

RATIONAL GROUPTHINK

MATAN HAREL¹, ELCHANAN MOSSEL², PHILIPP STRACK³, AND OMER TAMUZ⁴

ABSTRACT. We study how long-lived rational agents learn from repeatedly observing a private signal and each others' actions. With normal signals, a group of any size learns more slowly than just four agents who directly observe each others' private signals in each period. Similar results apply to general signal structures. We identify *rational groupthink*—in which agents ignore their private signals and choose the same action for long periods of time—as the cause of this failure of information aggregation.

JEL: D83.

1. INTRODUCTION

A key question in social learning is: How well do agents learn by observing each others' actions? As the analysis of the beliefs of long-lived Bayesian agents is challenging¹, most of the literature focuses either on short-lived agents² or on non-rational belief dynamics such as the DeGroot model or quasi-Bayesian agents³. By applying large deviation techniques we overcome the difficulty associated with the analysis of Bayesian beliefs and analyze social learning with long-lived rational agents.

Formally, we consider a group of myopic Bayesian agents who repeatedly observe private signals about a binary state, as well as each others' past actions. As every agent eventually learns the state, our main focus is on the speed of learning, i.e. the rate with which the agents converge to the correct action. Our main result is that information aggregation will fail for Bayesian agents in a large society: An arbitrarily large group of Bayesian agents

¹TEL AVIV UNIVERSITY

²MASSACHUSETTS INSTITUTE OF TECHNOLOGY

³YALE UNIVERSITY

⁴CALIFORNIA INSTITUTE OF TECHNOLOGY

We thank seminar audiences in Berkeley, Berlin, Bonn, Caltech, Chicago, Düsseldorf, Duke, Harvard, Medelín, Microsoft Research New England, MIT, Montreal, NYU, Penn State, Pittsburgh, Princeton, San Diego, UPenn, USC and Washington University, Yale, as well as Nageeb Ali, Ben Brooks, Dirk Bergemann, Kim Border, Federico Echenique, Wade Hann-Caruthers, Benjamin Golub, Rainie Heck, Paul Heidhues, Shachar Kariv, Navin Kartik, Steven Morris, Luciano Pomatto, Larry Samuelson, Lones Smith, Juuso Toikka, Leeat Yariv and others for insightful comments and discussions. Matan Harel was partially supported by the IDEX grant of Paris-Saclay. Elchanan Mossel is supported by ONR grant N00014-16-1-2227 and NSF grant CCF 1320105. Omer Tamuz was supported by a grant from the Simons Foundation (#419427).

¹See e.g. [Cripps, Ely, Mailath, and Samuelson \(2008\)](#).

²See e.g. [Dasaratha, Golub, and Hak \(2018\)](#); [Mueller-Frank and Arieli \(2018\)](#).

³See [Golub and Jackson \(2010\)](#); [Molavi, Tahbaz-Salehi, and Jadbabaie \(2018\)](#).

observing each others' actions will only learn as fast as a small group of agents observing each others' signals directly. For example, when signals are normal, four agents sharing their signals learn faster than a group of arbitrarily many agents who observe each others' actions, but not signals.

This failure of information aggregation is caused by endogenous correlation in the agents' actions, which reduces the amount of information that these actions reveal about the private signals.⁴ Whereas signals are independent, the agents' actions become correlated, as they are all heavily influenced by the past, observed actions. This correlation is an immediate consequence of the incentive to learn from each others' actions. For example, if agent 1 takes an action that is optimal in some state of the world, the other agents will infer that agent 1's private belief indicates that this state is relatively likely and will themselves take this action with greater probability. A greater number of agents increases this correlation, as agents share more common information. This decreases the amount of information that the agents' actions reveal about their private signals. The insight of our analysis is that as the number of agents grows, the correlation increases to an extent that completely out-weighs the gain of the additional independent private signals. We show that asymptotically this failure of information aggregation holds for any signal structure, any utility function and any number of agents.

What inference an agent draws from the actions of another agent depends on her belief about the other agent's belief. Thus, agents' actions may depend on their higher order beliefs. This poses a significant challenge for the exact characterization of behavior. We circumvent this problem by focusing on long-term probabilities, and by analyzing a phenomenon that we call "rational groupthink", conditional on which higher-order beliefs admit a tractable structure. We define rational groupthink to be the event that *all* agents take the wrong action for many periods. We show that when this event occurs, it is likely that all agents have private signals that indicate the correct action. Through a fixed-point argument we are able to estimate the asymptotic probability of rational groupthink, and find that rational groupthink occurs so often that agents in a large group learn almost as slowly as they do in autarky. Hence, in this sense, rational groupthink prevents almost all information aggregation.⁵

⁴As is well known, a large number of sufficiently correlated signals convey less information than a small number of independent signals.

⁵Our prediction seems to be in line with the findings in the empirical literature: [Da and Huang \(2020, page 5\)](#) find in a study on forecasters "that private information may be discarded when a user places weights on the prior forecasts [of others]. In particular, errors in earlier forecasts are more likely to persist and appear in the final consensus forecast, making it less efficient."

Groupthink arises in our model as a consequence of Bayesian updating, rather than being driven by an assumed desire for conformity.⁶ Rational groupthink occurs after a consensus on an action is formed in the initial periods, making it optimal for every agent to continue taking the consensus action, even when her private information indicates otherwise. Indeed, we show that typically, after a wrong consensus forms, all agents eventually observe private signals which provide strong evidence for choosing the correct action, and yet a long time may pass until any of them breaks the wrong consensus (Theorem 2). Thus a situation arises in which each agent’s private information indicates the correct action, and yet, because of the group dynamics, all agents choose the wrong action.

We study the effect of increasing the group size. On the one hand, with more agents, each individual agent is less likely to break a wrong consensus. On the other hand, the number of potential dissenters is larger, and so a priori it is not obvious whether rational groupthink becomes more or less likely. Our first main result shows that, even as the number of agents goes to infinity, the speed of learning from actions stays bounded by a constant (Theorem 1), whereas the speed of learning from the aggregated signals, which is proportional to the number of agents, goes to infinity (Fact 2). Thus, in a large group, almost no information is aggregated; the agents’ beliefs when observing only actions have the same precision as would result from observing a vanishingly small fraction of the available private signals. Specifically, for normal signals, a group of n agents observing each others’ actions learns asymptotically slower than a group of 4 agents who share their private signals; this holds for any number of agents! Hence, at most a fraction of $4/n$ of the private information is transmitted through actions (Corollary 1). We proceed beyond normal signals to show that for any signal distribution at most a fraction of c/n of the private information is transmitted through actions, for some constant c that depends only on the distribution of the private signals (Lemma 12).

As mentioned above, we quantify the speed of learning as the asymptotic rate at which agents converge to the correct action. An advantage of asymptotic rates is that they are independent of many details of the model, providing a measure that is robust to changes in model parameters such as the agents’ prior or the exact utility function. Furthermore, they are tractable. For similar reasons of tractability and robustness, many previous works have studied asymptotic (long run) rates of learning in various settings.⁷

⁶See [Angeletos and Pavan \(2007\)](#) for a setting in which payoff externalities lead to a desire for conformity, which in turn lead agents to discount their private signals.

⁷Examples of papers studying the rate of learning are [Vives \(1993\)](#); [Chamley \(2004\)](#); [Duffie and Manso \(2007\)](#); [Duffie, Malamud, and Manso \(2009\)](#); [Duffie, Giroux, and Manso \(2010\)](#). Asymptotic rates also have been studied in other settings in which it is difficult to analyze the short-term dynamics (e.g., [Hong and Shum, 2004](#); [Hörner and Takahashi, 2016](#)). [Molavi et al. \(2018\)](#) study the rate of learning in an almost identical setting, with boundedly rational agents.

As a robustness test, we complement our results on the asymptotic rates by an analysis of the probability with which the wrong action is chosen in early periods. We study a canonical setting of a large group of agents with normal private signals, where, as the size of the group is increased, the total precision of their aggregate signal is kept constant. This regime guarantees that the total amount of information available to society is independent of the number of agents, allowing us to study groups which can be very large, but still not learn immediately.

Using numerical simulations, we show that, for example, the probability with which an agent makes a mistake in period 10 equals roughly 14% with either 40 or 100 agents observing each others' actions, but equals 0.07% and is thus 200 times smaller, if signals are public. These and other simulation results indicate that our asymptotic results—which only have formal implications for late enough periods—sometimes already hold in the early periods.

We complement these simulations with our second main theorem, which shows that in this setting, as the number of agents goes to infinity, the probability that an agent chooses correctly in some given period tends to the (roughly constant) probability with which the majority of agents choose correctly in the first period (Theorem 3). This is because, after observing the first period actions, agents will tend to ignore their private signals for many periods. Thus, when the group is large, the private signals of period two and later periods are effectively lost, and information fails to aggregate not only asymptotically, but already after the first period.

A defining feature of our model is that information flows bidirectionally. In §6 we study a setting in which information flows only in one direction, and show that there, a non-vanishing fraction of information is aggregated. We consider a partial observation structure in which agent 1 observes the actions of all others in addition to his private signals, and each of the remaining agents observes his private signals only. In this setting, agent 1 will learn with a speed that increases linearly with the number of agents (Theorem 5), in sharp contrast to our main result. This highlights that the almost complete failure of information aggregation in our baseline setting occurs because of the bidirectional flow of information, and not just because agents observe actions rather than signals. This is in contrast to the herding literature where information aggregation fails even when information travels only unidirectionally.

We assume throughout that agents are Bayesian and myopic: they completely discount future payoffs, and thus at every time period choose the action that maximizes the expected payoff at that period. In a repeated action setting with non-myopic agents there may be a strategic incentive to change ones own action in order to gain more information from future actions of others. This effect does not exist for rational myopic agents, and we make this

assumption for tractability, as does most of the learning literature.⁸ A possible justification for this approach is that reasoning about the informational effect of one’s actions in such setups requires a level of sophistication that seems unrealistic in many applications.⁹

Related Literature. Most of the preceding literature studies situations where each agent observes a *single* signal and agents try to infer the others’ signals from repeatedly observing their actions. [Geanakoplos and Polemarchakis \(1982\)](#); [Sebenius and Geanakoplos \(1983\)](#); [Parikh and Krasucki \(1990\)](#); [Mossel, Sly, and Tamuz \(2015\)](#) give conditions under which agents actions agree in the long-run. [Rosenberg, Solan, and Vieille \(2009\)](#) show that, for a large class of social learning models, agents reach asymptotic agreement, provided the agents observe enough information about each others’ actions. The question of how well information is aggregated in such settings was considered in an important paper by [Vives \(1993\)](#), who studies the rate at which information is aggregated through noisy prices.

In contrast to this literature we allow for agents to repeatedly observe signals about the state of the world and the actions of others. The only other article which we are aware of that tackles this problem is [Molavi et al. \(2018\)](#), which studies asymptotic rates of learning under (non-rational) linear belief updating rules in complex observational networks. The focus of this paper differs from ours: they allow for complex network structures, but impose simple linear belief updating rules. In contrast, we study the complexities associated with Bayesian learning, but assume that all actions are commonly known. Interestingly, our results contrast their findings; while in their model information is quickly aggregated, in our model it is not. This is in part a consequence of the difference in the rationality assumptions.¹⁰

[Gale and Kariv \(2003\)](#) use numerical methods to characterize the asymptotic rates with which rational agents learn, and emphasize the importance of understanding the rates at which Bayesian agents learn from each other.¹¹

Our work is also related to models of rational herding as we use the same conditional i.i.d. structure of signals, and utilities depend only on one’s own actions and the state.¹²

⁸Indeed, the same choice is made in most of the learning literature (where signals are private and agents interact repeatedly) either explicitly (e.g., [Sebenius and Geanakoplos, 1983](#); [Parikh and Krasucki, 1990](#); [Bala and Goyal, 1998](#); [Keppo et al., 2008](#)), or implicitly, by assuming that there is a continuum of agents (e.g., [Vives, 1993](#); [Gale and Kariv, 2003](#); [Duffie and Manso, 2007](#); [Duffie et al., 2009, 2010](#)).

⁹We conjecture that all our results generalize to the case of non-myopic agents, but this extension requires substantial technical innovation, beyond the techniques developed in this paper.

¹⁰In §7 we present numerical simulations that indicate that non-Bayesian updating could lead to faster learning in our setting.

¹¹[Gale and Kariv \(2003, p.20\)](#): “Speeds of convergence can be established analytically in simple cases. For more complex cases, we have been forced to use numerical methods. The computational difficulty of solving the model is massive even in the case of three persons [...] This is an important subject for future research.”

¹²See the original papers of [Bikhchandani, Hirshleifer, and Welch \(1992\)](#) and [Banerjee \(1992\)](#), as well as [Smith and Sørensen \(2000\)](#); [Chamley \(2004\)](#); [Acemoglu, Dahleh, Lobel, and Ozdaglar \(2011\)](#); [Rosenberg and Vieille \(2019\)](#), and many others.

While in sequential models the number of agents is equal to the number of time periods, in our model these can be varied independently.¹³

A more significant difference is that in herding models each agent acts only once, and thus information is transmitted between agents only in one direction, which implies that higher order beliefs to play no role. A contribution of this paper is to show that the failure of information aggregation is not particular to sequential models, but more generally extends to situations of repeated interactions. Our main finding, the rational groupthink effect, has no analogue in sequential herding models, since, in these models, once a herd starts, it is not true that every agent’s private signal indicates the correct action.

Our work is also related to the literature that studies how groupthink arises from various other motives. [Bénabou \(2012\)](#) shows that anticipatory or Kreps-Porteus preferences can lead agents to willingly ignore freely available information if others chose to ignore it even absent any social learning. [Ottaviani and Sørensen \(2001\)](#) demonstrate how a desire to appear well informed can lead to herding and groupthink. [Angeletos and Pavan \(2007\)](#) show how a desire to coordinate actions can lead agents to discount their private signals in favor of public information.

Potential applications of our results appear in settings in which agents repeatedly learn from each other. These include the dissemination of information in developing countries (e.g., [Conley and Udry \(2010\)](#); [Banerjee et al. \(2013\)](#) among many studies), the adoption of opinions on social networks, and prediction markets where forecasters observe the forecasts of others (see [Da and Huang, 2020](#)).

2. LEADING EXAMPLE: AGGREGATING INFORMATION THROUGH PRICES

As a leading economic example, consider local monopolistic sellers who want to learn about the quality of a new product and the associated optimal price. For concreteness, imagine that each seller is the owner of a theater / shop in a different city, and has to decide how much to charge for a new movie, musical, book, toy or fashion item. Because sellers act in different markets there are no direct payoff externalities. Assume that the product is either good or bad, with corresponding demand either high or low. As the demand in other markets is informative about the product’s quality, it is also informative about demand in the seller’s home market. When marginal profits are not constant in the volume of sales, a seller will want to set one price if the demand is high, another price if the demand is low, and potentially intermediate prices when she is unsure about the state. Consequently, each seller wants to learn the state and can do so not only by observing her local demand, but also by observing the prices set by other sellers.

¹³As an exception, [Dasaratha and He \(2019\)](#) recently study the effect of changing population size in a sequential model with overlapping generations.

A second application is that of learning about the quality of a new government policy (e.g., Obamacare) via social media. A group of people who differ in their location, income, family status, etc., each receive private signals from their experience with the new policy, and share their coarse opinion of it on social media, where they also observe the opinions of others.¹⁴

To model these situations, we assume that each seller decides each period which price to charge, and then observes demand in her city, as well as the prices charged by sellers in other cities. To illustrate our main results in the cleanest possible setting, we make a number of simplifying assumptions in this section:

- (i) the sellers choose between only two prices;
- (ii) each seller receives a payoff of 1 if she charges the high price and the product is good or she charges a low price and the product is bad, and a payoff of 0 otherwise;
- (iii) the signal the seller observes at each period is normally distributed with mean -1 when demand is low, and mean $+1$ when demand is high;
- (iv) Finally, the variance of each seller’s signal equals n ; this ensures that the total information contained in all sellers combined signals is constant and allows us to compare outcomes for different numbers of sellers.¹⁵

We consider more than two possible prices, arbitrary payoff functions, and arbitrary signal distributions in our general results which we present in §3 and the following sections.¹⁶

2.1. Speed of Learning from Actions. For this setup we used Monte Carlo simulations to compute the error probabilities in each period, as well as other quantities of interest (see §F for details); we use the term “error” to describe the choice of an action that is not optimal, given the realized state, such as choosing the high price when the product is bad. Figure 2.1 displays the results of these simulations. The plot on the left side shows the error probability for different group sizes when observing actions, and for the case where all signals are public.

¹⁴A less economic example—which we nevertheless find compelling—is that of a group of agents who are interested in knowing whether or not there is a god. Every morning, each toasts bread for breakfast, and checks to see if a divine signal appears in the burn patterns. If there is no god, then the probability of this event is very low. If there is a god, then this probability is significantly higher, although still low. After breakfast, people declare to each other whether or not they believe in god, based on some common threshold of belief. The importance of minor miracles to the belief in god has a long history; for example, in Judaism and Christianity, the concept of *special providence* (Hebrew: *hashgacha pratit*, literally meaning “private monitoring”) refers to the idea that god frequently performs small miracles to his believers (e.g., Maimonides, 1904, part 3, chapters 17-18). Hume (1748) discusses the reliability of reports of miracles and their implications on beliefs. See Holder (1998) for a modern discussion of Bayesian updating of the belief in god following the reporting of miracles. For a recent example see, e.g., <http://news.bbc.co.uk/2/hi/americas/4019295.stm>. We thank the editor for suggesting this example.

