

Differences of Opinion

Dionissi Aliprantis

September 19, 2017

Abstract: A strong assumption of Bayesian social learning is that agents can solve complex inference problems to construct the likelihood function. Non-Bayesian social learning typically assumes agents use a simple rule of thumb for inference. This paper studies a non-Bayesian social learner whose goal is to solve a complex inference problem. I study how common obstacles to identifying causal effects can lead to ambiguity, a set of possible beliefs, and how an agent might use social learning to resolve that ambiguity. I allow the agent to impute missing data using information observed through her network in combination with a model of social learning. In some cases, the agent's belief formation reduces to DeGroot updating, and beliefs in a network reach a consensus. In other cases, the agent's updating can generate polarization and sustain clustered disagreement, even on a connected network where everyone observes the same data and processes that data with the same model.

Keywords: Belief Formation, Subjective Probability, Imprecise Probability, Social Learning, Partial Identification, Causal Inference, Network, DeGroot Learning Rule, Bounded Confidence

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 Siedlarek, Alireza Tahbaz-Salehi, Teddy Seidenfeld, Michael Smithson, seminar participants at the Cleveland Fed and Goethe University Frankfurt, 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

Transforming data into beliefs is not always straightforward. Data arriving between 2006 and 2012 from the United States’ [National Oceanic and Atmospheric Administration](#) show a clear continuation of the upward trend in average global temperature over previous decades. Yet these data actually moved us away from consensus; many more Americans believed there was an upward trend in 2006 than in 2012 (Pew (2012)). Is it possible to rationalize this type of belief updating?

To study disagreement, it can be useful to abstract from the precise statistical techniques used to interpret data, and to view inference through the lens of a coarse, binary model. Some cases of disagreement viewed through this lens are 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)); the US was at full employment in 2017.

The need for data with specific features makes disagreement especially common in causal inference. For example, there is no data available from an ideal experiment randomly allocating households across neighborhoods. In studies of neighborhood effects on employment, the best data that is available has been interpreted as evidence of both the absence (Ludwig et al. (2013)) and the presence (Aliprantis and Richter (2016)) of neighborhood effects.

This paper studies how common obstacles to identifying causal effects can lead to ambiguity, a set of possible beliefs, and how an agent might use social learning to resolve that ambiguity. When the agent does not know how to specify the likelihood function for this complex inference problem, she cannot engage in Bayesian social learning. She still might want to approach inference with more sophistication than the rules of thumb typically employed in non-Bayesian social learning.

I suppose the agent extrapolates from observed data under invariance assumptions, using disagreement over multiple issues to infer missing data from other people. This approach will often lead agents to agree. Sometimes, though, this approach will over-interpret differences of opinion, creating polarization in situations that might be unexpected.

To frame the issues, I first consider the case of inference when there are no problems of identification. An agent observes a sequence of data, where the sampling procedure is identical over time. At each point in time the agent uses her model to translate the observed data into a point-valued signal $\sigma_t \in [0, 1]$ about a binary state of the world $s \in \{0, 1\}$. Beliefs will converge if subjective beliefs $\lambda_t \in [0, 1]$ are formed by taking the average of signals over time.

In the general case when causal effects are not so cleanly identified, we might suppose that an agent observes iid data yielding coarse, or set-valued, signals. This circumstance is widespread in the social sciences, where we rarely observe data from an ideal experiment for identifying causal effects. Beliefs formed by averaging imprecise probabilities over discrete time will converge to a set (Artstein and Vitale (1975)), a scenario of partial identification.

How might a decision maker choose one belief from a set of possible beliefs? When an agent faces ambiguity, prominent decision rules instruct her to choose the single belief generating an extreme utility (Gilboa and Marinacci (2013)).¹ Alternatively, a common practice is to form beliefs after

¹For example, the maxmin expected utility decision rule maximizes expected utility after choosing the belief that

searching one’s social network for additional information.

I study the problem of an agent who directly observes data consistent with a set of beliefs, but who would like to choose the single belief she would have formed with access to data yielding a clear inference. Such an agent is committed to the scientific ideal of direct observation, but faces time, resource, or ethical constraints making it infeasible to personally verify the claims in question.

