Estimating True Beliefs from Declared Opinions

Jennifer Tang, Aviv Adler, Amir Ajorlou, and Ali Jadbabaie

***Abstract*—Social networks often exert social pressure, causing individuals to adapt their expressed opinions to conform to their peers. An agent in such systems can be modeled as having a (true and unchanging) *inherent belief* but broadcasts a *declared opinion* at each time step based on her inherent belief and the past declared opinions of her neighbors. An important question in this setting is *parameter estimation*: how to disentangle the effects of social pressure to estimate inherent beliefs from declared opinions. To address this, Jadbabaie et al. [1] formulated the *interacting Pólya urn model* of opinion dynamics under social pressure and studied it on complete-graph social networks using an aggregate estimator, and found that their estimator converges to the inherent beliefs unless majority pressure pushes the network to consensus. In this work, we study this model on arbitrary networks, providing an estimator which converges to the inherent beliefs even in consensus situations. Finally, we bound the convergence rate of our estimator in both consensus and non-consensus scenarios; to get the bound for consensus scenarios (which converge slower than non-consensus) we additionally found how quickly the system converges to consensus.**

I. Introduction

Opinion dynamics is the study of how people's opinions evolve over time as they interact with others on social networks. This can provide insights and predictions into how public opinion develops on a variety of political, social, commercial and cultural topics, as well as guide marketing and political campaign strategies. For instance, Ancona et al. [2] used opinion dynamics models to study the spread of vaccine hesitancy and to develop marketing strategies to help combat it. Many common opinion dynamics models assume that people are truthful in the opinions they share. However, in reality this is not always the case, as people often alter their expressed views to better fit in with their social environment, which in turn feeds back into the social environment. This social pressure feedback loop can cause publicly-expressed opinions to become arbitrarily uniform over time [3], which poses difficulties in estimating and studying the underlying true public opinion.

In this work, we study an *interacting Pólya urn model* for opinion dynamics under social pressure, originating from [1] and developed further in [4], which captures a system of agents with stochastic behaviors who alter their publicly-expressed opinions to conform to their neighbors. This model consists of n agents on a fixed network communicating on an issue with two basic sides, denoted 0 and 1. Each agent i has an *inherent* (true and unchanging) belief ϕ_i, which is either 0 or 1, and a *bias parameter* γ_i indicating the ratio of the strength of their attachment to opinions 1 and 0, with $\gamma_i = 1$ indicating a neutral position (equal preference for both, though we will assume that no agents are neutral) and $\phi_i = 1 \iff \gamma_i > 1$ (higher preference for 1 than 0). The agents communicate their *declared opinions* to their neighbors at discrete time steps: at each integer step t, each agent i (simultaneously) declares an opinion $\psi_{i,t} \in \{0, 1\}$; the declarations of the agents at any t are random and independent, and each agent's probability of declaring 1 is determined by her bias parameter, inherent belief, and the opinions declared by her neighbors in the past. These terms are fully defined in Section II-B.

This can represent both scenarios where agents alter their statements (contrary to their actual beliefs) to better fit in to their social environment and scenarios where the agents update their beliefs according to what they hear from others but retain a bias towards their original beliefs.

A. Background Literature

We refer the reader to [1] and [4] for in-depth discussion of prior work. Here, we discuss some relevant highlights.

A highly influential opinion dynamics model is the DeGroot model [5], where agents in a network average their neighbors' opinions (and their own) in an iterative manner. This causes the entire group to asymptotically approach a state where they all share a single opinion (which may be between 0 and 1), a phenomenon known as consensus. However, consensus is not always reached in real social networks, as people remain attached to their original or true beliefs, and other opinion dynamics models were created to reflect this. Among these is the Friedkin-Johnsen model [6], in which each agent updates her opinion at each step by averaging her neighbors' opinions (as in the DeGroot model) and then averaging the result with her initial opinion.

Besides the interacting Pólya urn model in [1], several other models also include agents that internally retain their initial opinions in some form [7], [8], [9]. Ye et al. [9] study a model in which each agent has both a private and expressed opinion, which evolve differently. Agents' private opinions evolve using the same update as in the Friedkin-Johnsen model, while their public opinions are updated as the average of their own private opinion and the average public opinion of their neighbors. Both [3] and [9] are very similar to [1], since agents' expressed opinions may not match their internal beliefs. However, unlike [1], [9] assumes opinions are precisely expressed on a continuous interval, which is unrealistic for certain applications. On the other hand [3] works with binary opinions like [1], though with a significantly more complex model that includes additional terms and parameters.

The analysis in [1] is primarily focused on studying whether inherent beliefs are recoverable using an aggregate estimator. This is carried out by establishing the convergence of the dynamics in the network and analyzing the equilibrium state,

The work was supported by ARO MURI W911 NF-19-1-0217 and a Vannevar Bush Fellowship from the Office of the Under Secretary of Defense.

though the analysis is limited to the complete graph and all agents having the same amount of resistance to social pressure. In [4], the authors generalize this result to allow for any network structure and for any bias parameters. They found that the proportion of declared opinions of each agent converges almost surely to an equilibrium point in any network configuration. They also determined necessary and sufficient conditions for a network to approach consensus. We note that the definition of consensus used for the interacting Pólya urn model is that all agents will declare a single opinion (all '0' or all '1') with probability tending to 1.

B. Contributions

In [1], the authors consider when it is possible to asymptotically determine the inherent beliefs of the agents based on their history of declared opinions and those of their neighbors. They study a simplified case in which the social network is an (unweighted) complete graph and all agents have the same, known, degree of bias towards their true beliefs, and consider a specific aggregate estimator which tries to first estimate the proportion of agents with true belief 1 and then determine which agents those are. In this setting, they show that the aggregate estimator estimates the proportion of agents with true belief 1 if and only if the agents do not asymptotically approach consensus (where a large majority causes all agents declare the same opinion with probability approaching 1).

In this work, we consider the general setting:

1) the social network is an arbitrary weighted (and connected) undirected graph, possibly with self-loops;
2) the agents can have heterogeneous bias parameters, indicating different levels of resistance to social pressure or certainty in their inherent beliefs, which are not known.

This greatly increases the applicability of the model, as real-life social networks have a variety of different structures and people often react very differently to social pressure.

In this setting, we study the maximum likelihood estimator (MLE), which estimates bias parameters from the history of declared opinions, rather than the aggregate estimator from [1]. We also derive a simplified estimator for inherent beliefs from the MLE, which takes a clean form with a low-dimensional sufficient statistic, consisting of two values which are simple to update at each step. We show that if the history of the agents' declared opinions is known, the MLE and the inherent belief estimator almost surely asymptotically converges to the correct inherent beliefs and bias parameters of all the agents in all such networks (even when the network approaches consensus). This resolves the fundamental question posed in [1] of whether such estimation is always possible. We also show bounds on the convergence rate of the inherent beliefs estimator. These bounds are slower when the system approaches consensus, reflecting the loss of information in the declared opinions.

II. Model Description

We use the model from [4] which is a slight generalization of the model from [1]. We refer the reader to [4, Sec II] for more details on the model and in particular [4, Sec II F] for intuition on the model. Some changes from [1] include the addition that each edge in the network has a (nonnegative) weight denoting how much the two agents' declared opinions influence each other and the use of bias parameters instead of honesty parameters (which are different, but mathematically equivalent, representations of the same process).

A. Graph Notation

Let undirected graph $G = (V, E)$ (possibly including self-loops) be a network of n agents (corresponding to the vertices) labeled $i = 1, 2, \ldots, n$. For each edge $(i, j) \in E$, there is a weight $a_{i,j} \geq 0$, where by convention $a_{i,j} = 0$ if $(i, j) \notin E$. We denote the matrix of these weights as $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. We denote the weighted degree of vertex i as $\mathsf{deg}(i) = \sum_j a_{i,j}$. The vector of degrees of all agents is denoted as

$$\boldsymbol{d} \triangleq [\mathsf{deg}(1), \mathsf{deg}(2), \ldots, \mathsf{deg}(n)] \tag{1}$$

and its diagonalization is denoted $\boldsymbol{D} = \mathsf{diag}(\boldsymbol{d})$, i.e. the diagonal matrix of the degrees. Let the *normalized adjacency matrix* be $\boldsymbol{W} = \boldsymbol{D}^{-1}\boldsymbol{A}$ whose entries are $w_{i,j}$. We assume that $\boldsymbol{W}$ is irreducible (G is connected). We denote the largest eigenvalue of a matrix by $\lambda_{\max}(\cdot)$ (the matrices we use this with have real eigenvalues). Let $\mathbb{I}\{\cdot\}$ be the indicator function.

B. Inherent Beliefs and Declared Opinions

Each agent i has an *inherent belief* $\phi_i \in \{0, 1\}$, which does not change. At each time step t, each agent i (simultaneously) announces a *declared opinion* $\psi_{i,t} \in \{0, 1\}$. We denote by $\mathcal{H}_t$ the *history* of the process, consisting of all $\psi_{i,\tau}$ for $\tau \leq t$. Each agent has a bias parameter $\gamma_i \in (0, \infty)$ where $\gamma_i \neq 1$ (every agent has a preference for either 1 or 0), which satisfies the relationship $\gamma_i > 1 \iff \phi_i = 1$.

The declarations $\psi_{i,t}$ are based on a probabilistic rule which we will give after the following definitions: for $t \in \mathbb{Z}_+$ let

$$M_i^0(t) = m_i^0 + \sum_{\tau=2}^{t}\sum_{j=1}^{n} a_{i,j}\mathbb{I}[\psi_{i,\tau} = 0] \tag{2}$$

$$M_i^1(t) = m_i^1 + \sum_{\tau=2}^{t}\sum_{j=1}^{n} a_{i,j}\mathbb{I}[\psi_{i,\tau} = 1] \tag{3}$$

where $m_i^0, m_i^1 > 0$ represent the initial settings of the model. (Initial settings are used in place of declared opinions at time 1. Some requirements for the initial settings are given shortly.) The quantity $M_i^0(t)$ is the (weighted) number of times agent i observed a neighbor declare opinion 0 up to step t (plus initial settings), and $M_i^1(t)$ is the total of observed 1's. If each $a_{i,j} \in \{0, 1\}$, then $M_i^0(t)$ and $M_i^1(t)$ represent counts of agent's neighbors' declarations (plus initial settings). The ratio $M_i^1(t)/M_i^0(t)$ can be viewed as the social pressure on agent i to choose opinion 1. Then for $t > 1$, let $M_i(t) \triangleq m_i^0 + m_i^1 + (t-1)\,\mathsf{deg}(i) = M_i^0(t) + M_i^1(t)$ and define

$$\mu_i^0(t) \triangleq M_i^0(t)/M_i(t) \text{ and } \mu_i^1(t) \triangleq M_i^1(t)/M_i(t)\,. \tag{4}$$

The parameter $\mu_i^1(t)$ is essentially the sufficient statistic that summarizes the proportion of (weighted) declared opinions observed by agent i up to time t. Since $\mu_i^0(t) = 1 - \mu_i^1(t)$, we simplify the notation to $\mu_i(t) \triangleq \mu_i^1(t)$.

We define the function (note that μ, γ are scalars)

$$f(\mu, \gamma) \triangleq \frac{\gamma\mu}{1+(\gamma-1)\mu} = \frac{1}{1+\frac{1}{\gamma}\left(\frac{1}{\mu}-1\right)}. \qquad (5)$$

The probabilistic rule for declared opinions $\psi_{i,t+1}$ is

$$\psi_{i,t+1} \triangleq \begin{cases} 1 & \text{with probability } f(\mu_i(t), \gamma_i) \\ 0 & \text{with probability } 1-f(\mu_i(t), \gamma_i) \end{cases}. \qquad (6)$$

Note that the bias parameter γ_i is always defined as agent i's bias towards opinion 1. However, the model is symmetric in the following way: a γ bias towards 1 is equivalent to a $1/\gamma$ bias towards 0,[1] which is captured by the equation

$$f(\mu_i^1(t), \gamma) = 1 - f(\mu_i^0(t), 1/\gamma). \qquad (7)$$

We also define a sufficient statistic that summarizes agent i's declarations. Let $b_i^0, b_i^1 > 0$ (the initialization) be such that $b_i^0 + b_i^1 = 1$ for each i. Let

$$\beta_i^0(t) = \frac{b_i^0}{t} + \frac{1}{t}\sum_{\tau=2}^{t}(1-\psi_{i,\tau}) \qquad (8)$$

$$\beta_i^1(t) = \frac{b_i^1}{t} + \frac{1}{t}\sum_{\tau=2}^{t}\psi_{i,\tau}. \qquad (9)$$

These are counts and proportions of declarations of each opinion (or "time-averaged declarations") for each agent (plus initial conditions). As before, $\beta_i(t) \triangleq \beta_i^1(t)$. We define

$$m_i^0 = \sum_{j=1}^{n} a_{i,j} b_j^0 \text{ and } m_i^1 = \sum_{j=1}^{n} a_{i,j} b_j^1, \qquad (10)$$

so it follows that

$$\mu_i(t) = \frac{1}{\mathsf{deg}(i)}\sum_{j=1}^{n} a_{i,j}\beta_j(t). \qquad (11)$$

We denote the corresponding vectors over the agents as:

$$\boldsymbol{\mu}(t) \triangleq [\mu_1(t), \ldots \mu_n(t)]^\top \text{ and } \boldsymbol{\beta}(t) \triangleq [\beta_1(t), \ldots \beta_n(t)]^\top \qquad (12)$$

C. Evolution of Declared Ratio

Given the history $\mathcal{H}_t$, we can write the expected value of the next $\boldsymbol{\beta}(t+1)$ as

$$\mathbb{E}[\boldsymbol{\beta}(t+1) - \boldsymbol{\beta}(t)] = \frac{1}{t+1}(F(\boldsymbol{\beta}(t), \boldsymbol{\gamma}) - \boldsymbol{\beta}(t)) \qquad (13)$$

where $F(\boldsymbol{\beta}(t), \boldsymbol{\gamma}) = [F_1(\boldsymbol{\beta}(t), \boldsymbol{\gamma}), \ldots, F_n(\boldsymbol{\beta}(t), \boldsymbol{\gamma})]$ and

$$F_i(\boldsymbol{\beta}(t), \boldsymbol{\gamma}) = f\left(\frac{1}{\mathsf{deg}(i)}\sum_{j=1}^{n} a_{i,j}\beta_j(t), \gamma_i\right). \qquad (14)$$

The equilibrium points of the expected dynamics are given by the solutions to the equations

$$0 = (\gamma_i - 1)\beta_i\mu_i + \beta_i - \gamma_i\mu_i \qquad (15)$$

where $i \in \{1, \ldots, n\}$ and $\mu_i = \frac{1}{\mathsf{deg}(i)}\sum_{j=1}^{n} a_{i,j}\beta_j$. (See [4, Section III] for more discussion.)

[1]The honesty parameter in [1] is equivalent to the bias towards the agent's true belief, i.e. a honesty parameter of γ with a true belief of 0 corresponds to a bias parameter of $1/\gamma$.

D. Consensus

An important term for this work is *consensus*, which needs to be defined appropriately for our stochastic system.