¹⁵The analogue of this assumption in a setting with Binary signals would be to keep the total number of signals fixed and have each agent privately observe an equal share of these signals.

¹⁶For a more realistic model of random demand within our framework, one could assume that the number of customers interested in buying the product is Poisson distributed, with parameter depending on the state.




FIGURE 2.1. The probability with which an agent takes the wrong action over time in absolute terms (on the left) and the number of agents sharing their signals which would lead to the same probability of error (on the right).

What immediately stands out is how much slower a group of agents learns by observing each others' actions relative to observing each others' signals. For example, the probability with which an agent makes a mistake in period 10 equals roughly 14% with either 40 or 100 agents observing each others' actions, but equals 0.07% and is thus 200 times smaller, if signals are public; note that by our choice of variance the probability of error with public signals is independent of the number of agents.

Another way one could measure how much information is lost due to the fact that agents observe actions is by looking at the smallest number of agents that could match the probability of error by sharing their signals. We draw this illustration in the right plot of Figure 2.1. For example, in period 30 we have the following comparisons: 3 agents sharing their signals are less likely to take the wrong action than 40 agents observing each others' actions. And 5 agents sharing their signals are less likely to take the wrong action than 100 agents observing each others' actions. Thus, in this example 92% (resp., 95%) of the information contained in the agents' private signals is lost when the agents learn only from actions. We establish in Theorem 1 that this phenomenon is not due to the specific set of parameters chosen in the example: whenever signals are normal, 4 agents sharing their signals are eventually more likely to make the correct choice than n agents observing each others' actions, for any n ! For non-normal signals the same result holds, with the number 4 replaced by a constant depending on the distribution, and given explicitly in Theorem 1.

2.2. Why Learning from Actions is Slow. Why is it that most information is lost when only actions are observed? To better understand this phenomenon it is instructive to study the correlation between an agent's actions and his private signals. We plot this correlation on the left in Figure 2.2. This correlation is a (rough) measure of the information that can be inferred from an agent's actions. In the first period each agent does not observe any




FIGURE 2.2. Correlation between the action taken by an agent, and the action the agent would take based on her private signals only (on the left). On the right is the correlation, conditioned on the agent choosing the wrong action.

information from other agents and thus chooses her action based solely on her first period signal. This leads to a correlation of 1 and the revelation of all private information in the first period. As agents' first period signals are independent conditional on the state, so are their actions. Thus, when observing the others' first period actions, an agent observes many conditionally independent signals.¹⁷ When the number of agents is large the information in the first period actions is likely to lead to a stronger signal than each agent's private signals in subsequent periods. As a consequence, agents are likely to follow the action taken by the majority in the first period. This in turn leads agents to not condition their actions on their own signals, making future actions uninformative and hindering information aggregation. The sudden drop in correlation between the agent's action and private signals in the second period shown in Figure 2.2 illustrates this effect. It is apparent from Figure 2.2 that the low correlation between the agents' actions and private signals prevails for many periods, and is significantly lower for 100 compared to 40 agents. This is formalized in Theorem 3, which shows that given a sufficiently large number of agents, in any given period after the second, all agents with high probability ignore their private signals, leading to a small correlation between actions and private signals.

2.3. Groupthink. The right plot of Figure 2.3 shows that the event in which all agents choose incorrectly does not have insignificant probability, but in fact happens often, conditional on an agent choosing incorrectly. Thus, a signal agent making a mistake is closely tied with the entire group making a mistake.

The right plot of Figure 2.2 shows that when the agents choose the incorrect action, their private signals are negatively correlated with their actions. In other words, conditioned on

¹⁷As a consequence, the probability of error drops considerably between the first and second period; see Figure 2.1.




FIGURE 2.3. Probability of correct private signals (on the left) and all agents taking the wrong action (on the right) conditional on taking the wrong action in a period.

making a mistake, an agent is likely to have a correct private signal. For example, as is shown on the left side of Fig. 2.3, conditioned on choosing the wrong action, an agent’s private signal indicates the correct action with probability 57%. This may be surprising, as one might have reasonably expected that when an agent chooses the wrong action, it is because of incorrect private signals.

This is formalized in Theorem 2, which captures the groupthink effect: agents take the incorrect action because of the group influence, and despite having the correct signal. Moreover, no agent needs to have an incorrect private signal for all agents to choose the incorrect action. In fact, as Theorem 2 shows, when all agents take the wrong action in late periods, they all, with high probability, have private signals indicating the correct action.

3. SETUP

Time is discrete and indexed by $t \in \{1, 2, \dots\}$. Each period, each agent $i \in \{1, 2, \dots, n\}$ first observes a signal (or shock) $s_t^i \in \mathbb{R}$, takes an action $a_t^i \in A$, and finally observes the actions taken by others this period. The set of possible actions is finite: $|A| < \infty$.

3.1. States and Signals. There is an unknown state

$$\Theta \in \{\mathbf{b}, \mathbf{g}\}$$

randomly chosen by nature, with probability $p_0 = \mathbb{P}[\Theta = \mathbf{g}] \in (0, 1)$. For ease of exposition, we call \mathbf{b} the bad state and \mathbf{g} the good state, even though the model is completely symmetric in the state. Signals s_t^i are i.i.d, across agents i and over time t , conditional on the state Θ , with distribution μ_Θ . Throughout, we denote by $\mathbb{E}_\theta[\cdot] := \mathbb{E}[\cdot | \Theta = \theta]$ and $\mathbb{P}_\theta[\cdot] := \mathbb{P}[\cdot | \Theta = \theta]$ the expectation and probability conditional on the state. The distributions $\mu_{\mathbf{g}}$

and $\mu_{\mathbf{b}}$ are mutually absolutely continuous¹⁸ and hence no signal perfectly reveals the state. As a consequence the log-likelihood ratio of every signal

$$\ell_t^i = \log \frac{d\mu_{\mathbf{g}}}{d\mu_{\mathbf{b}}}(s_t^i)$$

is well defined (i.e., $|\ell_t^i| < \infty$) and we assume that it has finite expectation $|\mathbb{E}[\ell_t^i]| < \infty$. We also assume that priors are generic¹⁹, so as to avoid the expository overhead of treating cases in which the agents are indifferent between actions; the results all hold even without this assumption.

Our signal structure allows for bounded as well as unbounded likelihoods.²⁰ Our main example is that of *normal signals* $s_t^i \sim \mathcal{N}(m_\theta, \sigma^2)$ with mean m_θ depending on the state and variance σ^2 . Another example is that of *binary signals* $s_t^i \in \{\mathbf{b}, \mathbf{g}\}$ which are equal to the state with constant probability $\mathbb{P}_\theta[s_t^i = \theta] = \phi > 1/2$.

3.2. Actions and Payoffs. Agent i 's payoff (or utility) in period t depends on her action a_t^i and next period's signal s_{t+1}^i , and is given by $u(s_{t+1}^i, a_t^i)$.²¹ Note that $u(\cdot, \cdot)$ does not depend on the agent's identity i or the time period t .²² We denote by a^θ the action that maximizes the flow payoff in state θ , which we assume is unique

$$\alpha^\theta := \arg \max_{\alpha \in A} \bar{u}(\theta, \alpha).$$

We call $\alpha^{\mathbf{g}}, \alpha^{\mathbf{b}}$ the *certainty actions* and assume that they are distinct (i.e., $\alpha^{\mathbf{g}} \neq \alpha^{\mathbf{b}}$), as otherwise the problem is trivial.

It is an important feature of this model that externalities are purely informational, i.e., each agent's utility is independent of the others' actions, and hence agents care about others' actions only because they may provide information. Furthermore, private signals are independent of actions, and so agents have no experimentation motive; they learn the same information from their signals, irregardless of the actions that they take.

¹⁸That is, every event with positive probability under one measure has positive probability under the other.

¹⁹That is, chosen from a Lebesgue measure one subset of $[0, 1]$.

²⁰In the herding literature agents either learn or do not learn the state, depending on whether private signals have bounded likelihood ratios (Smith and Sørensen, 2000). In our model, the distinction between unbounded and bounded private signals is not important, since the aggregate of each agent's private information suffices to learn the state.

²¹Note, that observing the utility $u(s_{t+1}^i, a_t^i)$ does not provide any information beyond the signal s_{t+1}^i and therefore past signals $(s_1^i, \dots, s_{t+1}^i)$ are a sufficient statistic for the private information available to agent i when taking an action in period $t + 1$.

²²This model is equivalent to a model where the agent's utility $\bar{u}(\Theta, a_t^i)$ is unobserved and depends directly on the state. Formally, we can translate the model where the utility depends on the signal into the model where it depends on the state by setting it equal to the expected payoff conditional on the state $\theta \in \{\mathbf{b}, \mathbf{g}\}$ $\bar{u}(\theta, \alpha) := \mathbb{E}_\theta[u(s_{t+1}^i, \alpha)]$.

3.3. Agents’ Behavior and Information. We denote by p_t^i the posterior probability that agent i assigns to the event $\Theta = \mathbf{g}$ after observing her private signal and before choosing her period t action.

We assume throughout that agents are myopic: at each period they choose the action that maximizes their stage utility, completely discounting the future. As an agent’s posterior belief p_t^i is a sufficient statistic for her expected payoff, her action a_t^i almost surely depends only on p_t^i .²³

Each agent observes only her own signals, and not the signals of others. To learn about the state, agents try to infer the signals of others from their actions. More precisely, at the end of each period an agent observes the actions taken by all other agents in this period.

3.3.1. Example: Matching the State. Our model allows for any (finite) number of actions, and any signal distribution. Nevertheless, a simple example which suffices to understand all the economic results of the paper is the case of two actions $A = \{\mathbf{b}, \mathbf{g}\}$ where the agent’s expected utility equals one if she matches the state, i.e.

$$\bar{u}(\theta, \alpha) = \begin{cases} 1 & \text{if } \alpha = \theta \\ 0 & \text{if } \alpha \neq \theta \end{cases}.$$

In this case the agent simply takes the action to which her posterior belief assigns higher probability:

$$a_t^i = \begin{cases} \mathbf{g} & \text{if } p_t^i > \frac{1}{2} \\ \mathbf{b} & \text{otherwise} \end{cases}.$$

4. RESULTS

In this section we describe our results; §5 derives the learning dynamics in detail and explains how they lead to the results of this section. We consider the probability with which an agent i takes a suboptimal action in period t :

$$a_t^i \neq \alpha^\Theta.$$

We refer to this event as agent i “making a mistake” by “choosing the wrong action”, even though she takes the action which is optimal given her information. As a benchmark we first briefly discuss the classical single agent case.

²³This statement holds, as for almost every prior the agent will not be indifferent.

4.1. **Autarky.** In the single agent case $n = 1$, the probability of a suboptimal action is known to decay exponentially, with a rate r_{aut} that can be calculated explicitly in terms of the cumulant generating functions²⁴ $\lambda_g(z) := -\log \mathbb{E}_g [e^{-z\ell}]$ and $\lambda_b(z) := -\log \mathbb{E}_b [e^{z\ell}]$.²⁵

Fact 1 (Speed of learning in autarky). *The probability that a single agent in autarky chooses the wrong action in period t satisfies*²⁶

$$(1) \quad \mathbb{P} [a_t \neq \alpha^\Theta] = e^{-r_{aut} \cdot t + o(t)},$$

where

$$r_{aut} := \sup_{z \geq 0} \lambda_g(z) = \sup_{z \geq 0} \lambda_b(z).$$

We provide a formal proof of this fact in the appendix, where we also explain why $\sup_{z \geq 0} \lambda_g(z) = \sup_{z \geq 0} \lambda_b(z)$. This type of autarky result is classical in the statistics literature and can be found, for example, in studies of Bayesian hypothesis testing; see, e.g. [Cover and Thomas \(2006, pages 314-316\)](#). For us it serves as a benchmark for the case when agents try to learn from the actions of others. We prove [Fact 1](#) in the Appendix, for the convenience of the reader.

Note, that the long-run probability of a mistake is independent of the set of actions and the utility function. It is also independent of the prior. Thus quantifying the speed of learning using the exponential rate has both advantages and disadvantages: the rate is independent of many details of the model and depends only on the private signal distributions. It is also tractable and can be explicitly calculated for many distributions. However, it is an asymptotic measure and in general does not say anything formally about what happens in early periods.

4.2. **Many agents.** We now turn to the case where there are $n \geq 2$ agents. We first consider the benchmark case where all signals are observed by all agents. Since there is no private information, all agents hold the same beliefs, and this case reduces to the single agent case, but where n signals are observed in every period. After t periods the agents will have observed $n \cdot t$ signals, and so, by [Fact 1](#), their probability of taking the wrong action will be the probability of error after $n \cdot t$ periods in the autarky setting.

Fact 2 (Speed of learning with public signals). *When signals are public, the probability that any agent i chooses the wrong action in period t satisfies*

²⁴Here ℓ is a random variable with a distribution that is equal to that of any of the log-likelihood ratios ℓ_i^j .

²⁵The signs used in this definition deviate from the standard definition $\log \mathbb{E}_\theta [e^{z\ell}]$ of the cumulant generating function of ℓ . Our choice allows for a convenient formulation of [Lemma 2](#) below, and reflects the fact that in the good state, high ℓ indicates a correct signal, while in the low state it indicates an incorrect one.

²⁶Here, and elsewhere, we write $o(t)$ to mean a lower order term. Formally a function $f: \mathbb{R} \rightarrow \mathbb{R}$ is in $o(t)$ if $\lim_{t \rightarrow \infty} f(t)/t = 0$.

$$\mathbb{P} [a_t^i \neq \alpha^\Theta] = e^{-nr_{aut}t+o(t)}.$$

Having considered this benchmark case, we turn to our model, in which $n \geq 2$ agents observe each others' actions, but signals are private. Our main result is that for any number of agents the speed of learning is bounded from above by a constant:

Theorem 1. *Suppose n agents all observe each others' past actions. Given the private signal distributions, there exists a constant $r_{bnd} > 0$ independent of the number of agents n , such that*

$$\mathbb{P} [a_t^i \neq \alpha^\Theta] \geq e^{-r_{bnd}t+o(t)}.$$

In particular, this holds for $r_{bnd} = \min \{\mathbb{E}_g [\ell], -\mathbb{E}_b [\ell]\}$. When private signals are normal then one can take $r_{bnd} = 4r_{aut}$.

Note that this theorem holds for all fixed signal distributions and all group sizes n , and does not require any assumptions about the relation between them, such as the ones we make in §2.

An immediate corollary from Theorem 1 and Fact 2 is the following result.

Corollary 1. *There exists a fixed group size k such that for any arbitrarily large group size n , the probability that any agent chooses the wrong action is eventually lower with k agents and public signals than with n agents who only observe actions. When signals are normal we can take $k = 4$.*

Thus, adding more agents (and with them more private signals and more information) cannot boost the speed of learning past some bound, and as n tends to infinity more and more of the information is lost: a vanishing fraction of the private signals would produce the the same error probabilities if observed directly.²⁷ In the case of normal signals $r_{bnd} = 4r_{aut}$, and thus, regardless of the number of agents, the probability of mistake is eventually higher than it would be if 4 agents shared their private signals. Thus for large groups almost all of the private signals are effectively lost, i.e. not aggregated in the decisions of others.

4.2.1. *Rational groupthink.* In the proof of this theorem we calculate the asymptotic probability of the event that *all* agents choose the wrong certainty action in *almost all time periods* up to time t . We call this event “*rational groupthink*” and show that its probability is already high, which implies that the probability that one particular agent errs at time t is also high.

²⁷Formally, Theorem 1 establishes that the upper bound on the rate of learning, r_{bnd} , is less than some constant times $1/n$ times the rate nr_{aut} of learning from observing n signals directly every period, i.e. $\frac{r_{bnd}}{nr_{aut}} \leq \frac{c}{n}$, and thus goes to zero for n tending to ∞ .

When a wrong consensus forms by chance in the beginning, it is hard to break and can last for a long time, with surprisingly high probability. To understand why this occurs, we observe that conditioned on a wrong consensus forming, each agent needs a stronger-than-indifferent signal to break the consensus. This is because the private signal needs to overcome what is learned by observing that the other agents have not broken the consensus. As periods progress, conditioned on the consensus not being broken, the required signal threshold rises and rises. Indeed, after a long time, the threshold will be arbitrarily high. As correct signals are, in the long-run, more likely than incorrect signals, it follows that conditioned on being below the threshold, the agents' signals will be close to it, and in particular will indicate the *correct* action. Thus the private signals of each agent, which initially indicated the wrong action, eventually strongly indicate the correct action, but are still ignored due to the overwhelming information provided by the actions of others. This intuition is formalized in the next result, which is based the large deviation principle that states that when an unlikely event occurs, it is very likely to occur in the most likely way.

Define α_t^{\min} to be the lowest action that is taken by any agent with positive probability at time t . In many cases $\alpha_t^{\min} = \alpha^b$; for example, this holds when the private signals are unbounded. For bounded signals this holds for all t large enough, but may not hold for initial t .²⁸

Denote by $\hat{p}_i^t = \mathbb{P}[\Theta = \mathbf{g} \mid s_1^i, \dots, s_t^i]$ the probability assigned to the good state given only agent i 's signals.

Theorem 2. *Condition on the state being good $\Theta = \mathbf{g}$. In the long run, conditional on all agents taking the eventually **incorrect** action α_t^{\min} in every period, the private signals of every agent strongly indicate the **correct** certainty action. That is, for every $\varepsilon > 0$ it holds that*

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\mathbf{g}} [\hat{p}_i^t > 1 - \varepsilon \text{ for all } i \mid a_\tau^j = \alpha_\tau^{\min} \text{ for all } \tau \leq t \text{ and all } j] = 1.$$

The analogous statement holds in the bad state.

