

Differences of Opinion

Dionissi Aliprantis

November 30, 2022

Abstract: This paper presents a generalization of DeGroot updating in which agents learn from, rather than just ignore, biased opinions. Agents observe data consistent with a set of beliefs and turn to social learning to resolve the ambiguity they face. The solution to each agent's missing data problem uses disagreement across multiple dimensions to predict the bias in others' beliefs. In some cases belief updating will reduce to DeGroot updating. In other cases belief updating will be non-constricting and generate polarization, even on a connected network where everyone observes the same data and processes that data with the same model.

Keywords: Partial Identification, Imprecise Probability, Social Learning, DeGroot Updating, Predicted Bias, Bounded Confidence

JEL Classification Numbers: D81, D83, D84

Contact: Federal Reserve Bank of Cleveland, +1(216)579-3021, dionissi.aliprantis@clev.frb.org.

Acknowledgements: Many of the ideas in this paper originated from discussions with Alon Bergman and Gregorio Caetano. I also thank Ben Craig, Michalis Haliassos, Charles Manski, Jan-Peter Siedlerek, Alireza Tahbaz-Salehi, Teddy Seidenfeld, Michael Smithson, Johannes Hörner, two anonymous referees, seminar participants at the Cleveland Fed, Goethe University Frankfurt, University of Lyon/GATE, and ISIPTA '17 conference participants for helpful comments.

The opinions expressed are those of the author and do not represent views of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

1 Introduction

Economics is full of disagreements so fundamental that they can be characterized in terms of a coarse, binary model. Researchers cannot agree on whether extending unemployment benefits increases unemployment (Hagedorn et al. (2013), Farber et al. (2015)); government spending stimulates the economy (Ramey (2011), Serrato and Wingender (2016)); neighborhood of residence affects employment (Ludwig et al. (2013), Aliprantis and Richter (2020)); or parenting affects children’s economic outcomes (Stephens-Davidowitz (2022), García and Heckman (2022)).

The difficulty of intervening on the data generating process, along with constraints on time and resources, make the scientific ideal of “seeing for one’s self” an infeasible approach to belief formation on all topics. How, then, does one discern the evidence behind the disagreements above when the researchers communicate their conclusions? In the examples above, individuals are interpreting the same data with different models. When data are open to multiple interpretations and time or resource constraints prevent us from investigating further for ourselves, how do we decide which sources of information have interpreted the original data appropriately?

Figure 1 shows global warming as an example of how disagreements can be driven by distrust in some sources of information (Gessen (2019), Lace-Evans (2016)) rather than limited exposure to information (Gentzkow and Shapiro (2011)). Data arriving between 2006 and 2012 showed a continuation of the upward trend in average global temperature (Figure 1a). While a statistical interpretation of these data would have had a monotonic effect on beliefs, the non-monotonic path of beliefs over this same time period, as shown in Figure 1b, reflects heterogeneous levels of skepticism when evidence on global warming is presented by a government agency like the United States’ National Oceanic and Atmospheric Administration, politicians (Levine (2017)), or public relations firms (Bramoullé and Orset (2018), Upin (2012)).

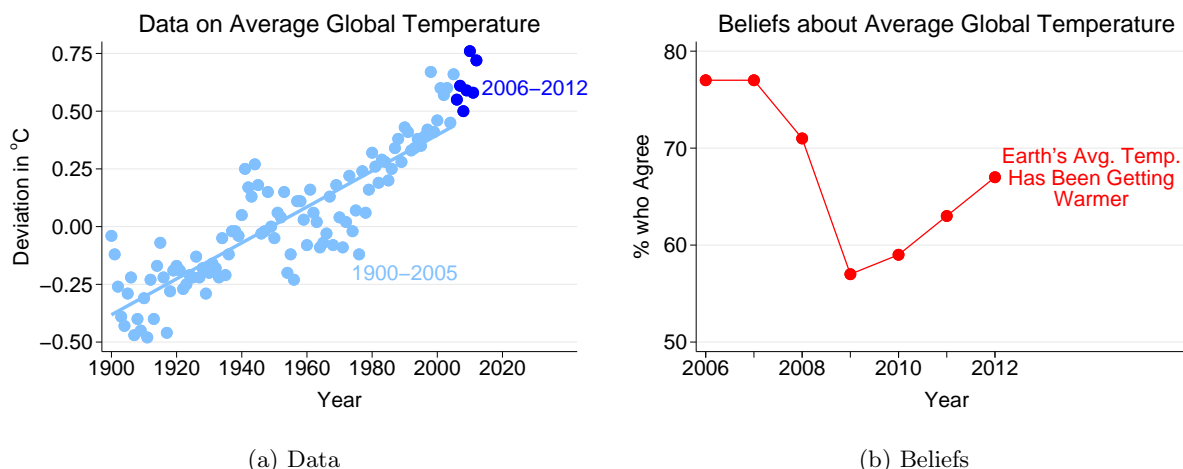


Figure 1: Average Global Temperature

Note: Data measuring average global temperature are from the United States’ National Oceanic and Atmospheric Administration (NOAA) and data on beliefs of US residents about average global temperature are from Pew surveys (Pew (2012)).

This paper studies belief formation when the data required to draw an unambiguous conclusion

are not observed and, instead, a set of beliefs is consistent with the data. When facing this ambiguity, I allow an agent’s belief formation to follow a common approach in decision theory, in that she chooses a single belief from her set of possible beliefs. One difference from the standard approach is that the agent’s choice maximizes an objective function based on solving her missing data problem, not based on generating an extreme utility.¹ Another difference from the standard approach is that by placing the agent in a social setting, she is able to use social learning to choose the one belief solving her missing data problem. The agent’s social learning problem can be stated in terms of predicting the bias of the model each individual in her network uses to interpret data relative to her own model.²

The first contribution of this paper is to show that DeGroot updating can be rationalized as a solution to the agent’s problem, and more generally to show that models of non-Bayesian social learning can be linked to the partial identification and ambiguity aversion literatures. In DeGroot updating a group of individuals holds subjective probability F_i , which denotes each person i ’s belief about a proposition. Updating is achieved via linear opinion pooling as $F_i^t = \sum_{j=1}^J p_{ij} F_j^{t-1}$ when $\{p_{ij}\}_{j=1}^J$ are nonnegative numbers such that $\sum_{j=1}^J p_{ij} = 1$ for each individual i , and can be written for the network as $\mathbf{F}^t = \mathbf{P}\mathbf{F}^{t-1}$.

DeGroot updating solves the agent’s problem when the agent places zero weight on biased individuals. In the context of the agent’s problem, this means that DeGroot updaters must choose between interpreting differences of opinion as resulting from either sampling error or bias. An agent interprets a difference of opinion as sampling error by placing positive weight $p_{ij} > 0$ on individual j ’s opinion. Placing no weight $p_{ij} = 0$ on individual j interprets their opinion as a result of bias. Thus, the agent’s problem offers a clear interpretation to the zero weights given in many non-Bayesian social learning rules.

The second contribution of this paper is to use the agent’s missing data problem to motivate the prediction of bias in bounded confidence models and to show how this prediction can lead to new updating rules. Choosing between positive and zero weight requires that the agent has some way of predicting an individual’s bias. In the bounded confidence models studied in Hegselmann and Krause (2002), for example, an individual is predicted to be biased, and therefore ignored, if their opinion is too far away from the agent’s opinion. The agent’s missing data problem suggests that predicted bias might not only allow the agent to choose which individuals to ignore, but might also allow the agent to learn from individuals whose bias can be accurately predicted. That is, when the agent considers an individual biased, she may move her beliefs away from their reported beliefs. I assume the agent partitions disagreement on a given proposition into sampling error and bias using differences of opinion across multiple other propositions. This mechanism allows for polarization through an updating rule that need not lead to constricting belief updating

¹For example, the maxmin expected utility decision rule maximizes expected utility after choosing the belief that would be set by a malevolent nature minimizing the agent’s utility for any decision (Gilboa and Schmeidler (1989)). The minimax regret decision rule maximizes expected utility after choosing the belief maximizing the agent’s lost utility from not knowing the true state of the world (Manski (2011)).

²Or, equivalently, the degree to which each individual samples unrepresentative data.

(Mueller-Frank (2015)).

The third contribution of this paper is to study a specific example of a predicted bias updating rule, motivated as a solution to the agent’s problem, that is capable of generating polarization and can sustain clustered disagreement, even on a connected network where everyone directly observes the same data and processes that data with the same model. This capacity of the predicted bias updating rule is a desideratum of the literature on non-Bayesian social learning for its ability to reconcile theory and evidence (Golub and Sadler (2016)).³ Two important factors for generating polarization are low-quality data and strong beliefs that different models are used to interpret data, which allow social learning to amplify disagreement beyond that predicted by individuals using different models to interpret data privately.

The generalization of bounded confidence updating rules presented in this paper helps to understand the contours of the learning environments capable of generating polarization. Some contrasts with related papers are particularly illustrative. Loh and Phelan (2019) study an environment where ambiguity arises due to the dimensionality of the learning, whereas here the dimensionality of the learning is used to overcome ambiguity. The assumptions in Sethi and Yildiz (2016) are in many ways the mirror to the assumptions here. For example, Sethi and Yildiz (2016) consider environments where learning about a proposition happens quickly and learning about perspectives (ie, bias) happens slowly. Here the opposite is true for both aspects of the environment, as learning about a proposition takes time but assessments of bias happen quickly. Sethi and Yildiz (2016) are focused on the tradeoff between well-understood and well-informed individuals, whereas here the inability to communicate how well-informed one is means that agents may update based on well-understood but poorly-informed individuals.

The paper proceeds as follows: Section 2 describes the agent’s belief formation under point and partial identification. Section 2 then states the agent’s problem as using social learning to infer the missing data that would have generated point identification. Section 3 states the assumptions under which DeGroot and predicted bias updating can solve the agent’s problem. Section 4 illustrates some of the belief dynamics possible under predicted bias updating. Section 5 concludes.