Definition 1. Consensus *is approached if*

$$\boldsymbol{\beta}(t) \to \mathbf{1} \textit{ or } \boldsymbol{\beta}(t) \to \mathbf{0} \quad \textit{as} \quad t \to \infty. \qquad (16)$$

Since $\beta_i(t)$ represents the ratio of 1's agent i has declared, consensus is approached when this ratio goes to 0 or 1.

We let the diagonal matrix with $\boldsymbol{\gamma}$ along the diagonal be

$$\boldsymbol{\Gamma} = \mathsf{diag}(\boldsymbol{\gamma}) \qquad (17)$$

and define $\boldsymbol{J_1} = \boldsymbol{\Gamma}^{-1}\boldsymbol{W}$ and $\boldsymbol{J_0} = \boldsymbol{\Gamma}\boldsymbol{W}$. In [4], it is shown that consensus $\boldsymbol{\beta}(t) \to \mathbf{1}$ occurs when $\lambda_{\max}(\boldsymbol{J_1}) \leq 1$ and $\boldsymbol{\beta}(t) \to \mathbf{0}$ occurs when $\lambda_{\max}(\boldsymbol{J_0}) \leq 1$.

III. Estimators for Inferring Inherent Beliefs and Bias Parameters

One of the key questions in [1] is whether it is possible to infer the inherent beliefs of agents from the history of declared opinions. The authors of [1] studied the interacting Pólya urn model on the complete graph using an aggregate estimator which keeps track of the fraction of declared opinions of all agents throughout time, and showed that this estimator may not converge to the inherent belief of all agents if they approach consensus. Consensus presents difficulties for estimators since asymptotically all agents approach the same behavior regardless of their inherent beliefs.

However, we show that estimators based maximum likelihood estimation (MLE) almost surely infer the inherent belief of any agent i in the limit, even when consensus is approached. This fact is connected to [4, Lemma 2] – each agent declares both opinions infinitely often, yielding sufficient information to determine inherent beliefs over time.

Additionally, unlike [1], our formulation also allows agents to have different bias parameters. Thus, it is natural to ask how to estimate the bias parameter of any agent. Intuitively, after enough time has passed, the values of $\mu_i(t)$ and $\beta_i(t)$ will converge to values close to the equilibrium point. In such a case, we can use (15) to estimate the bias parameter γ_i and inherent belief ϕ_i with

$$\widehat{\gamma}_i^{eq}(t) = \frac{\beta_i(t)}{1-\beta_i(t)}\frac{1-\mu_i(t)}{\mu_i(t)} \qquad (18)$$

$$\widehat{\phi}_i^{eq}(t) = \mathbb{I}\{\beta_i(t) < \mu_i(t)\} \qquad (19)$$

These estimators are asymptotically consistent, i.e.

$$\lim_{t\to\infty} \widehat{\gamma}_i^{eq}(t) = \gamma_i \text{ and } \lim_{t\to\infty} \widehat{\phi}_i^{eq}(t) = \phi_i \qquad (20)$$

when the dynamics converge to an interior equilibrium point (see [4] for details). However, plugging the equilibrium values into (18) is not well-defined if $\beta_i(t)$ and $\mu_i(t)$ both converge to either 0 or 1, i.e. when consensus is approached. This shows that more careful analysis needs to be done in order to estimate the bias parameters and inherent beliefs in all circumstances.

In the next sections, we develop estimators for both inherent beliefs and bias parameters which work even when agents in the network approach consensus. After that, we also discuss some results on the convergence rates of these parameters under consensus conditions. To determine this, we also determine the rate at which the system approaches consensus.

IV. Definition of Estimators

We assume at time t the estimator has at its disposal the history of agent i and agent i's neighbors' declarations up to and including time t (recall we denote as $\mathcal{H}_t$). Given $\mathcal{H}_{t-1}$, we can compute exactly the values of

$$\mathbb{P}\left[\psi_{i,t}=1|\mathcal{H}_{t-1}\right]=f(\mu_i(t-1),\gamma_i)\,. \tag{21}$$

Note that in general $\mathbb{P}\left[\psi_{i,t}=1\right]$ is a random variable dependent on $\mathcal{H}_{t-1}$, while $\mathbb{P}\left[\psi_{i,t}=1|\mathcal{H}_{t-1}\right]$ is constant. The sequence $\mathcal{H}_0\subseteq\mathcal{H}_1\subseteq\dots$ is also a filtration.

Our estimator to predict γ_i is based on the maximum log-likelihood estimator:

Definition 2. *The* single-step negative log-likelihood *for a given agent i at time $t>1$ and parameter γ is*

$$\begin{aligned}\ell_i(\gamma,t)\triangleq-\Bigg(&\mathbb{I}\{\psi_{i,t}=1\}\log(f(\mu_i(t-1),\gamma))\\&+\mathbb{I}\{\psi_{i,t}=0\}\log(1-f(\mu_i(t-1),\gamma))\Bigg)\end{aligned} \tag{22}$$

The negative log-likelihood for a given agent i at time t and parameter $\gamma\in(0,\infty)$ is

$$L_i(\gamma,t)\triangleq\sum_{\tau=2}^{t}\ell_i(\gamma,\tau)\,. \tag{23}$$

Note that γ_i is the actual bias parameter of agent i, whereas γ represents a proposed value whose loss we are measuring. The MLE for bias parameter γ_i gives the value of γ that maximizes the likelihood of agent i's declarations, which also minimizes the negative log-likelihood.

Definition 3 (Estimator for Bias Parameter)**.** *The* maximum likelihood estimator (MLE) *for the bias parameter γ at time t is given by*

$$\widehat{\gamma}_i(t)\triangleq\arg\min_{\gamma}L_i(\gamma,t) \tag{24}$$

Since the inherent belief of an agent is defined as whether the bias parameter is greater than or less than 1, given the MLE estimator, we can always predict the inherent belief of agent i by taking $\text{sign}(\log(\widehat{\gamma}_i(t)))$.

However, if we assume that $\gamma_i\neq1$, and are only interested estimating the inherent beliefs, this reduces to a simpler form. Let $\bar{\beta}_i(t)=\frac{1}{t-1}\sum_{\tau=2}^{t}\mathbb{I}[\psi_{i,\tau}=1]$, which a similar quantity to $\beta_i(t)$ except that the arbitrary initial conditions are not included. (If t is large, then difference between $\beta_i(t)$ and $\bar{\beta}_i(t)$ is negligible.)

Definition 4 (Inherent Belief Estimator)**.** *Let*

$$\widehat{\phi}_i(t)=\frac{1}{2}\text{sign}\left((t-1)\bar{\beta}_i(t)-\left(\sum_{\tau=1}^{t-1}\mu_i(\tau)\right)\right)+\frac{1}{2}\,. \tag{25}$$

We note that the transformation of multiplying by $1/2$ and adding $1/2$ is simply to map the output of the $\text{sign}(\cdot)$ operation to 0 and 1. Fundamentally, this estimator requires only comparing

$$\bar{\beta}_i(t)>\frac{1}{t-1}\sum_{\tau=1}^{t-1}\mu_i(\tau)\,. \tag{26}$$

Note that $\widehat{\phi}_i(t)$ does not depend on knowing the bias parameter, as it only assumes that $\gamma\neq1$, and the estimator is simple to compute as it only requires the aggregate count of an agent's declarations and her neighborhood's declarations.

Intuitively, this compares agent i's actual declarations against its expected declarations if $\gamma_i=1$ (i.e. if the agent were unbiased); however, the consistency of this estimator is derived from that of the MLE for the bias parameter given in Definition 3. We show this derivation in Section VII.

Lastly, note that while both the estimator in Definition 4 and the estimator in (19) have the same asymptotic values when the network does not approach consensus, only the estimator in Definition 4 is guaranteed to work when the network approaches consensus.

A. Preliminaries: Bounds on $\mu_i(t)$

The following property is essential to our the analysis:

- We have $M_i(t)=m_i^0+m_i^1+(t-1)\,\mathsf{deg}(i)$. This is because each of agent i's neighbors contributes a signal of weight $a_{i,j}$ to agent i at each time step.
- We know that

$$\mu_i(t)\in\left[\frac{1}{M_i(t)},1-\frac{1}{M_i(t)}\right]\,. \tag{27}$$

As an important consequence, for a constant κ,

$$\mu_i(t)\geq\frac{\kappa}{t}\text{ and }1-\mu_i(t)\geq\frac{\kappa}{t}\,. \tag{28}$$

The constant $\kappa=(\max_i\mathsf{deg}(i))^{-1}$ always works.

B. Negative Log-Likelihood Properties

We analyze in depth the MLE which is key to our analysis. We start by introducing an alternative representation for $\ell_i(\gamma,t)$. Let $\tilde{\psi}_{i,t}=2\psi_{i,t}-1$, which takes values -1 and $+1$, instead of 0 and 1, which gives a more symmetric representation of the process.

Since $f(\mu_i(t),\gamma)$ is still the probability of $\tilde{\psi}_{i,t}=1$,

$$\begin{aligned}\ell_i(\gamma,t)&=-\log\left(\frac{1}{1+e^{-\tilde{\psi}_{i,t}\log\left(\gamma\frac{\mu_i(t-1)}{1-\mu_i(t-1)}\right)}}\right)&(29)\\&=\log\left(1+e^{-\tilde{\psi}_{i,t}\log\left(\gamma\frac{\mu_i(t-1)}{1-\mu_i(t-1)}\right)}\right)\,.&(30)\end{aligned}$$

We reparameterize γ and $\mu_i(t)$ as follows:

$$\chi\triangleq\log\gamma\quad\text{and}\quad\nu_i(t)\triangleq\log\frac{\mu_i(t)}{1-\mu_i(t)}\,. \tag{31}$$

Using χ symmetrizes the bias parameter across $\mathbb{R}$ (so $\chi=0$ represents an unbiased agent).

We thus define some quantities which take $\chi=\log\gamma$ as the argument instead of γ and use them where convenient:

$$\tilde{\ell}_i(\chi,t)\triangleq\ell_i(\gamma,t)\text{ and }\tilde{L}_i(\chi,t)\triangleq L_i(\gamma,t)\,. \tag{32}$$

For this section to Section VI we will fix i and then use γ_1 and γ_2 to represent any two possible choices for γ_i. Define

$$Z(\gamma_1, \gamma_2, t) \triangleq L_i(\gamma_2, t) - L_i(\gamma_1, t)\,. \tag{33}$$

If $Z(\gamma_1, \gamma_2, t)$ is positive, intuitively, γ_1 is more likely than γ_2, so we expect γ_1 to be the true parameter. Indeed, if γ_1 is the true parameter, then

$$\mathbb{E}[Z(\gamma_1, \gamma_2, t)] = \sum_{\tau=2}^{t} \mathbb{E}\Bigg[\mathbb{I}\{\psi_{i,\tau} = 1\} \log \frac{f(\mu_i(\tau-1), \gamma_1)}{f(\mu_i(\tau-1), \gamma_2)} + \mathbb{I}\{\psi_{i,\tau} = 0\} \log \frac{1 - f(\mu_i(\tau-1), \gamma_1)}{1 - f(\mu_i(\tau-1), \gamma_2)} \Big| \mathcal{H}_{\tau-1}\Bigg] \tag{34}$$

$$= \sum_{\tau=2}^{t} D_{\text{KL}}(f(\mu_i(\tau-1), \gamma_1) \| f(\mu_i(\tau-1), \gamma_2)) \tag{35}$$

which is always a nonnegative quantity.

Proposition 1. *$L_i(\gamma, t)$ is a stochastic process which satisfies the following properties:*

(a) For fixed γ, $L_i(\gamma, t)$ (and $\tilde{L}_i(\chi, t)$) is an increasing function in t

(b) For fixed t, $\tilde{L}_i(\chi, t)$ is a strictly convex function in χ

(c) $\ell_i(\gamma, t) \in [0, \infty)$, and for a fixed t,

- *If $\tilde{\psi}_{i,t} = -1$, then $\ell_i(\gamma, t)$ is a decreasing function in γ (and $\tilde{\ell}_i(\chi, t)$ is decreasing in χ)*
- *If $\tilde{\psi}_{i,t} = 1$, then $\ell_i(\gamma, t)$ is an increasing function in γ (and $\tilde{\ell}_i(\chi, t)$ is increasing in χ)*

(d) If there exists $t_1, t_2 \le t$ where $\tilde{\psi}_{i,t_1} = 1$ and $\tilde{\psi}_{i,t_2} = -1$, then $\tilde{L}_i(\chi, t)$ has unique finite minimum as a function in χ. Also $L_i(\gamma, t)$ has the same minimum at $\gamma = e^{\chi}$.

(e) If γ^* is the true (bias) parameter, then for any $\gamma \neq \gamma^*$,

$$\mathbb{E}[\ell_i(\gamma, t) | \mathcal{H}_{t-1}] > \mathbb{E}[\ell_i(\gamma^*, t) | \mathcal{H}_{t-1}] \tag{36}$$

Proof.

(a) For each t, $\ell_i(\gamma, t)$ is nonnegative, so $L_i(\gamma, t)$ must be increasing. (Similar for $\tilde{L}_i(\chi, t)$).

(b)

$$\frac{d^2}{d\chi^2} \tilde{\ell}_i(\chi, t) = \frac{d^2}{d\chi^2} \log\left(1 + e^{-\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))}\right) \tag{37}$$

$$= \frac{d}{d\chi} \frac{-\tilde{\psi}_{i,t} e^{-\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))}}{1 + e^{-\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))}} \tag{38}$$

$$= -\tilde{\psi}_{i,t} \frac{d}{d\chi} \frac{1}{1 + e^{\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))}} \tag{39}$$

$$= \tilde{\psi}_{i,t}^2 \frac{e^{\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))}}{(1 + e^{\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))})^2} \tag{40}$$

$$= \frac{e^{\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))}}{(1 + e^{\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))})^2} \tag{41}$$

$$> 0\,. \tag{42}$$

Note that $\tilde{\psi}_{i,t}^2 = 1$. Thus $\tilde{\ell}_i(\chi, t)$ is convex for all t, and so $\tilde{L}_i(\chi, t) = \sum_{\tau=2}^{t} \tilde{\ell}_i(\chi, \tau)$ is also convex.

(c) Using (30), changing the sign of $\tilde{\psi}_{i,t}$ changes the sign on the exponent. If $\tilde{\psi}_{i,t}$ is positive, then the quantity in the exponent is decreasing as γ increases. The range is $[0, \infty)$ since the quantity in the log is greater than or equal to 1.

(d) Follows from (b) and (c). The function $\tilde{L}_i(\chi, t)$ must be convex and go to infinity at both ends. Since $L_i(\gamma, t) = \tilde{L}_i(\chi, t)$, it has the same minimum.

(e) Result follows from (35) setting $\gamma_1 = \gamma^*$ and $\gamma_2 = \gamma$. The KL divergence must always be nonnegative and equal to zero iff $\gamma_1 = \gamma_2$.