Note that Theorem 2 is *not* a consequence of the law of large numbers, as conditional on taking the wrong action the distribution of signals is not independent. Indeed, the result of Theorem 2 does not hold in the single agent case, where—in sharp contrast—conditional on choosing the wrong action the agent holds wrong beliefs. It shows that in a multi-agent learning problem agents will (with high probability) have received correct signals even conditioned on choosing the wrong action. This phenomenon, which does not have an analogue in sequential herding models, seems striking, as it does not involve irrationality, and

²⁸Recall that the action a_t^i is chosen according to the posterior belief p_t^i . By standard arguments, the set of beliefs in which each possible action is taken is an interval. This induces an order on the actions, from the lowest one, which must be α^b , and is taken for the lowest beliefs, to the highest one, α^g , which is taken for the highest beliefs. We discuss this technical issue in detail in §5.3.

yet results in a group taking an action which contradicts each and every member’s private information.

4.2.2. Early Period Mistake Probabilities. Theorem 1 is a statement about asymptotic rates. In fact, if one were to increase the number of agents while holding the private signal distributions fixed, the probability of the agents choosing correctly at any given period $t > 1$ approaches 1. Thus, a more interesting setting is the one studied numerically in §2, and which we analyze formally in this section. In this setting, as we increase the number of agents, we decrease the informativeness of each agent’s signal, while keeping fixed the amount of information available to all agents together.

We consider n agents who each receive normal private signals with fixed conditional means ± 1 and variance n . If such signals were publicly observable they would be informationally equivalent to a single normal signal with variance 1 each period. In this setting, Theorem 1 implies that the speed of learning would be inversely proportional to the size of society, and in particular would tend to zero as n tends to infinity.

To test the robustness of this asymptotic speed of learning result, we perform a detailed analysis of the early periods, showing that, as the number of agents increases, they learn less and less from each other’s actions. Thus, the asymptotic result of Theorem 1, which stated that the agents learn little from each other’s actions in the long run, “kicks in” early on (in fact, already in the *second* period), in the sense that with high probability the agents learn nothing from each other’s actions after the first period.

Theorem 3. *Suppose n agents have a uniform prior, normal private signals with conditional distributions $\mathcal{N}(\pm 1, n)$ and want to match the state, so that $\bar{u}(\theta, a) = \mathbf{1}_{\{a=\theta\}}$. Then, for every t , the probability that all agents in the periods $\{2, 3, \dots, t\}$ choose the action that the majority of the agents chose in period 1 converges to one as n goes to infinity.*

Note that the theorem statement also holds conditioned on the first period majority taking the wrong action, since this event occurs with probability that is bounded away from zero. Thus the private signals of periods $\{2, \dots, t\}$ are with high probability not strong enough to induce a deviation from the first period consensus. Consequently, the actions in these periods are correct only if the action taken by the majority in the first period is correct. This probability is bounded by $\Phi(1) \sim 0.84$ for any n . Of course, this probability can be arbitrarily close to 1/2 if the private signal distributions have a larger variance. The numerical simulations in §2 show that a lot of information is lost even for groups of moderate size such as 40 or 100 agents.

The intuition behind this result is the following: after observing the first round actions, the probability that a particular agent will have a strong enough signal to deviate from the majority opinion (action) is small. Increasing the number of agents yields two opposing

forces: with more agents and weaker signals for each agent, each particular agent is less likely to deviate from the consensus, but because there are more agents, it is more likely that some agent deviates. It follows from properties of the normal distribution that the probability of receiving a strong signal vanishes more quickly than $\frac{1}{n}$. Since the probability that at least one agent breaks the consensus grows linearly in n (by the union bound), the first effect dominates the second for large n . When agents observe that no one has deviated, it further strengthens (if not by much) their belief in the majority opinion, thus again delaying the breaking of the consensus. Of course, when the initial consensus is wrong, eventually it is broken.

5. LEARNING DYNAMICS

In this section we analyze the learning dynamics in detail and explain how we prove the results of §4. We discuss how agents interpret each other’s actions and how they choose their own. The analysis of these learning dynamics is related to questions in random walks and requires the application of large deviations techniques. We provide a self-contained introduction to large deviations in the appendix.

5.1. Preliminaries. As an agent’s expected utility for a given action is linear in her posterior belief p_t^i , the set of beliefs where she takes a given action is an interval. It will be convenient to define the agent’s log-likelihood ratio (LLR)

$$(2) \quad L_t^i := \log \frac{p_t^i}{1 - p_t^i}.$$

We define the *private LLR* R_t^i as the LLR calculated based only on an agent’s private signals. It follows from Bayes’ law that

$$(3) \quad R_t^i := L_0^i + \sum_{\tau=1}^t \ell_\tau^i.$$

As the LLR is a monotone transformation of the agent’s posterior belief, and as a myopic agent’s action is determined by her posterior, the same holds true in terms of LLRs. This can be summarized in the following lemma.

Lemma 1. *There exist disjoint intervals $(\underline{L}(\alpha), \bar{L}(\alpha)) \subset \mathbb{R} \cup \{-\infty, +\infty\}$, one for each action $\alpha \in A$, such that, with probability one, $a_t^i = \alpha$ if and only if $L_t^i \in (\underline{L}(\alpha), \bar{L}(\alpha))$.*

Note that we assumed the prior is generic, which rules out indifference. To characterize the agent’s actions it thus suffices to characterize her LLR. Note, that for the certainty action α^b it holds that $\underline{L}(\alpha^b) = -\infty$, and that analogously $\bar{L}(\alpha^g) = +\infty$.

5.2. **Autarky.** As a benchmark, we first describe the classical autarky setting where a single agent acts by himself. In this section we omit the superscript signifying the agent.

Probability of Mistakes. As a consequence of Lemma 1, the probability that the agent chooses the wrong action in period t when the state equals θ is given by

$$(4) \quad \mathbb{P}_\theta [a_t \neq \alpha^\theta] = \begin{cases} \mathbb{P}_\mathfrak{g} [L_t \leq \underline{L}(\alpha^\mathfrak{g})] & \text{if } \theta = \mathfrak{g} \\ \mathbb{P}_\mathfrak{b} [L_t \geq \bar{L}(\alpha^\mathfrak{b})] & \text{if } \theta = \mathfrak{b} \end{cases}.$$

Hence, to calculate the probability of a mistake one needs to calculate the probability that the LLR is in a given interval. In the single agent case the private signals are all the available information, so $L_t = R_t$. By (3) the LLR is the sum of increments which are i.i.d. conditional on the state, and hence $(L_t)_t$ is a random walk.

The short-run probability that a random walk is within a given interval is hard to calculate and depends very finely on the distribution of its increments.²⁹ As this makes it impossible—even in the single agent case—to obtain any general results on the probability that the agent makes a mistake, we focus on the long-run probability of mistakes, which can be analyzed for general signal structures.³⁰

Beliefs. As R_t is a random walk we can use large deviation theory to estimate the probability that the private LLR R_t deviates from its expectation, conditional on the state. To this end, recall that $\lambda_\theta : \mathbb{R} \rightarrow \mathbb{R}$ is the cumulant generating function of the increments of the LLR in state θ .³¹ Denote its Fenchel conjugate by

$$\lambda_\theta^*(\eta) := \sup_{z \geq 0} \lambda_\theta(z) - \eta \cdot z.$$

Given these definition, we are ready to state the basic classical large deviations estimate that we use in this paper.

²⁹The only exception are a few cases where the distribution of the LLR L_t is known in closed form for every t , such as the normal case. Even in the normal case it seems to us intractable to calculate in closed form the mistake probability in early periods in the multi-agent case.

³⁰The long-run behavior of random walks has been studied in *large deviations theory*, with one of the earliest results due to Cramér (1944), who studied these questions in the context of calculating premiums for insurers. We will use some of the ideas and tools from this theory in our analysis; a self-contained introduction is given in Appendix A for the convenience of the reader.

³¹Defined in §4.1 by $\lambda_\mathfrak{g}(z) = -\log \mathbb{E}_\mathfrak{g} [e^{-z\ell}]$ and $\lambda_\mathfrak{b}(z) = -\log \mathbb{E}_\mathfrak{b} [e^{z\ell}]$.

Lemma 2. For any $\mathbb{E}_{\mathbf{b}}[\ell] < \eta < \mathbb{E}_{\mathbf{g}}[\ell]$ it holds that³²

$$\begin{aligned}\mathbb{P}_{\mathbf{g}}[R_t \leq \eta \cdot t + o(t)] &= e^{-\lambda_{\mathbf{g}}^*(\eta) \cdot t + o(t)} \\ \mathbb{P}_{\mathbf{b}}[R_t \geq \eta \cdot t + o(t)] &= e^{-\lambda_{\mathbf{b}}^*(-\eta) \cdot t + o(t)}.\end{aligned}$$

This Lemma states that the probability that the random walk R_t deviates from its (conditional) expectation is exponentially small, and decays with a rate that can be calculated exactly in terms of $\lambda_{\mathbf{g}}^*$ or $\lambda_{\mathbf{b}}^*$. The proof of Lemma 2 in the Appendix uses the properties of λ_{θ} and λ_{θ}^* to verify that the increments of the LLR process in both states are such that large deviation theory results are applicable. Lemma 2 allows us to calculate the probability of a mistake conditional on each state, immediately implying Fact 1, which states that³³

$$\mathbb{P}[a_t \neq \alpha^{\Theta}] = e^{-r_{aut} \cdot t + o(t)},$$

where $r_{aut} = \lambda_{\mathbf{g}}^*(0) = \lambda_{\mathbf{b}}^*(0)$.

5.3. Many Agents and the Groupthink Effect. In this section we consider $n \geq 2$ agents. Each agent observes a sequence of private signals s_1^i, \dots, s_t^i , and the action taken by the other agents in previous periods $(a_{\tau}^j)_{\tau < t, j \neq i}$. In this setting we prove Theorem 1.

The Probability that All Agents Make a Mistake in Every Period. We define for each t the action α_t^{\min} to be the lowest action (i.e., having the lowest $\bar{L}(\alpha)$) that is taken by any agent with positive probability at time t , and observe that α_t^{\min} is equal to $\alpha^{\mathbf{b}}$ for all t large enough. To bound the probability of mistake, we consider the event G_t that all agents choose the action α_t^{\min} in all time periods up to t :

$$G_t = \{a_{\tau}^i = \alpha_{\tau}^{\min} \text{ for all } \tau \leq t \text{ and all } i\}.$$

To simplify the exposition we assume in the main text that $\alpha_t^{\min} = \alpha^{\mathbf{b}}$.³⁴ Conditioned on $\Theta = \mathbf{g}$, the event G_t is the event that all the agents are, and always have been, in unanimous agreement on the *wrong* action $\alpha^{\mathbf{b}}$. We thus call G_t the *rational groupthink* event. The event G_t implies that all agents made a mistake in period t , conditioned on $\Theta = \mathbf{g}$. Thus calculating

³²Here each $o(t)$ denotes a different function, so that the first line can be alternatively written as follows: For every f_t with $\lim_{t \rightarrow \infty} f_t/t = 0$ there exists a g_t with $\lim_{t \rightarrow \infty} g_t/t = 0$ such that $\mathbb{P}_{\mathbf{g}}[R_t \leq \eta \cdot t + f_t] = e^{-\lambda_{\mathbf{g}}^*(\eta) \cdot t + g_t}$.

³³We note that it is possible to strengthen this result by replacing the lower order $o(t)$ term by $O(\log(t))$ using the Bahadur-Rao exact asymptotics method (see Dembo and Zeitouni (1998, Pages 110-113) for a detailed derivation). However, such precision will provide little additional economic insight while significantly complicating the proofs, and thus we will not pursue it.

³⁴This is the case, for example, if the prior is not too extreme relative to the maximal possible private signal strength, or if the private signals are unbounded. Otherwise, it may be the case that agents never take the wrong certainty action in some initial periods, for example if the prior is extreme and the private signals are weak. In Appendix C we drop this assumption and formally show that all our results also hold in general.

the probability of G_t will provide a lower bound on the probability that a particular agent makes a mistake.

This event can be written as $G_t^1 \cap \dots \cap G_t^n$, where G_t^i is the event that agent i chooses the wrong action α^b in every period $\tau \leq t$. To calculate the probability of G_t , it would of course have been convenient if these n events were independent, conditioned on Θ . However, due to the fact that the agents' actions are strongly intertwined, these events are not independent; given that agent 1 played the action α^b —which is optimal in the bad state—in all previous time periods, agent 2 assigns a higher probability to the bad state and is more likely to also play the same action. This poses a difficulty for the analysis of this model, which is a direct consequence of the fact that the agents' actions are intricately dependent on their higher order beliefs.

Decomposition in Independent Events. Perhaps surprisingly, it turns out that G_t can nevertheless be written as the intersection of conditionally independent events, one for each agent. The event associated with agent i is the event that agent i 's private LLRs R_1^i, \dots, R_t^i stay below a time dependent threshold q_1, \dots, q_t (Lemma 3). This reduces the problem to characterizing the thresholds, which we do after stating this result.

Lemma 3. *There exists a sequence of thresholds $(q_\tau)_\tau$ such that the event G_t equals the event that no agent's private LLR R^i hits the threshold q_τ before period t*

$$G_t = \bigcap_{i=1}^n \{R_\tau^i \leq q_\tau \text{ for all } \tau \leq t\}.$$

Thus, if we denote

$$W_t^i := \{R_\tau^i \leq q_\tau \text{ for all } \tau \leq t\},$$

then we have written $G_t = \bigcap_i W_t^i$ as the intersection of independent events.

The proof of Lemma 3 in Appendix C shows this result recursively. Intuitively, whenever G_{t-1} occurs, all agents took the action α^b up to time $t-1$. By the induction hypothesis this implies that the private LLR of all other agents was below the threshold q_τ in all previous periods. As conditional on each state the private LLR's of different agents are independent, whether agent i takes the action α^b at time t conditional on G_{t-1} , depends only on her private LLR R_t^i . As α^b is the most extreme action it follows that the set of private LLRs where the agent takes the action α^b must be a half-infinite interval and is thus characterized by a threshold q_τ . By symmetry, this is the same threshold for all agents.

Calculating the Thresholds. We now provide a sketch of the argument which we use in the appendix to characterize the threshold q_t . The threshold q_t admits a simple interpretation: it determines how high a private LLR R_t^i an agent must have in order to break from the consensus, and not take action α^b at time t , after having seen everyone take it so far. To

calculate the q_{t+1} 's we consider agent j 's decision problem at time $t + 1$, conditioned on G_t . The information available to her is her own private signals (summarized in her private log-likelihood ratio R_{t+1}^j), and in addition the fact that all other agents have chosen α^b up to this point. But the latter observation is equivalent to knowing that all the other agent's private log-likelihood ratios have been under the thresholds q_τ in all previous time periods. Formally, for agent j to know that G_t has occurred, is equivalent to knowing that

$$W_t^i = \{R_\tau^i \leq q_\tau \text{ for all } \tau \leq t\}$$

has occurred for all agents $i \neq j$.

How does knowing that agent i 's private LLR has been below q_τ in all previous periods (i.e. W_t^i occurred) influence agent j 's posterior? To answer this question we consider the log-likelihood ratio induced by this event, and show that it is asymptotically equal to the logarithm of the probability of the event $R_t^i \leq q_t$, i.e., the event that agent i 's private LLR is below the threshold q_t at just the last period.³⁵

In Lemma 11 in the appendix we show that the threshold q_t is in fact asymptotically linear, i.e. the limit $\beta = \lim_{t \rightarrow \infty} q_t/t$ exists. We argue that $\mathbb{P}_b[W_t^i]$ is bounded away from zero. Combining this with $\log \mathbb{P}_g[W_t^i] \approx \log \mathbb{P}_g[R_t^i \leq q_t]$, the linearity of q , and the large deviations estimate given in Lemma 2 yields³⁶

$$(5) \quad \log \frac{\mathbb{P}_g[W_t^i]}{\mathbb{P}_b[W_t^i]} \approx \log \mathbb{P}_g[W_t^i] \approx \log \mathbb{P}_g[R_t^i \leq q_t] \approx \log \mathbb{P}_g[R_t^i \leq \beta \cdot t] \approx -\lambda_g^*(\beta) \cdot t.$$

Since $G_t = \bigcap_{i=1}^n W_t^i$, and since the events $(W_t^i)_i$ are conditionally independent, we get that when a agent j observes G_t , her likelihood ratio will be the sum of R_t^j and $n - 1$ times the likelihood ratio of W_t^i :

$$(6) \quad L_t^j \approx R_t^j - (n - 1) \cdot \lambda_g^*(\beta) \cdot t.$$

Thus, the threshold for the rational groupthink event at time $t + 1$ will satisfy

$$(7) \quad q_{t+1} \approx \beta \cdot t \approx (n - 1) \cdot \lambda_g^*(\beta) \cdot t.$$

Dividing by t and taking the limit as t tends to infinity yields the following fixed point equation for the slope β of the thresholds $(q_\tau)_\tau$ (Lemma 11)

$$(8) \quad \beta = (n - 1) \cdot \lambda_g^*(\beta).$$

³⁵This result is similar in spirit to the Ballot Theorem of [Bertrand \(1887\)](#), which implies that the probability that a random walk is below a constant threshold in all prior periods approximately equals (up to sub-exponential terms) the probability that the random walk is below this threshold in the last period.

³⁶Throughout the proof sketch we denote by \approx equality up to terms that are of the order $o(t)$.