I allow the agent to solve her missing data problem with information observed through her social network. I assume that communication is restricted to point-valued signals and beliefs. In the face of such imperfect communication, the agent needs a model of social learning to infer the data she does not observe from the signals she does observe.

I first show that DeGroot (1974) updating, the benchmark model of non-Bayesian social learning, solves a special case of the agent’s problem. If the agent addresses her problem of inference with linear opinion pooling of signals, a common method for combining forecasts and estimates, she will follow a DeGroot learning rule under a special case of observed data. Strong assumptions on the data and models in the agent’s network are required, however, for DeGroot updating to solve the agent’s problem.²

The major contribution of this paper is to show that by searching for a general solution to the agent’s problem, one can find generalizations of DeGroot updating capable of generating polarization. Such a learning rule is a desideratum of the literature on social learning for its ability to reconcile theory and evidence (Golub and Sadler (2016)). DeGroot learning and many of its generalizations converge to a degenerate distribution for connected networks (Jackson (2008)).³ Yet empirically we often observe the analogue of a connected network – individuals exposed to sources of information contradicting their beliefs – together with persistent disagreement.⁴

A general solution to the agent’s problem using linear opinion pooling requires, in contrast to standard DeGroot updating, a first stage in which signals are properly transformed. I present the selection of a model that properly interprets signals as a statistical learning problem, and show that this problem is not well-posed. That is, frictions from communication generate a fundamental problem of inference, in that signals do not convey the same information as directly observed data, and the agent cannot know whether she is properly interpreting signals without this information.

I show that there is an intermediate case, in which the agent can feasibly solve her problem while invoking more general assumptions than DeGroot updating. Because these assumptions will not always hold, the agent must make inductive inferences (outside the support of the data) about when the assumptions are most likely to hold. I study the implementation and belief dynamics of the updating rule when the agent makes an invariance assumption that differences in interpreting data

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)).

²Individuals can be justified in using different models to interpret the same data (Al-Najjar (2009)), the agent might observe new data over time (Jadbabaie et al. (2012)), and individuals might directly observe different data.

³Time to consensus, though, is not invariant across all connected network structures (Golub and Jackson (2012)).

⁴There is persistent disagreement over propositions like [Iraq had an active WMD program](#), [President Obama was born in the US](#), [vaccines cause autism](#), and [global warming is occurring](#) despite public debate. This disagreement persists despite exposure to opposing views (Gentzkow and Shapiro (2011)).

are consistent across propositions. The agent assesses the credibility of applying this assumption to each sender by using the relative entropy of their disagreement over all propositions. This assigns weight to the data inferred from a sender based not on agreement, but based on understanding how someone interprets data (Sethi and Yildiz (2016)).

Although the agent’s updating rule that is both feasible and general tends to generate consensus, I show that the rule is also 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. Polarization is possible because the agent can update her beliefs away from a signal, which contrasts with updating in DeGroot or bounded confidence models. In other words, the agent’s updating rule need not lead to constricting belief updating (Mueller-Frank (2015)). Two keys for generating polarization are low-quality data and perceptions about the distribution of models for interpreting directly observed data.

The paper proceeds as follows: Section 2 sets the stage for the agent’s problem, describing how she could arrive at a set of beliefs when directly observing data. Section 3 explores how the agent might try to resolve the ambiguity she faces by using the signals she receives from individuals in her social network. Section 4 characterizes solutions to the agent’s problem that are either easily implementable or fully general, and one that balances feasibility with generality. Section 5 illustrates some of the belief dynamics possible under the feasible and general updating rule. Section 6 concludes, with Appendix A listing notation to help the reader.

2 Belief Formation via Directly-Observed Data

Suppose there is a finite set of propositions $\{p^1, p^2, \dots, p^K\} = \mathcal{K}$, none of which can be written as a compound proposition using other propositions in the set.⁵ An agent must determine the truth value of the statements, $T(p^k) \in \{0, 1\}$, and agent i ’s beliefs at time t are denoted by $\lambda_{it}^k = \Pr(T(p^k) = 1)$.⁶ The agent directly observes data W_{it} sampled under identical procedures at each point in time.

Consider a classical (frequentist) setting. With high-quality data W_{it}^* , the agent would be able to use her model φ_i^k to translate the data into an independent and identically distributed (iid) sequence of signals $\{\sigma_{it}^{k*}\}_{t=1}^T$, where

$$\sigma_{it}^{k*} = \varphi_i^k(W_{it}^*) \in [0, 1].$$

The law of large numbers ensures convergence to the mean of the signal distribution, which I will

⁵This 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.