□

V. Log-Likelihood Ratios and Martingales

To properly analyze the quantity (33), we need the following definitions. Unless otherwise stated, γ_1 is the true parameter from which the random data is generated. The following definitions will be used starting from this section to Section VI. The *loss difference* is

$$Z(t) \triangleq Z(\gamma_1, \gamma_2, t) \tag{43}$$

$$z(t) \triangleq z(\gamma_1, \gamma_2, t) \triangleq \ell_i(\gamma_2, t) - \ell_i(\gamma_1, t)\,. \tag{44}$$

The *predictable expected value* is

$$X(t) \triangleq X(\gamma_1, \gamma_2, t) \triangleq \sum_{\tau=2}^{t} \mathbb{E}[z(\tau) | \mathcal{H}_{\tau-1}] \tag{45}$$

$$x(t) \triangleq x(\gamma_1, \gamma_2, t) \triangleq \mathbb{E}[z(t) | \mathcal{H}_{t-1}]\,. \tag{46}$$

The *loss martingale* is

$$Y(t) \triangleq Y(\gamma_1, \gamma_2, t) \triangleq X(t) - Z(t) \tag{47}$$

$$y(t) \triangleq y(\gamma_1, \gamma_2, t) \triangleq x(t) - z(t)\,. \tag{48}$$

The *predictable quadratic variation* is

$$W(t) \triangleq W(\gamma_1, \gamma_2, t) \triangleq \sum_{\tau=2}^{t} \text{Var}[z(\tau) | \mathcal{H}_{\tau-1}] \tag{49}$$

$$= \sum_{\tau=2}^{t} \text{Var}[y(\tau) | \mathcal{H}_{\tau-1}] \tag{50}$$

$$w(t) \triangleq w(\gamma_1, \gamma_2, t) \triangleq \text{Var}[z(t) | \mathcal{H}_{t-1}] \tag{51}$$

$$= \text{Var}[y(t) | \mathcal{H}_{t-1}]\,. \tag{52}$$

We also let $\chi_1 = \log \gamma_1$ and $\chi_2 = \log \gamma_2$. We give some preliminary results about these processes.

Proposition 2. *We have the following properties:*

(a) $Z(t)$ is a submartingale and $X(t)$ is strictly increasing

(b) $Y(t)$ is a martingale

(c) $W(t)$ is strictly increasing

Proof.

(a) The two statements are equivalent. From Proposition 1 (e), we have

$$x(t) = \mathbb{E}[\ell_i(\gamma_2, t) | \mathcal{H}_{t-1}] - \mathbb{E}[\ell_i(\gamma_1, t) | \mathcal{H}_{t-1}] > 0\,. \tag{53}$$

(b) This follows from the definitions of $Y(t)$.

$$\mathbb{E}[y(t) | \mathcal{H}_{t-1}] = \mathbb{E}[x(t) - z(t) | \mathcal{H}_{t-1}] \tag{54}$$

$$= \mathbb{E}[\mathbb{E}[z(t) | \mathcal{H}_{t-1}] - z(t) | \mathcal{H}_{t-1}] = 0 \tag{55}$$

and

$$\mathbb{E}[Y(t)|\mathcal{H}_{t-1}] = Y(t-1) + \mathbb{E}[y(t)|\mathcal{H}_{t-1}] = Y(t-1)\,. \tag{56}$$

(c) Because of the bounds on $\mu_i(t)$ given in (27), each $\ell_i(\gamma, t)$ must be non-constant so long as $\gamma \neq 1$. Then $z(t)$ is non-constant so long as $\gamma_1 \neq \gamma_2$. The quantity $w(t)$ is the conditional variance of $z(t)$ which therefore must always be positive. The quantity $W(t)$ is a sum of $w(t)$ so it must be increasing.

□

Next we determine bounds on our quantities. First we bound the predictable expected value $X(t)$. Since γ_1 is the true bias, like in (35), we can write

$$x(t) = D_{\text{KL}}\left(f(\mu_i(t), \gamma_1) \| f(\mu_i(t), \gamma_2)\right)\,. \tag{57}$$

Lemma 1. *For each time t, we can bound*

$$x(t) \geq \frac{(\sqrt{\gamma_1} - \sqrt{\gamma_2})^2 \mu_i(t)(1-\mu_i(t))}{\max\{\frac{\gamma_1+\gamma_2}{2}, 1\}^2}\,. \tag{58}$$

Proof. To start, since $x(t)$ can be expressed as a KL divergence, we will lower bound this KL divergence by using squared Hellinger distance, specifically,

$$D_{\text{KL}}(P\|Q) \geq 2H^2(P, Q) \tag{59}$$

which we can derive from [10, 7.3].

For discrete distributions p and q over set $[1, \dots, k]$,

$$H^2(p, q) = 1 - \sum_{i=1}^{k} \sqrt{p(i)q(i)}\,. \tag{60}$$

This gives that

$$H^2(f(\mu, \gamma_1), f(\mu, \gamma_2)) \tag{61}$$
$$= 1 - \sqrt{\frac{\gamma_1\mu\gamma_2\mu}{(\gamma_1\mu + (1-\mu))(\gamma_2\mu + (1-\mu))}} - \sqrt{\frac{(1-\mu)(1-\mu)}{(\gamma_1\mu + (1-\mu))(\gamma_2\mu + (1-\mu))}} \tag{62}$$
$$= 1 - \frac{\sqrt{\gamma_1\gamma_2}\mu + (1-\mu)}{\sqrt{(\gamma_1\mu + (1-\mu))(\gamma_2\mu + (1-\mu))}}\,. \tag{63}$$

Let

$$A = \sqrt{\gamma_1\gamma_2}\mu + (1-\mu) \tag{64}$$
$$B = \sqrt{(\gamma_1\mu + (1-\mu))(\gamma_2\mu + (1-\mu))}\,. \tag{65}$$

Note that $B > A$ since squared Hellinger distance is always between 0 and 1. Then

$$H^2(f(\mu, \gamma_1), f(\mu, \gamma_2)) = \frac{B-A}{B} \tag{66}$$
$$= \frac{B^2 - A^2}{B(B+A)} \geq \frac{B^2 - A^2}{2B^2}\,. \tag{67}$$

We can compute

$$B^2 - A^2 \tag{68}$$
$$= (\gamma_1\mu + (1-\mu))(\gamma_2\mu + (1-\mu)) - (\sqrt{\gamma_1\gamma_2}\mu + (1-\mu))^2 \tag{69}$$
$$= (\gamma_1 + \gamma_2 - 2\sqrt{\gamma_1\gamma_2})\mu(1-\mu) \tag{70}$$
$$= (\sqrt{\gamma_1} - \sqrt{\gamma_2})^2\mu(1-\mu) \tag{71}$$

and using AM-GM

$$B^2 \leq (\gamma_1\mu + (1-\mu))(\gamma_2\mu + (1-\mu)) \tag{72}$$
$$\leq \left(\frac{\gamma_1+\gamma_2}{2}\mu + (1-\mu)\right)^2 \leq \left(\max\left\{\frac{\gamma_1+\gamma_2}{2}, 1\right\}\right)^2. \tag{73}$$

This results in

$$H^2(f(\mu, \gamma_1), f(\mu, \gamma_1)) \geq \frac{(\sqrt{\gamma_1} - \sqrt{\gamma_2})^2\mu(1-\mu)}{2\max\left\{\frac{\gamma_1+\gamma_2}{2}, 1\right\}^2}\,. \tag{74}$$

and combining this with (59) and (57) completes the proof. □

Lemma 2. *If $\gamma_1 \neq \gamma_2$, then there is some $t_0 = t_0(\gamma_1, \gamma_2)$ and $c_0 = c_0(\gamma_1, \gamma_2) > 0$ such for all $t > t_0$*

$$x(t) \geq c_0(\kappa/t)\,. \tag{75}$$

Additionally, there are some constants k, t_1 (which depend on t_0, γ_1, γ_2) such that for all $t > t_1$,

$$X(t) > kc_0\kappa\log(t)\,. \tag{76}$$

Proof. Using Lemma 1, let

$$c_0 = \frac{1}{2}\frac{(\sqrt{\gamma_1} - \sqrt{\gamma_2})^2}{\max\{\frac{\gamma_1+\gamma_2}{2}, 1\}^2}\,. \tag{77}$$

Then using (27),

$$\mu_i(t)(1-\mu_i(t)) \geq \frac{1}{2}\min\{\mu_i(t), 1-\mu_i(t)\} \geq \frac{1}{2}\frac{\kappa}{t}\,. \tag{78}$$

This gives that $x(t) \geq c_0(\kappa/t)$ which implies

$$X(t) = \sum_{s=2}^{t} x(s) \geq \sum_{s=t_0}^{t} c_0\frac{\kappa}{s} \geq k_0c_0\kappa\log(t)\,. \tag{79}$$

Constant k_0 accounts for some loss which occurs since $X(t)$ is a sum of terms $x(t)$, and for small t, the results may not be exact. □

The stochastic process $Z(\gamma_1, \gamma_2, t)$ is a likelihood ratio test for determining whether γ_1 or γ_2 is the true parameter. Combining Lemma 2 and (35) gives that for $\gamma_1 \neq \gamma_2$,

$$\lim_{t\to\infty} \mathbb{E}[Z(\gamma_1, \gamma_2, t)] = \lim_{t\to\infty} X(\gamma_1, \gamma_2, t) = \infty\,. \tag{80}$$

We can use (27) to replace κ in the bound for $X(t)$.

The likelihood ratio $Z(\gamma_1, \gamma_2, t)$ on average is very large as t gets large. This means that $Z(\gamma_1, \gamma_2, t)$ can be used to distinguish which of the two parameters, γ_1 or γ_2, is the true parameter governing the data. If $Z(\gamma_1, \gamma_2, t)$ is very large (positive), then γ_1 is the true parameter. If $Z(\gamma_1, \gamma_2, t)$ is very small (negative), then γ_2 is the true parameter.

Remark 1. *The fact that* $\mathbb{E}[Z(t)] \to \infty$ *heavily relies on the fact that* $\mu_i(t) \in [\kappa/t, 1-\kappa/t]$*, as discussed in Section IV-A.*

If instead, $\mu_i(t)$ *scales as* $1/t^2$*, then the limit of* $\mathbb{E}[Z(t)]$ *would be finite. In such a scenario, randomness might make* $Z(t)$ *unreliable for distinguishing between* γ_1 *and* γ_2*.*

Some situations where the condition $\mu_i(t) \in [\kappa/t, 1-\kappa/t]$ *might not hold are when the full stochastic dynamics puts higher weight on previously declared opinions or if the network adds more agents at each time step* t*.*

Next we bound the loss martingale $Y(t)$.

Lemma 3. *For any* t*, we have*

$$|Z(t) - Z(t-1)| \leq \left|\log \frac{\gamma_1}{\gamma_2}\right| = |\chi_1 - \chi_2| \tag{81}$$

$$|Y(t) - Y(t-1)| \leq \left|\log \frac{\gamma_1}{\gamma_2}\right| = |\chi_1 - \chi_2|\,. \tag{82}$$

Proof. Let $\mu = \mu_i(t)$. Since $\gamma_2 \neq \gamma_1$, we know that either:

(i) $\log\left(\frac{f(\mu,\gamma_1)}{f(\mu,\gamma_2)}\right) < 0 < \log\left(\frac{1-f(\mu,\gamma_1)}{1-f(\mu,\gamma_2)}\right)$ (if $\gamma_1 < \gamma_2$),

(ii) $\log\left(\frac{1-f(\mu,\gamma_1)}{1-f(\mu,\gamma_2)}\right) < 0 < \log\left(\frac{f(\mu,\gamma_1)}{f(\mu,\gamma_2)}\right)$ (if $\gamma_1 > \gamma_2$).

Since

$$Z(t) = \begin{cases} Z(t-1) + \log\left(\frac{f(\mu,\gamma_1)}{f(\mu,\gamma_2)}\right) & \text{if } \psi_{i,t} = 1 \\ Z(t-1) + \log\left(\frac{1-f(\mu,\gamma_1)}{1-f(\mu,\gamma_2)}\right) & \text{if } \psi_{i,t} = 0 \end{cases} \tag{83}$$

and by

$$\left|\log\left(\frac{f(\mu,\gamma_1)}{f(\mu,\gamma_2)}\right) - \log\left(\frac{1-f(\mu,\gamma_1)}{1-f(\mu,\gamma_2)}\right)\right| \tag{84}$$

$$= \left|\log \frac{\gamma_1}{\gamma_2}\frac{\gamma_2\mu + (1-\mu)}{\gamma_1\mu + (1-\mu)} - \log \frac{\gamma_2\mu - (1-\mu)}{\gamma_1\mu + (1-\mu)}\right| \tag{85}$$

$$= \left|\log \frac{\gamma_1}{\gamma_2}\right| \tag{86}$$

this means that $Z(t-1)$ and $Z(t)$ are both in the same $\left|\log \frac{\gamma_1}{\gamma_2}\right|$-sized interval ($Z(t)$ is one of the endpoints and $Z(t-1)$ is somewhere in the middle). Additionally, $Y(t)$ (given history $\mathcal{H}_{t-1}$) is also a binary random variable whose possible outcomes are $\left|\log \frac{\gamma_1}{\gamma_2}\right|$ apart, and since $\mathbb{E}[Y(t)\,|\,\mathcal{H}_{t-1}] = Y(t-1)$ this interval also must contain $Y(t-1)$, and we are done. □

Finally, we bound the predictable quadratic variation $W(t)$. Since $W(t)$ is defined as a variance, the following standard identity is helpful for finding an upper bound. Suppose that variable X takes two values, a with probability p and b with probability $(1-p)$, then $\mathrm{Var}[X] = p(1-p)(a-b)^2$. Applied to $w(t)$, we get that

$$w(t) = \frac{\gamma_1\mu_i(t)(1-\mu_i(t))}{(1+(\gamma_1-1)\mu_i(t))^2}\left(\log\frac{\gamma_1}{\gamma_2}\right)^2. \tag{87}$$

Lemma 4. *When* $\gamma_1 \neq \gamma_2$*, there exists a constant* $c_1 = c_1(\gamma_1,\gamma_2) > 0$ *and* t_1 *such that for all* $t > t_1$

$$w(t) \leq c_1 x(t)\,. \tag{88}$$

This also implies

$$W(t) \leq c_1 X(t)\,. \tag{89}$$

Proof. Starting with (87), we have that

$$w(t) \leq \frac{\gamma_1\left(\log\frac{\gamma_1}{\gamma_2}\right)^2}{(\min\{1,\gamma_1\})^2}\mu_i(t)(1-\mu_i(t)) \tag{90}$$

$$\leq \frac{\gamma_1\left(\log\frac{\gamma_1}{\gamma_2}\right)^2}{(\min\{1,\gamma_1\})^2}\frac{\max\{\frac{\gamma_1+\gamma_2}{2},1\}^2}{(\sqrt{\gamma_1}-\sqrt{\gamma_2})^2}x(t) \tag{91}$$

and thus we can set

$$c_1 = \frac{\gamma_1\left(\log\frac{\gamma_1}{\gamma_2}\right)^2}{(\min\{1,\gamma_1\})^2}\frac{\max\{\frac{\gamma_1+\gamma_2}{2},1\}^2}{(\sqrt{\gamma_1}-\sqrt{\gamma_2})^2}\,. \tag{92}$$

Since $W(t)$ is a sum of $w(t)$ and $X(t)$ is a sum of $x(t)$, we naturally have $W(t) \leq c_1 X(t)$ for all t. □

VI. Concentration by Freedman's Inequality

We want to show that the test $Z(t) > 0$ works to distinguish whether γ_1 or γ_2 is the true parameter. We do this by showing that if γ_1 is the true parameter, then almost surely $Z(t) \leq 0$ (i.e. the test fails) for only finitely many t. We show this by applying Freedman's inquality (this formulation taken from [11] (Thm 1.1), but originally from [12] (Thm 1.6)):

Theorem 1 (Freedman's Martingale Inequality [11])**.** *If* $Y(t)$ *is a martingale with steps* $y(t) = Y(t) - Y(t-1)$ *such that* $|y(t)| \leq \alpha$ *almost surely (and* $Y(0) = 0$*), and the* predictable quadratic variation *of* $Y(t)$ *is*

$$W(t) = \sum_{\tau=2}^{t}\mathbb{E}[y(\tau)^2|Y(1),\ldots,Y(\tau-1)] \tag{93}$$

then for any $s, \sigma^2 > 0$*,*

$$\mathbb{P}[\exists\, t : Y(t) \geq s, W(t) \leq \sigma^2] \leq \exp\left(\frac{-s^2/2}{\sigma^2 + \alpha s/3}\right). \tag{94}$$

This inequality is an extension of Bernstein's inequality to martingales, where the variance of each step is not fixed but is itself a random variable dependent on the history.