Note that β depends only on the private signal distributions, through $\lambda_{\mathbf{g}}^*$. Since $\lambda_{\mathbf{g}}^*$ is non-negative and decreasing, this equation will always have a unique solution. We thus have calculated β as the solution of the fixed point equation (8).

This fixed point equation has a simple intuition: if the threshold is too high then it is likely that the others' private LLRs are below it, and so it is likely that they do not break the consensus. Thus, an agent gains little information from observing them agreeing with the consensus, and her threshold for breaking the consensus will be low. This contradicts the initial assumption that the threshold is high. Likewise, if the threshold is too low, then an agent learns a lot by observing the consensus endure, and thus sets a high threshold for breaking it. The fixed point of (8) is the value in which these effects are equal.

Given β , we can use (5) to determine the probability of the event W_t^i that agent i does not break the consensus. Using the facts that the rational groupthink event G_t satisfies $G_t = \bigcap_{i=1}^n W_t^i$ and that the W_t^i 's are conditionally independent, we thus have that

$$(9) \quad \log \mathbb{P}_{\mathbf{g}} [G_t] = \log (\mathbb{P}_{\mathbf{g}} [W_t^i]^n) = n \log \mathbb{P}_{\mathbf{g}} [W_t^i] \approx -\beta \cdot \frac{n}{n-1} \cdot t.$$

Consequently, the rate r_{grp} of the event G_t that all agents take the wrong action in all periods up to time t is

$$(10) \quad r_{grp} = \frac{n}{n-1} \beta.$$

Finally, a convexity argument yields that this rate is bounded by the expected log-likelihood ratio of a single signal: $r_{grp} < \mathbb{E}_{\mathbf{g}} [\ell]$ for any number of agents (Lemma 12). As the rational groupthink event implies that all agents make a mistake, this provides a bound on the speed of learning, conditioned on $\Theta = \mathbf{g}$:

$$\mathbb{P}_{\mathbf{g}} [a_t^i \neq \alpha^{\mathbf{g}}] \geq \mathbb{P}_{\mathbf{g}} [G_t] = e^{-r_{grp} \cdot t + o(t)}.$$

Performing the corresponding calculation when conditioning on the bad state, we have proven Theorem 1, for $r_{bnd} = \min \{ \mathbb{E}_{\mathbf{g}} [\ell], -\mathbb{E}_{\mathbf{b}} [\ell] \}$.

We note that r_{grp} can often be calculated explicitly. For example, for normal private signals a straightforward calculation shows that

$$r_{grp} = 4 \frac{(n - \sqrt{n})^2}{(n-1)^2} r_{aut}.$$

A tedious but straightforward calculation shows that $r_{bnd} = 4r_{aut}$.

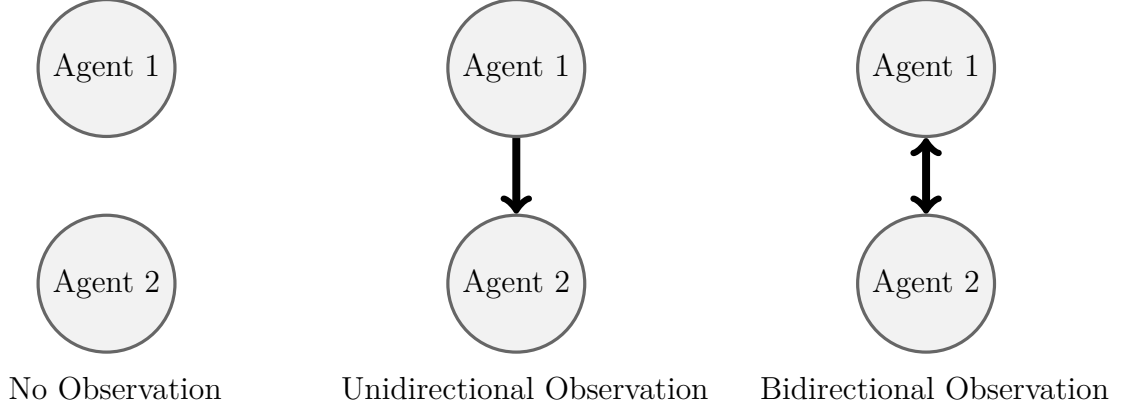


FIGURE 6.1. Different observability structures we analyze in this section. An arrow from agent i to agent j indicates that agent i can observe agent j 's actions.

6. INCOMPLETE OBSERVATION STRUCTURES

So far we have assumed that all agents observe the actions of all others in each period. It is natural to ask how the speed of learning changes when we relax this assumption. We consider two very simple cases, and leave the general case to future work.

First, we consider just two agents. There are three possible observation structures in this case: when neither observes the other, when both observe each other, and when one observes the other, but not vice versa. We have already treated the first two cases, and here we study the speed of learning in the third case. This speed will now depend on the agent. Of course, the agent who observes nothing but her own private signal will learn as in autarky, and so the new result is the speed of learning of the observing agent. Unsurprisingly, we show that the observing agent learns faster than she would in autarky, as she now has additional information in the form of the actions of the other. The less a priori obvious result is that the observing agent learns more quickly than she does under the bidirectional observation structure. Thus, in this case, adding another channel of communication between the agents reduces the speed of learning.

Theorem 4. *Consider 2 agents and two settings of observation structures: either (\leftrightarrow) both observe each others' past actions or (\rightarrow) agent 1 observes agent 2's past actions, but agent 2 does not observe agent 1's past actions. Denote by e_t^{\leftrightarrow} the probability that agent 1's action a_t^1 is not equal to α^\ominus in the first setting, any by e_t^{\rightarrow} the same probability, in the second setting. Then*

$$\frac{e_t^{\leftrightarrow}}{e_t^{\rightarrow}} \geq e^{rt+o(t)}$$

for some $r > 0$ that depends only on the private signal structure.

In Theorem 8 in the Appendix we compute the exact rate at which agent 1 learns in the unidirectional case. This result might be of independent interest. For example, in the case of normal signals it yields that agent 1 learns as fast as she would learn if she observed $\frac{9}{16} \approx 56\%$ of agent 2s private signals, instead of her actions (Corollary 2).

Next, we analyze another simple case: the case of a large group of agents, in which agent 1 can observe the actions of all others, but no other agent can observe any actions. In this case, we show that the speed of learning of agent 1 grows linearly with the number of agents she observes. While this result is rather straightforward to understand and prove (as agent 1 has access to $n - 1$ independent sources of information), it highlights the fact that the loss of information in the full observation setting is not due to the fact that agents observe actions rather than signals, but to the interdependence of these actions.

Theorem 5. *Suppose n agents all observe private signals only, except for agent 1, who additionally observes the others' past actions. Given the private signal distributions, there exists a constant $r > 0$, which depends only on the distribution of the private signals, such that for any number of agents*

$$\mathbb{P} [a_t^1 \neq \alpha^\Theta] \leq e^{-(n-1)r \cdot t + o(t)}.$$

7. NON-BAYESIAN BELIEFS AND OVER-PRECISION BIAS

We next relax the assumption that agents form beliefs using Bayes rule. The bias we consider is over-confidence about the precision of an agent's own signals as compared to the other agents' signals.³⁷ For tractability we focus on the case of normal signals, and, as in §2, each agent's signal is normally distributed with precision $1/n$ and mean $+1$ or -1 depending on the state. While the true precision of each agent's signal is $1/n$, each agent believes that their own signal has precision $c \cdot 1/n$ with $c > 1$ and all the other agents' signals have precision $1/n$. We consider the case where agents are sophisticated, that means, are aware of the over-precision bias of others and understand how other agents pick their actions.

7.1. The Effects of Over-Precision Bias. The over-precision bias of the agents has a direct as well as an indirect effect on the agent's ability to learn the state. The direct effect is straightforward: As agents make a mistake when updating their beliefs they are less likely to chose the correct action. The indirect effect is more subtle: as agents (erroneously) attribute a higher precision to their own signal, they put a higher weight on it when picking their action. As an agent's signals are now more likely to influence his actions, his actions reveal more of his private information. Intuitively, this benefits all other agents and allows

³⁷This bias seems especially relevant in the context of social learning, where it distorts information aggregation. Its importance has been suggested, for example, by Vives (2010), exercises 4.7 and 6.7. See Moore et al. (2015) and the references therein.




FIGURE 7.1. The probability with which an agent takes the wrong action over time in absolute terms for a biased and unbiased agent (on the left) and the error probability in period $t = 30$ for various degrees of bias and $n = 40$ agents (on the right).

them to learn faster. The right graph of Figure 7.1 displays the error probability for various degrees of over-precision in period 30 in an example with 40 agents. Maybe surprisingly, agents are less likely to take the wrong action for intermediate biases (in the range between 1 and 4). This means that the indirect positive effect coming from the fact that other agents reveal more of their private information dominates the direct effect caused by the deviation from Bayes rule, which leads to wrong beliefs and thus sub-optimal actions. The left graph of Figure 7.1 shows how this comparison evolves over time. In the first period, the agent observed only his own signal and the over-precision bias thus has no effect. In the second period there is no positive effect of the over-precision bias as other agents reveal exactly the same information in the first period independent of the bias. The error probability is thus higher—but only slightly—with the bias (21.7% vs 21.8% for $c = 1$ vs $c = 2$). However, already from the third period the indirect benefit is larger than the direct loss, leading to lower error probabilities for the biased agents (18% vs 16% in period 3, and, for example, 13% vs 9% in period 10).

We conjecture that for an appropriate choice of the bias parameter c , the asymptotic error probability is smaller for biased agents than it is for rational agents, as suggested by Figure 7.1. Indeed, a straightforward modification of the proof of Theorem 1 shows that the rate of the groupthink event indeed can increase (implying lower probability of the groupthink event) for biased agents. However, the probability of groupthink provides only a lower bound on the error probability, which we do not know how to explicitly calculate. We thus leave this conjecture for future work.

7.2. A Social Planners Perspective. An interesting question is which strategies a social planner would pick for the agents in order to maximize the long-run probability with which

they pick the right action. The main trade-off faced by such a social planner is that taking a sub-optimal action today increases the mistake probability today, but potentially leads the agent to reveal more information which benefits other agents in the future. While in equilibrium agents do not take this positive informational externality into account, a forward-looking social planner would, and thus potentially has an incentive to intervene with the agents' actions. While solving for the optimal policy is beyond the scope of this paper, the numerical simulations of the previous section already provide some insight into this question: The simulations indicate that agents learn faster when they over-weigh their own signals, which suggests that a social planner could improve welfare by instructing agents to use the *non-Bayesian* biased updating rule described in the previous section. Thus, while biased learning is suboptimal for myopic individuals, it might be socially beneficial.

8. CONCLUSION

We show that rational groupthink occurs in a complex environment of agents who observe each other and take actions repeatedly. As a result, almost all information is lost when the group of agents is large. We use asymptotic rates as a measure of the speed of learning. As a robustness test, we show that the same effect holds also in the early periods, for the case of normal signals.

This article leaves many open questions which could potentially be analyzed using our approach.

- (1) We think that it may be feasible to extend our methods beyond the two state case to an arbitrary, finite number of states. This will require the use of high-dimensional large deviation techniques, as the beliefs are now multi-dimensional.
- (2) The extension to the case in which different agents have different signals seems more straightforward. We conjecture that the methodology we developed could be used in this setting and will lead to similar conclusions.
- (3) What happens when the state changes over time? This setting is potentially very interesting, as one could derive steady-state results instead of asymptotic results. We conjecture that results that are similar in spirit will hold, with large groups not performing significantly better than single agents. A major challenge in the analysis is that, since the probability of taking the wrong action does not vanish over time, large deviation techniques no longer apply. Social learning with a changing state and short-lived agents has been studied by [Moscarini et al. \(1998\)](#), [Frongillo et al. \(2011\)](#) and recently [Dasaratha et al. \(2018\)](#).
- (4) What happens with payoff externalities, for example when agents have incentive to coordinate?

- (5) What is the optimal policy of a forward-looking social planner who cannot transfer information between the agents? It is unclear to us how one could approach this problem.
- (6) Of particular interest is the study of more complex societal structures: how fast do agents learn for a given arbitrary network of observation, which is not the complete network? We briefly tackle some particularly simple examples in §6, but our techniques break down in the general case, as they rely on the fact that the groupthink event is common knowledge.

REFERENCES

- Daron Acemoglu, Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011.
- George-Marios Angeletos and Alessandro Pavan. Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142, 2007.
- Venkatesh Bala and Sanjeev Goyal. Learning from neighbours. *The Review of Economic Studies*, 65(3):595–621, 1998.
- Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. The diffusion of microfinance. *Science*, 341(6144):1236–1238, 2013.
- Abhijit V Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, pages 797–817, 1992.
- Roland Bénabou. Groupthink: Collective delusions in organizations and markets. *Review of Economic Studies*, 80(2):429–462, 2012.
- Joseph Bertrand. Solution d’un problème. *Comptes Rendus de l’Académie des Sciences, Paris*, 105:369, 1887.
- Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, pages 992–1026, 1992.
- Christophe Chamley. *Rational herds: Economic models of social learning*. Cambridge University Press, 2004.
- Timothy G Conley and Christopher R Udry. Learning about a new technology: Pineapple in Ghana. *The American Economic Review*, 100(1):35–69, 2010.
- Thomas M Cover and Joy A Thomas. *Elements of information theory*. John Wiley & Sons, 2006.
- Harald Cramér. On a new limit theorem of the theory of probability. *Uspekhi Mat. Nauk*, 10:166–178, 1944.
- Martin W Cripps, Jeffrey C Ely, George J Mailath, and Larry Samuelson. Common learning. *Econometrica*, 76(4):909–933, 2008.

- Zhi Da and Xing Huang. Harnessing the wisdom of crowds. *Management Science*, 66(5): 1847–1867, 2020.
- Krishna Dasaratha and Kevin He. Speed of rational social learning in networks with Gaussian information, 2019.
- Krishna Dasaratha, Benjamin Golub, and Nir Hak. Social learning in a dynamic environment. 2018.
- Amir Dembo and Ofer Zeitouni. *Large deviations techniques and applications*. Springer, second edition, 1998.
- Darrell Duffie and Gustavo Manso. Information percolation in large markets. *The American Economic Review*, pages 203–209, 2007.
- Darrell Duffie, Semyon Malamud, and Gustavo Manso. Information percolation with equilibrium search dynamics. *Econometrica*, 77(5):1513–1574, 2009.
- Darrell Duffie, Gaston Giroux, and Gustavo Manso. Information percolation. *American Economic Journal: Microeconomics*, pages 100–111, 2010.
- Rick Durrett. *Probability: theory and examples*. Cambridge University Press, 1996.
- Rafael M Frongillo, Grant Schoenebeck, and Omer Tamuz. Social learning in a changing world. In *International Workshop on Internet and Network Economics*, pages 146–157. Springer, 2011.
- Douglas Gale and Shachar Kariv. Bayesian learning in social networks. *Games and Economic Behavior*, 45(2):329–346, 2003.
- John D Geanakoplos and Heraklis M Polemarchakis. We can’t disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982.
- Benjamin Golub and Matthew O Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–49, 2010.
- Jan Hązła, Ali Jadbabaie, Elchanan Mossel, and M. Amin Rahimian. Reasoning in Bayesian opinion exchange networks is PSPACE-hard. In Alina Beygelzimer and Daniel Hsu, editors, *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 1614–1648, Phoenix, USA, 25–28 Jun 2019. PMLR.
- Rodney D. Holder. Hume on Miracles: Bayesian Interpretation, Multiple Testimony, and the Existence of God. *The British Journal for the Philosophy of Science*, 49(1):49–65, 03 1998. ISSN 0007-0882. doi: 10.1093/bjps/49.1.49.
- Han Hong and Matthew Shum. Rates of information aggregation in common value auctions. *Journal of Economic Theory*, 116(1):1–40, 2004.
- Johannes Hörner and Satoru Takahashi. How fast do equilibrium payoff sets converge in repeated games? *Journal of Economic Theory*, 165:332–359, 2016.
- David Hume. Section X: Of miracles. In *An Enquiry Concerning Human Understanding*. A. Millar, London, 1748.

- Jussi Keppo, Lones Smith, and Dmitry Davydov. Optimal electoral timing: Exercise wisely and you may live longer. *The Review of Economic Studies*, 75(2):597–628, 2008.
- Moses Maimonides. *The Guide for the Perplexed. Translated by M. Friedländer*. George Routledge & Sons, London, 1904. Original manuscript published circa 1190.
- Pooya Molavi, Alireza Tahbaz-Salehi, and Ali Jadbabaie. A theory of non-Bayesian social learning. *Econometrica*, 86(2):445–490, 2018.
- Don A Moore, Elizabeth R Tenney, and Uriel Haran. Overprecision in judgment. *The Wiley Blackwell handbook of judgment and decision making*, pages 182–209, 2015.
- Giuseppe Moscarini, Marco Ottaviani, and Lones Smith. Social learning in a changing world. *Economic Theory*, 11(3):657–665, 1998.
- Elchanan Mossel, Allan Sly, and Omer Tamuz. Strategic learning and the topology of social networks. *Econometrica*, 83(5):1755–1794, 2015.
- Manuel Mueller-Frank and Itai Arieli. Multi-dimensional social learning. *The Review of Economic Studies*, forthcoming, 2018.
- Marco Ottaviani and Peter Sørensen. Information aggregation in debate: who should speak first? *Journal of Public Economics*, 81(3):393–421, 2001.
- Rohit Parikh and Paul Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52(1):178–189, 1990.
- Dinah Rosenberg and Nicolas Vieille. On the efficiency of social learning. *Econometrica*, 87(6):2141–2168, 2019.
- Dinah Rosenberg, Eilon Solan, and Nicolas Vieille. Informational externalities and emergence of consensus. *Games and Economic Behavior*, 66(2):979–994, 2009.
- James K Sebenius and John Geanakoplos. Don’t bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424–426, 1983.
- Lones Smith and Peter Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2):371–398, 2000.
- Daniel W Stroock. *Mathematics of probability*, volume 149. American Mathematical Society, 2013.
- Xavier Vives. How fast do rational agents learn? *The Review of Economic Studies*, 60(2):329–347, 1993.
- Xavier Vives. *Information and learning in markets: the impact of market microstructure*. Princeton University Press, 2010.