2 Belief Formation as a Missing Data Problem

Suppose there is a set of $k = 1, \dots, K$ propositions, none of which can be written as a compound proposition using other propositions in the set.⁴ Agent i ’s beliefs at time t are $F_i^{k,t}$, and the agent directly observes data $W_i^{k,t} \in \mathbb{R}^q$ sampled under an identical procedure at each point in discrete time.

³While DeGroot learning and many of its generalizations converge to a degenerate distribution for connected networks (Jackson (2008)), we often observe the analogue of a connected network – individuals exposed to sources of information contradicting their beliefs – together with persistent disagreement.

⁴This greatly simplifies the analysis. Al-Najjar (2009) studies the statistical implications of relaxing this assumption, and Paris and Vencovská (1990) and Wilmers (2010) study propositional calculus without this assumption.

2.1 No Missing Data: The Ideal Case

A key scientific ideal is that belief formation be based on direct observation.⁵ We can state this ideal of “seeing for one’s self” in terms of the agent forming her beliefs based on directly-observed data. In this case, the agent uses the model φ_i^k to interpret data as the signal $\sigma_i^{k,t} = \varphi_i^k(W_i^{k,t}) \in [0, 1]$. We can think of $\sigma_i^{k,t}$ as a parameter of the agent’s model estimated on the data they observed. Assuming $\{W_i^{k,t}\}$ is an independent and identically distributed (iid) sequence of random variables and φ_i^k is continuous, $\{\sigma_i^{k,t}\}$ will also be an iid sequence (White (2000), Proposition 3.2), with distribution denoted by Γ_i^k . The law of large numbers ensures convergence to the mean of the signal distribution, or $F_i^{k,t} \xrightarrow{a.s.} \mu_i^k \equiv \mathbb{E}[\sigma_i^{k,t}]$, for beliefs formed as

$$F_i^{k,t} = \frac{1}{t} \sum_{n=1}^t \sigma_i^{k,n} = \beta_t \sigma_i^{k,t} + (1 - \beta_t) F_i^{k,t-1} \quad \text{where} \quad \beta_t = 1/t. \quad (1)$$

2.2 Missing Data: The Case with Ambiguity

There is an abundance of economic applications in which the observed data do not allow for clear identification of signals. The Introduction discussed several examples of causal propositions where the data necessary for drawing unambiguous conclusions is missing.

Inspired by the literature on partial identification (Manski (2007), Tamer (2010)), now consider a setting in which the agent’s directly-observed data $W_i^{k,t}$ only allows her to determine a set containing the true signal, $\sigma_i^{k,t} \in [\underline{\sigma}_i^{k,t}, \overline{\sigma}_i^{k,t}]$. The agent can determine a point estimate $\widehat{\sigma}_{ii}^{k,t}$ from the data, but doing so requires less than credible assumptions, so the agent cannot be sure that $\mathbb{E}[\widehat{\sigma}_{ii}^{k,t}] = \mu_i^k$. We will refer to the quality of a signal as $\theta_i^{k,t}$ where $(\widehat{\sigma}_{ii}^{k,t}, \theta_i^{k,t}) = \varphi_i^k(W_i^{k,t}) \in [0, 1] \times [0, 1]$. In this case the agent holds a set-valued belief

$$\bar{F}_i^{k,t} = \left[\frac{1}{t} \sum_{n=1}^t \underline{\sigma}_i^{k,n}, \frac{1}{t} \sum_{n=1}^t \overline{\sigma}_i^{k,n} \right]$$

often referred to as an “imprecise probability” (Coolen et al. (2011)). The set $[\underline{\sigma}_i^{k,t}, \overline{\sigma}_i^{k,t}]$ is what can be learned about proposition k from the directly-observed data under the most credible assumptions, where $[\underline{\sigma}_i^{k,t}, \overline{\sigma}_i^{k,t}] \equiv [\max\{0, \widehat{\sigma}_{ii}^{k,t} - (1 - \theta_i^{k,t})\}, \min\{\widehat{\sigma}_{ii}^{k,t} + (1 - \theta_i^{k,t}), 1\}]$. Note that $\widehat{\sigma}_{ii}^{k,t} = \sigma_i^{k,t}$ when $\theta_i^{k,t} = 1$.

While the Data Generating Process could generate partial identification, preventing the observation of all relevant states of the world, set-valued beliefs could also be driven by time or resource constraints. In a model of rational inattention, beliefs taking the unit interval, or $[\underline{\sigma}_i^{k,t}, \overline{\sigma}_i^{k,t}] = [0, 1]$, could be interpreted as the default value of beliefs that spontaneously comes to mind with no think-

⁵The Royal Society, for example, expresses this scientific ideal through their motto ‘Nullius in verba,’ which “is taken to mean ‘take nobody’s word for it’. It is an expression of the determination . . . to withstand the domination of authority and to verify all statements by an appeal to facts determined by experiment.” The agent’s approach to belief formation does not meet [all scientific ideals for belief formation](#); updating beliefs with a fixed model φ_i^k is one reason.

ing (Gabaix (2019)) or as the result of proposition k being outside of the agent’s consideration set (Caplin et al. (2019)).

2.3 Solving the Missing Data Problem Via Social Learning

When holding beliefs represented by an imprecise probability, several approaches to decision making can be interpreted as picking one belief from the set

$$\widehat{F}_i^{k,t} \in \bar{F}_i^{k,t}$$

and then using this probability as a subjective belief with which to make decisions following the Savage axioms. The chosen probability is typically pessimistic, assuming the worst case in some sense of utility. For example, the Γ -maxmin utility decision rule maximizes expected utility after choosing the belief that would be set by a malevolent nature minimizing the agent’s utility for any decision (Gilboa and Schmeidler (1989)). Similarly, the Γ -minimax regret decision rule chooses the single belief that maximizes the loss from making decisions with the chosen belief rather than the true probability when the agent makes decisions to minimize this loss (Manski (2011)).⁶

The subsequent model explores belief formation when the agent chooses one belief $\widehat{F}_i^{k,t} \in \bar{F}_i^{k,t}$ using information from her social network. Decision makers often infer missing data from social observations when they must form beliefs with incomplete data. The literature provides many examples of this behavior, but it can also be observed everyday in internet forums providing information on [investments](#), [real estate](#), [consumer goods](#), [computer programs](#), or [open-ended questions](#).⁷

Suppose the agent is a member of an exogenous network of J individuals from which she might gather information, and we denote the set of others in the agent’s network as \mathcal{J} . Imperfect communication means that individuals in the agent’s social network can only communicate their previous beliefs and their interpretation of the data they observe, $\widehat{\sigma}_{jj}^{k,t}$. Individuals cannot directly report the data they observe, the model they use to interpret data, or any measure of ambiguity like $\bar{F}_j^{k,t-1}$ or $\theta_j^{k,t}$. This means that the agent observes information $\mathcal{I}_i^t \equiv (\mathcal{I}_{ii}^t, \mathcal{I}_{i\mathcal{J}}^t)$ where

$$\mathcal{I}_{ii}^t \equiv \left\{ \left(\widehat{F}_i^{k,t-1}, W_i^{k,t}, \varphi_i^k \right) \right\}_{k=1}^K = \left\{ \left(\widehat{F}_i^{k,t-1}, \widehat{\sigma}_{ii}^{k,t}, \theta_i^{k,t} \right) \right\}_{k=1}^K \quad \text{and} \quad \mathcal{I}_{i\mathcal{J}}^t \equiv \left\{ \left\{ \widehat{F}_j^{k,t-1}, \widehat{\sigma}_{jj}^{k,t} \right\}_{j \in \mathcal{J}^k} \right\}_{k=1}^K.$$

I assume that the agent uses socially-observed information in an effort to replicate the classical inference in Equation 1. This approach to belief formation is based in the scientific ideal of seeing for one’s self and questioning authority, and can be seen as an expression of epistemic vigilance (?). The agent’s problem is to estimate $\sigma_i^{k,t}$ as $\widehat{\sigma}_i^{k,t}$ each period so that the resulting inferred beliefs $\widehat{F}_i^{k,t}$ are identical to the beliefs $F_i^{k,t}$ formed with the ideal data used to generate $\sigma_i^{k,t}$. Given a loss

⁶Appendix B illustrates these decision rules in a stylized binary model.

⁷Examples from the literature include neighborhood and school choice (de Souza Briggs et al. (2008)); residential mobility (Brown and Moore (1970), Krysan (2008)); adoption of a new technology (Conley and Udry (2010), Foster and Rosenzweig (1995)); consumption of a new good (Moretti (2011)); investment decisions (Bursztyn et al. (2014)); mortgage refinance decisions (McCartney and Shah (2022)) retirement savings decisions (Beshears et al. (2015)); and choice of health insurance plans (Sorensen (2006)).

function \mathcal{L} , the agent's problem is to choose functions g^k from some set \mathcal{G} to solve the problem

$$\begin{aligned} & \min_{g^1, \dots, g^K \in \mathcal{G}} \sum_{k=1}^K \mathcal{L} \left(\mathbb{E} \left[\widehat{F}_i^{k,t} - \mu_i^k \right] \right) & (2) \\ \text{s.t. } & \widehat{F}_i^{k,t} \in \bar{F}_i^{k,t} \\ & \widehat{\sigma}_i^{k,t} = g^k(\mathcal{I}_i^t) \\ & \widehat{F}_i^{k,t} = \beta_t \widehat{\sigma}_i^{k,t} + (1 - \beta_t) \widehat{F}_i^{k,t-1}. \end{aligned}$$

3 Two Solutions to the Agent's Problem

A natural restriction on \mathcal{G} is for the agent to infer the signal $\widehat{\sigma}_i^{k,t}$ using linear opinion pooling (LOP):

$$\widehat{\sigma}_i^{k,t} = \sum_{j=1}^J p_{ij}^{k,t} \widehat{\sigma}_{ij}^{k,t} \quad \text{with } p_{ij}^{k,t} \geq 0 \quad \forall j \in \mathcal{J}, \quad \sum_{j=1}^J p_{ij}^{k,t} = 1 \quad (\text{R1})$$

Although undesirable properties of LOP have been documented (Seidenfeld et al. (2010), Bradley (2017), Ranjan and Gneiting (2010), Lichtendahl et al. (2017)), I investigate LOP for two reasons. First, LOP is commonly used when facing problems like the agent's problem (where fully Bayesian updating may not be feasible). And second, I investigate LOP because it results in DeGroot updating in a special case of data observation.