Using Theorem 1, we get our result for our test:

Theorem 2. *Let* γ^* *be the true parameter. The likelihood ratio test* $Z(t) = L_i(\gamma_2,t) - L_i(\gamma_1,t)$ *is such that*

$$Z(t)\begin{cases} > 0 \text{ if } \gamma^* = \gamma_1 \\ < 0 \text{ if } \gamma^* = \gamma_2 \end{cases} \tag{95}$$

for all but finitely many t*.*

Proof. For the proof, suppose that $\gamma^* = \gamma_1$. If $\gamma^* = \gamma_2$, the same proof holds except with $-Z(t)$.

For the purposes of finding a contradiction, assume that $Z(t) \leq 0$ infinity often. Since $Z(t) = X(t) - Y(t)$, we have

$$Z(t) \leq 0 \iff X(t) \leq Y(t)\,. \tag{96}$$

Using (89) from Lemma 4, we have

$$X(t) \leq Y(t) \iff W(t) \leq c_1 X(t) \leq c_1 Y(t)\,. \tag{97}$$

Now suppose there are infinite values of t where $Z(t) \leq 0$. This means there are infinite values where $W(t) \leq c_1 Y(t)$.

From (80), we know that $X(t) \to \infty$ as $t \to \infty$. This means that there is a t where $X(t) > 2$. In that case, we have

$$Y(t) - \frac{1}{2c_1} W(t) \geq X(t) - \frac{1}{2} X(t) \geq 1 \,. \tag{98}$$

This means that so long as t is large enough so that $X(t) > 2$, we have that

$$\begin{aligned} & Y(t) - \frac{1}{2c_1} W(t) \geq 1 \\ \implies & \exists s \in \mathbb{Z}_{>0} \text{ such that } \frac{1}{2c_1} W(t) \leq s \leq Y(t) \,. \end{aligned} \tag{99}$$

Let us define a set of *bad times*

$$\mathcal{T} \triangleq \{t > 1 : Z(t) \leq 0, X(t) > 2\} = \{\tilde{t}_1, \tilde{t}_2, \dots\} \tag{100}$$

where they are ordered $\tilde{t}_1 < \tilde{t}_2 < \dots$. We assume that $\mathcal{T}$ is infinite and then derive a contradiction. We define

$$\tilde{s}_k \triangleq \max_s \left\{ s \in \mathbb{Z}_{>0} : \frac{1}{2c_1} W(\tilde{t}_k) \leq s \leq Y(\tilde{t}_k) \right\} . \tag{101}$$

From (99), we know such an $\tilde{s}_k$ exists for each k. In fact, we know that $\tilde{s}_k = \lfloor Y(\tilde{t}_k) \rfloor$. If $\mathcal{T}$ is infinite, then this produces an infinite sequence of integers $\tilde{s}_1, \tilde{s}_2, \dots$. Because $X(t) \to \infty$, which implies that $Y(\tilde{t}_k) \to \infty$, we have also that

$$\lim_{k \to \infty} \tilde{s}_k = \infty \,. \tag{102}$$

While the sequence $\tilde{s}_1, \tilde{s}_2, \dots$ could have many copies of the same integer, the set $\{s : s = \tilde{s}_k \text{ for some } k\}$ must be infinite since $\lim_{k\to\infty} \tilde{s}_k = \infty$. In other words, if $\mathcal{T}$ is infinite there must be infinitely many positive integers s such that

$$Y(t) \geq s \text{ and } W(t) \leq 2c_1 s \,. \tag{103}$$

Applying Theorem 1, we get that there are infinitely many s such that

$$\begin{aligned} \mathbb{P}[\exists\, t : Y(t) \geq s, W(t) \leq 2c_1 s] &\leq \exp\left(\frac{-s^2/2}{2c_1 s + \alpha s/3} \right) \\ &= \exp\left(\frac{s}{4c_1 + 2\alpha/3} \right) = \exp(-\xi s) \end{aligned} \tag{104}$$

where $\xi = 1/(4c_1 + 2\alpha/3)$. Lemma 3 then gives the bound

$$y(t) \leq \left| \log \frac{\gamma_1}{\gamma_2} \right| \triangleq \alpha \,. \tag{105}$$

For each $s \in \mathbb{Z}_{>0}$, let event A_s be

$$A_s = \{\exists t : Y(t) > s, W(t) \leq 2c_1 s\} \,. \tag{106}$$

Then,

$$\sum_{s=1}^{\infty} \mathbb{P}[A_s] \leq \sum_{s=1}^{\infty} \exp(-\xi s) < \infty \tag{107}$$

and therefore by Borel-Cantelli, almost surely only finitely many events A_s can occur. This is a contradiction and proves our result. □

VII. Consistency of Estimators

The above concentration results show that the MLE (24) $\widehat{\gamma}_i(t)$ converges asymptotically to the true γ_i.

Theorem 3. *For any agent i, almost surely,*

$$\lim_{t \to \infty} \widehat{\gamma}_i(t) = \gamma_i \,. \tag{108}$$

Proof. We take advantage of the alternative parameterization of $\chi = \log \gamma$. Let the MLE for χ be

$$\widehat{\chi}_i(t) = \arg\min_{\chi} \tilde{L}_i(\chi, t) \,. \tag{109}$$

We will show that $\widehat{\chi}_i(t)$ converges to the true parameter, which we call χ_i, so $\widehat{\gamma}_i(t)$ converges to the true γ_i.

For any fixed $\epsilon > 0$, let $a = \chi_i - \epsilon$ and $b = \chi_i + \epsilon$. From Theorem 2, there exists some time t_a so that $\tilde{L}_i(a, t) > \tilde{L}_i(\chi_i, t)$ for all $t > t_a$, and there exists some time t_b so that $\tilde{L}_i(b, t) > \tilde{L}_i(\chi_i, t)$ for all $t > t_b$.

At all times $t > \max\{t_a, t_b\} \triangleq t(\epsilon)$, the value of $\tilde{L}_i(\chi_i, t)$ is less than both $\tilde{L}_i(a, t)$ and $\tilde{L}_i(b, t)$. By Proposition 1(b), the function $\tilde{L}_i(\chi, t)$ is convex in χ, and thus the minimum of $\tilde{L}_i(\chi, t)$ at any $t > t(\epsilon)$ must be in $[a, b] = [\chi_i - \epsilon, \chi_i + \epsilon]$.

Thus, for every $\epsilon > 0$, we can always find a $t(\epsilon)$ where for all $t > t(\epsilon)$ we have that $\widehat{\chi}_i(t)$ is within ϵ of χ_i, and thus $\lim_{t\to\infty} \widehat{\chi}_i(t) = \chi_i$ completing the proof. □

This also shows that the inherent belief estimator from Definition 4 almost surely converges to the correct result.

Theorem 4. *Almost surely, if $\gamma_i \neq 1$, then*

$$\lim_{t \to \infty} \widehat{\phi}_i(t) = \phi_i \tag{110}$$

Proof. This is equivalent to

$$\lim_{t \to \infty} \left(\sum_{\tau=1}^{t-1} \mu_i(\tau) \right) - (t-1)\bar{\beta}_i(t) \begin{cases} < 0 & \text{if } \phi_i = 1 \\ > 0 & \text{if } \phi_i = 0 \end{cases} \tag{111}$$

The result follows from three facts:

(i) letting $\chi_i = \log(\gamma_i)$ and $\widehat{\chi}_i(t) = \arg\min_{\chi} \tilde{L}_i(\chi, t) = \log(\widehat{\gamma}_i(t))$ be the maximum likelihood estimator of χ_i, then $\lim_{t\to\infty} \widehat{\chi}_i(t) = \chi_i$;
(ii) for any t, $\tilde{L}_i(\chi, t)$ is strictly convex in χ;
(iii) $\left(\sum_{\tau=1}^{t-1} \mu_i(\tau) \right) - (t-1)\bar{\beta}_i(t) = \frac{\partial}{\partial \chi} \tilde{L}_i(\chi, t) \Big|_{\chi=0}$.

Fact (i) follows directly from Theorem 3 and (ii) is Proposition 1(b). Fact (ii) also shows that

$$\widehat{\chi}_i(t) > 0 \iff \frac{\partial}{\partial \chi} \tilde{L}_i(\chi, t) \Big|_{\chi=0} < 0 \,; \tag{112}$$

and (assuming $\chi_i \neq 0$), $\phi_i = 1 \iff \chi_i > 0$. Thus facts (i) and (ii) show that (almost surely)

$$\phi_i = 1 \implies \widehat{\chi}_i(t) > 0 \text{ for all sufficiently large } t \tag{113}$$

$$\implies \frac{\partial}{\partial \chi} \tilde{L}_i(\chi, t) \Big|_{\chi=0} < 0 \text{ for all sufficiently large } t \tag{114}$$

Thus, only fact (iii) remains to be shown.

Using $\tilde{L}_i(\chi, t) = \sum_t \tilde{\ell}_i(\chi, t)$ and

$$\tilde{\ell}_i(\chi, t) = \log\left(1 + e^{-\tilde{\psi}_{i,t}(\chi + \nu_i(t-1))} \right) \tag{115}$$

to evaluate the derivative (as in (38)) at $\chi = 0$, we get

$$\frac{\partial}{\partial \chi}\tilde{\ell}_i(\chi, t)\bigg|_{\chi=0} = \frac{-\tilde{\psi}_{i,t}}{e^{\tilde{\psi}_{i,t}\nu_i(t-1)} + 1} \tag{116}$$

$$= \begin{cases} \frac{1}{\frac{1-\mu_i(t-1)}{\mu_i(t-1)}+1} & \text{if } \tilde{\psi}_{i,t} = -1 \\ \frac{-1}{\frac{\mu_i(t-1)}{1-\mu_i(t-1)}+1} & \text{if } \tilde{\psi}_{i,t} = 1 \end{cases} \tag{117}$$

$$= \begin{cases} \mu_i(t-1) & \text{if } \tilde{\psi}_{i,t} = -1 \\ \mu_i(t-1) - 1 & \text{if } \tilde{\psi}_{i,t} = 1 \end{cases} \tag{118}$$

$$= \mu_i(t-1) - \mathbb{I}\{\psi_{i,t} = 1\}. \tag{119}$$

And thus the derivative of the entire negative log-likelihood evaluated at 0 is given by

$$\frac{\partial}{\partial \chi}\tilde{L}_i(\chi, t)\bigg|_{\chi=0} = \sum_{\tau=2}^{t} \mu_i(\tau-1) - \mathbb{I}\{\psi_{i,t} = 1\} \tag{120}$$

$$= \left(\sum_{\tau=1}^{t-1} \mu_i(\tau)\right) - (t-1)\bar{\beta}_i(t). \tag{121}$$

This shows (iii) and completes the proof. □

VIII. Convergence Rates for Inferring Inherent Beliefs

While we have shown that the estimator for inherent beliefs in Definition 4 will eventually correctly converge to an agent's inherent belief, we are also interested in how fast it converges. Since the estimator $\widehat{\phi}_i(t)$ only takes values of 0 and 1 (and is therefore always exactly correct or exactly wrong) we say the estimator *converges by time* t^* if

$$\widehat{\phi}_i(t) = \phi_i \text{ for all } t \geq t^*. \tag{122}$$

The question is: for any $\delta > 0$, how many steps t^* does it take for the estimator to have a $1 - \delta$ probability of converging? Or, in other words, for what t^* do we have

$$\mathbb{P}\left[\exists t \geq t^* : \widehat{\phi}_i(t) \neq \phi_i\right] \leq \delta\,? \tag{123}$$

In this section, we give bounds for the worst-case convergence rate. As in Section VI we fix the agent i and omit it from the notation. Also as in Section VI, we assume that agent i has inherent belief $\phi_i = 1$ (so $\gamma_i > 1$).

A. Analysis for Estimator

The analysis for the convergence rate of the inherent belief estimator is similar to the analysis above for showing the convergence of the MLE. We use many of the same symbols in this proof as we did for the proof in the previous section (such as $X(t)$, $Y(t)$, etc.) but these will represent different (though analogous) quantities.

Note that $\widehat{\phi}_i(t) = \phi_i = 1$ if and only if

$$Z(t) \triangleq (t-1)\bar{\beta}_i(t) - \sum_{\tau=1}^{t-1} \mu_i(\tau) > 0. \tag{124}$$

Then $Z(t)$ is a stochastic process with differences

$$z(t) \triangleq Z(t) - Z(t-1) = \mathbb{I}\{\psi_{i,t} = 1\} - \mu_i(t-1). \tag{125}$$

We then make a martingale $Y(t)$ as in Section VI: first, the expected updates and (cumulative) predictable expected value

$$x(t) \triangleq \mathbb{E}[z(t)|\mathcal{H}_{t-1}] \text{ and } X(t) \triangleq \sum_{\tau=2}^{t} x(\tau). \tag{126}$$

We then derive $x(t)$ as

$$x(t) = f(\mu_i(t-1), \gamma_i) - \mu_i(t-1) \tag{127}$$

$$= \frac{(\gamma_i - 1)\mu_i(t-1)(1 - \mu_i(t-1))}{1 + (\gamma_i - 1)\mu_i(t-1)} \tag{128}$$

Note that since $\gamma_i > 1$ and $\mu_i(t-1) \in (0, 1)$, we have $x(t) > 0$ for all t, and hence $X(t)$ is increasing and $Z(t)$ is a submartingale.

Proposition 3. *If $X(t) > g(t)$, and $X(t_0) > 2$ for some t_0, then*

$$\mathbb{P}\left[\exists t \geq t^* : \widehat{\phi}_i(t) \neq \phi_i\right] \leq \delta \tag{129}$$

holds for

$$t^* \geq \max\left\{g^{-1}\left(\frac{1}{\xi_i}\log\frac{1}{\delta(e^{\xi_i} - 1)}\right), t_0\right\} \tag{130}$$

$$\xi_i = \frac{1}{4c_i + 2/3} \tag{131}$$

and $c_i = \frac{\gamma_i}{\gamma_i - 1}$.