APPENDIX A. THE CUMULANT GENERATING FUNCTIONS, THEIR FENCHEL
CONJUGATES, AND LARGE DEVIATIONS ESTIMATES

The long-run behavior of random walks has been studied in large deviations theory. In this section we first introduce some well known tools from this literature, which will be crucial to understanding the long-run behavior of agents. In the end of this section we derive a sample path large deviation theorem which will be the main tool in our analysis. The proof of this theorem follows well known techniques (see [Dembo and Zeitouni, 1998](#), Chapter 5).

Large Deviations of Random Walks. Let X_1, X_2, \dots be i.i.d random variables with $\mathbb{E}[X_t] = \mu$ and $Y_t = \sum_{\tau=1}^t X_\tau$ the associated random walk with steps X_t . By the law of large numbers we know that Y_t should approximately equal $\mu \cdot t$. Large deviation theory characterizes the probability that Y_t is much lower, and in particular smaller than $\eta \cdot t$, for some $\eta < \mu$. Under some technical conditions, this probability is exponentially small, with a rate $\lambda^*(\eta)$:

$$\mathbb{P}[Y_t < \eta \cdot t + o(t)] = e^{-\lambda^*(\eta) \cdot t + o(t)},$$

or equivalently stated

$$\lim_{t \rightarrow \infty} -\frac{1}{t} \log \mathbb{P}[Y_t < \eta \cdot t + o(t)] = \lambda^*(\eta).$$

The rate λ^* can be calculated explicitly and is the *Fenchel Conjugate* of the *cumulant generating function* of the increments

$$\lambda^*(\eta) := \sup_{z \geq 0} (-\log \mathbb{E}[e^{-z X_1}] - \eta \cdot z).$$

The first proof of a “large deviation” result of this flavor is due to [Cramér \(1944\)](#), who studied these questions in the context of calculating premiums for insurers. A standard textbook on large deviations theory is [Dembo and Zeitouni \(1998\)](#).

In this section we provide an independent proof of this classical large deviations result, and prove a more specialized one suited to our needs. We consider a very general setting: we make no assumptions on the distribution of each step X_t , and in particular do not need to assume that it has an expectation.

Denoting $X = X_1$, The cumulant generating function λ is (up to sign, as compared to the usual definition) given by

$$\lambda(z) = -\log \mathbb{E}[e^{-z X}].$$

Note that when the right hand side is not finite it can only equal $-\infty$ (and never $+\infty$).

Lemma 4. *λ is finite on an interval I , on which it is concave and on whose interior it is smooth (that is, having continuous derivatives of all orders).*

Proof. Note that I contains 0, since $\lambda(0) = 0$ by definition. Assume $\lambda(a)$ and $\lambda(b)$ are both finite. Then for any $r \in (0, 1)$

$$\lambda(r \cdot a + (1 - r) \cdot b) = -\log \mathbb{E} \left[e^{-(r \cdot a + (1-r) \cdot b) \cdot X} \right] = -\log \mathbb{E} \left[(e^{-a \cdot X})^r \cdot (e^{-b \cdot X})^{1-r} \right],$$

which by Hölder's inequality is at least $r \cdot \lambda(a) + (1 - r) \cdot \lambda(b)$. Hence λ is finite and concave on a convex subset of \mathbb{R} , or an interval. We omit here the technical proof of smoothness; it can be found, for example, in [Stroock \(2013, Theorem 1.4.16\)](#). \square

It also follows that unless the distribution of X is a point mass (which is a trivial case), λ is strictly concave on I . We assume this henceforth. Note that it could be that I is simply the singleton $[0, 0]$. This is not an interesting case, and we will show later that in our setting I is larger than that.

The Fenchel conjugate of λ is given by

$$\lambda^*(\eta) = \sup_{z \geq 0} \lambda(z) - \eta \cdot z.$$

We note a few properties of λ^* . First, since $\lambda(0) = 0$ and $\lambda(z) < \infty$, λ^* is well defined and non-negative (but perhaps equal to infinity for some η). Second, since λ is equal to $-\infty$ whenever it is not finite, the supremum is attained on I , unless it is infinity. Third, since λ is strictly concave on I , $\lambda(z) - \eta \cdot z$ is also strictly concave there, and so the supremum is a maximum and is attained at a single point $z^* \in I$ whenever it is finite. Additionally, since λ is smooth on I , this single point z^* satisfies $\lambda'(z^*) = \eta$ if $z^* > 0$ (equivalently, if $\lambda^*(\eta) > 0$). I.e., if $\lambda'(z^*) = \eta$ for some z^* in the interior of I then

$$(11) \quad \lambda^*(\eta) = \lambda(z^*) - \eta \cdot z^*.$$

Finally, it is immediate from the definition that λ^* is weakly decreasing, and it is likewise easy to see that it is continuous. This, together with (11) and the fact that λ' is decreasing, yields that $\lambda^*(\eta) = \lambda(0) = 0$ whenever $\eta \geq \sup_{z \geq 0} \lambda'(z)$. We summarize this in the following lemma.

Lemma 5. *Let I be the interval on which λ is finite, and let $I^* = \{\eta : \exists z \in \text{int} I \text{ s.t. } \lambda'(z) = \eta\}$. Then*

- (1) λ^* is continuous, non-negative and weakly decreasing. It is positive and strictly decreasing on I^* .
- (2) $\lambda^*(\eta) = 0$ whenever $\eta \geq \sup_{z \geq 0} \lambda'(z)$.
- (3) If $\eta \in I^*$ and $\lambda'(z^*) = \eta$ then $\lambda^*(\eta) = \lambda(z^*) - \eta \cdot z^*$.

Given all this, we are ready to state and prove our first large deviations theorem.

Theorem 6 ([Cramér, 1944](#)). *For every η such that $\eta > \inf_{z \in I} \lambda'(z)$ it holds that*

$$\mathbb{P}[Y_t \leq \eta \cdot t + o(t)] = e^{-\lambda^*(\eta) \cdot t + o(t)}.$$

Proof. For the upper bound, we use a Chernoff bound strategy: for any $z \geq 0$

$$\mathbb{P}[Y_t \leq \eta \cdot t + o(t)] = \mathbb{P}[e^{-zY_t} \geq e^{-z \cdot (\eta \cdot t + o(t))}],$$

and so by Markov's inequality

$$\mathbb{P}[Y_t \leq \eta \cdot t + o(t)] \leq \frac{\mathbb{E}[e^{-zY_t}]}{e^{-z \cdot (\eta \cdot t + o(t))}}.$$

Now, note that $\mathbb{E}[e^{-zY_t}] = e^{-\lambda(z) \cdot t}$, and so

$$\mathbb{P}[Y_t \leq \eta \cdot t + o(t)] \leq e^{-(\lambda(z) - z \cdot \eta) \cdot t + z \cdot o(t)}.$$

Choosing $z \geq 0$ to maximize the coefficient of t yields

$$\mathbb{P}[Y_t \leq \eta \cdot t + o(t)] \leq e^{-\lambda^*(\eta) \cdot t + o(t)},$$

which is the desired upper bound.

We now turn to proving the lower bound. Denote by ν the law of X , and for some fixed z in the interior of I (to be determined later) define the probability measure $\tilde{\nu}$ by

$$\frac{d\tilde{\nu}}{d\nu}(x) = \frac{e^{-zx}}{\mathbb{E}[e^{-zX}]} = e^{\lambda(z) - zx},$$

and let \tilde{X}_t be i.i.d. random variables with law $\tilde{\nu}$. Note that

$$\mathbb{E}[\tilde{X}] = \frac{\mathbb{E}[Xe^{-zX}]}{\mathbb{E}[e^{-zX}]} = \lambda'(z).$$

Now, fix any η_1, η_2 such that $\eta_1 < \eta_2 < \eta$ and $\lambda'(z) = \eta_2$ for some z in the interior of I ; this is possible since $\eta > \inf_{z \in I} \lambda'(z)$. This is the z we choose to take in the definition of $\tilde{\nu}$. If we think of η_2 as being close to η then the expectation of \tilde{X} , which is equal to η_2 , is close to η . We have thus “tilted” the random variable X , which had expectation μ , to a new random variable with expectation close to η .

We can bound

$$\mathbb{P}[Y_t \leq \eta \cdot t + o(t)] \geq \mathbb{P}[\eta_1 \cdot t \leq Y_t \leq \eta \cdot t + o(t)] = \int_{\eta_1 t}^{\eta t + o(t)} 1 \, d\nu^{(t)},$$

where $\nu^{(t)}$ is the t -fold convolution of ν with itself, and hence the law of Y_t . It is easy to verify³⁸ that $d\nu^{(t)}(y) = e^{zy - \lambda(z) \cdot t} d\tilde{\nu}^{(t)}(y)$, and so

$$\int_{\eta_1 t}^{\eta t + o(t)} 1 d\nu^{(t)} = e^{-\lambda(z) \cdot t} \int_{\eta_1 t}^{\eta t + o(t)} e^{zy} d\tilde{\nu}^{(t)}(y),$$

which we can bound by taking the integrand out of the integral and replacing y with the lower integration limit:

$$e^{-\lambda(z) \cdot t} \int_{\eta_1 t}^{\eta t + o(t)} e^{zy} d\tilde{\nu}^{(t)}(y) \geq e^{(\eta_1 z - \lambda(z)) \cdot t} \int_{\eta_1 t}^{\eta t + o(t)} 1 d\tilde{\nu}^{(t)}.$$

Since the law of $\tilde{Y}_t = \sum_{\tau=1}^t \tilde{X}_\tau$ is $\tilde{\nu}^{(t)}$, this implies that

$$e^{(\eta_1 z - \lambda(z)) \cdot t} \int_{\eta_1 t}^{\eta t + o(t)} 1 d\tilde{\nu}^{(t)} = e^{(\eta_1 z - \lambda(z)) \cdot t} \mathbb{P} \left[\eta_1 \cdot t \leq \tilde{Y}_t \leq \eta \cdot t + o(t) \right].$$

Since $\eta_1 < \mathbb{E}[\tilde{X}] < \eta$ we have that $\lim_{t \rightarrow \infty} \mathbb{P} \left[\eta_1 \cdot t \leq \tilde{Y}_t \leq \eta \cdot t + o(t) \right] = 1$, by the law of large numbers. Hence

$$(12) \quad \liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} [Y_t \leq \eta \cdot t + o(t)] \geq \eta_1 z - \lambda(z)$$

which, by (11), and recalling that $z = (\lambda')^{-1}(\eta_2)$, can be written as

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} [Y_t \leq \eta \cdot t + o(t)] \geq -\lambda^*(\eta_2) - (\eta_2 - \eta_1) \cdot (\lambda')^{-1}(\eta_2).$$

Taking the limit as η_1 approaches η_2 yields

$$(13) \quad \liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} [Y_t \leq \eta \cdot t + o(t)] \geq -\lambda^*(\eta_2).$$

We now consider two cases. First, assume that $\eta \leq \sup_{z \geq 0} \lambda'(z)$. In this case we can choose η_2 arbitrarily close to η , and by the continuity of λ^* we get that

$$\liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P} [Y_t \leq \eta \cdot t + o(t)] \geq -\lambda^*(\eta),$$

or equivalently

$$\mathbb{P} [Y_t \leq \eta \cdot t + o(t)] \geq e^{-\lambda^*(\eta) \cdot t + o(t)}.$$

The second case is that $\eta > \sup_{z \geq 0} \lambda'(z)$. In this case $\lambda^*(\eta) = 0$ (Lemma 5). Also, (13) holds for any $\eta_2 < \sup_z \lambda'(z)$ and thus it holds for $\eta_2 = \sup_{z \geq 0} \lambda'(z)$. But then $\lambda^*(\eta_2) = 0 = \lambda^*(\eta)$, and so we again arrive at the same conclusion. \square

A.1. Sample Path Large Deviation Bounds. In this section we prove a large deviation result that is similar in spirit, and in some sense is stronger than Theorem 6, as it shows that

³⁸See, e.g., Durrett (1996, Page 74) or note that the Radon-Nikodym derivative between the law of X and \tilde{X} is $e^{zx - \lambda(z)}$, and so the derivative between the laws of (X_1, \dots, X_t) and $(\tilde{X}_1, \dots, \tilde{X}_t)$ is $e^{z(x_1 + \dots + x_t) - \lambda(z) \cdot t}$.

the same rate applies to the event that the sum is below the threshold at all time periods prior to t , rather than just at period t . It furthermore does not require the threshold to be linear, but only asymptotically and from one direction; both of these generalizations are important. This theorem is similar in spirit to other sample path large deviation results (see, e.g., [Dembo and Zeitouni, 1998](#), Chapter 5).

Theorem 7. *For every η such that $\eta > \inf_{z \in I} \lambda'(z)$, and every sequence $(y_t)_{t \in \mathbb{N}}$ with $\liminf_{t \rightarrow \infty} y_t/t = \eta$ and $\mathbb{P}[Y_t \leq y_t] > 0$ it holds that*

$$\mathbb{P} \left[\bigcap_{\tau=1}^t \{Y_\tau \leq y_\tau\} \right] = e^{-\lambda^*(\eta) \cdot t + o(t)}.$$

Proof. Let E_t be the event $\bigcap_{\tau=1}^t \{Y_\tau \leq y_\tau\}$. Let (t_k) be a sequence such that $\lim_{k \rightarrow \infty} y_{t_k}/t_k = \eta$. For every t let t' be the largest t_k with $t_k \leq t$. Then by inclusion we have that

$$\frac{1}{t} \log \mathbb{P}[E_t] \leq \frac{1}{t'} \log \mathbb{P}[Y_{t'} \leq y_{t'}].$$

Using the same Chernoff bound strategy of the proof of [Theorem 6](#), we get that

$$\frac{1}{t} \log \mathbb{P}[E_t] \leq -\lambda^*(y_{t'}/t').$$

The continuity of λ implies that taking the limit superior of both sides yields

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}[E_t] \leq -\lambda^*(\eta),$$

or

$$\mathbb{P}[E_t] \leq e^{-\lambda^*(\eta) \cdot t + o(t)}.$$

To show the other direction, define (as in the proof of [Theorem 6](#)) \tilde{X}_t to be i.i.d. random variables with law $\tilde{\nu}$ given by

$$\frac{d\tilde{\nu}}{d\nu}(x) = e^{\lambda(z) - zx},$$

where ν is the law of X , and $z \in I$ is chosen so that $\lambda'(z) = \eta_2$ for some $\eta_1 < \eta_2 < \eta$, so that the expectation of \tilde{X}_t is η_2 . It follows from inclusion that

$$\mathbb{P}[E_t] \geq \mathbb{P}[E_t \cap \{Y_t \geq \eta_1 \cdot t\}].$$

Now, the Radon-Nikodym derivative between the laws of (X_1, \dots, X_t) and $(\tilde{X}_1, \dots, \tilde{X}_t)$ is $e^{z(x_1 + \dots + x_t) - \lambda(z) \cdot t}$. Hence

$$\mathbb{P}[E_t] \geq \mathbb{E}[1_{E_t} \cdot 1_{Y_t \geq \eta_1 \cdot t}] = \mathbb{E}\left[1_{\tilde{E}_t} \cdot 1_{\tilde{Y}_t \geq \eta_1 \cdot t} \cdot e^{z\tilde{Y}_t - \lambda(z) \cdot t}\right],$$

where \tilde{E}_t is the event $\bigcap_{\tau=1}^t \{\tilde{Y}_\tau \leq y_\tau\}$. We can bound this expression by taking $e^{z\tilde{Y}_t - \lambda(z) \cdot t}$ out of the integral and replacing it with the lower bound $\eta_1 \cdot t$. This yields

$$(14) \quad \mathbb{P}[E_t] \geq e^{z(\eta_1 \cdot t) - \lambda(z) \cdot t} \cdot \mathbb{P}\left[\tilde{E}_t \cap \left\{\tilde{Y}_t \geq \eta_1 \cdot t\right\}\right].$$

Since the expectation of \tilde{Y}_t/t is strictly higher than η_1 , we have that $\lim_{t \rightarrow \infty} \mathbb{P}[\tilde{Y}_t \geq \eta_1 \cdot t] = 1$ by the weak law of large numbers. We claim that $\lim_{t \rightarrow \infty} \mathbb{P}[\tilde{E}_t] > 0$, and show this below. Thus $\lim_t \mathbb{P}[\tilde{E}_t \cap \{\tilde{Y}_t \geq \eta_1 \cdot t\}] > 0$. We can therefore deduce from (14) that

$$\begin{aligned} \liminf_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}[E_t] &\geq z \cdot \eta_1 - \lambda(z) + \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}[\tilde{E}_t \cap \{\tilde{Y}_t \geq \eta_1 \cdot t\}] \\ &\geq z \cdot \eta_1 - \lambda(z). \end{aligned}$$

Proceeding as in the proof of Theorem 6 following equation (12) yields that

$$\mathbb{P}[E_t] \geq e^{-\lambda^*(\eta) \cdot t + o(t)},$$

which is what we set out to prove.