⁶A proposition is a statement that is either true ($T(p^k) = 1$) or false ($T(p^k) = 0$).

denote by μ_i^{k*} , for beliefs formed as

$$\begin{aligned}\lambda_{it+1}^{k*} &= \frac{1}{t} \sum_{n=1}^t \sigma_{in}^{k*} \\ &= \beta_t \sigma_{it}^{k*} + (1 - \beta_t) \lambda_{it}^{k*} \quad \text{where} \quad \beta_t = 1/t.\end{aligned}\tag{1}$$

Now consider a setting in which the agent’s directly-observed data W_{it} only allows her to determine a set containing the true signal σ_{it}^{k*} . Inspired by the literature on partial identification (Manski (2007), Tamer (2010)), suppose the agent’s model and data allow her to determine a signal σ_{it}^k and its quality θ_{it}^k ,

$$(\sigma_{it}^k, \theta_{it}^k) = \varphi_i^k(W_{it}) \in [0, 1]^2,$$

where the true signal is related to the observed signal by

$$\sigma_{it}^{k*} \in [\max\{0, \sigma_{it}^k - (1 - \theta_{it}^k)\} , \min\{\sigma_{it}^k + (1 - \theta_{it}^k), 1\}] \equiv [\underline{\sigma}_{it}^{k*} , \overline{\sigma}_{it}^{k*}].\tag{2}$$

The agent then knows from her signals of imperfect quality that the average

$$\lambda_{it+1}^{k*} = \frac{1}{t} \sum_{n=1}^t \sigma_{in}^{k*} \in \Lambda_{it+1}^{k*} = \left[\frac{1}{t} \sum_{n=1}^t \underline{\sigma}_{in}^{k*} , \frac{1}{t} \sum_{n=1}^t \overline{\sigma}_{in}^{k*} \right],$$

where the sets $[\underline{\sigma}_{it}^{k*}, \overline{\sigma}_{it}^{k*}]$ and Λ_{it+1}^{k*} are often referred to as “imprecise probabilities” (Coolen et al. (2011)). The set $[\underline{\sigma}_{it}^{k*}, \overline{\sigma}_{it}^{k*}]$ is what can be learned about p^k from the directly-observed data under the most credible assumptions. While the agent can also determine a point estimate σ_{it}^k , doing so requires less credible assumptions, so the agent cannot be sure that $\mathbb{E}[\sigma_{it}^k] = \mu_i^{k*}$ unless $\theta_{it}^k = 1$. Manski (2016) discusses how the length of an interval estimate can be reduced, even as far as a point estimate, by adopting assumptions of decreasing credibility.

In order to give some interpretation to the above notation, suppose that two researchers had access to the same sequence of data generated by many states randomly increasing and decreasing their state budgets. The first researcher is interested in learning about the proposition $p^1 =$ “Increasing state spending stimulates the state economy.” Suppose that this proposition were true ($T(p^1) = 1$). In the case of high-quality data where $\theta_{it}^1 = 1$ for all t , $\sigma_{it}^1 = \sigma_{it}^{1*}$, and so λ_{it+1}^{1*} can be calculated from Equation (1) as the mean signal. Supposing that the true mean of the signal distribution were 0.9, the Law of Large Numbers ensures that the first researcher’s belief will converge to 0.9 as $t \rightarrow \infty$.

Now suppose that the second researcher is interested in using the same data to learn the truth of proposition $p^2 =$ “Increasing federal spending stimulates the national economy.” Given the differences between the state and national economies, as well as differences in state and federal government purchases, the researcher judges her state-level data to be of low-quality, mapping into signals represented by $\theta_{it}^2 = 0.2$ for all t . A stylized example would be that the signal distribution took on a discrete support, with probability 0.25 that $\sigma_{it}^2 = 0$, probability 0.25 that $\sigma_{it}^2 = 0.5$, and

probability 0.5 that $\sigma_{it}^2 = 1$.