Proof. First, note that g must be a monotonic function. We define

$$y(t) \triangleq x(t) - z(t) \text{ and } Y(t) \triangleq \sum_{\tau=2}^{t} y(\tau) = X(t) - Z(t). \tag{132}$$

Then $Y(t)$ is a martingale, as $x(t) = \mathbb{E}[z(t)|\mathcal{H}_{t-1}] \implies \mathbb{E}[y(t)|\mathcal{H}_{t-1}] = 0$. Furthermore, the martingale $Y(t)$ has bounded step sizes. We define the predictable quadratic variation as:

$$w(t) \triangleq \mathrm{Var}[y(t)|\mathcal{H}_{t-1}] \text{ and } W(t) \triangleq \sum_{\tau=2}^{t} w(\tau) \tag{133}$$

Noting that (given $\mathcal{H}_{t-1}$) the value $\mu_i(t)$ is fixed yields

$$w(t) = \mathrm{Var}[\psi_{i,t}|\mathcal{H}_{t-1}] = \frac{\gamma_i\mu_i(t-1)(1 - \mu_i(t-1))}{(1 + (\gamma_i - 1)\mu_i(t-1))^2} \tag{134}$$

This then implies bounds on the ratio between $w(t)$ and $x(t)$:

$$\frac{1}{\gamma_i - 1}x(t) \leq w(t) \leq \frac{\gamma_i}{\gamma_i - 1}x(t) \tag{135}$$

Since $c_i = \frac{\gamma_i}{\gamma_i - 1}$, this gives

$$W(t) \leq c_i \sum_{\tau=2}^{t} x(\tau) \leq c_i X(t) \tag{136}$$

Rewriting as $W(t)/c_i \leq X(t)$, we note that $X(t) \geq 2$ implies $W(t)/(2c_i) \leq X(t) - 1$; this means that when $X(t) \geq 2$, if $Y(t) > X(t)$ there must be some integer s such that $W(t)/(2c_i) < s < Y(t)$.

Let t_0 be when $X(t_0) \geq 2$. For any $t > t_0$, the estimator being wrong then implies

$$Z(t) < 0 \iff X(t) - Y(t) < 0 \tag{137}$$
$$\iff Y(t) > X(t) \tag{138}$$
$$\implies \exists s \in \mathbb{Z}_+ : X(t) - 1 < s < Y(t) \tag{139}$$
$$\implies \exists s \in \mathbb{Z}_+ : W(t)/(2c_i) < s < Y(t) \tag{140}$$

Let event A_s for integer s be defined as

$$A_s = \{\exists t > t_0 : X(t) - 1 < s < Y(t)\} \tag{141}$$

If A_s occurs, we call s a *separator*; by the above, for any $t > t_0$, if the estimator is wrong at time t there must be a separator corresponding to t (note however that one separator s can work for multiple t). We now apply Freedman's Inequality (Theorem 1) to bound the probability that any given s is a separator:

$$\mathbb{P}[A_s] = \mathbb{P}[\exists t > t_0 : X(t) - 1 < s < Y(t)] \tag{142}$$
$$\leq \mathbb{P}[\exists t : Y(t) \geq s, W(t) \leq 2c_i s] \tag{143}$$
$$\leq \exp\left(\frac{-s}{2(2c_i + 1/3)}\right). \tag{144}$$

However, given some s_0, we want to bound the probability that *any* integer $s > s_0$ is a separator. We thus define

$$B_{s_0} = \{\nexists\ s \in \mathbb{Z} > s_0 \text{ such that } A_s \text{ holds}\}. \tag{145}$$

Using (131) gives:

$$1 - \mathbb{P}[B_{s_0}] = \sum_{s=s_0+1}^{\infty} \mathbb{P}[A_s] \leq \sum_{s=s_0+1}^{\infty} \exp(-\xi_i s) \tag{146}$$
$$= \frac{e^{-\xi_i(s_0+1)}}{e^{\xi_i} - 1}. \tag{147}$$

If t is such that

$$s_0 + 1 \leq g(t) \leq X(t) \implies t \geq g^{-1}(s_0 + 1) \tag{148}$$

and B_{s_0} holds, then $Z(t) > 0$ and the estimator is correct. Thus, if (148) holds, then

$$\mathbb{P}[Z(t) < 0] \leq 1 - \mathbb{P}[B_{s_0}] \leq \frac{e^{-\xi_i(s_0+1)}}{e^{\xi_i} - 1}. \tag{149}$$

If we want

$$\mathbb{P}[Z(t) < 0] \leq \frac{e^{-\xi_i(s_0+1)}}{e^{\xi_i} - 1} \leq \delta \tag{150}$$

then

$$e^{-\xi_i(s_0+1)} \leq \delta(e^{\xi_i} - 1) \tag{151}$$
$$\implies s_0 + 1 \geq \frac{1}{\xi_i} \log \frac{1}{\delta(e^{\xi_i} - 1)}. \tag{152}$$

Combining this with (148) gives

$$g(t) \geq \frac{1}{\xi_i} \log \frac{1}{\delta(e^{\xi_i} - 1)}. \tag{153}$$

□

To use Proposition 3, we need to determine how quickly $\mu_i(t)$ or $\beta_i(t)$ approaches 0. Before doing a calculation of this, we show what happens if we use a worst-case bound on $\mu_i(t)$. Since $\mu_i(t-1) \in [\kappa/t, 1 - \kappa/t]$, we have

$$x(t) \geq \frac{\gamma_i - 1}{\gamma_i} \frac{1}{2} \frac{\kappa}{t} \tag{154}$$

as $\mu_i(t)(1 - \mu_i(t)) \geq \frac{1}{2} \min(\mu_i(t), 1 - \mu_i(t))$. This implies that for $t \geq 3$,

$$X(t) \geq \frac{\gamma_i - 1}{\gamma_i} \frac{1}{4} \kappa \log(t). \tag{155}$$

By (155), we know that $X(t) \geq 2$ when

$$t \geq e^{8 \frac{\gamma_i}{\gamma_i - 1} \kappa^{-1}} \triangleq t_0. \tag{156}$$

Then the estimator has converged with probability $\geq 1 - \delta$ for all t such that

$$\frac{\gamma_i - 1}{4\gamma_i} \kappa \log(t) \geq \frac{1}{\xi_i} \log \frac{1}{\delta(e^{\xi_i} - 1)} \tag{157}$$
$$\implies t \geq \left(\frac{1}{\delta(e^{\xi_i} - 1)}\right)^{\frac{4}{\kappa} \frac{\gamma_i}{\gamma_i - 1}\left(4 \frac{\gamma_i}{\gamma_i - 1} + \frac{2}{3}\right)} = \Theta\left((1/\delta)^c\right) \tag{158}$$

where c is a constant depending on κ and γ_i. Thus, we know that for some t^* on the order of $(1/\delta)^c$, the inherent belief estimator converges by time t^* with probability at least $1 - \delta$.

However, since (155) is a lower bound, corresponding to using the lower bound of $\Theta(1/t)$ for $\mu_i(t)$, the computed convergence rate (158) is too low. This raises the question of improving it by using better bounds on $\mu_i(t)$, thus yielding a better bound of $X(t)$ in (155). This can be divided into two cases: consensus and non-consensus.

If the system does not approach consensus (i.e. it converges to an interior equilibrium point) then $X(t)$ is linear since for large enough t, $x(t)$ will be very close to a constant, and thus $X(t) \geq Kt$ for some constant K. Using Proposition 3 then gives that we converge for all t such that

$$t \geq \max\left\{\frac{2}{K}, \frac{1}{K} \frac{1}{\xi_i} \log \frac{1}{\delta(e^{\xi_i} - 1)}\right\} = \Theta\left(\log \frac{1}{\delta}\right) \tag{159}$$

(where $\xi_i = \frac{1}{4c_i + 2/3}$ as defined in proposition 3).

When the system approaches consensus $X(t)$ will be sublinear as $\mu_i(t) \to 0$ as $t \to \infty$; however, by analyzing the rate of convergence to consensus, we will obtain a better bound on $\mu_i(t)$ than $\Theta(1/t)$, which will yield a more precise convergence rate for the inherent belief estimator.

IX. Bounds on Rate of Convergence to Consensus

In this section, we look at what the rate of convergence to consensus is. Consensus occurs when either $\lambda_{\max}(\boldsymbol{J_0})$ or $\lambda_{\max}(\boldsymbol{J_1})$ is less than 1 (see [4]). For this section, suppose that $\lambda_{\max}(\boldsymbol{J_0}) < 1$ meaning that consensus to $\mathbf{0}$ occurs.

We define that $\beta_i(t)$ for agent i converges to 0 at a rate of t^r if for every constant $\epsilon > 0$, we have that

$$\lim_{t \to \infty} \frac{\beta_i(t)}{t^{r+\epsilon}} = 0 \quad \text{and} \quad \lim_{t \to \infty} \frac{\beta_i(t)}{t^{r-\epsilon}} = \infty. \tag{160}$$

(We note that this definition does exclude subpolynomial terms, for instance, $t^r \log t$ and $t^r / \log t$ will both satisfy the conditions.)

Let $\lambda = \lambda_{\max}(\boldsymbol{J_0})$. Let $\boldsymbol{v}$ be the associated (left) eigenvector of eigenvalue λ. Since $\boldsymbol{J_0}$ is irreducible (since $\boldsymbol{W}$ is irreducible) and a nonnegative matrix, the Perron-Frobenius theorem [13] implies that $\boldsymbol{v}$ is a positive eigenvector. Scale $\boldsymbol{v}$ so that $\boldsymbol{v}^\top \mathbf{1} = 1$. Let

$$V(\boldsymbol{\beta}) = \boldsymbol{v}^\top \boldsymbol{\beta}\,. \tag{161}$$

For convenience, we represent the declared opinions at a given step t as a (random) vector $\boldsymbol{\psi}(\boldsymbol{\beta}(t))$ depending on the state $\boldsymbol{\beta}(t)$. Thus the dynamics follow the update

$$\boldsymbol{\beta}(t+1) = \frac{t}{t+1}\boldsymbol{\beta}(t) + \frac{1}{t+1}\boldsymbol{\psi}(\boldsymbol{\beta}(t)) \tag{162}$$

and $\boldsymbol{\psi}(\boldsymbol{\beta}(t))$ is a vector with ith component is given by

$$\psi_i(\boldsymbol{\beta}(t)) = \begin{cases} 1 & \text{w.p. } f(\mu_i(t), \gamma_i) \\ 0 & \text{w.p. } 1 - f(\mu_i(t), \gamma_i) \end{cases}. \tag{163}$$

Since consensus to $\mathbf{0}$ occurs, this implies that $\lim_{t\to\infty} V(\boldsymbol{\beta}(t)) = 0$. We will show that the rate at which $V(\boldsymbol{\beta}(t))$ approaches 0 is also the rate at which each $\beta_i(t) \to 0$ up to a constant factor.

To analyze the convergence of $V(\boldsymbol{\beta}(t))$, we will define linear functions which we can use to upper and lower bound the value of $V(\boldsymbol{\beta}(t)$ when $\boldsymbol{\beta}(t)$ is in a certain region.

For $\zeta > 0$, let random vector $\bar{\boldsymbol{\psi}}(\boldsymbol{\beta}(t), \zeta)$ be such that the ith component is given by

$$\bar{\psi}_i(\boldsymbol{\beta}(t), \zeta) = \begin{cases} \begin{cases} 1 & \text{w.p. } \zeta\gamma_i\mu_i(t) \\ 0 & \text{w.p. } 1 - \zeta\gamma_i\mu_i(t) \end{cases} & \text{if } \zeta\gamma_i\mu_i(t) \le 1 \\ \zeta\gamma_i\mu_i(t) & \text{if } \zeta\gamma_i\mu_i(t) > 1 \end{cases} \tag{164}$$

(The case where $\zeta\gamma_i\mu_i(t) > 1$ is a technicality we need to consider) where $\bar{\boldsymbol{\psi}}(\boldsymbol{\beta}(t), \zeta)$ is also maximally coupled with $\boldsymbol{\psi}(\boldsymbol{\beta}(t))$.This means that the joint distribution between $\bar{\boldsymbol{\psi}}(\boldsymbol{\beta}(t), \zeta)$ and $\boldsymbol{\psi}(\boldsymbol{\beta}(t))$ is such that

$$\mathbb{P}[\bar{\psi}_i(\boldsymbol{\beta}(t), \zeta) = \psi_i(\boldsymbol{\beta}(t))] \tag{165}$$

is maximized. Note that by (164),

$$\mathbb{E}[\bar{\psi}_i(\boldsymbol{\beta}(t), \zeta) \,|\, \mu_i(t)] = \zeta\gamma_i\mu_i(t) \quad \text{and} \tag{166}$$

$$\mathrm{Var}[\bar{\psi}_i(\boldsymbol{\beta}(t), \zeta) \,|\, \mu_i(t)] \le \zeta\gamma_i\mu_i(t)\,. \tag{167}$$

where the variance term follows from

$$\mathrm{Var}[\bar{\psi}_i(\boldsymbol{\beta}(t), \zeta) \,|\, \mu_i(t)] = \max(\zeta\gamma_i\mu_i(t)(1 - \zeta\gamma_i\mu_i(t)), 0)\,. \tag{168}$$

Definition 5 (Linearized Process)**.** *Given t_0 and $\boldsymbol{\beta}(t_0)$, for each constant value $\alpha > 0$ define the stochastic process $h^\alpha(t)$ for $t \ge t_0$ to have the following joint distribution with random stochastic process $\boldsymbol{\beta}(t)$:*

$$h^\alpha(t_0) = V(\boldsymbol{\beta}(t_0)) \tag{169}$$

$$h^\alpha(t+1) = \frac{t}{t+1}h^\alpha(t) + \frac{1}{t+1}\boldsymbol{v}^\top \bar{\boldsymbol{\psi}}\left(\boldsymbol{\beta}(t), \alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\right). \tag{170}$$

While in the above $h^\alpha(t)$ is defined with respect to random process $\boldsymbol{\beta}(t)$, we will show shortly that the marginal distribution of $h^\alpha(t)$ will have expectation and a variance bound that does not depend on $\boldsymbol{\beta}(t)$. The processes $h^\alpha(t)$ will be used as linear upper and lower bounds to $V(\boldsymbol{\beta}(t))$. Let us define

$$R(t+1, \eta) = \prod_{\tau=0}^{t} \frac{\tau + \eta}{\tau + 1} \tag{171}$$

for $\eta \in (0, 1)$ and time t. We then get the following result:

Lemma 5.

$$\frac{1}{\Gamma(\eta)(t+1)^{1-\eta}} \le R(t, \eta) \le \frac{1}{\Gamma(\eta)(t)^{1-\eta}} \tag{172}$$

Proof. This follows from Gautschi's inequality. □

Lemma 6. *Some properties of $h^\alpha(\cdot)$ are:*

(a) For any $\epsilon > 0$, there is a time t_0, where for $t > t_0$, $\boldsymbol{\beta}(t)$ remains in a disc of radius $r(\epsilon)$ from $\mathbf{0}$. Then the trajectory $V(\boldsymbol{\beta}(t))$ is such that there exists an $1 < \alpha_+ < 1 + \epsilon$ and $1 - \epsilon < \alpha_- < 1$ such that

$$h^{\alpha_-}(t) \le V(\boldsymbol{\beta}(t)) \le h^{\alpha_+}(t) \tag{173}$$

for $t > t_0$.