Under Restriction R1, the agent's problem in Equation 2 can be restated as predicting the bias in a sender's signal (choosing g_i^k) and assessing one's confidence in that prediction (choosing $p_{ij}^{k,t}$):

$$\begin{aligned} & \min_{\{g_i^k\}_{k=1}^K, \{p_{ij}^{k,t}\}_{k=1}^K} \sum_{k=1}^K \mathcal{L} \left(\mathbb{E} \left[\widehat{F}_i^{k,t} - \mu_i^k \right] \right) & (3) \\ \text{s.t. } & \widehat{\sigma}_{ij}^{k,t} = g_i^k(\mathcal{I}_i^t) \in [\underline{\sigma}_i^{k,t}, \overline{\sigma}_i^{k,t}] \\ & \widehat{\sigma}_i^{k,t} = \sum_{j=1}^J p_{ij}^{k,t}(\mathcal{I}_i^t) \widehat{\sigma}_{ij}^{k,t} \\ & \widehat{F}_i^{k,t} = \beta_t \widehat{\sigma}_i^{k,t} + (1 - \beta_t) \widehat{F}_i^{k,t-1}. \end{aligned}$$

3.1 DeGroot Updating

Consider an additional restriction on \mathcal{Q} so that agent i must interpret j 's signal as being unbiased

$$\widehat{\sigma}_{ij}^{k,t} = \widehat{\sigma}_{jj}^{k,t} \quad (\text{R2})$$

along with the following assumptions:

A0a Data are only observed at time $t = 0$ and beliefs at time $t = 1$ are based only on directly-observed data;

A0b Beliefs and signals are interchangeable, or $\hat{\sigma}_{ii}^{k,t} = \hat{F}_i^{k,t-1}$ and $\hat{\sigma}_{jj}^{k,t} = \hat{F}_j^{k,t-1}$;

A1 The agent's point estimates are unbiased, $\mathbb{E}[\hat{\sigma}_{ii}^{k,0}] = \mu_i^k$;

A2 The agent can predict bias, or $\mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$.

Proposition 1 in Appendix A shows that under Restrictions R1-R2 and Assumption A0, the agent's updating can be written as DeGroot updating where $\hat{\mathbf{F}}^{k,t} = \mathbf{\Omega}^{k,t} \hat{\mathbf{F}}^{k,t-1}$. Proposition 2 in Appendix A shows that under Restrictions R1-R2 and Assumptions A0-A2, this DeGroot updating solves the agent's problem. Note that an implication of Restriction R2 and Assumption A2 together is that the agent will choose $p_{ij}^{k,t}(\mathcal{I}_i^t) > 0 \iff \mu_i^k = \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$.

3.2 Predicted Bias Updating: A Generalization of DeGroot Updating

Assumption A2 requires that the agent has some way of predicting the bias $\mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$ of each individual's signal. This assumption is made in the bounded confidence generalization of DeGroot updating studied in Hegselmann and Krause (2002). In the context of the agent's problem, the assumption in those models is that $\mu_i^k \neq \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$, and therefore $p_{ij}^{k,t}(\mathcal{I}_i^t) = 0$, if $|\hat{F}_i^{k,t} - \hat{F}_j^{k,t}| \geq \epsilon$ for some $\epsilon > 0$.

Having the ability to predict the bias $\mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$ would suggest weakening restriction R2 to allow the agent to interpret an individual's signal in terms of its expected bias

$$\hat{\sigma}_{ij}^{k,t} = \hat{\sigma}_{jj}^{k,t} + \mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]. \quad (\text{R2}^*)$$

Proposition 3 in Appendix A shows that under Restrictions R1-R2* and Assumptions A1-A2, linear opinion pooling using R2* solves the agent's problem in Equation 3 for any choice of $p_{ij}^{k,t}(\mathcal{I}^t)$ generating a convex combination.

4 Opinion Dynamics in Predicted Bias Updating

4.1 Choosing g_i^k (Predicted Bias)

The justification for adopting Assumption A2 is often tenuous, whether adopted to rationalize DeGroot updating or its generalization developed in R2*. While similarly tenuous assumptions have been required in previous efforts to provide microeconomic foundations to DeGroot updating (Golub and Sadler (2016)), the arbitrariness must nevertheless be noted here. Here I show how one approach to satisfying Assumption A2 can create polarization under the generalized version of DeGroot updating that solves the agent's problem.

How might the agent predict the bias $\mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$? If each individual j aside from the agent observes ideal data $W_j^{k,t}$ allowing them to point-identify signals with distributions Γ_j^k , then $\mu_i^k \neq \mathbb{E}[\hat{\sigma}_{jj}^{k,t}]$ either because of biased models ($\varphi_i^k \neq \varphi_j^k$) or biased sampling ($\Gamma_i^k \neq \Gamma_j^k$). In this case,

$$\mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,t}] = \mathbb{E}[F_i^{k,t-1}] - \mathbb{E}[\hat{F}_j^{k,t-1}].$$

The agent might therefore interpret signals as

$$\hat{\sigma}_{ij}^{k,t} = \begin{cases} \hat{\sigma}_{jj}^{k,t} + (\hat{F}_i^{k,t-1} - \hat{F}_j^{k,t-1}) & \text{if } \hat{\sigma}_{jj}^{k,t} + (\hat{F}_i^{k,t-1} - \hat{F}_j^{k,t-1}) \in [0, 1]; \\ 0 & \text{if } \hat{\sigma}_{jj}^{k,t} + (\hat{F}_i^{k,t-1} - \hat{F}_j^{k,t-1}) < 0; \\ 1 & \text{if } \hat{\sigma}_{jj}^{k,t} + (\hat{F}_i^{k,t-1} - \hat{F}_j^{k,t-1}) > 1. \end{cases} \quad (\text{R2}^*)$$

Note that R1-R2* allows for moving beliefs away from a signal, which violates the Monotonicity assumption in Molavi et al. (2017), and characterizes the difference between this model's heterogeneous confidence learning rule and bounded confidence models like those developed in Hegselmann and Krause (2002) or Sotiropoulos et al. (2015).

4.2 Choosing $p_{ij}^{k,t}$ (Assigning Confidence to Predicted Bias)

Given R2* as the choice of $g_i^k(\mathcal{I}_i^t)$ for interpreting signals from other individuals j in her network, how should the agent make the choice of weights $p_{ij}^{k,t}$ indicating the strength of the agent's confidence in the resulting interpretation $\hat{\sigma}_{ij}^{k,t}$? Define Δ_{ij}^t as the credibility that the agent i deems to R2* as applied to individual j 's signal, with $\sum_{j \in \mathcal{J}^k} \Delta_{ij}^t = 1$. Noting that $j \in \mathcal{J}^k$ implies $i \neq j$, the agent would assign weights to interpreted signals as

$$p_{ii}^{k,t} = \theta_i^{k,t} \quad (\text{R3}^*\text{-i})$$

$$p_{ij}^{k,t} = (1 - \theta_i^{k,t}) \Delta_{ij}^t \quad (\text{R3}^*\text{-j})$$

The final hurdle to empirically implementing the agent's model of social learning is empirically determining the credibility the agent gives to R2* as a means of adjusting a signal from sender j , Δ_{ij}^t . Any inductive inference requires invariance assumptions that may not be true; this is the problem of induction. An invariance assumption the agent could invoke to assess credibility would pertain to differences of beliefs across propositions:⁸

$$\mathbf{(A3)} \quad \mu_i^{k,t} - \mu_j^{k,t} = \mu_i^{k',t} - \mu_j^{k',t} \quad \text{for all } t \in \mathbb{N}, j \in \mathcal{J}, \text{ and } k, k' \in \{1, \dots, K\}.$$

Under A3, the agent might use the distribution of $\delta_{ij}^{k,t} \equiv \hat{F}_i^{k,t} - \hat{F}_j^{k,t}$ over the proposition space as a means of assigning credibility to R2* applied to individual j . Many measures could be used to characterize this distribution, like root-mean-squared-error, and here I will suppose the agent uses relative entropy as developed in information theory for measuring uncertainty.

Figure 2 helps to illustrate the idea of (relative) entropy using the notion of Shannon entropy from information theory (Shannon (1948)). The agent would be most informed about individual j if the distribution of $\delta_{ij}^{k,t}$ were a degenerate distribution, and would be least informed were $\delta_{ij}^{k,t}$ to follow a uniform distribution. In the Figure, $f(\delta_{i \max}^{k,t}) = U[-1, 1]$ has the maximum entropy (representing the least informative sender), senders $j = 1, 2, 3$ have high entropy (representing low

⁸We could imagine other invariance assumptions, including ones in which disagreement over any proposition is proportional to disagreement over a set of predictive propositions. Such a generalization of A3 could help to explain why so much attention is paid to seemingly unimportant propositions in the US' "culture wars."

information), senders $j = 4, 5, 6$ have medium entropy (representing moderate information), and senders $j = 7, 8, 9$ have low entropy (representing high information).

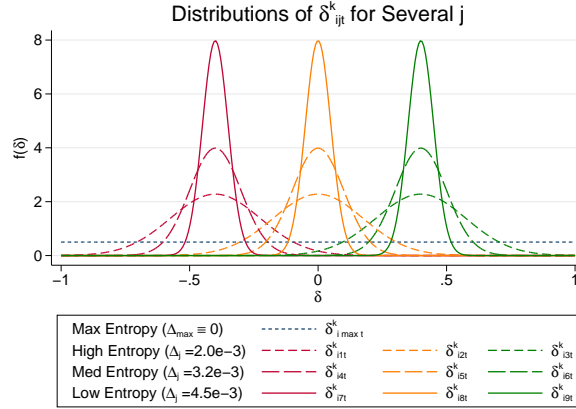


Figure 2: Entropy of Sender j

Although senders $j = 2$ and $j = 5$ have the same average disagreement across propositions with the agent, the agent will give more weight to interpreted signals from sender $j = 5$ because their disagreement has lower entropy than that of sender $j = 2$. That is, the agent is more certain about how she will disagree with sender $j = 5$. On the other hand, note that while the agent expects to disagree differently with senders $j = 1$, $j = 2$, and $j = 3$, the agent deems interpreted signals from these senders to be equally credible because they all have the same entropy. This helps to illustrate that what matters to the agent’s judgment about credibility is not her average disagreement with a sender, but the signal to noise ratio in her disagreement with a sender (Sethi and Yildiz (2016)).