It thus remains to be shown that $\lim_{t \rightarrow \infty} \mathbb{P}[\tilde{E}_t] > 0$. Recall that $\tilde{E}_t = \cap_{\tau=1}^t \{\tilde{Y}_\tau \leq y_\tau\}$ is the event that \tilde{Y}_t is under the threshold y_t up to time t . Since $\tilde{E}_{t+1} \subseteq \tilde{E}_t$, we need to show that the probability of $\tilde{E} := \cap_{t=1}^\infty \tilde{E}_t$ is positive. This is the event that \tilde{Y}_t is always under the threshold y_t .

Denote $\tilde{F}_t = \cap_{\tau=t+1}^\infty \{\tilde{Y}_\tau \leq y_\tau\}$. This is the event that Y_t is under the threshold y_t after time t . Thus, for any t , $\tilde{E} = \tilde{E}_t \cap \tilde{F}_t$. Recall also that $\liminf_t y_t/t = \eta$, and that $\mathbb{E}[\tilde{Y}_t/t] = \eta_1 < \eta$. It thus follows from the strong law of large numbers that $Y_t \leq y_t$ eventually: $\lim_t \mathbb{P}[\tilde{F}_t] = 1$. In particular there is some t_0 such that $\mathbb{P}[\tilde{F}_{t_0}] > 0$.

An hypothesis of this theorem is that $\mathbb{P}[\tilde{Y}_t \leq y_t] > 0$ for all t . As these events are all positively correlated, it is easy to show that the intersection of any finite number of them has positive probability, and in particular that $\mathbb{P}[\tilde{E}_{t_0}] > 0$. By the same reasoning \tilde{E}_{t_0} is positively correlated with \tilde{F}_{t_0} , and therefore

$$\lim_{t \rightarrow \infty} \mathbb{P}[\tilde{E}_t] = \mathbb{P}[\tilde{E}] = \mathbb{P}[\tilde{E}_{t_0} \cap \tilde{F}_{t_0}] \geq \mathbb{P}[\tilde{E}_{t_0}] \cdot \mathbb{P}[\tilde{F}_{t_0}] > 0.$$

□

APPENDIX B. APPLICATION OF LARGE DEVIATION ESTIMATES

In this section we prove a number of claims regarding the functions λ_θ and λ_θ^* . Recall that for $\theta \in \{h, l\}$

$$\lambda_{\mathfrak{g}}(z) := -\log \mathbb{E}_{\mathfrak{g}}[e^{-z\ell}] \quad \lambda_{\mathfrak{b}}(z) := -\log \mathbb{E}_{\mathfrak{b}}[e^{z\ell}],$$

where ℓ is a random variable with the same law as any ℓ_t^i , and

$$\lambda_\theta^*(\eta) = \max_z \lambda_\theta(z) - \eta \cdot z.$$

We first note that by the definition of λ_θ we have that

$$(15) \quad \lambda_{\mathfrak{g}}(z) = -\log \int \exp \left(-z \cdot \log \frac{d\mu_{\mathfrak{g}}}{d\mu_{\mathfrak{b}}}(s) \right) d\mu_{\mathfrak{g}}(s) = -\log \int \left(\frac{d\mu_{\mathfrak{b}}}{d\mu_{\mathfrak{g}}}(s) \right)^z d\mu_{\mathfrak{g}}(s).$$

It follows immediately that there is a simple connection between $\lambda_{\mathfrak{g}}$ and $\lambda_{\mathfrak{b}}$

$$\lambda_{\mathfrak{b}}(z) = \lambda_{\mathfrak{g}}(1 - z).$$

Furthermore, as for every η between $\mathbb{E}_{\mathfrak{g}}[\ell]$ and $\mathbb{E}_{\mathfrak{b}}[\ell]$ the maximum in the definition of $\lambda_{\mathfrak{g}}^*$ is achieved for some $z \in (0, 1)$, it follows that there is also a simple connection between $\lambda_{\mathfrak{g}}^*$ and $\lambda_{\mathfrak{b}}^*$:

$$(16) \quad \lambda_{\mathfrak{b}}^*(\eta) = \lambda_{\mathfrak{g}}^*(-\eta) - \eta.$$

We will accordingly state some results in terms of $\lambda_{\mathfrak{g}}$ and $\lambda_{\mathfrak{g}}^*$ only. It also follows from (15) that the interval I on which $\lambda_{\mathfrak{g}}$ is finite contains $[0, 1]$. Since from the definitions we have that $\lambda'_{\mathfrak{g}}(0) = \mathbb{E}_{\mathfrak{g}}[\ell]$, and since $\lambda'_{\mathfrak{g}}(1) = \mathbb{E}_{\mathfrak{b}}[\ell]$ by the relation between $\lambda_{\mathfrak{g}}$ and $\lambda_{\mathfrak{b}}$, we have shown the following lemma.

Lemma 6. $\lambda_{\theta}(z)$ and $\lambda_{\theta}^*(\eta)$ are finite for all $z \in [0, 1]$ and $\eta \in (\mathbb{E}_{\mathfrak{b}}[\ell], \mathbb{E}_{\mathfrak{g}}[\ell])$. Furthermore,

$$(17) \quad \lambda_{\theta}(z) = \lambda_{\mathfrak{b}}(1 - z) \text{ and } \lambda_{\theta}^*(\eta) = \lambda_{\mathfrak{b}}^*(-\eta) - \eta.$$

Proof of Lemma 2. Given Lemma 6, Lemma 2 is an immediate corollary of Theorem 6. \square

The following simple observation will be useful on several occasions:

Lemma 7. Let $r_{aut} = \lambda_{\mathfrak{g}}^*(0)$. Then $r_{aut} = \max_{z \in (0, 1)} \lambda_{\mathfrak{g}}(z) = \max_{z \in (0, 1)} \lambda_{\mathfrak{b}}(z) = \lambda_{\mathfrak{b}}^*(0)$, $r_{aut} < \min \{ \mathbb{E}_{\mathfrak{g}}[\ell], -\mathbb{E}_{\mathfrak{b}}[\ell] \}$, and $\min \{ \lambda_{\mathfrak{g}}^*(r_{aut}), \lambda_{\mathfrak{b}}^*(r_{aut}) \} > 0$.

Proof. That $r_{aut} = \max_{z \in (0, 1)} \lambda_{\mathfrak{g}}(z) = \max_{z \in (0, 1)} \lambda_{\mathfrak{b}}(z) = \lambda_{\mathfrak{b}}^*(0)$ follows immediately from the definitions. Now, note that $\mathbb{E}_{\mathfrak{g}}[\ell] = \lambda'_{\mathfrak{g}}(0)$. Thus $r_{aut} < \mathbb{E}_{\mathfrak{g}}[\ell]$ is a simple consequence of the fact that $r_{aut} = \lambda_{\mathfrak{g}}^*(0) = \max_{z \geq 0} \lambda(z)$, that this maximum is obtained in $(0, 1)$, and that $\lambda_{\mathfrak{g}}$ is strictly concave. It follows from the same considerations that $r_{aut} < -\mathbb{E}_{\mathfrak{b}}[\ell]$. Finally, by Lemma 5, $\lambda_{\mathfrak{g}}^*(r_{aut}) > 0$ as $\lambda'_{\mathfrak{g}}(0) < r_{aut} < \lambda'_{\mathfrak{g}}(1)$. The same arguments show that $r_{aut} < -\mathbb{E}_{\mathfrak{b}}[\ell]$ and $\lambda_{\mathfrak{b}}^*(r_{aut}) > 0$. \square

Proof of Fact 1. Consider the case $\Theta = \mathfrak{g}$. As shown in Lemma 1 the probability that the agent makes a mistake is equal to the probability that the LLR is below $\underline{L}(\alpha^{\mathfrak{g}})$. Thus, Lemma 2 allows us to characterize this probability explicitly:

$$\mathbb{P}_{\mathfrak{g}} [a_t^i \neq \alpha^{\theta}] = \mathbb{P}_{\mathfrak{g}} [R_t^i \leq \underline{L}(\alpha^{\mathfrak{g}})] = \mathbb{P}_{\mathfrak{g}} [R_t^i \leq o(t)] = e^{-\lambda_{\mathfrak{g}}^*(0) \cdot t + o(t)}.$$

An analogous argument yields that $\mathbb{P}_{\mathfrak{b}} [a_t^i \neq \alpha^{\theta}] = e^{-\lambda_{\mathfrak{b}}^*(0) \cdot t + o(t)}$. By (17) $\lambda_{\mathfrak{g}}^*(0) = \lambda_{\mathfrak{b}}^*(0)$. \square

APPENDIX C. MANY AGENTS

Recall that we define for each t the action α_t^{\min} to be the lowest action (i.e., having the lowest $\bar{L}(\alpha)$) that is taken by any agent with positive probability at time t , and observe that α_t^{\min} is equal to α^b for all t large enough. We define

$$G_t = \bigcap_{i=1}^n \bigcap_{\tau=1}^t \{a_\tau^i = \alpha_\tau^{\min}\}.$$

Proof of Lemma 3. Note first, that each agent chooses action α_1^{\min} in the first period if the likelihood ratio she infers from her first private signal is at most $\bar{L}(\alpha_1^{\min})$. Hence

$$G_1 = \bigcap_{1 \leq i \leq n} \{a_1^i = \alpha_1^{\min}\} = \bigcap_{1 \leq i \leq n} \{R_1^i \leq \bar{L}(\alpha_1^{\min})\}.$$

Thus the claim holds for $t = 1$. Assume now that all agents choose the action α_τ^{\min} up to period $t - 1$; that is, that G_{t-1} has occurred, which is a necessary condition for G_t . What would cause any one of them to again choose α_t^{\min} at period t ? It is easy to see that there will be some threshold q_t^i such that, given G_{t-1} , agent i will choose α_t^{\min} if and only if her private likelihood ratio R_t^i is lower than q_t^i . By the symmetry of the equilibrium, q_t^i is independent of i , and so we will simply write it as q_t . It follows that

$$G_t = G_{t-1} \cap \bigcap_{1 \leq i \leq n} \{R_t^i \leq q_t\}.$$

Therefore, by induction, and if we denote $q_1 = \bar{L}(\alpha_\tau^{\min})$, we have that

$$G_t = \bigcap_{\substack{\tau \leq t \\ 1 \leq i \leq n}} \{R_\tau^i \leq q_\tau\}. \quad \square$$

Lemma 8. *The threshold q_t is characterized by the recursive relation*

$$(18) \quad q_t = \bar{L}(\alpha_t^{\min}) - (n-1) \cdot \log \frac{\mathbb{P}_g [W_{t-1}^1]}{\mathbb{P}_b [W_{t-1}^1]} \quad \text{and} \quad W_t^i = \bigcap_{1 \leq \tau \leq t} \{R_\tau^i \leq q_\tau\}.$$

Proof. Agent 1's log-likelihood ratio conditional on $\bigcap_{i=1}^n W_{t-1}^i$ at time t equals b

$$L_t^1 = R_t^1 + \log \frac{\mathbb{P}_g [\bigcap_{i=2}^n W_{t-1}^i]}{\mathbb{P}_b [\bigcap_{i=2}^n W_{t-1}^i]}.$$

Since the W_{t-1}^i 's are conditionally independent, we have that

$$L_t^1 = R_t^1 + \sum_{i=2}^n \log \frac{\mathbb{P}_g [W_{t-1}^i]}{\mathbb{P}_b [W_{t-1}^i]}.$$

Finally, by symmetry, all the numbers in the sum are equal, and

$$L_t^1 = R_t^1 + (n-1) \cdot \log \frac{\mathbb{P}_{\mathfrak{g}} [W_{t-1}^1]}{\mathbb{P}_{\mathfrak{b}} [W_{t-1}^1]}.$$

Now, the last addend is just a number. Therefore, if we denote

$$(19) \quad q_t = \bar{L}(\alpha_t^{\min}) - (n-1) \cdot \log \frac{\mathbb{P}_{\mathfrak{g}} [W_{t-1}^1]}{\mathbb{P}_{\mathfrak{b}} [W_{t-1}^1]},$$

then

$$L_t^1 = R_t^1 - q_t + \bar{L}(\alpha_t^{\min}),$$

and $L_t^1 \leq \bar{L}(\alpha_t^{\min})$ (and thus $a_t^1 = \alpha_t^{\min}$) whenever $R_t^1 \leq q_t$. \square

Lemma 9. $q_t \geq \bar{L}(\alpha_t^{\min})$ for all t .

Proof. Let $F_{\mathfrak{g}}$ and $F_{\mathfrak{b}}$ be the cumulative distribution functions of a private log-likelihood ratio ℓ , conditioned on $\Theta = \mathfrak{g}$ and $\Theta = \mathfrak{b}$, respectively. Then it is easy to see that $F_{\mathfrak{g}}$ stochastically dominates $F_{\mathfrak{b}}$, in the sense that $F_{\mathfrak{b}}(x) \geq F_{\mathfrak{g}}(x)$ for all $x \in \mathbb{R}$.³⁹ It follows that the joint distribution of $(R_{\tau}^i)_{\tau \leq t}$ conditioned on $\Theta = \mathfrak{g}$ dominates the same distribution conditioned on $\Theta = \mathfrak{b}$, and so $\mathbb{P}_{\mathfrak{g}} [W_t^1] \leq \mathbb{P}_{\mathfrak{b}} [W_t^1]$. Hence $q_t \geq \bar{L}(\alpha_t^{\min})$. \square

Lemma 10. There is a constant $C > 0$ such that $\mathbb{P}_{\mathfrak{b}} [W_t^1] \geq C$ for all t .

Proof. Since the events W_t^1 are decreasing, i.e. $W_t^1 \subseteq W_{t-1}^1$, we will prove the lemma by showing that

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\mathfrak{b}} [W_t^1] > 0,$$

which by definition is equivalent to

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\mathfrak{b}} [\cap_{\tau \leq t} \{R_{\tau}^i \leq q_{\tau}\}] > 0.$$

Since $q_t \geq \bar{L}(\alpha_t^{\min})$, it suffices to prove that

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\mathfrak{b}} [\cap_{\tau \leq t} \{R_{\tau}^i \leq \bar{L}(\alpha_{\tau}^{\min})\}] > 0.$$

To prove the above, note that agents eventually learn Θ , since the private signals are informative. Therefore, conditioned on $\Theta = \mathfrak{b}$, the limit of R_t^i as t tends to infinity must be $-\infty$. Thus, with probability 1, for all t large enough it does hold that $R_t^1 \leq \bar{L}(\alpha_t^{\min})$. Since each of the events W_t^1 has positive probability, and by the Markov property of the random walk R_t^1 , it follows that the event $\cap_{\tau} \{R_{\tau}^i \leq \bar{L}(\alpha_{\tau}^{\min})\}$ has positive probability. \square

³⁹To see this observe that for any non-decreasing function h , we have that $\mathbb{E}_{\mathfrak{g}}[h(\ell)] = \mathbb{E}_{\mathfrak{b}}[h(\ell) \frac{d\mu_{\mathfrak{g}}}{d\mu_{\mathfrak{b}}}(\ell)] = \mathbb{E}_{\mathfrak{b}}[h(\ell)e^{\ell}] \geq \mathbb{E}_{\mathfrak{b}}[h(\ell)] \cdot \mathbb{E}_{\mathfrak{b}}[e^{\ell}] = \mathbb{E}_{\mathfrak{b}}[h(\ell)] \cdot \mathbb{E}_{\mathfrak{b}}[\frac{d\mu_{\mathfrak{g}}}{d\mu_{\mathfrak{b}}}(\ell)] = \mathbb{E}_{\mathfrak{b}}[h(\ell)]$, where the inequality follows from Chebyshev's sum inequality.

Lemma 11. *The limit $\beta = \lim_{t \rightarrow \infty} q_t/t$ exists, and*

$$\beta = (n - 1)\lambda_{\mathfrak{g}}^*(\beta).$$

Proof. It follows immediately from Lemma 10 and Lemma 8 that

$$(20) \quad \beta := \lim_{t \rightarrow \infty} \frac{q_t}{t} = -(n - 1) \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}_{\mathfrak{g}} [W_{t-1}^1],$$

provided that the limit exists. To show that this limit exists and to calculate it, let $\underline{\beta} = \liminf_{t \rightarrow \infty} q_t/t$. Since $W_t^i = \cap_{\tau=1}^t \{R_{\tau}^i \leq q_{\tau}\}$, it follows from Theorem 7 that

$$\lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{P}_{\mathfrak{g}} [W_t^i] = -\lambda_{\mathfrak{g}}^*(\underline{\beta}),$$

provided that $\underline{\beta} > \inf_z \lambda_{\mathfrak{g}}'(z)$. But $\underline{\beta} \geq 0$ (Lemma 9), and so this indeed holds. The claim now follows from (20). \square

Lemma 12. *For any number of agents n it holds that $r_{grp} < \mathbb{E}_{\mathfrak{g}}[\ell]$.*

Proof. Recall that $\lambda_{\mathfrak{g}}^*$ is strictly convex, and that $\lambda_{\mathfrak{g}}^*(\mathbb{E}_{\mathfrak{g}}[\ell]) = 0$. Hence

$$\begin{aligned} \lambda_{\mathfrak{g}}^*(\beta) &< \frac{\beta}{\mathbb{E}_{\mathfrak{g}}[\ell]} \lambda_{\mathfrak{g}}^*(\mathbb{E}_{\mathfrak{g}}[\ell]) + \frac{\mathbb{E}_{\mathfrak{g}}[\ell] - \beta}{\mathbb{E}_{\mathfrak{g}}[\ell]} \lambda_{\mathfrak{g}}^*(0) \\ &= \frac{\mathbb{E}_{\mathfrak{g}}[\ell] - \beta}{\mathbb{E}_{\mathfrak{g}}[\ell]} \lambda_{\mathfrak{g}}^*(0). \end{aligned}$$

Substituting $\beta/(n - 1)$ for $\lambda_{\mathfrak{g}}^*(\beta)$ (which we can do by Lemma 11) yields

$$\frac{1}{n - 1} \beta < \frac{\mathbb{E}_{\mathfrak{g}}[\ell] - \beta}{\mathbb{E}_{\mathfrak{g}}[\ell]} \lambda_{\mathfrak{g}}^*(0).$$

Since $\lambda_{\mathfrak{g}}^*(0) < \mathbb{E}_{\mathfrak{g}}[\ell]$ (Lemma 7),

$$\frac{1}{n - 1} \beta < \mathbb{E}_{\mathfrak{g}}[\ell] - \beta,$$

or

$$\frac{n}{n - 1} \beta < \mathbb{E}_{\mathfrak{g}}[\ell].$$

Since, by (10), $r_{grp} = \frac{n}{n-1}\beta$, our proof is complete. \square

We now turn to proving Theorem 2, which states that conditioned on rational groupthink—that is, conditioned on the event G_t —all agents have, with high probability, a private LLR R_t^i that strongly indicates the correct action. In fact, we prove a stronger statement, which implies Theorem 2: the private LLR is arbitrarily close to $\beta \cdot t$, the asymptotic threshold for R_t^i above which rational groupthink ends.