(b) For $t \ge t_1$ and any $\alpha > 0$ where $\alpha\lambda < 1$,

$$\mathbb{E}[h^\alpha(t) | h^\alpha(t_1)] = \frac{R(t, \alpha\lambda)}{R(t_1, \alpha\lambda)} h(t_1) \tag{174}$$

$$= (1 + o(1))\frac{t^{\alpha\lambda - 1}}{t_1^{\alpha\lambda - 1}} h(t_1) \tag{175}$$

(c) For sufficiently large t_1 there is a constant c where for $t > t_1$ and any $\alpha > 0$,

$$\mathrm{Var}[h^\alpha(t) | h^\alpha(t_1)] \le c\frac{t^{2\alpha\lambda - 2}}{t_1^{2\alpha\lambda - 1}} h^\alpha(t_1) \tag{176}$$

Proof. Proof of Part (a):

From [4, Theorem 2], we know that $\boldsymbol{\beta}(t)$ almost surely converges to an equilibrium point. In this case, that equilibrium point is $\mathbf{0}$ and thus there must be some time t_0 after which which $\boldsymbol{\beta}(t)$ lies in a disc around $\mathbf{0}$. When $\boldsymbol{\beta}(t)$ is within a disc of radius r of $\mathbf{0}$, each $\mu_i(t)$ must also be within r of 0.

There are constants α_- and α_+ so that for each agent i,

$$\alpha_-\gamma_i\mu_i \le f(\mu_i, \gamma_i) \le \alpha_+\gamma_i\mu_i \tag{177}$$

for all $\mu_i \le r$. (The values α_- and α_+ get closer to 1 as r decreases.) We focus on showing $h^{\alpha_-}(t) \le V(\boldsymbol{\beta}(t))$ (showing the other case is symmetric).

First, by definition $h^{\alpha_-}(t_0) = V(\boldsymbol{\beta}(t_0))$. If we assume that $h^{\alpha_-}(t) \le V(\boldsymbol{\beta}(t))$, then

$$h^{\alpha_-}(t+1) \tag{178}$$

$$= \frac{t}{t+1}h^{\alpha_-}(t) + \frac{1}{t+1}\boldsymbol{v}^\top \bar{\boldsymbol{\psi}}\left(\boldsymbol{\beta}(t), \alpha_-\frac{h^{\alpha_-}(t)}{V(\boldsymbol{\beta}(t))}\right) \tag{179}$$

$$\le \frac{t}{t+1}V(\boldsymbol{\beta}(t)) + \frac{1}{t+1}\boldsymbol{v}^\top \bar{\boldsymbol{\psi}}\left(\boldsymbol{\beta}(t), \alpha_-\frac{h^{\alpha_-}(t)}{V(\boldsymbol{\beta}(t))}\right) \tag{180}$$

Since $h^{\alpha_-}(t) \le V(\boldsymbol{\beta}(t))$, random variable

$$\bar{\psi}_i\left(\boldsymbol{\beta}(t), \alpha_-\frac{h^{\alpha_-}(t)}{V(\boldsymbol{\beta}(t))}\right) \tag{181}$$

is 1 with probability

$$\alpha_{-}\frac{h^{\alpha_{-}}(t)}{V(\boldsymbol{\beta}(t))}\gamma_i\mu_i(t) \le \alpha_{-}\gamma_i\mu_i(t) \le f(\mu_i,\gamma_i) \tag{182}$$

where $f(\mu_i,\gamma_i)$ is the probability of $\psi_i(\boldsymbol{\beta}(t))=1$. Because the two variables are maximally coupled, every instance when (181) is 1, it must also be that $\psi_i(\boldsymbol{\beta}(t))=1$. Then

$$\bar{\psi}_i\left(\boldsymbol{\beta}(t),\alpha_{-}\frac{h^{\alpha_{-}}(t)}{V(\boldsymbol{\beta}(t))}\right) \le \psi_i(\boldsymbol{\beta}(t)) \tag{183}$$

and thus

$$h^{\alpha_{-}}(t+1) \le \frac{t}{t+1}V(\boldsymbol{\beta}(t)) + \frac{1}{t+1}\boldsymbol{v}^\top\boldsymbol{\psi}(\boldsymbol{\beta}(t)) \tag{184}$$

$$= V(\boldsymbol{\beta}(t+1))\,. \tag{185}$$

Proof of Part (b): We will use proof by induction. For the base case we trivially have

$$\mathbb{E}[h^\alpha(t_1)|h^\alpha(t_1)] = h^\alpha(t_1) = \frac{R(t_1,\alpha\lambda)}{R(t_1,\alpha\lambda)}h^\alpha(t_1)\,. \tag{186}$$

For the inductive step, we can assume that

$$\mathbb{E}[h^\alpha(t)|h^\alpha(t_1)] = \frac{R(t,\alpha\lambda)}{R(t_1,\alpha\lambda)}h^\alpha(t_1)\,. \tag{187}$$

Then we consider $t+1$:

$$\mathbb{E}[h^\alpha(t+1)|h^\alpha(t_1)] = \frac{t}{t+1}\mathbb{E}[h^a(t)|h^\alpha(t_1)] + \frac{1}{t+1}\mathbb{E}\left[\boldsymbol{v}^\top\bar{\boldsymbol{\psi}}\left(\boldsymbol{\beta}(t),\alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\right)\middle| h^\alpha(t_1)\right]\,. \tag{188}$$

The expectation in the second term is equivalent to

$$\mathbb{E}\left[\boldsymbol{v}^\top\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\alpha\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\beta}(t)\middle| h^\alpha(t_1)\right] \tag{189}$$

$$= \mathbb{E}\left[\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\alpha\lambda\boldsymbol{v}^\top\boldsymbol{\beta}(t)\middle| h^\alpha(t_1)\right] \tag{190}$$

$$= \alpha\lambda\mathbb{E}[h^\alpha(t)|h^\alpha(t_1)] \tag{191}$$

which gives that

$$\mathbb{E}[h^\alpha(t+1)|h^\alpha(t_1)] \tag{192}$$

$$= \frac{t}{t+1}\mathbb{E}[h^\alpha(t)|h^\alpha(t_1)] + \frac{1}{t+1}\alpha\lambda\mathbb{E}[h^\alpha(t)|h^\alpha(t_1)] \tag{193}$$

$$= \frac{R(t+1,\alpha\lambda)}{R(t_1,\alpha\lambda)}h(t_1)\,. \tag{194}$$

Then by Lemma 5 we have for any $\epsilon > 0$, there is a sufficiently large t_1 such that for all $t \ge t_1$,

$$(1-\epsilon)\frac{t^{\lambda-1}}{t_1^{\lambda-1}}h^\alpha(t_1) \le \mathbb{E}[h^\alpha(t)\,|\,h^\alpha(t_1)] \le (1+\epsilon)\frac{t^{\lambda-1}}{t_1^{\lambda-1}}h^\alpha(t_1)\,. \tag{195}$$

Proof of Part (c):

When we apply the law of total variance, we get that

$$\mathrm{Var}[h^\alpha(t+1)|h^\alpha(t_1)] \tag{196}$$

$$= \mathbb{E}\left[\mathrm{Var}[h^\alpha(t+1)|\mathcal{H}_t,h^\alpha(t_1)]\middle| h^\alpha(t_1)\right] + \mathrm{Var}\left[\mathbb{E}[h^\alpha(t+1)|\mathcal{H}_t,h^\alpha(t_1)]\middle| h^\alpha(t_1)\right]\,. \tag{197}$$

We compute the first term in (197) by starting with

$$\mathrm{Var}\left[\bar{\psi}_i\left(\boldsymbol{\beta}(t),\alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\right)\middle|\mathcal{H}_t,h^\alpha(t_1)\right] \tag{198}$$

$$\le \alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\gamma_i\mu_i(t)\,. \tag{199}$$

where we used (167).

Let $c_0 = \max_i\{v_i\}$. Then we can compute

$$\mathrm{Var}[h^\alpha(t+1)|\mathcal{H}_t,h^\alpha(t_1)] \tag{200}$$

$$= \mathrm{Var}\left[\frac{1}{t+1}\boldsymbol{v}^\top\bar{\boldsymbol{\psi}}\left(\boldsymbol{\beta}(t),\alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\right)\middle|\mathcal{H}_t,h^\alpha(t_1)\right] \tag{201}$$

$$= \frac{1}{(t+1)^2}\sum_{i=1}^n v_i^2\mathrm{Var}\left[\bar{\psi}_i\left(\boldsymbol{\beta}(t),\alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\right)\middle|\mathcal{H}_t,h^\alpha(t_1)\right] \tag{202}$$

$$\le \frac{1}{(t+1)^2}\sum_{i=1}^n v_i^2\alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}\gamma_i\mu_i(t) \tag{203}$$

$$\le \frac{1}{(t+1)^2}\alpha\frac{h^\alpha(t)}{V(\boldsymbol{\beta}(t))}c_0\lambda\boldsymbol{v}^\top\boldsymbol{\beta}(t) = \frac{c_0\alpha\lambda}{(t+1)^2}h^\alpha(t)\,. \tag{204}$$

To finish computing the first term in (197), we have

$$\mathbb{E}\left[\mathrm{Var}[h^\alpha(t+1)|\mathcal{H}_t,h^\alpha(t_1)]\middle| h^\alpha(t_1)\right] \tag{205}$$

$$\le \mathbb{E}\left[\frac{c_0\alpha\lambda}{(t+1)^2}h^\alpha(t_1)\middle| h^\alpha(t_1)\right] \tag{206}$$

$$\le c_1\alpha\lambda\frac{1}{t^2}\frac{t^{\alpha\lambda-1}}{t_1^{\alpha\lambda-1}}h^\alpha(t_1) \tag{207}$$

where we used the approximation result from Lemma 6(b) which is a correct upper bound for some constant c_1 and sufficiently large t_1. For the second term in (197), the expectation inside the variance is

$$\mathbb{E}[h^\alpha(t+1)|\mathcal{H}_t,h^\alpha(t_1)] = \frac{R(t+1,\alpha\lambda)}{R(t,\alpha\lambda)}h^\alpha(t) \tag{208}$$

and thus

$$\mathrm{Var}\left[\mathbb{E}[h^\alpha(t+1)|\mathcal{H}_t,h^\alpha(t_1)]\middle| h^\alpha(t_1)\right] \tag{209}$$

$$= \left(\frac{R(t+1,\alpha\lambda)}{R(t,\alpha\lambda)}\right)^2\mathrm{Var}[h^\alpha(t)|h^\alpha(t_1)]\,. \tag{210}$$

Putting this together gets

$$\mathrm{Var}[h^\alpha(t+1)|h^\alpha(t_1)] \le c_1\alpha\lambda\frac{1}{t^2}\frac{t^{\alpha\lambda-1}}{t_1^{\alpha\lambda-1}}h^\alpha(t_1) + \left(\frac{R(t+1,\alpha\lambda)}{R(t,\alpha\lambda)}\right)^2\mathrm{Var}[h^\alpha(t)|h^\alpha(t_1)]\,. \tag{211}$$

Telescoping the variance terms yields:

$$\mathrm{Var}[h^\alpha(t+1)|h^\alpha(t_1)] \tag{212}$$

$$\le c_2\alpha\lambda\sum_{\tau=t_1}^t\left(\frac{(t+1)^{\alpha\lambda-1}}{\tau^{\alpha\lambda-1}}\right)^2\frac{1}{\tau^2}\frac{\tau^{\alpha\lambda-1}}{t_1^{\alpha\lambda-1}}h^\alpha(t_1) \tag{213}$$

$$= c_2\alpha\lambda\frac{(t+1)^{2\alpha\lambda-2}}{t_1^{\alpha\lambda-1}}h^\alpha(t_1)\sum_{\tau=t_1}^t\frac{1}{\tau^{\alpha\lambda+1}}\,. \tag{214}$$

Approximating the sum using an integral then gives:

$$\sum_{\tau=t_1}^{t} \frac{1}{\tau^{\alpha\lambda+1}} \le \int_{t_1}^{\infty} \frac{1}{\tau^{\alpha\lambda+1}} d\tau = \frac{1}{\alpha\lambda} t_1^{-\alpha\lambda} \tag{215}$$

which results in

$$\mathrm{Var}[h^\alpha(t)|h^\alpha(t_1)] \le c_2 \alpha\lambda \frac{t^{2\alpha(\lambda-1)}}{t_1^{\alpha\lambda-1}} h^\alpha(t_1) \frac{1}{\alpha\lambda} t_1^{-\alpha\lambda} \tag{216}$$

$$= c_2 \frac{t^{2(\alpha\lambda-1)}}{t_1^{\alpha\lambda-1}} h^\alpha(t_1) t_1^{-\alpha\lambda}. \tag{217}$$

where c_2 is a constant not depending on α, λ, or t (so long as t_1 is sufficiently large). □

We will use the results from the above lemma to get a bound on the variance with regard to the square of the expected value. For this bound, we need to use that at any time t, we expect that $V(\boldsymbol{\beta}(t)) > c\kappa/t \triangleq \kappa^*/t$ (this is discussed in (IV-A)).

Lemma 7. *For any sufficiently large t_0 and any $t > t_0$,*

$$\frac{\mathrm{Var}[h^\alpha(t) \mid h^\alpha(t_0)]}{\mathbb{E}[h^\alpha(t) \mid h^\alpha(t_0)]^2} \le \frac{c}{h^\alpha(t_0) t_0} \le c^* \tag{218}$$

for some constant c^* which does not depend on t or t_0.

Proof.

$$\frac{\mathrm{Var}[h^\alpha(t) \mid h^\alpha(t_0)]}{\mathbb{E}[h^\alpha(t) \mid h^\alpha(t_0)]^2} \le \frac{c\frac{t^{(\alpha\lambda-1)2}}{t_0^{2\alpha\lambda-1}} h^\alpha(t_0)}{\left(\frac{t^{\alpha\lambda-1}}{t_0^{\alpha\lambda-1}} h^\alpha(t_0)\right)^2} \tag{219}$$

$$= \frac{c\, t_0^{\alpha\lambda-1}}{h^\alpha(t_0)} \frac{1}{t_0^{\alpha\lambda}} = \frac{c}{h^\alpha(t_0) t_0} \le c^* \tag{220}$$

where the last step uses $h^\alpha(t_0) = V(\boldsymbol{\beta}(t_0)) = \Omega(t_0^{-1})$. □

Note that $h^\alpha(t_0) \propto t_0^{-1}$ is a (guaranteed, not probabilistic) worst-case bound, and significantly worse than the expected $h^\alpha(t_0) \propto t_0^{\alpha\lambda-1}$. Note also that replacing $h^\alpha(t)$ on the left hand side by a scaled version $\rho\, h^\alpha(t)$ (where ρ can depend on t, t_0 but not on the value of $h^\alpha(t)$ itself) will not change the bound, as it multiplies both the numerator and denominator by ρ^2. We therefore define a martingale $\bar{h}^\alpha(t)$ as follows:

Definition 6. *Given a t_0, consider the process $\bar{h}^\alpha(\cdot)$ starting from time t_0: $\bar{h}^\alpha(t_0) = h^\alpha(t_0)$; then for any $t > t_0$ we define the normalized convergence process as*

$$\bar{h}^\alpha(t) = \frac{R(t_0, \alpha\lambda)}{R(t, \alpha\lambda)} h^\alpha(t). \tag{221}$$

We also define the random variable $\bar{h}_\alpha$ as follows:

$$\bar{h}_\alpha = \liminf_{t\to\infty} \bar{h}^\alpha(t). \tag{222}$$

Then $\bar{h}^\alpha(t)$ is nonnegative, uniformly integrable, and is a martingale.