Assuming that the distribution of $\delta_{ij}^{k,t}$ across propositions $k \in \{1, \dots, K\}$ has a probability density function Q_{ij}^t , the informational content of sender j ’s signal (or certainty about $\delta_{ij}^{k,t}$) can be defined as the Kullback-Leibler divergence from the uniform distribution over $[-1, 1]$:

$$D_{KL}(Q_{ij}^t : U) = \int_{-1}^1 Q_{ij}^t(\delta) \log \left(\frac{Q_{ij}^t(\delta)}{1/2} \right) d\delta,$$

which is a measure of the difference in entropy of Q_{ij}^t relative to the maximum entropy (uncertainty) distribution.⁹ I assume that the credibility the agent assigns to interpreted signals from sender j is

$$\Delta_{ij}^t \equiv \rho(D_{KL}(Q_{ij}^t : U)) = [\gamma_1 D_{KL}(Q_{ij}^t : U)]^{\gamma_2} \quad (\text{R3}^*-\Delta)$$

where $(\gamma_1, \gamma_2) \in [0, \infty) \times [0, \infty)$ can be thought of as distrust parameters.

⁹In another setting, Zanardo (2017) shows that Kullback-Leibler relative entropy satisfies desirable axioms as a measure of disagreement.

4.3 Simulations for 1 Proposition and a 2-Individual Network

To illustrate how a predicted bias updating rule can generate polarization, consider here a simple example given in DeGroot (1974) of two individuals learning about one proposition. This example was originally chosen to illustrate the conditions on the weighting matrix \mathbf{P} under which individuals' beliefs will reach a consensus. DeGroot (1974) notes that if beliefs are updated as $\mathbf{F}^t = \mathbf{P}\mathbf{F}^{t-1}$, beliefs will ultimately coincide if $\mathbf{P} = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix}$ but will remain different when

$$\mathbf{P} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

To see how predicted bias updating can generate polarization, now suppose that each element of \mathbf{P} is $\frac{1}{2}$, but that agent i judges any difference of opinion with individual j as bias. Recalling Equation R2*, this implies that i would interpret j 's beliefs as $\hat{F}_{ij} = F_j + (F_i - F_j)$ to account for their bias. If two individuals were to update beliefs as $\mathbf{F}^t = \mathbf{P}\hat{\mathbf{F}}^{t-1}$, Figure 3 shows that even small initial differences of opinion would lead beliefs to diverge.

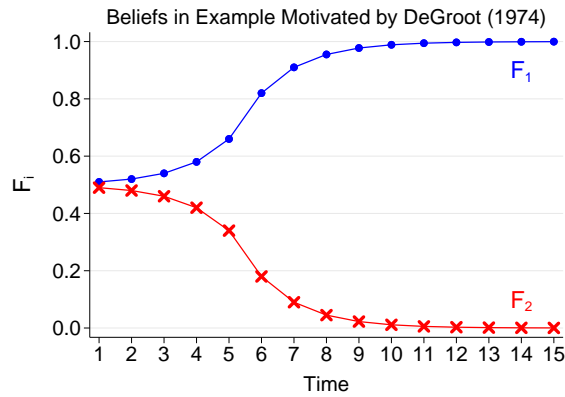


Figure 3

Even though predicted bias updating tends to generate consensus, the possibility of polarization is opened because the updating rule need not lead to constricting belief updating (Mueller-Frank (2015)).

4.4 Simulations for 30 Propositions and a 300-Individual Network

I now show simulations illustrating some of the belief dynamics that solving the agent's problem in Equation 3 via R1, R2*, and R3* can generate. I first show that beliefs can polarize and clusters can be sustained in a relatively simply setting. I then show that this result is not an artifact of the simple setting by replicating it in more nuanced settings.

In each numerical experiment, I consider a network of $J = 300$ individuals learning about $K = 30$ propositions. A key parameter tilting the simulations toward polarization is setting the quality of the signal from directly-observed data to be low, so that $\theta_i^{k,t} = 0.1$ for all k, t , and

i. Each agent assesses the credibility of R2* applied to sender j 's signal according to $\Delta_{ij}^t = [\gamma_1 D_{KL}(Q_{ij}^t : U)]^{\gamma_2}$ where the distrust parameters are $(\gamma_1, \gamma_2) = (100, 8)$. While the ambiguity of each individual's directly-observed data is high, tilting the simulations toward agreement, or convergence to a degenerate distribution, is an assumption that all individuals receive the same signal over and over again. I assume all individuals directly-observe data generating the same signals at all times for all propositions:

$$\hat{\sigma}_{ii}^{k,t} = 0.5 \quad \forall \quad k = 1, \dots, K; \quad t = 2, 3, 4, \dots; \quad \text{and} \quad i = 1, \dots, J.$$

What varies across Experiments 1, 2, and 3 is the initial distribution of beliefs at $t = 1$. In Experiment 1, I assume there are two clusters of individuals, \mathcal{C}_1 and \mathcal{C}_2 , with $\text{card}(\mathcal{C}_1) = 100$ and $\text{card}(\mathcal{C}_2) = 200$. Initial beliefs $\hat{F}_i^{k,1}$ are generated as follows:

$$\ddot{F}_i^k \sim \begin{cases} \mathcal{N}(0.8, 0.1) & \text{if } i \in \mathcal{C}_1 \quad \forall k = 1, \dots, K; \\ \mathcal{N}(0.2, 0.01) & \text{if } i \in \mathcal{C}_2 \quad \forall k = 1, \dots, K. \end{cases}$$

with

$$\hat{F}_i^{k,1} = \min \left\{ \max \left\{ \ddot{F}_i^k, 0 \right\}, 1 \right\} \quad \forall k = 1, \dots, K.$$

Figure 4a shows the initial distribution of beliefs by clusters for proposition p^1 .

Figure 4b shows that as individuals in the network update using R1, R2*, and R3*, the updating maintains the clustering, with \mathcal{C}_1 and \mathcal{C}_2 still clearly distinguishable from one another. The basic idea is that if a given agent in Cluster 1, \mathcal{C}_1 , tends to agree with those in this widely-distributed cluster (unbiased but imprecise), but tends to disagree with those in the tightly-distributed Cluster 2, \mathcal{C}_2 (biased but precise), that agent will rely more on interpreted signals from the disagreeing Cluster 2, and this can cause her to overcompensate when they provide her with unbiased signals. This appears as the Cluster 1 steady state distribution being further from the (common and repeated) directly-observed signal of 0.5 than the Cluster 2 steady state.

This is what happens to agents in cluster \mathcal{C}_1 in Experiment 1 (Figure 4c). Given their disagreement, for an individual in \mathcal{C}_1 the interpretation is that any signal from someone in \mathcal{C}_2 must, on average, be an understatement of $\sigma_i^{1,t}$, so the interpreted signal $\hat{\sigma}_{ij}^{1,t}$ adjusts the received signal upwards. Since the disagreement for individuals within \mathcal{C}_1 is more uncertain than the disagreement across individuals in \mathcal{C}_1 and \mathcal{C}_2 , individuals in \mathcal{C}_1 give more credibility to the interpreted signals from those in \mathcal{C}_2 . As a result, individuals in \mathcal{C}_1 overcompensate and move their beliefs away from 0.5. Note that while all agents form beliefs in the same way, initial conditions act as a mechanism for generating behavior like the expectation types documented in Dominitz and Manski (2011).

Appendix C shows two further experiments illustrating that this type of clustered polarization is not an artifact of having two clusters, and that there are possibilities for interesting questions about which clusters polarize and which ones converge. In these experiments as well we see the repelling force of the perception of bias.

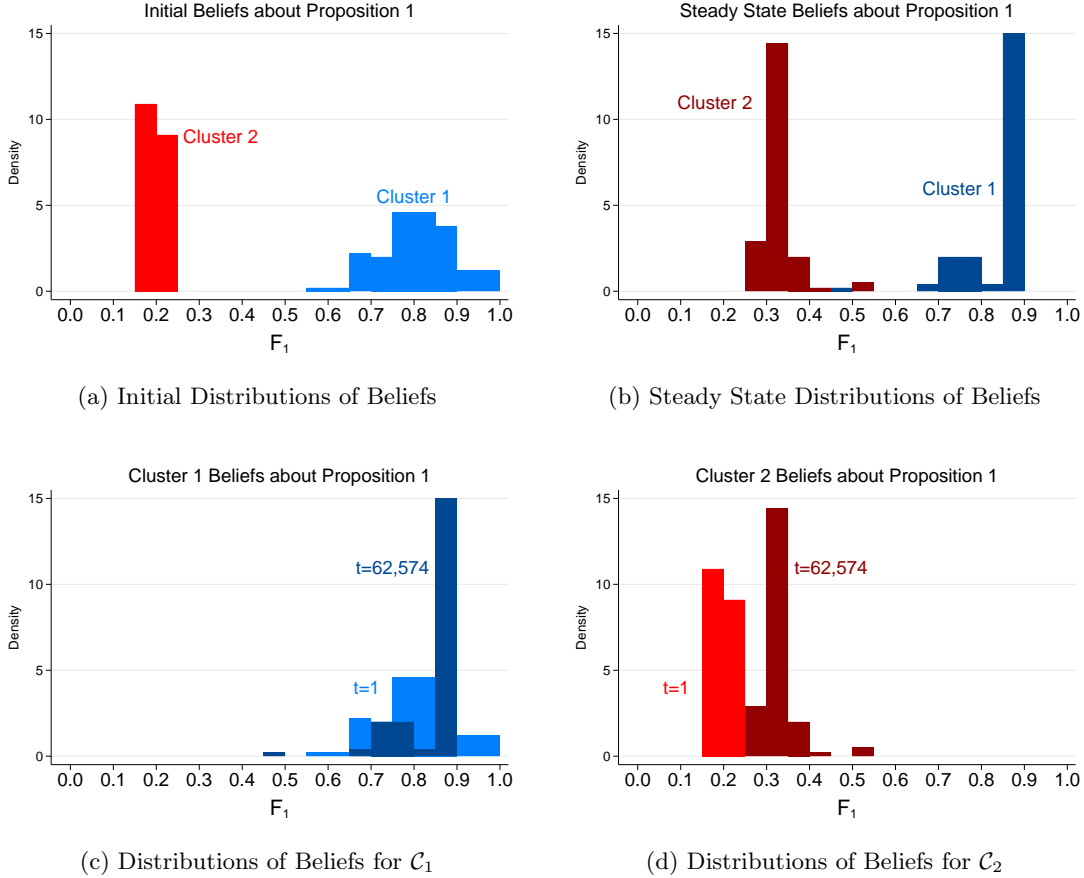


Figure 4: Beliefs in Experiment 1

Steady state belief distributions are those for which $\max_{k,i} \|\widehat{F}_i^{k,t} - \widehat{F}_i^{k,t-1}\| < 1e-6$, which occurs in Experiment 1 at $t = 62,574$.