In this case, the second researcher will be subject to ambiguity in addition to risk.⁷ If the observed signal is $\sigma_{it}^2 = 0, 0.5, \text{ or } 1$, then based on Equation 2 the researcher can bound the true signal to be within, respectively, $[0, 0.8]$, $[0.1, 0.9]$, or $[0.2, 1]$ (See Figure 1.). Thus, as $t \rightarrow \infty$, the second researcher will infer that the mean of the true signals is $\mu_i^{2*} \in \Lambda_i^{2*} = [0.125, 0.925]$.⁸




Figure 1: $\varphi_i^k(W_{it}) = (\sigma_{it}^k, \theta_{it}^k) \Rightarrow [\underline{\sigma}_{it}^{k*}, \bar{\sigma}_{it}^{k*}]$

Non-causal propositions can also have low-quality signals for reasons like survey non-response (Manski (2015)). The above model of belief revision can apply to any proposition for which the agent faces a missing data problem.

In addition to describing signals, throughout the analysis I will use “high-quality” (relative to the agent’s model) to describe data yielding point-identified signals ($\theta_{it}^k = 1$), and “low-quality” to describe data yielding set-identified signals ($\theta_{it}^k < 1$). Another interpretation of an extremely low-quality signal, $\theta_{it}^k = 0$, is that the agent does not directly observe any data for a given proposition

⁷In this context a point-valued belief $\lambda_{it}^k \in (0, 1)$ represents risk, while a set-valued belief $\Lambda_{it}^k \subseteq [0, 1]$ represents Knightian uncertainty or ambiguity.

⁸Confidence intervals for the identified set Λ_i^{k*} are studied in Imbens and Manski (2004) and Stoye (2009), more generally as confidence regions in Chernozhukov et al. (2007) and Romano and Shaikh (2010), and using Bayesian methods in Moon and Schorfheide (2012) and Bollinger and van Hasselt (2017).

p^k , so that $\varphi_i^k(\emptyset) = (\sigma_{it}^k, 0) \Rightarrow \sigma_{it}^{k*} \in [0, 1]$. It could also be the case that the agent’s model φ_i^k is not capable of extracting information from data. For example, an agent ignorant of genetics and molecular biology would likely have a model incapable of interpreting data on the human genome. In such cases, one could assign $\varphi_i^k(W_{it}^*) = \varphi_i^k(W_{it}) = (\sigma_{it}^k, 0) \Rightarrow \sigma_{it}^{k*} \in [0, 1]$ for any data set. For this analysis I will assume that for each proposition in question, the agent’s model produces a point-identified signal given a high-quality data set.

3 Belief Formation via Social Learning

A criticism of Bayesian decision theory is that in some circumstances, it might not be possible for the agent to express her beliefs using a distribution over the set Λ_{it}^{k*} . Bayesian decision theory is difficult to apply to these circumstances, since an imprecise probability cannot be used to make decisions according to the standard Savage axioms (Gilboa and Marinacci (2013)).

When holding beliefs represented by an imprecise probability Λ_{it}^{k*} , several approaches to decision making can be interpreted as picking one belief from the set Λ_{it}^{k*} , 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)).⁹

When a decision maker does not observe all of the relevant data to form beliefs, there is empirical evidence that such a decision maker will often try to infer the relevant information from social observations. Examples are as diverse as neighborhood and school choice (de Souza Briggs et al. (2008)); adoption of a new technology (Conley and Udry (2010), Foster and Rosenzweig (1995)); consumption of a new good (Moretti (2011)); investment decisions (Bursztyrn et al. (2014)); retirement savings decisions (Beshears et al. (2015)); and choice of health insurance plans (Sorensen (2006)).¹⁰ The subsequent model explores belief formation when the agent chooses one belief from Λ_{it}^{k*} using information from her social network.

3.1 The Agent’s Problem

Suppose the agent is a member of a network of $J + 1$ individuals from which she might gather information. The agent directly-observes the information

$$\mathcal{I}_{it} \equiv \left\{ (\lambda_{it}^1, \sigma_{it}^1, \theta_{it}^1) , \dots , (\lambda_{it}^K, \sigma_{it}^K, \theta_{it}^K) \right\}.$$

⁹Appendix B illustrates these decision rules in the stylized binary model considered in this paper.

¹⁰This type of social learning is not restricted to causal propositions. As an example, the Federal Reserve looks at the beliefs of others when forming beliefs about whether the economy is in a recession or is at full employment.