Lemma 8. *The sequence $\{\bar{h}^\alpha(t)\}_{t\ge t_0}$ is a uniformly integrable martingale and $\lim_{t\to\infty} \bar{h}^\alpha(t) = \bar{h}_\alpha$ almost surely. Furthermore, for any t, we have $\bar{h}^\alpha(t) = \mathbb{E}[\bar{h}_\alpha \mid \mathcal{H}(t)]$.*

Proof. The process $\{\bar{h}^\alpha(t)\}_{t\ge t_0}$ is a martingale due to Lemma 6(b). The Martingale Convergence Theorem shows that it converges to a well-defined (random) limit almost surely, since $\bar{h}^\alpha(t) \ge 0$ and by definition

$$\mathbb{E}[|\bar{h}^\alpha(t)|] = \mathbb{E}[\bar{h}^\alpha(t)] = \bar{h}^\alpha(t_0) < \infty \tag{223}$$

Note that this means that almost surely,

$$\bar{h}_\alpha = \lim_{t\to\infty} \bar{h}^\alpha(t) \tag{224}$$

as the limit almost surely exists.

Finally, Lemma 7 (and the fact that $\{\bar{h}^\alpha(t)\}_{t\ge t_0}$ is a martingale) shows that $\mathbb{E}[|\bar{h}^\alpha(t)|^2]$ is bounded for all t, which implies uniform integrability by [14, Section 13.3]. □

Then Lemma 7 yields the following:

Corollary 1. *For any sufficiently large t^*, for any $t > t^*$*

$$\mathbb{P}\left[\bar{h}^\alpha(t) \le \frac{\bar{h}^\alpha(t^*)}{2} \,\middle|\, \mathcal{H}(t^*)\right] \le \frac{c^*}{c^* + 1/4} \tag{225}$$

This also implies that for any sufficiently large t^*,

$$\mathbb{P}\left[\bar{h}_\alpha \le \frac{\bar{h}^\alpha(t^*)}{4} \,\middle|\, \mathcal{H}(t^*)\right] \le \frac{c^*}{c^* + 1/8} \tag{226}$$

where c^* is the constant used in Lemma 7.

Proof. Note that if $\bar{h}^\alpha(t) \le \frac{\bar{h}^\alpha(t^*)}{2}$ then $\bar{h}^\alpha(t^*) - \bar{h}^\alpha(t) \ge \frac{\bar{h}^\alpha(t^*)}{2}$; since $\bar{h}^\alpha(t^*) - \bar{h}^\alpha(t)$ has mean 0 (conditioned on $\mathcal{H}(t^*)$) and variance $\le c^*$ given by Lemma 7, the Chebyshev-Cantelli inequality states that

$$\mathbb{P}\left[\bar{h}^\alpha(t^*) - \bar{h}^\alpha(t) \ge \frac{\bar{h}^\alpha(t^*)}{2} \middle| \mathcal{H}(t^*)\right] \tag{227}$$

$$\le \frac{\mathrm{Var}[\bar{h}^\alpha(t)]}{\mathrm{Var}[\bar{h}^\alpha(t)] + (\frac{\bar{h}^\alpha(t^*)}{2})^2} \tag{228}$$

$$\le \frac{c^* \bar{h}^\alpha(t^*)^2}{c^* \bar{h}^\alpha(t^*)^2 + (\frac{\bar{h}^\alpha(t^*)}{2})^2} = \frac{c^*}{c^* + 1/4} \tag{229}$$

Note that the function $\frac{x}{x+y}$ is increasing in x if y is positive, so we appropriately get an upperbound when applying $\mathrm{Var}[\bar{h}^\alpha(t)] \le c^* \bar{h}^\alpha(t^*)^2$.

To prove (226) (note that the bound is now $\frac{\bar{h}^\alpha(t^*)}{4}$ rather than $\frac{\bar{h}^\alpha(t^*)}{2}$) we have that by Lemma 8, $\bar{h}^\alpha(t) \to \bar{h}_\alpha$ almost surely; this means that $\bar{h}^\alpha(t) \to \bar{h}_\alpha$ in probability as well, so for any $\delta_1, \delta_2 > 0$, and for sufficiently large t,

$$\mathbb{P}[|\bar{h}_\alpha - \bar{h}^\alpha(t)| > \delta_1] \le \delta_2 \tag{230}$$

Set $\delta_1 = \frac{\bar{h}^\alpha(t^*)}{4}$ and $\delta_2 = \frac{c^*}{c^*+1/8} - \frac{c^*}{c^*+1/4}$. Then we assume to the contrary that

$$\mathbb{P}\left[\bar{h}_\alpha \le \frac{\bar{h}^\alpha(t^*)}{4} \,\middle|\, \mathcal{H}(t^*)\right] > \frac{c^*}{c^* + 1/8} \tag{231}$$

This means that, for any sufficiently large t,

$$\mathbb{P}\left[\bar{h}_\alpha \le \frac{\bar{h}^\alpha(t^*)}{4} \text{ and } \bar{h}^\alpha(t) > \frac{\bar{h}^\alpha(t^*)}{2} \middle| \mathcal{H}(t^*)\right] \tag{232}$$

$$> \frac{c^*}{c^* + 1/8} - \frac{c^*}{c^* + 1/4} \tag{233}$$

which yields the desired contradiction given by (230). □

Lemma 9. *For any α, almost surely $\bar{h}_\alpha > 0$*

Proof. Therefore we have established that for any α:

- $\bar{h}^\alpha(t)$ is a uniformly integrable martingale;
- $\bar{h}^\alpha(t) \to \bar{h}_\alpha$ almost surely;
- for any sufficiently large t^*,

$$\mathbb{P}\left[\bar{h}_\alpha \le \frac{\bar{h}^\alpha(t^*)}{4} \,\middle|\, \mathcal{H}(t^*)\right] \le \frac{c^*}{c^* + 1/8} \tag{234}$$

which implies that for all sufficiently large t^*,

$$\mathbb{P}[\bar{h}_\alpha > 0 \,|\, \mathcal{H}(t^*)] \ge 1 - \frac{c^*}{c^* + 1/8} > 0 \tag{235}$$

We define the process $\eta(t)$ as

$$\eta(t) = \mathbb{P}[\bar{h}_\alpha > 0 \,|\, \mathcal{H}(t)] \tag{236}$$

which is a martingale due to the tower property (and uniformly integrable because it is bounded). We likewise define

$$\eta = \mathbf{1}\{\bar{h}_\alpha > 0\} \tag{237}$$

and (due to uniform integrability) almost surely $\eta(t) \to \eta$ because of to Levy's 0-1 Law [15, Theorem 5.5.8]. But $\lim_{t\to\infty} \eta(t) \ge 1 - \frac{c^*}{c^*+1/8} > 0$, thus showing that 1 is the only possible limit out of $\{0, 1\}$. Thus, $\eta = 1$ almost surely, so $\bar{h}_\alpha > 0$ almost surely. □

Theorem 5. *For any $\epsilon > 0$, we have that*

$$\lim_{t\to\infty} \frac{V(\boldsymbol{\beta}(t))}{t^{\lambda-1+\epsilon}} = 0 \text{ and } \lim_{t\to\infty} \frac{V(\boldsymbol{\beta}(t))}{t^{\lambda-1-\epsilon}} = \infty \tag{238}$$

almost surely.

Proof. For a given ϵ, using Lemma 6(a), we can find a $\delta(\epsilon)$ with corresponding t_0 large enough and an α_+ and α_- which are such that

$$\alpha_+ < 1 + \epsilon/\lambda \quad \text{and} \quad \alpha_- > 1 - \epsilon/\lambda \tag{239}$$

Then for any trajectory $V(\boldsymbol{\beta}(t))$, there exists some trajectory $h^{\alpha_-}(t)$ and $h^{\alpha_+}(t)$ such that

$$h^{\alpha_-}(t) \le V(\boldsymbol{\beta}(t)) \le h^{\alpha_+}(t) \tag{240}$$

Using Lemma 9, any trajectory of $h^{\alpha_-}(t)$ has a corresponding martingale $\bar{h}^{\alpha_-}(t)$ which converges to a constant. Thus $h^{\alpha_-}(t)$ converges to zero at a rate of $\Omega(t^{\alpha_- \lambda - 1})$ almost surely. Similarly, $h^{\alpha_+}(t)$ converges to zero at a rate of $\Omega(t^{\alpha_+ \lambda - 1})$ almost surely.

Since the value of t_0 does not affect the asymptotic rate after t_0, we have

$$\lim_{t\to\infty} \frac{V(\boldsymbol{\beta}(t))}{t^{\lambda-1+\epsilon}} \le \lim_{t\to\infty} \frac{h^{\alpha_+}(t)}{t^{\lambda-1+\epsilon}} \le \lim_{t\to\infty} \frac{c_+ t^{\alpha_+\lambda-1}}{t^{\lambda-1+\epsilon}} = 0 \tag{241}$$

$$\lim_{t\to\infty} \frac{V(\boldsymbol{\beta}(t))}{t^{\lambda-1+\epsilon}} \ge \lim_{t\to\infty} \frac{h^{\alpha_-}(t)}{t^{\lambda-1+\epsilon}} \ge \lim_{t\to\infty} \frac{c_- t^{\alpha_-\lambda-1}}{t^{\lambda-1-\epsilon}} = \infty \tag{242}$$

□

Next we need the above result on $V(\boldsymbol{\beta}(t))$ to imply a result on all $\beta_i(t)$. We prove a few lemmas first.

Lemma 10. *Given any $\epsilon > 0$, there exists some time t_0 and some c, such that for all $t > t_0$, there exists some i such that*

$$\beta_i(t) \ge c t^{\lambda-1-\epsilon} \tag{243}$$

almost surely. Additionally, for any $t_2 > t_0$, there exists some i such that $\beta_i(t) \ge c_i t^{\lambda-1-\epsilon}$ holds at least $1/n$ of the times t in $t_2 < t \le 2t_2$.

Proof. First, $V(\boldsymbol{\beta}(t)$ converges at a rate at least $\Omega(t^{\lambda-1-\epsilon})$. Thus, there is some t_0, where for $t > t_0$,

$$V(\boldsymbol{\beta}(t)) \ge c t^{\lambda-1-\epsilon} . \tag{244}$$

At time $t > t_0$, let k be such that $\beta_k(t) \ge \beta_i(t)$ for all i. Since $\boldsymbol{v}$ is a vector so that $\boldsymbol{v}^\top \mathbf{1} = 1$, we have that

$$\beta_k(t) = (\boldsymbol{v}^\top \mathbf{1}) \beta_k(t) \ge \boldsymbol{v}^\top \boldsymbol{\beta}(t) = V(\boldsymbol{\beta}(t)) . \tag{245}$$

And thus this shows (243) for each t.

For the second statement, we know that there must be some i which satisfies (243). Since there are only n candidates for i, at least one i must occur the most often, which means this i occurs at least $1/n$ of the time in a certain interval. □

Definition 7. *For any agent i and constants $c, \rho > 0$, we say that time $t_1 > 0$ is (c, ρ)-good for agent i if*

$$\beta_i(t) \ge c t^{\lambda-1-\epsilon} \tag{246}$$

for at least ρ fraction of the times $t \in [t_1/2, t_1]$. We denote the set of times where this holds as

$$\mathcal{T}_i(c, \rho) = \{t_1 : t_1 \text{ is } (c, \rho)\text{-good for } i\} \tag{247}$$

For shorthand, once c, ρ are fixed, we denote the set of good times for i as $\mathcal{T}_i$. The key observation is that the set of good times for an agent i eventually becomes good for all her neighbors j, which then eventually become good for all their neighbors, and so forth until the set of good times for i must be good for all agents.

Lemma 11. *Fix an agent i and constants $c_i, \rho_i > 0$, and let $\mathcal{T}_i := \mathcal{T}_i(c_i, \rho_i)$, and let j be adjacent to i. Then there is some $c_j, \rho_j > 0$ such that, almost surely, there is some t^* for which*

$$t_1 > t^* \text{ and } t_1 \in \mathcal{T}_i \implies t_1 \in \mathcal{T}_j \tag{248}$$

where $\mathcal{T}_j := \mathcal{T}_j(c_j, \rho_j)$.

Note that ρ_j can be less than ρ_i, meaning that in the argument where good times spread from a source i, the ρ's diminish as the process gets further from i.

Proof. Recall that $t_1 \in \mathcal{T}_i$ means that $\beta_i \ge c_i t^{\lambda-1-\epsilon}$ for at least a ρ_i fraction of $t \in [t_1/2, t_1]$; we say a time t is c_i-*enough* (for agent i) if $\beta_i \ge c_i t^{\lambda-1-\epsilon}$ (unlike *good* times t_1, this only depends on the value of β_i at time t, not at any previous time).

First, note that at least $(\rho_i/4)t_1$ different t in $[t_1/2, (1 - \rho_i/4)t_1]$ are c_i-enough (since there are at least $(\rho_i/2)t_1$ c_i-enough times in total). For any $t \in [t_1/2, t_1]$,

$$\beta_i(t) \ge \begin{cases} c_i t_1^{\lambda-1-\epsilon} & \text{if } t \text{ is } c_i\text{-enough} \\ 0 & \text{otherwise} \end{cases} \tag{249}$$

since $t \le t_1$, and therefore (as $w_{j,i} = \frac{a_{j,i}}{\deg(j)}$)

$$\mu_j(t) \ge \begin{cases} c_i w_{j,i} t_1^{\lambda-1-\epsilon} & \text{if } t \text{ is } c_i\text{-enough} \\ 0 & \text{otherwise} \end{cases} . \tag{250}$$

So, when t is c_i-enough, we get

$$f(\mu_j(t),\gamma_j) \geq f(w_{j,i}c_i t_1^{\lambda-1-\epsilon},\gamma_j) \tag{251}$$
$$= \frac{\gamma_j w_{j,i}c_i t_1^{\lambda-1-\epsilon}}{\gamma_j w_{j,i}c_i t_1^{\lambda-1-\epsilon} + 1 - w_{j,i}c_i t_1^{\lambda-1-\epsilon}} \tag{252}$$
$$\geq c'_j t_1^{\lambda-1-\epsilon} \tag{253}$$

where $c'_j = \min(\gamma_j, 1)w_{j,i}c_i$, which is a lower bound on the probability of agent j declaring 1 at any c_i-enough time. Consider the $\geq \rho_i/4$ such times in $[t_1/2,(1-\rho_i/4)t_1]$; the number of 1's declared by j in the range $[t_1/2,(1-\rho_i/4)t_1]$ thus stochastically dominates the sum of $(\rho_i/4)t_1$ independent Bernoulli random variables with probability $c'_j t_1^{\lambda-1-\epsilon}$ each (whose sum has expected value $(\rho_i/4)c'_j t_1^{\lambda-\epsilon}$). By the Chernoff bound, this then yields that

$$\mathbb{P}\left[\sum_{t=t_1/2}^{(1-\rho_i/4)t_1} \psi_{j,t} \leq \frac{\rho_i c'_j}{8} t_1^{\lambda-\epsilon}\right] \leq e^{-\frac{\rho_i}{32}c'_j t_1^{\lambda-\epsilon}} \tag{254}$$

meaning that there is a very high probability of getting at least $\frac{\rho_i c'_j}{8} t_1^{\lambda-\epsilon}$ declarations of 1 from agent j by time $(1-\rho_i/4)t_1$. But then for the $(\rho_i/4)t_1$ times $t \in [(1-\rho_i/4)t_1, t_1]$, we have

$$\beta_j(t) \geq \frac{\rho_i c'_j}{8} t_1^{\lambda-1-\epsilon} \tag{255}$$

so $t_1 \in \mathcal{T}_j(c''_j,\rho_j)$ where $c''_j = \frac{\rho_i c'_j}{8}$ and $\rho_j = \rho_i/2$ (since $(\rho_i/4)t_1$ needs to be a ρ_j proportion of $t_1/2$).