5 Conclusion

Models of non-Bayesian social learning typically assume that agents ignore evidence when provided by biased sources. But the choice of whether to ignore a source suggests that agents have a way to predict that source’s bias. This paper developed a model of non-Bayesian social learning in which agents use predicted bias to learn from biased sources rather than just to ignore them.

The model of social learning is motivated by an agent who must solve a missing data problem. The agent faces ambiguity because her directly-observed data only allow her to partially identify a signal about the truth of a proposition. I proposed social learning as a way that the agent might choose a single subjective probability from a set of possible probabilities.

I showed that DeGroot updating can solve a special case of the agent’s problem. I also showed that a generalization of DeGroot updating, which I call predicted bias updating, can solve the agent’s problem. Predicted bias is based on differences of opinion across multiple propositions under an assumption of invariant bias across propositions. I showed that although predicted bias updating will often lead to consensus, this updating rule can lead to polarization and clustered, permanent disagreement on a connected network where everyone observes the same data and processes that

data with the same model. In these cases social learning implies greater polarization than that implied by individuals using different models to interpret data in isolation.

Topics for future investigation include understanding when a network is wise under various definitions, whether beliefs must necessarily become unidimensional as in DeMarzo et al. (2003), how one might endogenize the agent’s network along the lines in Sethi and Yildiz (2012) so as to generate a model of rational inattention (Sims (2003), Caplin et al. (2019)), and how one might endogenize φ_i^k in order to satisfy scientific ideals in addition to “seeing for one’s self.” We might also be interested in whether the agent’s problem of inference can address some of the concerns raised in Al-Najjar and Weinstein (2009), and how belief dynamics change as specific assumptions are changed to move the agent closer to Bayesian social learning (Molavi et al. (2017)). Finally, we might be curious about how the aggregation procedure in this paper would behave if it were generalized to account not only for disagreement but also for imprecision along the lines studied in Smithson (1999), Cabantous (2007), and Gajdos and Vergnaud (2013). Adding socially-acquired information to the empirical studies on belief revision surveyed in Manski (2017) would help to discipline all of this theoretical work.

References

- Al-Najjar, N. I. (2009). Decision makers as statisticians: Diversity, ambiguity, and learning. *Econometrica* 77(5), 1371–1401.
- Al-Najjar, N. I. and J. Weinstein (2009). The ambiguity aversion literature: A critical assessment. *Economics and Philosophy* 25(03), 249–284.
- Aliprantis, D. and F. G.-C. Richter (2020). Evidence of neighborhood effects from Moving to Opportunity: LATEs of neighborhood quality. *The Review of Economics and Statistics* 102(4), 633–647. DOI: 10.1162/rest_a_00933.
- Beshears, J., J. J. Choi, D. Laibson, B. C. Madrian, and K. L. Milkman (2015). The effect of providing peer information on retirement savings decisions. *The Journal of Finance* 70(3), 1161–1201.
- Bradley, R. (2017, Jul). Learning from others: conditioning versus averaging. *Theory and Decision*.
- Bramoullé, Y. and C. Orset (2018). Manufacturing doubt. *Journal of Environmental Economics and Management* 90, 119–133.
- Brown, L. A. and E. G. Moore (1970). The intra-urban migration process: A perspective. *Geografiska Annaler: Series B, Human Geography* 52(1), 1–13.
- Bursztyn, L., F. Ederer, B. Ferman, and N. Yuchtman (2014). Understanding mechanisms underlying peer effects: Evidence from a field experiment on financial decisions. *Econometrica* 82(4), 1273–1301.

- Cabantous, L. (2007). Ambiguity aversion in the field of insurance: Insurers' attitude to imprecise and conflicting probability estimates. *Theory and Decision* 62(3), 219–240.
- Caplin, A., M. Dean, and J. Leahy (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies* 86(3), 1061–1094.
- Conley, T. G. and C. R. Udry (2010). Learning about a new technology: Pineapple in Ghana. *The American Economic Review* 100(1), 35–69.
- Coolen, F. P., M. C. Troffaes, and T. Augustin (2011). Imprecise probability. In M. Lovric (Ed.), *International Encyclopedia of Statistical Science*. Springer.
- de Souza Briggs, X., K. S. Ferryman, S. J. Popkin, and M. Rendon (2008). Why did the Moving to Opportunity experiment not get young people into better schools? *Housing Policy Debate* 19(1), 53–91.
- DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association* 69(345), 118–121.
- DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003). Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics* 118(3), 909–968.
- Dominitz, J. and C. F. Manski (2011). Measuring and interpreting expectations of equity returns. *Journal of Applied Econometrics* 26(3), 352–370.
- Farber, H. S., J. Rothstein, and R. G. Valletta (2015). The effect of extended unemployment insurance benefits: Evidence from the 2012–2013 phase-out. *American Economic Review* 105(5), 171–176.
- Foster, A. D. and M. R. Rosenzweig (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy* 103(6), 1176–1209.
- Gabaix, X. (2019). Chapter 4 - behavioral inattention. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics - Foundations and Applications 2*, Volume 2 of *Handbook of Behavioral Economics: Applications and Foundations 1*, pp. 261–343. North-Holland.
- Gajdos, T. and J.-C. Vergnaud (2013). Decisions with conflicting and imprecise information. *Social Choice and Welfare* 41(2), 427–452.
- García, J. L. and J. J. Heckman (2022). Parenting promotes social mobility within and across generations. *NBER Working Paper 30610*.
- Gentzkow, M. and J. M. Shapiro (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics* 126(4), 1799–1839.

- Gessen, M. (2019, March 2). Why measles is a quintessential political issue of our time. *The New Yorker*.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Golub, B. and E. Sadler (2016). Learning in social networks. In Y. Bramoullé, A. Galeotti, and B. Rogers (Eds.), *The Oxford Handbook of the Economics of Networks*. Oxford University Press.
- Hagedorn, M., F. Karahan, I. Manovskii, and K. Mitman (2013). Unemployment benefits and unemployment in the Great Recession: The role of macro effects. *NBER Working Paper 19499*.
- Hegselmann, R. and U. Krause (2002). Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation* 5(3).
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton: Princeton University Press.
- Krysan, M. (2008). Does race matter in the search for housing? An exploratory study of search strategies, experiences, and locations. *Social Science Research* 37(2), 581–603.
- Lace-Evans, O. (2016, December 5). ‘No truth in journalism’: Trump voters cast wide net for news. *BBC News*. Retrieved from <http://www.bbc.com/news/av/world-us-canada-38191313/no-truth-in-journalism-trump-voters-cast-wide-net-for-news>.
- Levine, D. K. (2017, August 14). Confirmatory bias, the Bayesian detective and an inconvenient truth. *Economic and Game Theory Blog*. <http://www.dklevine.com/general/bayes.08.pdf>.
- Lichtendahl, Jr., K. C., Y. Grushka-Cockayne, V. R. R. Jose, and R. L. Winkler (2017). Bayesian ensembles of binary-event forecasts. *arXiv:1705.02391 [stat.ME]*.
- Loh, I. and G. Phelan (2019). Dimensionality and disagreement: Asymptotic belief divergence in response to common information. *International Economic Review* 60(4), 1861–1876.
- Ludwig, J., G. J. Duncan, L. A. Gennetian, L. F. Katz, R. C. Kessler, J. R. Kling, and L. Sanbonmatsu (2013). Long-term neighborhood effects on low-income families: Evidence from Moving to Opportunity. *American Economic Review* 103(3), 226–231.
- Manski, C. F. (2007). *Identification for Prediction and Decision*. Harvard University Press.
- Manski, C. F. (2011). Choosing treatment policies under ambiguity. *Annual Review of Economics* 3, 25–49.
- Manski, C. F. (2017). Survey measurement of probabilistic macroeconomic expectations: Progress and promise. *NBER WP 23418*.
- McCartney, W. B. and A. M. Shah (2022). Household mortgage refinancing decisions are neighbor influenced, especially along racial lines. *Journal of Urban Economics* 128, 103409.

- Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2017). Foundations of non-Bayesian social learning. *Mimeo., Northwestern University.*
- Moretti, E. (2011). Social learning and peer effects in consumption: Evidence from movie sales. *The Review of Economic Studies* 78(1), 356–393.
- Mueller-Frank, M. (2015). Reaching consensus in social networks. *IESE Working Paper 1116-E.*
- Paris, J. B. and A. Vencovská (1990). A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning* 4(3), 183–223.
- Pew (2012, October 15). More say there is solid evidence of global warming. *The Pew Research Center.* Retrieved from <http://www.people-press.org/2012/10/15/more-say-there-is-solid-evidence-of-global-warming/>.
- Ramey, V. A. (2011). Can government purchases stimulate the economy? *Journal of Economic Literature* 49(3), 673–685.
- Ranjan, R. and T. Gneiting (2010). Combining probability forecasts. *Journal of the Royal Statistical Society. Series B* 72(1), 71–91.
- Seidenfeld, T., M. J. Schervish, and J. B. Kadane (2010). Coherent choice functions under uncertainty. *Synthese* 172(1), 157–176.
- Serrato, J. C. S. and P. Wingender (2016). Estimating local fiscal multipliers. *NBER WP 22425.*
- Sethi, R. and M. Yildiz (2012). Public disagreement. *American Economic Journal: Microeconomics* 4(3), 57–95.
- Sethi, R. and M. Yildiz (2016). Communication with unknown perspectives. *Econometrica* 84(6), 2029–2069.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell Systems Technical Journal* 27, 379–423.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3), 665–690.
- Smithson, M. (1999). Conflict aversion: preference for ambiguity vs conflict in sources and evidence. *Organizational Behavior and Human Decision Processes* 79(3), 179–198.
- Sorensen, A. T. (2006). Social learning and health plan choice. *The Rand Journal of Economics* 37(4), 929–945.
- Sotiropoulos, D. N., C. Bilanakos, and G. M. Giaglis (2015). A time-variant and non-linear model of opinion formation in social networks. *Proceedings of the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 872–879.

- Stephens-Davidowitz, S. (2022, May 7). The one parenting decision that really matters. *The Atlantic*.
- Tamer, E. (2010). Partial identification in econometrics. *Annual Review of Economics* 2(1), 167–195.
- Upin, C. (2012, October 23). *Climate of Doubt*. WGBH/Boston: Frontline.
- White, H. (2000). *Asymptotic Theory for Econometricians*. Emerald Publishing Limited.
- Wilmers, G. (2010). The social entropy process: Axiomatising the aggregation of probabilistic beliefs. *Mimeo., Manchester Institute for Mathematical Sciences*.
- Zanardo, E. (2017). How to measure disagreement? *Mimeo., Columbia University*.

A Proofs

Proposition 1. *Under the specification of the agent's problem in Equation 2, Assumption A0, and Restrictions R1 and R2, the agent's updating rule is equivalent to DeGroot updating with time-varying weights.*

Proof. Then the agent's updating is

$$\begin{aligned}\widehat{F}_i^{k,t} &= \beta_t \widehat{\sigma}_i^{k,t} + (1 - \beta_t) \widehat{F}_i^{k,t-1} \\ &= \beta_t \sum_{j=1}^J p_{ij}^{k,t} \widehat{\sigma}_{ij}^{k,t} + (1 - \beta_t) \widehat{F}_i^{k,t-1}\end{aligned}\tag{4}$$

$$= \beta_t \sum_{j=1}^J p_{ij}^{k,t} \widehat{\sigma}_{ij}^{k,t} + (1 - \beta_t) \widehat{\sigma}_{ii}^{k,t}\tag{5}$$

$$= [\beta_t p_{ii}^{k,t} + (1 - \beta_t)] \widehat{\sigma}_{ii}^{k,t} + (1 - \beta_t) \sum_{j \in \mathcal{J}^k} p_{ij}^{k,t} \widehat{\sigma}_{ij}^{k,t}\tag{6}$$

$$= [\beta_t p_{ii}^{k,t} + (1 - \beta_t)] \widehat{\sigma}_{ii}^{k,t} + (1 - \beta_t) \sum_{j \in \mathcal{J}^k} p_{ij}^{k,t} \widehat{\sigma}_{jj}^{k,t}\tag{7}$$

$$= [\beta_t p_{ii}^{k,t} + (1 - \beta_t)] \widehat{F}_i^{k,t-1} + (1 - \beta_t) \sum_{j \in \mathcal{J}^k} p_{ij}^{k,t} \widehat{F}_j^{k,t-1}\tag{8}$$

$$= \sum_{j=1}^J \omega_{ij}^{k,t} \widehat{F}_j^{k,t-1}.\tag{9}$$

where

$$\omega_{ii}^{k,t} = \beta_t p_{ii}^{k,t} + (1 - \beta_t) \quad \text{and} \quad \omega_{ij}^{k,t} = (1 - \beta_t) p_{ij}^{k,t};$$

Equation 4 follows from R1; Equation 5 follows from A0b; Equation 6 is rearranging terms; Equation 7 follows from R2; and Equation 8 follows from A0b.

Note that we could re-write Equation 5 as:

$$\beta_t \sum_{j=1}^J p_{ij}^{k,t} \widehat{F}_j^{k,t-1} + (1 - \beta_t) \widehat{F}_i^{k,t-1}.$$

Since $\sum_{j=1}^J p_{ij}^{k,t} = 1$, the sum of weights on previous beliefs is $\beta_t + (1 - \beta_t)$, showing that $\sum_{j=1}^J \omega_{ij}^{k,t} = 1$. Similarly, the fact that $\omega_{ij}^{k,t} \geq 0$ for all i, j follows from the similar condition on $p_{ij}^{k,t}$ and the definition of β_t .

In matrix form, the agents' updating in Equation 9 can be written as $\widehat{\mathbf{F}}^{k,t} = \mathbf{\Omega}^{k,t} \widehat{\mathbf{F}}^{k,t-1}$, which is equivalent to DeGroot updating with time-varying weights. Note that due to A0a we could assume $p_{ij}^{k,t} = p_{ij}^{k,1}$ for all $t \geq 2$, so that the weights in $\omega_{ij}^{k,t}$ only change over time due to β_t . Thus Assumption A0 and Restrictions R1-R2 generate DeGroot updating with time-varying weights. \square

Proposition 2. *Under Assumptions A1 and A2, in addition to the assumptions in Proposition 1, updating via DeGroot updating solves the agent's problem.*

Proof. The proof is by induction. For $t = 1$ we have that

$$\begin{aligned}\mathbb{E}[\widehat{F}_i^{k,1}] &= \mathbb{E}\left[\sum_{j=1}^J p_{ij}^{k,1} \widehat{\sigma}_{jj}^{k,0}\right] \\ &= \sum_{j=1}^J p_{ij}^{k,1} \mathbb{E}[\widehat{\sigma}_{jj}^{k,0}]\end{aligned}\tag{10}$$

$$= \sum_{j \in \mathcal{J}^{k*}} p_{ij}^{k,1} \mu_i^k\tag{11}$$

$$= \mu_i^k\tag{12}$$

where Equation 10 is due to the linearity of expectation. Equation 11 follows from A1 and A2 in combination with the definition of \mathcal{J}^{k*} being those j such that the agent assigns $p_{ij}^{k,1} > 0 \iff \mathbb{E}[\widehat{\sigma}_{jj}^{k,t}] = \mu_i^k$.

Now, suppose by way of induction that $\mathbb{E}[\widehat{F}_i^{k,t-1}] = \mu_i^k$. Then

$$\mathbb{E}[\widehat{F}_i^{k,t}] = \mathbb{E}[\beta_t \widehat{\sigma}_{ii}^{k,t} + (1 - \beta_t) \widehat{F}_i^{k,t-1}]\tag{13}$$

$$= \beta_t \mathbb{E}[\widehat{F}_{ii}^{k,t-1}] + (1 - \beta_t) \mathbb{E}[\widehat{F}_i^{k,t-1}]\tag{14}$$

$$= \mathbb{E}[\widehat{F}_i^{k,t-1}]\tag{15}$$

$$= \mu_i^k\tag{16}$$

where Equation 14 follows from A0b and the linearity of expectation. □

Proposition 3. Assume that the $\hat{\sigma}_{jj}^{k,n}$ are each sampled from the distribution $\hat{\Gamma}_j^k$ with expected value μ_j^k . Under the specification of the agent's problem in Equation 2, and Assumptions A1-A2, the predicted bias updating rule in Restrictions R1-R2* solves the agent's problem.

Proof.

$$\begin{aligned} \mathbb{E} \left[\lim_{t \rightarrow \infty} \hat{F}_i^{k,t} \right] &= \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \hat{\sigma}_i^{k,n} \right] \\ &= \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \left(\sum_{j=1}^J p_{ij}^{k,n} \hat{\sigma}_{ij}^{k,n} \right) \right] \end{aligned} \quad (17)$$

$$= \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \left(\sum_{j=1}^J p_{ij}^{k,n} (\hat{\sigma}_{jj}^{k,n} + \mu_i^k - \mathbb{E}[\hat{\sigma}_{jj}^{k,n}]) \right) \right] \quad (18)$$

$$= \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \left(\sum_{j=1}^J p_{ij}^{k,n} \mu_i^k \right) \right] + \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \left(\sum_{j=1}^J p_{ij}^{k,n} (\hat{\sigma}_{jj}^{k,n} - \mathbb{E}[\hat{\sigma}_{jj}^{k,n}]) \right) \right] \quad (19)$$

$$= \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \mu_i^k \right] + \sum_{j=1}^J \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t p_{ij}^{k,n} (\hat{\sigma}_{jj}^{k,n} - \mathbb{E}[\hat{\sigma}_{jj}^{k,n}]) \right] \quad (20)$$

$$= \mu_i^k + \sum_{j=1}^J \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t p_{ij}^{k,n} (\hat{\sigma}_{jj}^{k,n} - \mu_j^k) \right] \quad (21)$$

$$= \mu_i^k + \sum_{j=1}^J \mathbb{E} \left[\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t p_{ij}^{k,n} \hat{\sigma}_{jj}^{k,n} \right] - \mu_j^k \quad (22)$$

$$= \mu_i^k. \quad (23)$$

Equation 17 follows from R1; Equation 18 follows from R2*; Equation 19 is rearranging terms; Equation 20 follows from the convex combination portion of R1 and the linearity of expectation; Equation 21 follows from A2 and the assumption that the $\hat{\sigma}_{jj}^{k,n}$ are sampled from the distribution $\hat{\Gamma}_j^k$ with expected value μ_j^k ; Equation 22 again follows from the convex combination portion of R1; and Equation 23 is a consequence of the law of large numbers. \square

B Appendix: Neighborhood Effects Decision Problem

Consider the following stylized decision problem: A Decision Maker (DM) is the head of a low-income household that lives in a low quality neighborhood ($D = 0$), and must decide whether to move to a high quality neighborhood ($D = 1$). If the household moves, they must pay a higher rent and a moving cost, which I normalize to c . However, the DM's probability of being employed at wage w (p_D) may be higher if $D = 1$ than if $D = 0$. The expected utility of remaining in the low quality neighborhood is $U_0 = p_0w$ and for moving it is $U_1 = p_1w - c$.