To initialize the process we might let $\lambda_{i1}^k = \sigma_{i1}^k$; assume that the agent observes point identified signals from $t = -T$ until $t = 1$ and then set identified signals for $t > 1$; or else assume that the agent has just randomly reset $t = 1$ (as a random mutation in an evolutionary algorithm). The agent also observes information in her social network about the truth of propositions. We denote the set of others in the agent's network as \mathcal{J} . However, the agent does not directly observe the data individuals in her network ($j \in \mathcal{J}$) directly observe. Instead, the agent observes individuals' beliefs and their interpreted data in the form of their signals. Thus, the socially-observed information available to the agent is

$$\mathcal{I}_{Jt} \equiv \left\{ \{\lambda_{jt}^1, \sigma_{jt}^1\}_{j \in \mathcal{J}^1}, \dots, \{\lambda_{jt}^K, \sigma_{jt}^K\}_{j \in \mathcal{J}^K} \right\},$$

where the agent receives information about proposition p^k from individuals in $\mathcal{J}^k \subseteq \mathcal{J}$.

The agent might try updating according to Bayes' rule:

$$Pr(T(p^k) = 1 | \sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k}) = \frac{Pr(\sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k} | T(p^k) = 1) Pr(T(p^k) = 1)}{Pr(\sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k})}$$

Using beliefs λ_{it}^k as the agent's prior, this would imply updating as

$$\lambda_{it+1}^k = \frac{f(\sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k} | T(p^k) = 1) \lambda_{it}^k}{f(\sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k} | T(p^k) = 1) \lambda_{it}^k + f(\sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k} | T(p^k) = 0) (1 - \lambda_{it}^k)}.$$

In a related setting, Acemoglu et al. (2016) document the restrictions that would be required on the conditional pdfs $f(\cdot | T(p^k))$ for there to be asymptotic agreement across agents. More fundamentally, correctly specifying the likelihood function $f(\sigma_{it}^k, \{\sigma_{jt}^k\}_{j \in \mathcal{J}^k} | T(p^k))$ can require unrealistic assumptions about the information and computation available to the agent (Acemoglu and Ozdaglar (2011)).¹¹ Weakening these assumptions is a key motivation of the literature on non-Bayesian social learning (Molavi et al. (2017)).

Correctly specifying the likelihood function $f(\varphi_{it}^k(W_{it}^*), \{\varphi_j^k(W_{jt})\}_{j \in \mathcal{J}^k} | T(p^k))$ would require not only that the agent know the sampling processes for W_{it} and W_{jt} conditional on $T(p^k)$, but also the models $\{\varphi_j^k\}_{j \in \mathcal{J}^k}$. I rule out Bayesian social learning by restricting social information to beliefs and signals, assuming that the agent does not observe the additional information required to specify the likelihood function:

(A1) Imperfect Communication: Agent i can only observe point estimates λ_{jt}^k and σ_{jt}^k . She cannot observe measures of the sender's ambiguity $\Lambda_{jt}^{k*}, \theta_{jt}^k$ or their model $\varphi_j^k \forall j, t, k$

The issue captured by A1 is that data must be transformed into information using a model, and it is difficult for individuals to communicate this process. Therefore, valuable details are lost

¹¹Benoît and Dubra (2015) and Andreoni and Mylovanov (2012) study polarization under private learning when agents disagree about $f(\sigma_{it}^{k*} | T(p^k))$. Alternatively, in this context their analyses could be interpreted as agents having different models for private learning φ_i^k , each proposition p^k being a conjunction of simple propositions $p^k = p^{k'} \wedge p^{k''}$, and W_{it}^* being revealed at different subperiods of t for $p^{k'}$ and $p^{k''}$.

relative to directly observing the data when information is obtained socially.¹² This assumption is appealing because there is a well-documented tendency for researchers to focus on communicating their point estimates σ_{it}^k without communicating about their models φ_i^k or measures of uncertainty θ_{it}^k (Manski (2007)). This practice also extends to official statistics (Manski (2015)) and surveys like the Survey of Professional Forecasters (Manski (2017)). Nevertheless, there is clear evidence that decision makers treat conflict and imprecision differently (Smithson (1999), Smithson (2013), Cabantous (2007)). A natural extension for future work would be to relax A1 to allow for agents to communicate Λ_{jt}^{k*} and θ_{jt}^k in addition to point estimates.

