}_{t^*+1})$, the previous condition is equivalent to $H(\hat{\omega}_{t^*+1}) > 0$. Since $H(\hat{\omega}) > 0$ for all $\hat{\omega} > \omega^*$, this will hold so long as $\hat{\omega}_{t^*+1} > \omega^*$. To prove that this final inequality must hold, note that $D_A(\hat{\omega}_{t^*+1}) = D_F(\hat{\omega}_{t^*})$. Since the induction hypotheses implies that $D_F(\hat{\omega}_{t^*}) > D_A(\hat{\omega}_{t^*})$, we therefore have $D_A(\hat{\omega}_{t^*+1}) > D_A(\hat{\omega}_{t^*}) \Leftrightarrow \hat{\omega}_{t^*+1} > \hat{\omega}_{t^*}$. This implies that $\hat{\omega}_{t^*+1} > \omega^*$ given that we must have $\hat{\omega}_{t^*} > \omega^*$ due to the induction hypothesis and the fact that $\hat{\omega}_3 > \omega > \omega^*$. Finally, note that $\Delta_\omega(t) > 0$ implies that $\hat{\omega}_t > \hat{\omega}_{t-1}$, and thus we have shown that for $\hat{\omega}_t$ strictly increases in t so long as $\hat{\omega}_{t-1} \in \Omega^I$ whenever $\omega > \omega^*$. An entirely analogous argument shows that $\omega < \omega^*$ implies that $\Delta_\omega(t) < 0$ for all $t > 2$, and hence $\hat{\omega}_t$ decreases in t when $\omega < \omega^*$. The details are omitted given that the argument is nearly identical to the one above. Hence, the proof of Part 1 is complete.

Part 2. We now show that ω^* is increasing in $\bar{\omega}$ and that $\omega^* < \bar{\omega}$. Since H is strictly increasing in ω on the domain of interest, the implicit function theorem implies that, for any parameter y , we have $\text{sgn}\left(\frac{\partial \omega^*}{\partial y}\right) = -\text{sgn}\left(\frac{\partial H}{\partial y}\Big|_{\omega=\omega^*}\right)$.

ω^* is strictly increasing in $\bar{\omega}$. since D_F is independent of $\bar{\omega}$ and the only component of D_A that depends on $\bar{\omega}$ is μ , we have

$$\frac{\partial H}{\partial \bar{\omega}} = \sigma \left[\Phi\left(\frac{1-\mu}{\sigma}\right) - \Phi\left(\frac{-\mu}{\sigma}\right) \right] \left(\frac{-\partial \mu / \partial \bar{\omega}}{\sigma} \right) \quad (\text{B.22})$$

$$= - \left[\Phi\left(\frac{1-\mu}{\sigma}\right) - \Phi\left(\frac{-\mu}{\sigma}\right) \right] \frac{1}{2(\rho_s + \rho_\omega) + \rho_\eta} \left(\frac{\rho_\omega \rho_\eta}{\alpha W_0} \right) \quad (\text{B.23})$$

$$< 0, \quad (\text{B.24})$$

and thus ω^* is increasing in $\bar{\omega}$. Furthermore, it is straightforward to verify that $\lim_{\bar{\omega} \rightarrow \infty} D_A(\omega) = 1$ and $\lim_{\bar{\omega} \rightarrow -\infty} D_A(\omega) = 0$.

The threshold ω^ lies below $\bar{\omega}$.* The expression for $\frac{\partial \omega^*}{\partial \bar{\omega}}$ derived above can be used to show that $\frac{\partial \omega^*}{\partial \bar{\omega}} \in (0, 1)$. Hence, to show that $\omega^* < \bar{\omega}$, it suffices to show that $\omega^* < 0$ when $\bar{\omega} = 0$. Note that we have $\omega^* < 0$ if $D_F(0) > D_A(0)$, and note that $D_F(0) = 1/2$. Thus, we need only show that $D_A(0) < 1/2$ when $\bar{\omega} = 0$. First notice that $D_A(\omega) = 1/2$ if and only if $\mu = 1/2$. To see this, Equation B.21 implies that if $\mu = 1/2$, then

$$D_A(\omega) = 1 - \sigma[b\Phi(b) - a\Phi(a) + \phi(b) - \phi(a)], \quad (\text{B.25})$$

where $b = 1/2\sigma$ and $a = -b$. Substituting these values of a and b into the expression above and using the symmetry properties of ϕ and Φ , we have

$$D_A(\omega) = 1 - \frac{1}{2}[\Phi(b) + \Phi(-b)] = \frac{1}{2}. \quad (\text{B.26})$$

Thus, to show $D_A(0) < 1/2$ when $\bar{\omega} = 0$, it suffices to show $\mu < 1/2$ at $\omega = 0$ when $\bar{\omega} = 0$. In this case, Equation B.17 implies that $\nu_0 = \frac{1}{2 + \rho_\eta / (\rho_s + \rho_\omega)} < 1/2$. Hence, $\omega^* < \bar{\omega}$.