Proof of Theorem 2. We prove the theorem by showing a stronger statement. Namely, that for every $\epsilon > 0$ it holds that

$$\lim_{t \rightarrow \infty} \mathbb{P}_{\mathfrak{g}} [R_t^i > t \cdot (\beta - \epsilon) \text{ for all } i \mid G_t] = 1,$$

where, as above, β is the solution to $\beta = (n - 1)\lambda_{\mathfrak{g}}^*(\beta)$.

By Theorem 6 we know that

$$\lim_{t \rightarrow \infty} -\frac{1}{t} \log \mathbb{P}_{\mathfrak{g}} [R_t^i \leq t \cdot (\beta - \epsilon)] = \lambda_{\mathfrak{g}}^*(\beta - \epsilon).$$

Since $\lambda_{\mathfrak{g}}^*(\beta - \epsilon) > \lambda_{\mathfrak{g}}^*(\beta)$ it follows that

$$\lim_{t \rightarrow \infty} -\frac{1}{t} \log \mathbb{P}_{\mathfrak{g}} [A_t] = n \cdot \lambda_{\mathfrak{g}}^*(\beta - \epsilon) > n \cdot \lambda_{\mathfrak{g}}^*(\beta),$$

where A_t is the event $\{R_t^i \leq t \cdot (\beta - \epsilon) \text{ for all } i\}$. Since for t high enough the event A_t is included in G_t , and since, by Lemma 11,

$$\lim_{t \rightarrow \infty} -\frac{1}{t} \log \mathbb{P}_{\mathfrak{g}} [G_t] = n \cdot \lambda_{\mathfrak{g}}^*(\beta),$$

it follows that $\mathbb{P}_{\mathfrak{g}} [A_t \mid G_t]$ decays exponentially with t . Hence $\mathbb{P}_{\mathfrak{g}} [A_t^c \mid G_t] \rightarrow_t 1$, which is the claim we set to prove. \square

APPENDIX D. EARLY PERIOD MISTAKE PROBABILITIES

We now prove Theorem 3. We assume that each agent i observes a normal signal $s_t^i \sim \mathcal{N}(m_{\Theta}, n)$ with mean

$$m_{\Theta} = \begin{cases} +1 & \text{if } \Theta = \mathfrak{g} \\ -1 & \text{if } \Theta = \mathfrak{b} \end{cases}$$

and variance n . Note, that for any number of agents the precision of the joint signal equals 1, and thus the total information the group receives every period is fixed, independent of n .

We assume that the prior belief assigns probability one-half to each state $p_0 = 1/2$ and that there are two actions $A = \{\mathfrak{b}, \mathfrak{g}\}$ and each agent wants to match the state, as in the “matching the state” example (§3.3.1). As in the first period each agent bases her decision only on her own private signal, she takes the action \mathfrak{g} whenever her signal s_1^i is greater than 0 and the action \mathfrak{b} otherwise:

$$a_1^i = \begin{cases} h & s_1^i > 0 \\ l & s_1^i \leq 0 \end{cases}.$$

The private likelihood of each agent after observing the first t signals is given by

$$R_t^i = \log \frac{\prod_{\tau=1}^t \exp\left(-\frac{(s_\tau^i - 1)^2}{2n}\right)}{\prod_{\tau=1}^t \exp\left(-\frac{(s_\tau^i + 1)^2}{2n}\right)} = \frac{2}{n} \sum_{\tau=1}^t s_\tau^i.$$

The probability that an agent takes the correct action Θ in period 1 (conditional only on her own first period signal) is thus given by

$$\mathbb{P}_{\mathbf{g}}[\Theta = a_1^i] = \mathbb{P}_{\mathbf{g}}[s_1^i \geq 0] = 1 - \Phi\left(\frac{-m_{\mathbf{g}}}{\sqrt{n}}\right) = \Phi\left(\frac{1}{\sqrt{n}}\right).$$

By symmetry, $\mathbb{P}_{\mathbf{b}}[a_1^i = \Theta] = \Phi(1/\sqrt{n})$ as well. Denote $\pi_n = \Phi\left(\frac{1}{\sqrt{n}}\right)$ and by $N_1 = |\{i: a_1^i = \mathbf{g}\}|$ the number of agents taking the action $a_1^i = \mathbf{g}$. Let $\kappa_n = \log(\pi_n/(1 - \pi_n))$, and note that $2/\sqrt{n} \geq \kappa_n \geq 1/\sqrt{n}$.

As the action of each agent is independent, the LLR of agent i at the beginning of period 2 is given by

$$L_2^i = \frac{2}{n}(s_1^i + s_2^i) - (2N_1 - n)\kappa_n - \text{sgn}(s_1^i)\kappa_n.$$

We define the public part of the LLR at the beginning of period 2 as

$$L_2^p = (2N_1 - n)\kappa_n.$$

This is the LLR of an outside observer. We define the private part of the LLR as the remainder

$$\hat{R}_2^i = L_2^i - L_2^p = \frac{2}{n}(s_1^i + s_2^i) - \text{sgn}(s_1^i)\kappa_n.$$

Let α_m be the action that the majority of the agents chose in the first period (with $\alpha_m = \mathbf{b}$ in case of a tie). Note that $\alpha_m = \mathbf{g}$ iff $L_2^p > 0$. Let E_t be the event that all agents take the first period majority action α_m in all subsequent periods up to time t , i.e., $a_\tau^i = \alpha_m$ for all $1 < \tau \leq t$.

Proof of Theorem 3. We prove the theorem by showing that the probability of E_t goes to one as the number of agents goes to infinity, i.e.,

$$\lim_{n \rightarrow \infty} \mathbb{P}[E_t] = 1.$$

We in fact provide a quantitative statement and prove that $\mathbb{P}[E_t] \geq 1 - 20 \cdot t \cdot \sqrt{\frac{\log n}{n}}$ for all $n \geq 3$.

We first show that the the probability of the event E_2 that all agents take the same action in period 2 goes to one. The LLR of agent i at the beginning of period 2 is given by

$$L_2^i = \frac{2}{n} \sum_{\tau=1}^2 s_\tau^i + (2N_1 - n) \kappa_n - \text{sgn}(s_1^i) \kappa_n = \hat{R}_2^i + L_2^p.$$

To show that E_2 has high probability we show that with high probability it holds that L_2^p , the public LLR induced by the first period actions, is large (in absolute value) and that the private beliefs are all small. Intuitively, this holds since both are (approximately) zero mean normal, with L_2^p having constant variance and \hat{R}_2^i having variance of order $1/\sqrt{n}$. It will then follow that with high probability the signs of L_2^p and L_2^i are equal for all i , which is a rephrasing of the definition of E_2 .

Let A_t be the event that all of the private signals in the first t periods have absolute values at most $M = 4\sqrt{n \log n}$. Using the union bound (over the agents and time periods), this happens except with probability at most

$$\mathbb{P}[A_t^c] \leq t \cdot n \cdot \mathbb{P}[|s_t^i| > M] \leq t \cdot n \cdot 2 \cdot \Phi\left(-\frac{1}{2}M/\sqrt{n}\right);$$

the $1/2$ factor in the argument of Φ is taken to account for the fact that the private signals do not have zero mean. Since $\Phi(-x) < e^{-\frac{x^2}{2}}$ for all $x < -1$, we have that

$$\mathbb{P}[A_t^c] \leq \frac{2 \cdot t}{n}.$$

Let

$$\hat{R}_t^i = \frac{2}{n} \sum_{\tau=1}^t s_\tau^i - \text{sgn}(s_1^i) \kappa_n.$$

Thus the event A_t implies that for all $\tau \leq t$

$$|\hat{R}_\tau^i| \leq \frac{2}{n} \cdot \tau \cdot M + \kappa_n \leq 8 \cdot \tau \cdot \sqrt{\frac{\log n}{n}} + \frac{2}{\sqrt{n}} \leq 9 \cdot \tau \cdot \sqrt{\frac{\log n}{n}}.$$

Let B_t be the event that the absolute value of the public LLR L_2^p is at least $9 \cdot t \cdot \sqrt{\frac{\log n}{n}}$; this is chosen so that the intersection of A_t and B_t implies E_t . Conditioned on $\Theta = \mathbf{g}$, the random variable N_1 has the unimodal binomial distribution $\mathcal{B}(n, \pi_n)$, which has mode $\lfloor (n+1) \cdot \pi_n \rfloor$. The probability at this mode is easily shown to be at most $1/\sqrt{n}$. The same applies conditioned on $\Theta = \mathbf{b}$. It follows that the probability of B_t^c , which by definition is equal to the probability that $|N_1 - n/2| \leq \frac{1}{\kappa_n} 9 \cdot t \cdot \sqrt{\frac{\log n}{n}}$, is at most $\frac{2}{\kappa_n} 9 \cdot t \cdot \sqrt{\frac{\log n}{n}}$ times the probability of the mode, or

$$\mathbb{P}[B^c] \leq \frac{2}{\kappa_n} 9 \cdot t \cdot \sqrt{\frac{\log n}{n}} \cdot \frac{1}{\sqrt{n}} \leq 18 \cdot t \cdot \sqrt{\frac{\log n}{n}}.$$

Together with the bound on the probability of A , we have that

$$\mathbb{P}[A_t \text{ and } B_t] \geq 1 - 20 \cdot t \cdot \sqrt{\frac{\log n}{n}},$$

and in particular

$$\mathbb{P}[E_2] \geq 1 - 40 \cdot \sqrt{\frac{\log n}{n}}.$$

We now claim that $A_t \cap B_t$ implies E_t . To see this, note that as $A_t \cap B_t$ implies E_2 , the agents all observe at period 2 that no other agent has a strong enough signal to dissent with the first period majority. This only strengthens their belief in the first period majority, requiring them an even higher (in absolute value) threshold than L_2^p to choose another action; the formal proof of this statement is identical to the proof of Lemma 9. But since, under the event $A_t \cap B_t$, each of their private LLRs \hat{R}_τ^i is weaker than L_2^p for all $\tau \leq t$, they will not do so at period 3, or, by induction, in any of the periods prior to period t . This completes the proof. \square

APPENDIX E. INCOMPLETE OBSERVATIONS STRUCTURES

In this section we study the case that agent 1 observes agent 2's actions, but not vice versa. We prove Theorem 4, and moreover precisely calculate the error rate of agent 1, which allows us to compare it to the error rate in the bidirectional case. We think that this result is of independent interest.

Theorem 8. *The probability that agent 1 makes a mistake if she observes agent 2's actions unidirectionally satisfies*

$$e_t^{\rightarrow} = \mathbb{P}_\theta [a_t^1 \neq \alpha^g] = e^{-r_{uni} \cdot t + o(t)},$$

where $r_{uni} := r_{aut} + \min \{ \lambda_g^*(r_{aut}), \lambda_b^*(r_{aut}) \} = \min \{ \lambda_b^*(-r_{aut}), \lambda_g^*(-r_{aut}) \}$.

In the case of normal signals we can calculate r_{uni} exactly:

Corollary 2. *Let μ_θ be the normal distribution with mean m_θ and variance $\sigma^2 > 0$. In this case $r_{uni} = \frac{25}{16} r_{aut}$.*

This implies that agent 1 learns as fast as she would learn if she observed $9/16 \approx 56\%$ of agent 2's private signals, instead of her actions.

Observing the last action. To gain some intuition into unidirectional observations, let us first assume that agent 1 observes only agent 2's last action a_{t-1}^2 , rather than the entire history of 2's actions. That is, at time t the information available to agent 1 is $s_1^1, \dots, s_t^1, a_{t-1}^2$, and the information available to agent 2 is only s_1^2, \dots, s_t^2 .

Bayes rule yields that the LLR of agent 1 when agent 2 takes the action α is given by

$$(21) \quad L_t^1 = R_t^1 + I_t(a_{t-1}^2),$$

where $I_t(a_{t-1}^2)$ is the amount by which agent 1's log-likelihood is shifted when she observes agent 2 take action a_{t-1}^2 in period $t - 1$:

$$I_t(\alpha) := \log \frac{\mathbb{P}_g [a_{t-1}^2 = \alpha]}{\mathbb{P}_b [a_{t-1}^2 = \alpha]}.$$

The next claim shows that there are three different types of inference $I_t(\alpha)$ agent 1 can draw from agent 2's behavior.

Lemma 13. *The function $I_t(\alpha)$ satisfies*

$$I_t(\alpha) = \begin{cases} -r_{aut} \cdot t + o(t) & \text{if } \alpha = \alpha^b \\ +r_{aut} \cdot t + o(t) & \text{if } \alpha = \alpha^g \\ o(1) & \text{if } \alpha \notin \{\alpha^b, \alpha^g\} \end{cases}.$$

This lemma follows simply from Fact 1, which characterizes agent 2's autarky behavior: When agent 2 takes a certainty action $\alpha \in \{\alpha^b, \alpha^g\}$ agent 1 believes that agent 2 has strong evidence for the state in which agent 2's action is optimal. If agent 2 does not take a certainty action $\alpha \notin \{\alpha^b, \alpha^g\}$ agent 1 believes that agent 2 must have gotten a sequence of very uninformative signals as she knows that agent 2's belief is bounded away from certainty. As a consequence the influence that agent 2's action has on agent 1's LLR $I_t(\alpha)$ vanishes for large t in this case.

The fact, that the amount by which a full certainty action of agent 2 shifts agent 1's belief is asymptotically linear in the period t , with slope equal to the rate r_{aut} , follows as, by Fact 1, the probability of a mistake in autarky vanishes at the rate r_{aut} :

$$\begin{aligned} I_t(\alpha^b) &= \log \frac{\mathbb{P}_g [a_{t-1}^2 = \alpha^b]}{\mathbb{P}_b [a_{t-1}^2 = \alpha^b]} = \log \mathbb{P}_g [a_{t-1}^2 = \alpha^b] - \log \mathbb{P}_b [a_{t-1}^2 = \alpha^b] \\ &= \log (e^{-r_{aut} \cdot t + o(t)}) - o(1) \\ &= -r_{aut} \cdot t + o(t). \end{aligned}$$

Intuitively, as agent 1 knows that agent 2, who acts in autarky, will take a suboptimal action approximately with probability $e^{-r_{aut} \cdot t}$, agent 1 shifts her LLR by approximately $-r_{aut} \cdot t$ when she sees that agent 2 chose α^b , and shifts by $+r_{aut} \cdot t$ when she sees agent 2 chose α^g . When agent 1 sees agent 2 take an action that is not optimal in either state she concludes that agent 2 is uninformed and ignores her action.

To calculate the probability of a mistake by agent 1, let us first consider the case of the good state. Recall that the LLR of agent 1 is the sum of the LLRs of her private signals R_t^1

as well as the inference $I_t(a_{t-1}^2)$ she draws from agent 2's action

$$(22) \quad L_t^1 = R_t^1 + I_t(a_{t-1}^2) = \begin{cases} R_t^1 - r_{aut} \cdot t + o(t) & \text{if } a_{t-1}^2 = \alpha^b \\ R_t^1 + r_{aut} \cdot t + o(t) & \text{if } a_{t-1}^2 = \alpha^g \\ R_t^1 + o(1) & \text{if } a_{t-1}^2 \notin \{\alpha^b, \alpha^g\}, \end{cases}$$

where the second equality follows from Lemma 13. As shown in Lemma 1, agent 1 makes a mistake in the good state (i.e., does not choose α^g) whenever her likelihood is below $\underline{L}(a^g)$. Thus, when $a_{t-1}^2 = \alpha^b$, agent 1 does not choose α^g whenever $R_t^1 \leq r_{aut} \cdot t + o(t)$. We can estimate the probability of this event using Lemma 2: it is $e^{-\lambda_g^*(r_{aut}) \cdot t + o(t)}$. A similar calculation for the other two cases yields

$$(23) \quad \begin{aligned} \mathbb{P}_g [a_t^1 \neq \alpha^g \mid a_{t-1}^2 = \alpha] &= \mathbb{P}_g [L_t^1 \leq \underline{L}(\alpha^g) \mid a_{t-1}^2 = \alpha] \\ &= \begin{cases} e^{-\lambda_g^*(+r_{aut}) \cdot t + o(t)} & \text{if } \alpha = \alpha^b \\ e^{-\lambda_g^*(-r_{aut}) \cdot t + o(t)} & \text{if } \alpha = \alpha^g \\ e^{-\lambda_g^*(0) \cdot t + o(t)} & \text{if } \alpha \notin \{\alpha^b, \alpha^g\} \end{cases} . \end{aligned}$$

To calculate the overall probability of a mistake in state \mathbf{g} we calculate the probability with which the three above cases occur.