With A1 ruling out Bayesian social learning, I assume that the agent uses signals in an effort to replicate classical inference. Given a loss function \mathcal{L} , the agent's problem is to choose functions f^k from some set \mathcal{F} to solve the problem

$$\begin{aligned} \min_{f^1, \dots, f^K \in \mathcal{F}} \quad & \sum_{k=1}^K \mathcal{L} \left(\mathbb{E} \left[\mu_i^{k*} - \lim_{t \rightarrow \infty} \lambda_{it+1}^k \right] \right) \\ \text{s.t. } \quad & (\mathcal{I}_{it}, \mathcal{I}_{Jt}) \\ & \hat{\sigma}_{it}^k = f^k(\mathcal{I}_{it}, \mathcal{I}_{Jt}) \quad \text{for } k = 1, \dots, K \\ & \lambda_{it+1}^k = \beta_t \hat{\sigma}_{it}^k + (1 - \beta_t) \lambda_{it}^k \quad \text{for } k = 1, \dots, K \end{aligned} \quad (3)$$

I will refer to the agent's construction of her unobserved, high-quality signals $\hat{\sigma}_{it}^k$ as her inferred signals. A natural restriction on \mathcal{F} is to make inferred signals a weighted average of directly- and socially-observed signals. In this case, f^k can be written as

$$\hat{\sigma}_{it}^k = \underbrace{\theta_i^k}_{\substack{\text{share of signal} \\ \text{directly-observed}}} \sigma_{it}^k + \underbrace{(1 - \theta_{it}^k)}_{\substack{\text{share of signal} \\ \text{socially-observed}}} \sigma_{Jt}^k.$$

This restriction reframes the choice of f^k as the choice of σ_{Jt}^k .¹³ Posing the inferred signals as weighted averages also gives an interpretation to θ_{it}^k as the agent's subjective judgment about the credibility of her modeling assumptions and/or a measure of the quality of her data.

4 Solving the Agent's Problem

4.1 A Feasible Solution Requiring Strong Assumptions

I begin by further restricting \mathcal{F} so that the agent infers the signal σ_{Jt}^k using linear opinion pooling (LOP) of social signals. 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

¹²A1 imposes a version of word-of-mouth learning (Ellison and Fudenberg (1995), Banerjee and Fudenberg (2004)).

¹³Assuming that $\{W_{it}\}_{t=1}^{\infty}$ and $\{\varphi_i^k\}_{k=1}^K$ are exogenous, both $\{\sigma_{it}^k\}_{t=1}^{\infty}$ and $\{\theta_{it}^k\}_{k=1, t=1}^{K, \infty}$ are given. Thus, in an abuse of notation, I will refer to f^k both as the function determining $\hat{\sigma}_{it}^k$ and as the function determining σ_{Jt}^k .

agent's problem (where fully Bayesian updating may not be feasible). And second, I investigate LOP because repeated linear opinion pooling results in DeGroot updating if data are only observed in the first period, and signals continue to be sent in later periods.

Proposition 1 (DeGroot). *If data are only observed once at $t = 1$, the agent sets $\lambda_{i1}^k = \sigma_{i1}^k$, $\theta_{it}^k = \theta_{i1}^k$ for all $t > 1$, and subsequent signals are interchangeable with beliefs ($\sigma_{it}^k = \lambda_{it}^k$ and $\sigma_{jt}^k = \lambda_{jt}^k$ for $j \geq 2$), then linear opinion pooling where the agent constructs her inferred signals for $t \geq 2$ as*

$$\widehat{\sigma}_{it}^k = \theta_i^k \sigma_{it}^k + (1 - \theta_i^k) \sigma_{Jt}^k \quad \text{where} \quad (4)$$

$$\sigma_{Jt}^k = \sum_{j \in \mathcal{J}^k} \underbrace{w_j^k}_{\substack{\text{share of social signal} \\ \text{from individual } j}} \sigma_j^k \quad \text{with } w_j^k \geq 0 \quad \forall j \in \mathcal{J}^k, \quad \sum_{j \in \mathcal{J}^k} w_j^k = 1 \quad (5)$$

is equivalent to DeGroot updating where $\lambda_{t+1}^k = \Omega_t^k \lambda_t^k$ and the entries of Ω_t^k are

$$\begin{aligned} \omega_{iit}^k &= \beta_t \theta_i^k + (1 - \beta_t) \\ \omega_{ijt}^k &= \beta_t (1 - \theta_i^k) w_j^k. \end{aligned}$$