Finally, we need to show that this probabilistic bound then implies that almost surely there are only finitely many t_1 which are in $\mathcal{T}_i(c_i,\rho_i)$ but not $\mathcal{T}_j(c''_j,\rho_j)$. This case only happens when $\sum_{t=t_1/2}^{(1-\rho_i/4)t_1} \psi_{j,t} \leq \frac{\rho_i c'_j}{8} t_1^{\lambda-\epsilon}$, and the probability of this (by (254)) decreases faster than any inverse polynomial of t_1, and hence has a finite sum over all $t_1 \in \mathcal{T}_i$ (notably, $\sum_{t_1=1}^{\infty} e^{-\frac{\rho_i}{32}c'_j t_1^{\lambda-\epsilon}} < \infty$, so summing only over $t_1 \in \mathcal{T}_i$ must also be finite). Thus, by the Borel-Cantelli Lemma, almost surely it happens only finitely many times, and we are done. □

Proposition 4. *For any agent i and constants $c_i, \rho_i > 0$, let $\mathcal{T}_i := \mathcal{T}_i(c_i,\rho_i)$. Then there is a set of constants $c_j, \rho_j > 0$ for all $j \neq i$ such that, almost surely, there is some $t^* > 0$ such that*

$$t_1 > t^* \text{ and } t_1 \in \mathcal{T}_i \implies t_1 \in \mathcal{T}_j \text{ for all } j \tag{256}$$

where $\mathcal{T}_j := \mathcal{T}_j(c_j,\rho_j)$.

Proof. This follows from Lemma 11 by induction over distance to i. Let $\text{dist}(i,j)$ denote the distance of vertex j from i, and $N_i(k) := \{j : \text{dist}(i,j) \leq k\}$. We can show that if the condition holds for all $j \in N_i(k)$, it holds for all $j \in N_i(k+1)$ as well, and therefore it holds for all $j \in N_i(n)$ (i.e. for all j, since no vertex can be more than n distance from i). □

It is important that the graph is finite, since then we only have a finite number of induction steps.

We also need that t_1 being a good time for agent i means that $\beta_i(t_1)$ also obeys a constant factor lower bound of order $t_1^{\lambda-1-\epsilon}$ (for all sufficiently large t_1):

Lemma 12. *For any $t_1 \geq 2/\rho_i$, if $t_1 \in \mathcal{T}_i(c_i,\rho_i)$, then*

$$\beta_i(t_1) \geq (c_i/2)t_1^{\lambda-1-\epsilon}. \tag{257}$$

Proof. This follows since $t_1 \in \mathcal{T}_i(c_i,\rho_i)$ and $t_1 \geq 2/\rho_i$ means that there is at least one $t \in [t_1/2, t_1]$ such that

$$\beta_i(t) \geq c_i t^{\lambda-1-\epsilon} \geq c_i t_1^{\lambda-1-\epsilon}. \tag{258}$$

But this means that

$$\beta_i(t_1) = \frac{\sum_{\tau=1}^{t_1}\psi_{i,\tau}}{t_1} \geq \frac{1}{2}\frac{\sum_{\tau=1}^{t_1}\psi_{i,\tau}}{t_1/2} \tag{259}$$
$$\geq \frac{1}{2}\frac{\sum_{\tau=1}^{t}\psi_{i,\tau}}{t} = \frac{1}{2}\beta_i(t) \geq (c_i/2)t_1^{\lambda-1-\epsilon}. \tag{260}$$

□

Finally, this yields the result that, almost surely, as $t \to \infty$, for all i, $\beta_i(t) = \tilde{O}(t^{\lambda-1})$ Formally, this is:

Proposition 5. *For any $i \in [n]$ and any $\epsilon > 0$, almost surely*

$$\lim_{t\to\infty}\frac{\beta_i(t)}{t^{\lambda-1+\epsilon}} = 0 \tag{261}$$
$$\lim_{t\to\infty}\frac{\beta_i(t)}{t^{\lambda-1-\epsilon}} = \infty \tag{262}$$

Proof. First, we can show (261) as a corollary of Theorem 5. Since $V(\boldsymbol{\beta}(t)) = \boldsymbol{v}^\top\boldsymbol{\beta}(t) = \sum_{i=1}^n v_i\beta_i(t)$ and each $v_i > 0$ is constant, for each i there is some c_i such that $\beta_i(t) \leq c_i V(\boldsymbol{\beta}(t))$ and therefore

$$\lim_{t\to\infty}\frac{\beta_i(t)}{t^{\lambda-1+\epsilon}} \leq \lim_{t\to\infty}\frac{c_i V(\boldsymbol{\beta}(t))}{t^{\lambda-1+\epsilon}} = 0\,. \tag{263}$$

Next, Lemma 10 shows that there is some $c > 0$ and time t_0 such that for all $t > t_0$, there is some i such that $t \in \mathcal{T}_i(c, 1/n)$. By Proposition 4, there is (almost surely) some time t^* and constants $c'_j, \rho_j > 0$ such that for any i, any $t > t^*$ and $j \in [n]$, we have $t \in \mathcal{T}_i(c,1/n) \implies t \in \mathcal{T}_j(c'_j,\rho_j)$. Thus, for all $t > \max(t_0,t^*)$, we have $t \in \mathcal{T}_j(c'_j,\rho_j)$ for all j. Finally, Lemma 12 shows this implies that there are constants $c_j > 0$ such that $\beta_j(t) \geq c_j t^{\lambda-1-\epsilon}$ for all $t > \max(t^*, t_0)$ and we are done. □

Finally, we can use this consensus rate to bound the convergence rate of the inherent belief estimator. If for each i, if $\beta_i(t) \geq ct^{\lambda-1-\epsilon}$, then $\mu_i(t) \geq ct^{\lambda-1-\epsilon}$. Using (128), we have that $x(t) \geq (\gamma_i - 1)\mu_i(t) \geq (\gamma_i-1)ct^{\lambda-1-\epsilon}$. Then

$$X(t) \geq (\gamma_i - 1)c\sum_{\tau=t_0}^{t}\tau^{\lambda-1-\epsilon} \approx c_1 t^{\lambda-\epsilon}. \tag{264}$$

Proposition 3 yields that $\mathbb{P}\left[\exists t \geq t^* : \widehat{\phi}_i(t) \neq \phi_i\right] \leq \delta$ if

$$t^* \geq \left(\frac{1}{c_1}\frac{1}{\xi_i}\log\frac{1}{\delta(e^{\xi_i}-1)}\right)^{1/(\lambda-\epsilon)} \tag{265}$$

(assuming that t^* is such that $c_1(t^*)^{\lambda-\epsilon} > 2$).

Compared to (158), we see that instead of a rate which is $1/\delta$ to some power, we get $\tilde{\Theta}(\log(1/\delta)^{1/\lambda})$ (since (265) holds for all $\epsilon > 0$), which is a big improvement. This, along with (159), yields the following theorem, suggesting that estimating inherent beliefs in consensus is more difficult:

Theorem 6. *For the inherent belief estimator $\widehat{\phi}_i(t)$ given in Definition 4, let t^* be defined as the time of convergence:*

$$t^* := \max(t : \widehat{\phi}_i(t) \neq \phi_i) \tag{266}$$

i.e. the first time such that $\widehat{\phi}_i(t) = \phi_i$ for all $t > t^*$. Then:

$$t^* = \begin{cases} O(\log(1/\delta)) & \text{if no consensus} \\ O(\log(1/\delta)^{\frac{1}{\lambda}+\epsilon}) \text{ for any } \epsilon > 0 & \text{if consensus} \end{cases} \tag{267}$$

with probability $\geq 1 - \delta$, where λ is the largest eigenvalue of $\mathbf{\Gamma W}$ if the consensus is to $\mathbf{0}$, and of $\mathbf{\Gamma}^{-1}\mathbf{W}$ if to $\mathbf{1}$.

Proof. This follows directly from (159) (for the non-consensus case) and from (265) (for the consensus case). Note that if something holds for an exponent of $\frac{1}{\lambda - \epsilon}$ for all $\epsilon > 0$, this is equivalent to holding for an exponent of $\frac{1}{\lambda} + \epsilon$ for all $\epsilon > 0$, so we can make the substitution. $\square$

X. Conclusion

In this work, we study the interacting Pólya urn model of opinion dynamics model under social pressure. We expanded upon [1] by showing there exists an estimator for bias parameters and inherent beliefs. Specifically, we showed that the history of any agent and their neighbors' declarations is sufficient in the limit to determine an agent's inherent belief and bias parameter for any network structure, using estimators based on maximum likelihoods. We also analyzed the rate at which the inherent belief estimator converges.

References

[1] Ali Jadbabaie, Anuran Makur, Elchanan Mossel, and Rabih Salhab, "Inference in opinion dynamics under social pressure," *IEEE Transactions on Automatic Control*, vol. 68, no. 6, pp. 3377–3392, 2023.

[2] Camilla Ancona, Francesco Lo Iudice, Franco Garofalo, and Pietro De Lellis, "A model-based opinion dynamics approach to tackle vaccine hesitancy," *Scientific Reports*, vol. 12, no. 1, pp. 11835, 2022.

[3] Damon Centola, Robb Willer, and Michael Macy, "The emperor's dilemma: A computational model of self-enforcing norms," *American Journal of Sociology*, vol. 110, no. 4, pp. 1009–1040, 2005.

[4] Jennifer Tang, Aviv Adler, Amir Ajorlou, and Ali Jadbabaie, "Stochastic opinion dynamics under social pressure in arbitrary networks," *arXiv preprint arXiv:2308.09275*, 2023.

[5] Morris H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[6] Noah E Friedkin and Eugene C Johnsen, "Social influence and opinions," *Journal of Mathematical Sociology*, vol. 15, no. 3-4, pp. 193–206, 1990.

[7] Daron Acemoğlu, Giacomo Como, Fabio Fagnani, and Asuman Ozdaglar, "Opinion fluctuations and disagreement in social networks," *Mathematics of Operations Research*, vol. 38, no. 1, pp. 1–27, 2013.

[8] Jason Gaitonde, Jon Kleinberg, and Eva Tardos, "Adversarial perturbations of opinion dynamics in networks," in *Proceedings of the 21st ACM Conference on Economics and Computation*, 2020, pp. 471–472.

[9] Mengbin Ye, Yuzhen Qin, Alain Govaert, Brian D.O. Anderson, and Ming Cao, "An influence network model to study discrepancies in expressed and private opinions," *Automatica*, vol. 107, pp. 371–381, 2019.

[10] Y. Polyanskiy and Y. Wu, "Lecture notes on information theory," *MIT (6.441), UIUC (ECE 563)*, 2013-2016.

[11] Joel Tropp, "Freedman's inequality for matrix martingales," *Electronic Communications in Probability*, vol. 16, 01 2011.

[12] David A. Freedman, "On Tail Probabilities for Martingales," *The Annals of Probability*, vol. 3, no. 1, pp. 100 – 118, 1975.

[13] Abraham Berman and Robert J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Society for Industrial and Applied Mathematics, 1994.

[14] David Williams, *Probability with Martingales*, Cambridge University Press, 1991.

[15] Rick Durrett, "Probability: Theory and examples, 4th edition," 2010, Cambridge University Press.

Jennifer Tang is a Postdoctoral Associate at the Institute of Data, Systems and Society (IDSS) and the Laboratory for Information and Decision Systems (LIDS) at MIT. She received her Ph.D and S.M in Electrical Engineering and Computer Science at MIT, and a B.S.E in Electrical Engineering from Princeton University. Her research interests include information theory, quantization and data compression, high-dimensional statistics and models for social dynamics and inference.

Aviv Adler is a Postdoctoral Scholar at AUTOLab at UC Berkeley. He received his Ph.D. and S.M. in the Department of Electrical Engineering and Computer Science at MIT as a member of LIDS, and his A.B. in Mathematics at Princeton University. Aviv's research interests include robotics, computational geometry, motion planning, optimization, complexity theory, and information theory.

Amir Ajorlou is a research scientist at the Laboratory for Information and Decision Systems (LIDS) at MIT. Previously, he was a postdoctoral research fellow at MIT and Penn from 2013 to 2016. He received his BS from Sharif University of Technology in Tehran, Iran and his PhD in electrical and computer engineering from Concordia University in Montreal, Canada, in 2013. He has been the recipient of several awards, including two gold medals in the International Mathematical Olympiad (IMO), Concordia University Doctoral Prize in Engineering and Computer Science, Governor General of Canada Academic Gold Medal, and the NSERC Postdoctoral Fellowship. He has served as an Associate Editor for the Conference Editorial Board of IEEE Control Systems Society from 2017 to 2022, and as Chair of the Control Systems Chapter of the IEEE Montreal Section in 2011. His current research uses a collection of tools and techniques from game theory, networks, microeconomics, optimization, applied probability theory, and control theory to model and analyze strategic decision making under uncertainty in complex social and economic networks.

Ali Jadbabaie is the JR East Professor and Head of the Department of Civil and Environmental Engineering at Massachusetts Institute of Technology (MIT), where he is also a core faculty in the Institute for Data, Systems, and Society (IDSS) and a Principal Investigator with the Laboratory for Information and Decision Systems. Previously, he served as the Director of the Sociotechnical Systems Research Center and as the Associate Director of IDSS which he helped found in 2015.

He received a B.S. degree with High Honors in electrical engineering with a focus on control systems from the Sharif University of Technology, an M.S. degree in electrical and computer engineering from the University of New Mexico, and a Ph.D. degree in control and dynamical systems from the California Institute of Technology. He was a Postdoctoral Scholar at Yale University before joining the faculty at the University of Pennsylvania, where he was subsequently promoted through the ranks and held the Alfred Fitler Moore Professorship in network science in the Department of Electrical and Systems Engineering. He is a recipient of a National Science Foundation Career Development Award, an US Office of Naval Research Young Investigator Award, the O. Hugo Schuck Best Paper Award from the American Automatic Control Council, and the George S. Axelby Best Paper Award from the IEEE Control Systems Society. He has been a senior author of several student best paper awards, in several conferences including ACC, IEEE CDC and IEEE ICASSP. He is an IEEE fellow, and the recipient of a Vannevar Bush Fellowship from the Office of Secretary of Defense. His research interests are broadly in systems theory, decision theory and control, optimization theory, as well as multiagent systems, collective decision making in social and economic networks, and computational social science.