The DM knows that $p_1 \in \{p_0, p_0 + \theta\}$.¹⁰ The DM faces risk when her beliefs are the point $Pr(p_1 = p_0 + \theta) = \pi \in [0, 1]$, and ambiguity/uncertainty when she holds a set of possible beliefs where $Pr(p_1 = p_0 + \theta) \in [\underline{\pi}, \bar{\pi}] \subseteq [0, 1]$.

To illustrate the importance of belief formation for choices, consider the DM's decision under the following 3 decision rules: Expected utility maximization under risk (EU), Γ -Maximin ($\Gamma - M$), and Γ -Minimax Regret ($\Gamma - MR$). I show below that all decision rules compare a single belief with a combination of the cost of moving c , the neighborhood effect θ , and the wage w :

$$D = 1 \iff \pi^{DR} \geq \frac{c}{\theta w},$$

where $\pi^{EU} = \pi$, $\pi^{\Gamma-M} = \underline{\pi}$, and $\pi^{\Gamma-MR} = (\underline{\pi} + \bar{\pi})/2$.

B.1 Expected Utility under Risk

Suppose the DM's belief is $\pi \in [0, 1]$. The DM chooses

$$\begin{aligned} D = 1 &\iff U_1 \geq U_0 \iff p_1w - c \geq p_0c \\ &\iff \pi(p_0 + \theta)w + (1 - \pi)p_0w - c \geq \pi p_0w + (1 - \pi)p_0w \\ &\iff \pi\theta w \geq c \end{aligned} \tag{24}$$

$$\iff \pi \geq \frac{c}{\theta w}. \tag{25}$$

Equation 24 shows that this comparison in Equation 25 is the same as whether the anticipated gains from moving outweigh the costs.

B.2 Γ -Maximin

Suppose the DM's belief is $\pi \in [\underline{\pi}, \bar{\pi}] \subseteq [0, 1]$. Under the Γ -Maximin decision rule the DM chooses

$$D = 1 \iff \underline{U}_1 \geq \underline{U}_0 \tag{26}$$

¹⁰Analogous to this binary state of the world, a binary Average Treatment Effect (ATE) would allow for $p_1 \in \{p_0, p_0 + \theta\}$, where θ is some positive constant. If potential outcomes $Y_i(D)$ represent employment under various neighborhood treatments, then the ATE is defined as $\mathbb{E}[Y_i(1) - Y_i(0)] \equiv p_1 - p_0$.

where

$$\begin{aligned}\underline{U}_1 &= \min_{\pi \in [\underline{\pi}, \bar{\pi}]} U_1 = \min_{\pi \in [\underline{\pi}, \bar{\pi}]} \pi(p_0 + \theta)w + (1 - \pi)p_0w - c \\ &= \underline{\pi}(p_0 + \theta)w + (1 - \underline{\pi})p_0w - c\end{aligned}$$

and

$$\begin{aligned}\underline{U}_0 &= \min_{\pi \in [\underline{\pi}, \bar{\pi}]} U_0 = \min_{\pi \in [\underline{\pi}, \bar{\pi}]} p_0w \\ &= p_0w.\end{aligned}$$

Thus Equation 26 can be stated as

$$\begin{aligned}D = 1 &\iff \underline{\pi}(p_0 + \theta)w + (1 - \underline{\pi})p_0w - c \geq \underline{\pi}p_0w + (1 - \underline{\pi})p_0w \\ &\iff \underline{\pi}\theta w \geq c\end{aligned}\tag{27}$$

$$\iff \underline{\pi} \geq \frac{c}{\theta w}.\tag{28}$$

Note the similarity between Equations 27 and 28 and Equations 24 and 25. Now the condition for moving is that the expected benefit of moving must be higher than the cost in the “worst-case” scenario of moving.

B.3 Γ -Minimax Regret

Suppose the DM’s belief is $\pi \in [\underline{\pi}, \bar{\pi}] \subseteq [0, 1]$. Under the Γ -Minimax Regret decision rule the DM chooses based on a comparison between

$$\begin{aligned}\bar{R}_1 &= \max_{\pi \in [\underline{\pi}, \bar{\pi}]} [U_0 - U_1] & \text{and} & & \bar{R}_0 &= \max_{\pi \in [\underline{\pi}, \bar{\pi}]} [U_1 - U_0] \\ &= \max_{\pi \in [\underline{\pi}, \bar{\pi}]} -\pi\theta w + c & & & &= \max_{\pi \in [\underline{\pi}, \bar{\pi}]} \pi\theta w - c \\ &= c - \underline{\pi}\theta w & & & &= \bar{\pi}\theta w - c.\end{aligned}$$

Thus the decision rule is

$$\begin{aligned}D = 1 &\iff \bar{R}_1 \leq \bar{R}_0 \\ &\iff c - \underline{\pi}\theta w \leq \bar{\pi}\theta w - c \\ &\iff (\underline{\pi} + \bar{\pi})\theta w \geq 2c\end{aligned}\tag{29}$$

$$\iff \frac{\underline{\pi} + \bar{\pi}}{2} \geq \frac{c}{\theta w}.\tag{30}$$

C Appendix: Experiments 2 and 3

In Experiment 2, I assume there are four clusters of individuals, \mathcal{C}_1 , \mathcal{C}_2 , \mathcal{C}_3 , and \mathcal{C}_4 with $\text{card}(\mathcal{C}_1) = 50$, $\text{card}(\mathcal{C}_2) = 100$, $\text{card}(\mathcal{C}_3) = 50$, and $\text{card}(\mathcal{C}_4) = 100$. Initial beliefs $\widehat{F}_i^{k,1}$ are generated as follows:

$$\ddot{F}_i^k \sim \begin{cases} \mathcal{N}(0.8, 0.1) & \text{if } i \in \mathcal{C}_1 \ \forall \ k = 1, \dots, K; \\ \mathcal{N}(0.2, 0.01) & \text{if } i \in \mathcal{C}_2 \ \forall \ k = 1, \dots, K; \\ \mathcal{N}(0.8, 0.1) & \text{if } i \in \mathcal{C}_3 \ \text{and } k \text{ even}; \\ \mathcal{N}(0.5, 0.1) & \text{if } i \in \mathcal{C}_3 \ \text{and } k \text{ odd}; \\ \mathcal{N}(0.2, 0.01) & \text{if } i \in \mathcal{C}_4 \ \text{and } k \text{ even}; \\ \mathcal{N}(0.5, 0.01) & \text{if } i \in \mathcal{C}_4 \ \text{and } k \text{ odd}; \end{cases}$$

with

$$\widehat{F}_i^{k,1} = \min \left\{ \max \left\{ \ddot{F}_i^k, 0 \right\}, 1 \right\} \quad \forall k = 1, \dots, K.$$

Figures 5a and 6a shows the initial distribution of beliefs by clusters for propositions p^1 and p^2 , respectively. Figures 5 and 6 show how beliefs evolve by cluster.

In Experiment 3, I also assume there are four clusters of individuals, \mathcal{C}_1 , \mathcal{C}_2 , \mathcal{C}_3 , and \mathcal{C}_4 with $\text{card}(\mathcal{C}_1) = 50$, $\text{card}(\mathcal{C}_2) = 100$, $\text{card}(\mathcal{C}_3) = 50$, and $\text{card}(\mathcal{C}_4) = 100$. Initial beliefs \widehat{F}_{i1}^k are generated as follows:

$$\ddot{F}_i^k \sim \begin{cases} \mathcal{N}(0.8, 0.1) & \text{if } i \in \mathcal{C}_1 \ \forall \ k = 1, \dots, K; \\ \mathcal{N}(0.2, 0.01) & \text{if } i \in \mathcal{C}_2 \ \forall \ k = 1, \dots, K; \\ \mathcal{N}(0.9, 0.1) & \text{if } i \in \mathcal{C}_3 \ \forall \ k = 1, \dots, K; \\ \mathcal{N}(0, 0) & \text{if } i \in \mathcal{C}_4 \ \forall \ k = 1, \dots, K; \end{cases}$$

with

$$\widehat{F}_i^{k,1} = \min \left\{ \max \left\{ \ddot{F}_i^k, 0 \right\}, 1 \right\} \quad \forall k = 1, \dots, K.$$

Figures 7a and 8a shows the initial distribution of beliefs by clusters for propositions p^1 and p^2 , respectively. Figures 7 and 8 show how beliefs evolve by cluster.