First, consider the case where agent 1 chooses a wrong action and agent 2 chooses the correct action α^g . By Fact 1 the probability that agent 2 chooses the correct action $a_{t-1}^2 = \alpha^g$ satisfies

$$\mathbb{P}_g [a_{t-1}^2 = \alpha^g] = 1 - e^{-r_{aut} + o(t)}$$

As a consequence the probability that agent 1 chooses a wrong action and agent 2 chooses the correct action equals

$$(24) \quad \begin{aligned} \mathbb{P}_g [a_t^1 \neq \alpha^g \text{ and } a_{t-1}^2 = \alpha^g] &= \mathbb{P}_g [a_t^1 \neq \alpha^g \mid a_{t-1}^2 = \alpha^g] \times \mathbb{P}_g [a_{t-1}^2 = \alpha^g] \\ &= e^{-\lambda_g^*(-r_{aut}) \cdot t + o(t)} (1 - e^{-r_{aut} \cdot t + o(t)}) \\ &= e^{-\lambda_g^*(-r_{aut}) \cdot t + o(t)} . \end{aligned}$$

The analysis of the other two cases (i.e., when agent 2 chooses $\alpha = \alpha^b$ or $\alpha \notin \{\alpha^b, \alpha^g\}$) are completed in the proof of the following lemma.

Lemma 14. *The probability that agent 1 makes a mistake if she observes agent 2's last action unidirectionally satisfies*

$$\mathbb{P} [a_t^1 \neq \alpha^\Theta] = e^{-r_{uni} \cdot t + o(t)} ,$$

where $r_{uni} := r_{aut} + \min \{ \lambda_g^*(r_{aut}), \lambda_b^*(r_{aut}) \} = \min \{ \lambda_b^*(-r_{aut}), \lambda_g^*(-r_{aut}) \}$.

Proof. Assuming that agent 1 only observes the last action of agent 2, we would like to calculate $\mathbb{P}_{\mathfrak{g}} [a_1^t \neq \alpha^{\mathfrak{g}}]$. We can write this as

$$(25) \quad \begin{aligned} & \mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}] \\ &= \mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}, a_{t-1}^2 = \alpha^{\mathfrak{g}}] + \mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}, a_{t-1}^2 = \alpha^{\mathfrak{b}}] + \mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}, a_{t-1}^2 \notin \{\alpha^{\mathfrak{g}}, \alpha^{\mathfrak{b}}\}]. \end{aligned}$$

We already calculated the first term in (24): it is equal to $e^{-\lambda_{\mathfrak{g}}^*(-r_{aut}) \cdot t + o(t)}$. To calculate the second term we write

$$\begin{aligned} \mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}} \text{ and } a_{t-1}^2 = \alpha^{\mathfrak{b}}] &= \mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}} \mid a_{t-1}^2 = \alpha^{\mathfrak{b}}] \times \mathbb{P}_{\mathfrak{g}} [a_{t-1}^2 = \alpha^{\mathfrak{b}}] \\ &= e^{-\lambda_{\mathfrak{g}}^*(+r_{aut}) \cdot t + o(t)} \times \mathbb{P}_{\mathfrak{g}} [a_{t-1}^2 = \alpha^{\mathfrak{b}}], \end{aligned}$$

where the second equality is an application of (23). To estimate $\mathbb{P}_{\mathfrak{g}} [a_{t-1}^2 = \alpha^{\mathfrak{b}}]$ we note that agent 2 acts as in autarky, and therefore, by Lemma 2, $\mathbb{P}_{\mathfrak{g}} [a_{t-1}^2 = \alpha^{\mathfrak{b}}] = e^{-\lambda^*(0) \cdot t + o(t)} = e^{-r_{aut} \cdot t + o(t)}$. Hence

$$\mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}} \text{ and } a_{t-1}^2 = \alpha^{\mathfrak{b}}] = e^{-(\lambda_{\mathfrak{g}}^*(r_{aut}) + r_{aut}) \cdot t + o(t)}.$$

We are thus left with the estimation of the last addend, $\mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}} \text{ and } a_{t-1}^2 \notin \{\alpha^{\mathfrak{g}}, \alpha^{\mathfrak{b}}\}]$. To this end we note that

$$\mathbb{P}_{\mathfrak{g}} [a_{t-1}^2 \notin \{\alpha^{\mathfrak{g}}, \alpha^{\mathfrak{b}}\}] \leq \mathbb{P}_{\mathfrak{g}} [R_t^2 \leq \underline{L}(\alpha^{\mathfrak{g}})] = e^{-r_{aut} \cdot t + o(t)},$$

where the last equality is another consequence of Lemma 2. Therefore, by (23),

$$\mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}} \text{ and } a_{t-1}^2 \notin \{\alpha^{\mathfrak{g}}, \alpha^{\mathfrak{b}}\}] = e^{-2r_{aut} \cdot t + o(t)}.$$

We thus have that

$$\mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}] = e^{-\lambda_{\mathfrak{g}}^*(-r_{aut}) \cdot t + o(t)} + e^{-(\lambda_{\mathfrak{g}}^*(r_{aut}) + r_{aut}) \cdot t + o(t)} + e^{-2r_{aut} \cdot t + o(t)}.$$

Recall that $\lambda_{\mathfrak{b}}^*(\eta) = \lambda_{\mathfrak{g}}^*(-\eta) - \eta$ (by (16)) and so $\lambda_{\mathfrak{g}}^*(r_{aut}) + r_{aut} = \lambda_{\mathfrak{b}}^*(-r_{aut})$. Hence

$$\mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}] = e^{-\lambda_{\mathfrak{g}}^*(-r_{aut}) \cdot t + o(t)} + e^{-\lambda_{\mathfrak{b}}^*(-r_{aut}) \cdot t + o(t)} + e^{-2r_{aut} \cdot t + o(t)}.$$

We show in Lemma 15 below that $\lambda_{\mathfrak{g}}^*(-r_{aut}) < 2r_{aut}$, and likewise $\lambda_{\mathfrak{b}}^*(-r_{aut}) < 2r_{aut}$. Given this, the last addend can be absorbed into the $o(t)$ term, and we have that

$$\mathbb{P}_{\mathfrak{g}} [a_t^1 \neq \alpha^{\mathfrak{g}}] = e^{-r_{uni} \cdot t + o(t)},$$

where

$$r_{uni} = \min \{ \lambda_{\mathfrak{g}}^*(-r_{aut}), \lambda_{\mathfrak{b}}^*(-r_{aut}) \} = r_{aut} + \min \{ \lambda_{\mathfrak{b}}^*(r_{aut}), \lambda_{\mathfrak{g}}^*(r_{aut}) \}.$$

By symmetry the same holds conditioned on $\Theta = \mathfrak{b}$, and so we have shown that

$$\mathbb{P} [a_t^1 \neq \alpha^{\theta}] = e^{-r_{uni} \cdot t + o(t)}.$$

This concludes the proof of Lemma 14. \square

The next claim is used in the proof of the lemma above. Moreover, it shows that $r_{uni} < 2r_{aut}$: that is, for agent 1 in this setting, learning from actions is slower than learning from signals.

Lemma 15. $\lambda_{\mathbf{g}}^*(-r_{aut}) < 2r_{aut}$ and $\lambda_{\mathbf{b}}^*(-r_{aut}) < 2r_{aut}$.

Proof. We show the former; the proof of the latter is identical. To this end, we first note that $-r_{aut} > \lambda'_{\mathbf{g}}(1)$ (Claim 7). It thus follows that the maximum in

$$\lambda_{\mathbf{g}}^*(-r_{aut}) = \max_{z \geq 0} \lambda_{\mathbf{g}}(z) + r_{aut}z$$

is also obtained in $(0, 1)$, since the z in which it is obtained is the solution to $\lambda'(z) = -r_{aut}$. Thus

$$\lambda_{\mathbf{g}}^*(-r_{aut}) = \max_{z \in (0,1)} \lambda_{\mathbf{g}}(z) + r_{aut}z < \max_{z \in (0,1)} \lambda_{\mathbf{g}}(z) + r_{aut} = 2r_{aut}. \quad \square$$

Observing all Actions Unidirectionally and the Proof of Theorem 8. We now return to the case that agent 1 observes all of agent 2's past actions, while agent 2 only observes her own signals. We show that in this case the speed of learning is identical to the speed in the case that she observes only the last action:

$$\mathbb{P}[a_t^1 \neq \alpha^\Theta] = e^{-r_{uni} \cdot t + o(t)}.$$

One direction is immediate: observing all actions can only reduce the probability of error relative to observing the last action, and so we know that

$$\mathbb{P}[a_t^1 \neq \alpha^\Theta] \leq e^{-r_{uni} \cdot t + o(t)}.$$

It thus remains to be shown that

$$\mathbb{P}[a_t^1 \neq \alpha^\Theta] \geq e^{-r_{uni} \cdot t + o(t)}.$$

To show this we show that the probability of a smaller event already satisfies this inequality. Specifically, we condition (without loss of generality) on $\Theta = \mathbf{g}$ and would like to consider the case that agent 2 chooses the wrong action $\alpha^{\mathbf{b}}$ at all time periods up to time t . As in Appendix C, we define for each t the action α_τ^{\min} to lowest (i.e., having the lowest \bar{L}) that is taken by agent 2 with positive probability at time t . By the above, α_τ^{\min} is equal to $\alpha^{\mathbf{b}}$ for all t large enough. We then prove the claim by showing that

$$(26) \quad \mathbb{P}_{\mathbf{g}}[a_t^1 \neq \alpha^{\mathbf{g}}, \cap_{1 \leq \tau \leq t} \{a_\tau^2 = \alpha_\tau^{\min}\}] = e^{-r_{uni} \cdot t + o(t)}.$$

That is, we show that even when agent 1 observes agent 2 take the wrong action at every period in which this is possible - even then agent 1 gets it wrong with probability that is

comparable to the probability of mistake when observing only the last action. Denote by E_t the event

$$E_t = \cap_{1 \leq \tau \leq t} \{a_\tau^2 = \alpha_\tau^{\min}\}.$$

We first claim that

$$(27) \quad \mathbb{P}_g[E_t] = e^{-r_{aut} \cdot t + o(t)}$$

and that

$$(28) \quad \mathbb{P}_b[E_t] = e^{-o(t)},$$

so that asymptotically this event has the same rate as the event $a_t^2 = \alpha^b$, for both possible values of Θ . Given this, the analysis is identical to the one carried out for the case of observing the last action only, and likewise yields (26). It thus remains to calculate the conditional rates of E_t , and in particular to show that they are the same at the rates of the event $a_t^2 = \alpha^b$. The key insight from which this follows is the classical Ballot Theorem (Bertrand, 1887). It states that if (X_1, X_2, \dots) are i.i.d. random variables, and if $Y_t = \sum_{\tau=1}^t X_\tau$ then

$$\frac{1}{t} \mathbb{P}[Y_t \leq 0] \leq \mathbb{P}[\cap_{\tau=1}^t \{Y_\tau \leq 0\}] \leq \mathbb{P}[Y_t \leq 0],$$

and so in particular the event that $Y_t \leq 0$ has the same rate as the event that $Y_\tau \leq 0$ for all $\tau \leq t$. Instead of using the Ballot Theorem, we use our Theorem 7.

Indeed, noting that the event E_t can be written as

$$E_t = \cap_{1 \leq \tau < t} \{R_\tau^2 \leq \bar{L}(\alpha_\tau^{\min})\}.$$

Thus, if we define $X_t = \ell_t$ and $y_t = \bar{L}(\alpha_\tau^{\min}) - L_0$ then $\lim_t y_t/t = 0$ and Theorem 7 yields the desired rates. This completes the proof of Theorem 8.

Proof of Theorem 4. Theorem 8 yields

$$e_t^{\rightarrow} = e^{-r_{uni} \cdot t + o(t)}.$$

Thus, to complete the proof of the Theorem, we need to analyze e_t^{\leftrightarrow} , the probability of error in the bidirectional case. Indeed, it suffices to lower bound it, and show that its rate is lower than r_{uni} .

We already lower bound e_t^{\leftrightarrow} in our main results. There, we show that it is at least $e^{-r_{grp} \cdot t + o(t)}$, where r_{grp} , the rate of the groupthink event, conditioned on the good state, is given in (10) by $r_{grp} = \frac{n}{n-1} \beta$, and where β is the solution of the fixed point equation (8). In the case of $n = 2$ agents, we get that

$$\mathbb{P}_g[a_t^1 \neq \alpha^\Theta] \geq e^{-2\beta \cdot t + o(t)}.$$

and that β is given by the fixed point equation $\beta = \lambda_g^*(\beta)$.

In the normal signal case this rate is about $1.37r_{aut}$, and is in particular less than $r_{uni} = \frac{25}{16}r_{aut} \approx 1.56r_{aut}$. To complete the proof we need to show that for every signal distribution it still holds that $2\beta < r_{uni}$.

Since $r_{uni} = \min \{ \lambda_{\mathfrak{b}}^*(-r_{aut}), \lambda_{\mathfrak{g}}^*(-r_{aut}) \}$, in order to prove the claim we need to show that $2 \cdot \beta < \lambda_{\mathfrak{g}}^*(-r_{aut})$; the corresponding condition for the bad state will follow by the same argument.

We consider two cases. If $\beta \geq r_{aut}$ then $\lambda_{\mathfrak{g}}^*(\beta) \geq \lambda_{\mathfrak{g}}^*(0)$, since $\beta = \lambda_{\mathfrak{g}}^*(\beta)$ and $r_{aut} = \lambda_{\mathfrak{g}}^*(0)$. By the monotonicity of $\lambda_{\mathfrak{g}}^*$ (Lemma 5) it then follows that $\beta \leq 0$. But this is false since it implies that $\beta = \lambda_{\mathfrak{g}}^*(\beta) \geq \lambda_{\mathfrak{g}}^*(0) > 0$, and so we have reached a contradiction.

Hence $\beta < r_{aut}$, in which case $\lambda_{\mathfrak{g}}^*(-\beta) < \lambda_{\mathfrak{g}}^*(-r_{aut})$, since $\lambda_{\mathfrak{g}}^*$ is strictly decreasing (Lemma 5). Now, since $\beta = \lambda_{\mathfrak{g}}^*(\beta)$, we have that $2 \cdot \beta = \beta + \lambda_{\mathfrak{g}}^*(\beta) = \lambda_{\mathfrak{b}}^*(-\beta)$, where the last equality follows from the general fact (see Appendix B) that $\lambda_{\mathfrak{b}}^*(\eta) = \lambda_{\mathfrak{g}}^*(-\eta) - \eta$. \square

Proof of Theorem 5. Condition on $\Theta = \mathfrak{g}$, with the other case admitting identical analysis. The probability that each of the agents $2, \dots, n$ does not choose the correct action $\alpha^{\mathfrak{g}}$ is $\gamma_t = e^{-r_{aut} \cdot t + o(t)}$, by Fact 1. Denote by M_t the event that the majority of these agents did not take the action $\alpha^{\mathfrak{g}}$. Since these actions are conditionally independent, by the Chernoff-Hoeffding bound

$$\mathbb{P}_{\mathfrak{g}}[M_t] \leq (4\gamma_t(1 - \gamma_t))^{(n-1)/2}$$

and hence

$$\mathbb{P}_{\mathfrak{g}}[M_t] \leq e^{-\frac{r_{aut}t + o(t) - \log 4}{2}(n-1)}.$$

Thus, if we take any $r < r_{aut}/2$, we have that

$$\mathbb{P}_{\mathfrak{g}}[M_t] \leq e^{-(n-1)r \cdot t + o(t)}.$$

The claim follows from the fact that agent 1's probability of mistake is lower than it would be if she were to (suboptimally) choose an action based on the majority of the actions of the others. \square

APPENDIX F. A DESCRIPTION OF THE NUMERICAL PROCEDURE

A major challenge in computing behaviour numerically is that for n agents the set of private histories at time t equals $\mathbb{R} \times A^{t \times n}$, where the first component equals the private log-likelihood ratio and the second component equals the public history of actions taken by every agent until period t . To simplify the exposition we focus on the case of two actions $|A| = 2$. In this case, for every public history of actions there exists a cut-off on the private LLR such that above the cut-off one action is taken and below the other action is taken. A strategy

that specifies behaviour at time t thus corresponds to a vector of cut-offs with dimension $2^{(t-1)\times n}$. For example, for 10 agents in period 10 this space has dimension $2^{90} \approx 1.2 \times 10^{27}$. This “curse of dimensionality” makes the computation of equilibria challenging. Indeed, it has been shown in related settings that calculating the actions of Bayesian agents is computationally hard (Hajła et al., 2019).

We overcome this challenge by computing the agents’ best response and beliefs for each realized path of actions using a nested Monte-Carlo simulation, rather than computing the whole equilibrium strategy. The main idea of our approach is to keep track of the agents’ private LLRs, and the beliefs an outside observer would have about each agent’s private LLR. As each agent’s belief can be computed from these statistics, it suffices to keep track of these n numbers and n beliefs, independent of the time period. The caveat here is that the beliefs of an outside observer are distributions, which are infinite dimensional. To solve this issue we simply approximate these distributions by large, but finite samples.