Proof. As hypothesized, set $\lambda_{i1}^k = \sigma_{i1}^k$. For $t \geq 2$, the equality of beliefs and signals, together with the updating equation in the agent's problem (3) imply that

$$\begin{aligned} \sigma_{it+1}^k &= \beta_t \widehat{\sigma}_{it}^k + (1 - \beta_t) \sigma_{it}^k \\ &= \beta_t \theta_i^k \sigma_{it}^k + (1 - \beta_t) \sigma_{it}^k + \beta_t (1 - \theta_i^k) \sum_{j \in \mathcal{J}^k} w_j^k \sigma_{jt}^k. \end{aligned}$$

□

Furthermore, when the data observed in $t = 1$ generate unbiased point-estimates of signals, repeated linear opinion pooling/DeGroot updating solves the agent's problem.

Proposition 2 (Unbiased Social Signals). *Maintain assumption A1 and assume again, as we did in the case of private learning, that*

$$(A2) \text{ Averaging Signals: } \beta_t = 1/t, \text{ so that } \beta_t \widehat{\sigma}_{it}^k + (1 - \beta_t) \lambda_{it}^k = \frac{1}{t} \sum_{n=1}^t \widehat{\sigma}_{in}^k$$

If the observed data yield unbiased signals

$$(A3) \text{ Private signals are iid with } \mathbb{E}[\sigma_{it}^{k*}] \equiv \mu_i^{k*} = \mu_i^k \equiv \mathbb{E}[\sigma_i^k], \text{ and}$$

$$(A4a) \text{ Social signals are iid for each } j \in \mathcal{J}^k \text{ with } \mathbb{E}[\sigma_{jt}^{k*}] \equiv \mu_j^{k*} = \mu_j^k \equiv \mathbb{E}[\sigma_j^k] \quad \forall j \in \mathcal{J}^k,$$

then repeated linear opinion pooling/DeGroot updating following Equations 4 and 5 solves the agent's problem.

Proof. Proposition 6 in Golub and Sadler (2016) states that as long as Ω^k is strongly connected and primitive, then

$$\lim_{t \rightarrow \infty} \sigma_{it+1}^k = \sum_{n=1}^{J+1} \pi_n^k \sigma_{n1}^k$$

where π_n^k is n 's left-hand eigenvector centrality in Ω^k .¹⁴ Additionally assuming that each weight w_i is strictly positive and that $\theta^k \in (0, 1)$ to ensure Ω^k is primitive, then since $\sum_{n=1}^{J+1} \pi_n^k = 1$ and $\mathbb{E}[\sigma_{n1}^k] = \mu_i^{k*}$ for all n , we know that

$$\mathbb{E}[\mu_i^{k*} - \lim_{t \rightarrow \infty} \lambda_{it+1}^k] = \mathbb{E}[\mu_i^{k*} - \sum_{n=1}^{J+1} \pi_n \sigma_{n1}^k] = \mu_i^{k*} - \mu_i^{k*} = 0.$$

□

4.2 A Fully General Solution that Is Not Feasible

The assumptions justifying DeGroot updating as a solution to the agent's problem are very restrictive, and need not hold in the general case of the agent's problem. For example, the agent might observe data and signals in each period, and social signals could potentially be biased. In this case, the agent can still solve her problem if she has a model capable of accurately interpreting the social signals she receives.