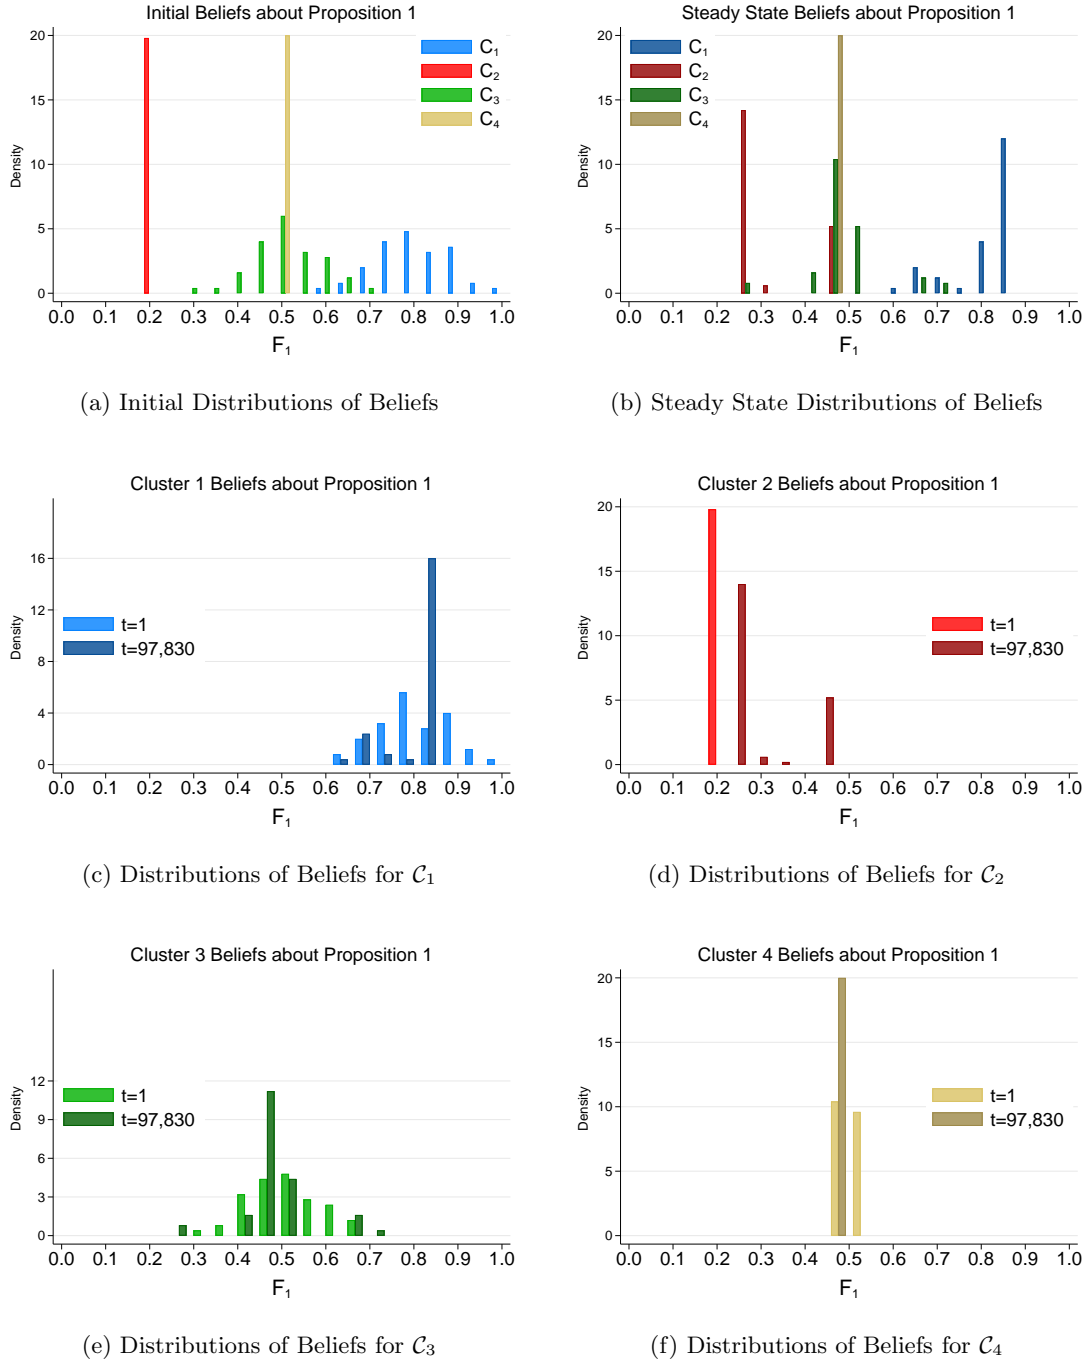
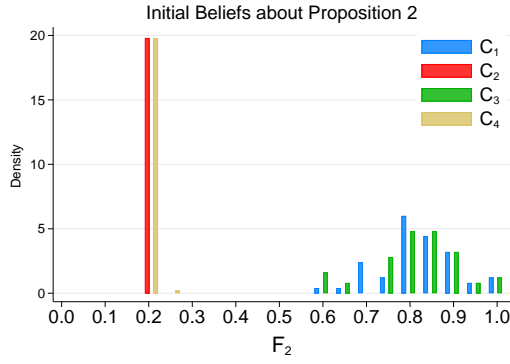



Figure 5: Beliefs in Experiment 2


Steady state belief distributions are those for which $\max_{k,i} \|\widehat{F}_i^{k,t} - \widehat{F}_i^{k,t-1}\| < 1e-6$, which occurs in Experiment 2 at $t = 97,830$.



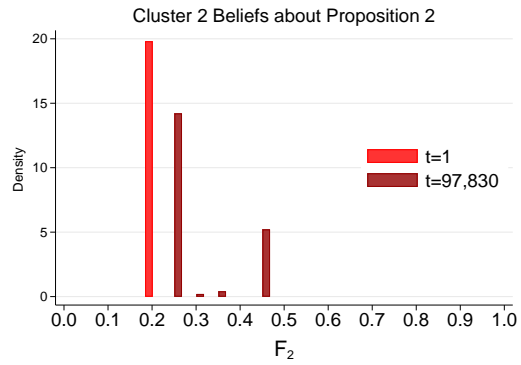
(a) Initial Distributions of Beliefs




(b) Steady State Distributions of Beliefs




(c) Distributions of Beliefs for C_1



(d) Distributions of Beliefs for C_2




(e) Distributions of Beliefs for C_3



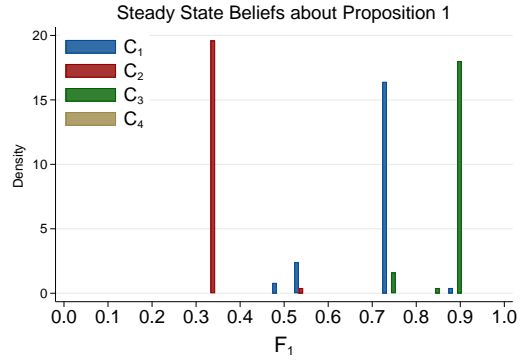
(f) Distributions of Beliefs for C_4

Figure 6: Beliefs in Experiment 2


Steady state belief distributions are those for which $\max_{k,i} \|\widehat{F}_i^{k,t} - \widehat{F}_i^{k,t-1}\| < 1e-6$, which occurs in Experiment 2 at $t = 97,830$.



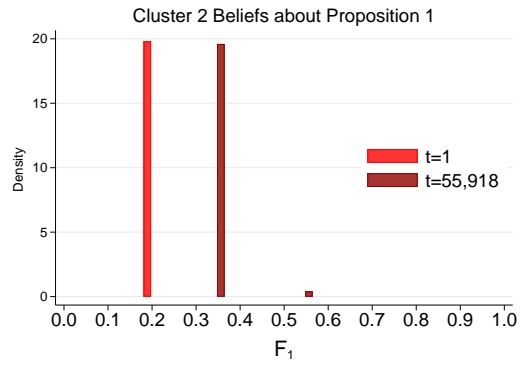
(a) Initial Distributions of Beliefs



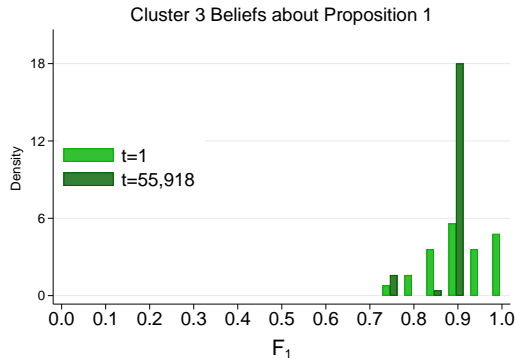
(b) Steady State Distributions of Beliefs



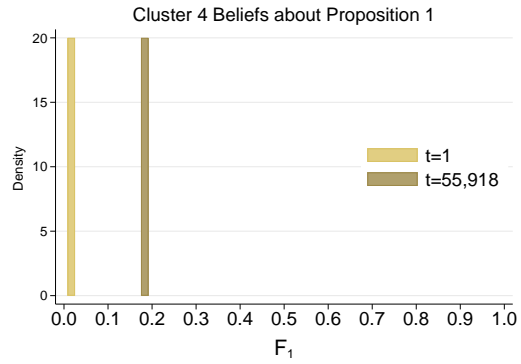
(c) Distributions of Beliefs for C_1



(d) Distributions of Beliefs for C_2




(e) Distributions of Beliefs for C_3




(f) Distributions of Beliefs for C_4

Figure 7: Beliefs in Experiment 3


Steady state belief distributions are those for which $\max_{k,i} \|\widehat{F}_i^{k,t} - \widehat{F}_i^{k,t-1}\| < 1e-6$, which occurs in Experiment 3 at $t = 55,918$.



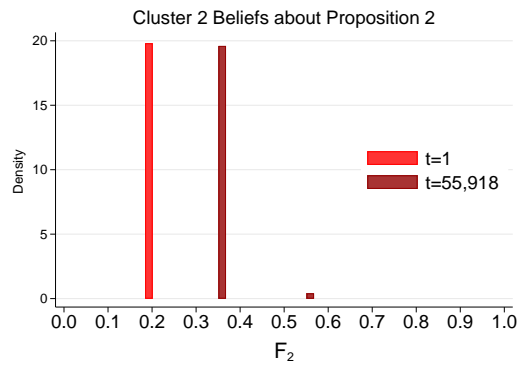
(a) Initial Distributions of Beliefs



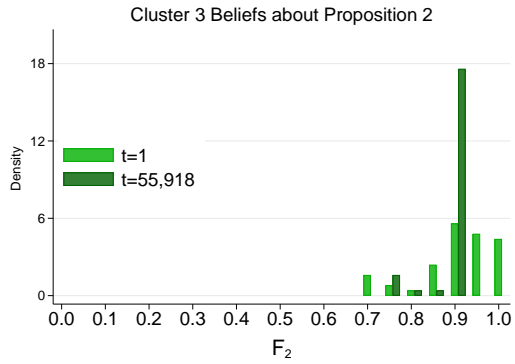
(b) Steady State Distributions of Beliefs




(c) Distributions of Beliefs for C_1



(d) Distributions of Beliefs for C_2



(e) Distributions of Beliefs for C_3



(f) Distributions of Beliefs for C_4

Figure 8: Beliefs in Experiment 3

Steady state belief distributions are those for which $\max_{k,i} \|\widehat{F}_i^{k,t} - \widehat{F}_i^{k,t-1}\| < 1e-6$, which occurs in Experiment 3 at $t = 55,918$.