Part 3. We show that the change in beliefs across periods, $\Delta_\omega(t)$, is increasing in magnitude in t so long as $\hat{\omega}_{t-1} \in \Omega^I$. For simplicity, suppose $\omega > \omega^*$ so that $\Delta_\omega(t)$ is positive for $\hat{\omega}_{t-1} \in \Omega^I$ (as in Part 1); the case for $\omega < \omega^*$ is analogous aside from changes in beliefs being negative rather than positive, and hence is omitted. Since D_F is linear in $\hat{\omega}$, notice that $\Delta_\omega(t) > \Delta_\omega(t-1) \Leftrightarrow$

$D_F(\hat{\omega}_t) - D_F(\hat{\omega}_{t-1}) > D_F(\hat{\omega}_{t-1}) - D_F(\hat{\omega}_{t-2})$. Now notice that $D_A(\hat{\omega}_t) = D_F(\hat{\omega}_{t-1})$ and $D_A(\hat{\omega}_{t-1}) = D_F(\hat{\omega}_{t-2})$. Substituting these values into the previous condition yields the equivalent condition of $\Delta_\omega(t) > \Delta_\omega(t-1) \Leftrightarrow D_F(\hat{\omega}_t) - D_A(\hat{\omega}_t) > D_F(\hat{\omega}_{t-1}) - D_A(\hat{\omega}_{t-1}) \Leftrightarrow H(\hat{\omega}_t) > H(\hat{\omega}_{t-1})$. Since Part 1 established that $\Delta_\omega(t) > 0$, we have $\hat{\omega}_t > \hat{\omega}_{t-1}$ and thus $H(\hat{\omega}_t) > H(\hat{\omega}_{t-1})$ since H is strictly increasing. Since the sequence $\langle \hat{\omega}_t \rangle$ is monotonically increasing and exhibits increasing differences while in Ω^I , then if $\omega \in \Omega^I$, there must exist a finite t^* such that $\hat{\omega}_{t^*} \in \Omega^I$ yet $\hat{\omega}_{t^*+1} \notin \Omega^I$. Thus, $\bar{x}_t = 1$ for all $t \geq t^* + 1$. Moreover, we will argue that $\hat{\omega}_t = \infty$ for all $t > t^* + 1$. Note $\bar{x}_t = 1$ implies each player (n, t) selects $x_{(n,t)} = 1$. From Equation 4, if acting in autarky, $x_{(n,t)} = 1 \Leftrightarrow s_{(n,t)} > c$, where c is the value of $s_{(n,1)}$ in Equation B.15 such that $z_{(n,1)} = 1$. Hence, $\mathbb{P}_\omega(x = 1) = 1 - \Phi((c - \omega)/\sqrt{\rho_s^{-1}})$. Since Lemma 1 assumes a finite action and state space, we cannot directly invoke it here. However, the logic naturally extends. The continuous analog of the cross entropy between $\mathbb{P}_\omega(x)$ and $\mathbb{T}_{\hat{\omega}}(x)$ is $-\int_0^1 T_{\hat{\omega}}(x) \log \mathbb{P}_\omega(x) dx$. Note that for $t > t^* + 1$, $\mathbb{T}_{\hat{\omega}_t}(x) = \delta(x - 1)$, where $\delta(\cdot)$ is the Dirac delta function. Hence $-\int_0^1 T_{\hat{\omega}_t}(x) \log \mathbb{P}_\omega(x) dx = -\log \mathbb{P}_\omega(1)$, which is minimized at the state in which $\mathbb{P}_\omega(1) = 1 - \Phi((c - \omega)/\sqrt{\rho_s^{-1}})$ is maximized. Thus $\hat{\omega}_{t+1} = \sup \Omega = \infty$. Furthermore, given this belief, $\mathbb{T}_{\hat{\omega}_{t+1}}(x) = \delta(x - 1)$, implying again that $\hat{\omega}_{t+2} = \infty$. Extending this logic, we have $\hat{\omega}_\tau = \infty$ for all $\tau > t$. As noted above, the proof that $\hat{\omega}_\tau \rightarrow -\infty$ when the system reaches its lower boundary ($\bar{x}_t = 0$) is analogous and omitted. ■

Proof of Corollary 1. Since the aggregate shocks are identically distributed for both assets, Asset r dominates Asset s under full information if and only if $\mathbb{E}[d_t^r|\omega] > \mathbb{E}[d_t^s|\omega] \Leftrightarrow \omega > 0$. There are two cases to consider: $\omega^* > 0$ and $\omega^* < 0$. First, suppose $\bar{\omega} > 0$ is high enough that $\omega^* > 0$. Define $\Omega' = (0, \omega^*)$. From Proposition 8, for any value of $\omega \in \Omega'$, $\langle \hat{\omega}_t \rangle$ diverges to $-\infty$. Hence, whenever $\omega \in \Omega'$, investors believe Asset s dominates Asset r even though the opposite is true. Now suppose $\bar{\omega} < 0$ is low enough that $\omega^* < 0$. Define $\Omega' = (\omega^*, 0)$. For any realization of $\omega \in \Omega'$, $\langle \hat{\omega}_t \rangle$ diverges to $+\infty$. Hence, investors wrongly believe Asset r dominates Asset s . ■