Proposition 3 (Biased Social Signals). *Maintain assumptions A1-A3. Now suppose that the agent receives biased signals in the sense that $\mathbb{E}[\sigma_{jt}^k] \neq \mu_{it}^{k*}$, but that the agent has successfully engaged in statistical learning in the following sense:*

(A4b) *The agent has a model of social learning g^k that interprets social signals as $s_{jt}^k = g^k(\mathcal{I}_{it}, \mathcal{I}_{Jt})$. The s_{jt}^k are iid for each $j \in \mathcal{J}^k$ with $\mathbb{E}[\sigma_{it}^{k*}] \equiv \mu_i^{k*} = \mathbb{E}[s_{jt}^k] \forall j \in \mathcal{J}^k$.*

Then linear opinion pooling where the agent constructs unobserved high-quality signals with her model as

$$\hat{\sigma}_{it}^k = \theta_{it}^k \sigma_{it}^k + (1 - \theta_{it}^k) \sigma_{Jt}^k \quad \text{where} \quad (6)$$

$$\sigma_{Jt}^k = \sum_{j \in \mathcal{J}^k} w_{jt}^k s_{jt}^k \quad \text{with } w_{jt}^k \geq 0 \quad \forall j \in \mathcal{J}^k, \quad \sum_{j \in \mathcal{J}^k} w_{jt}^k = 1 \quad (7)$$

$$s_{jt}^k = g^k(\mathcal{I}_{it}, \mathcal{I}_{Jt}) \quad (8)$$

solves the agent's problem.

Proof. By A2 we know that $\lim_{t \rightarrow \infty} \lambda_{it+1}^k = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \hat{\sigma}_{in}^k$. If the signals are iid, then by the law of large numbers we know that $\lim_{t \rightarrow \infty} \lambda_{it+1}^k = \mathbb{E}[\hat{\sigma}_{it}^k]$. After repeatedly applying the linearity

¹⁴A network is strongly connected if any agent i has a directed path in the network to any agent j . A strongly connected network is primitive if each agent attaches a non-zero weight to each agent.

of the expectations operator, A3 and A4b imply that

$$\begin{aligned}
\lim_{t \rightarrow \infty} \lambda_{it+1}^k &= \mathbb{E}[\hat{\sigma}_{it}^k] = \mathbb{E}[\bar{\theta}_i^k \sigma_{it}^k + (1 - \bar{\theta}_i^k) \sigma_{jt}^k] = \bar{\theta}_i^k \mathbb{E}[\sigma_{it}^k] + (1 - \bar{\theta}_i^k) \mathbb{E}[\sigma_{jt}^k] \\
&= \bar{\theta}_i^k \mathbb{E}[\sigma_{it}^k] + (1 - \bar{\theta}_i^k) \mathbb{E}\left[\sum_{j \in \mathcal{J}^k} w_{jt}^k s_{jt}^k\right] = \bar{\theta}_i^k \mathbb{E}[\sigma_{it}^k] + (1 - \bar{\theta}_i^k) \sum_{j \in \mathcal{J}^k} w_{jt}^k \mathbb{E}[s_{jt}^k] \\
&= \bar{\theta}_i^k \mu_{it}^k + (1 - \bar{\theta}_i^k) \sum_{j \in \mathcal{J}^k} w_{jt}^k \mu_{jt}^k \\
&= \mu_i^{k*}.
\end{aligned} \tag{9}$$

□

Unfortunately, Proposition 3 is just a theoretical result, since the agent does not know the correct g^k . Moreover, the agent cannot find the correct g^k .

Determining the correct model for interpreting signals could be viewed as a statistical learning problem. To put the problem in the notation of statistical learning, here we adapt some notation from Chapter 1 of Schölkopf and Smola (2002). Suppose that at time t the agent observes outcomes of high-quality data for herself and everyone in her network,

$$(x_i^1, y_i^1), \dots, (x_i^K, y_i^K) \in \mathcal{X} \times [0, 1]$$

where $x_i^k = x_i$ for $k = 1, \dots, K$ and

$$\begin{aligned}
(x_i^k, y_i^k) &= \left([\mathcal{I}_i, \mathcal{I}_j], \sigma_i^{k*} \right) \\
&\equiv \left(\left[\begin{array}{ccc|ccc} \lambda_{it}^1 & \sigma_{it}^1 & \theta_{it}^1 & \lambda_{1t}^1 & \dots & \lambda_{Jt}^1 & \sigma_{1t}^1 & \dots & \sigma_{Jt}^1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \lambda_{it}^K & \sigma_{it}^K & \theta_{it}^K & \lambda_{1t}^K & \dots & \lambda_{Jt}^K & \sigma_{1t}^K & \dots & \sigma_{Jt}^K \end{array} \right], \sigma_{it}^{k*} \right).
\end{aligned} \tag{10}$$

The agent constructs her high-quality signal as:

$$y_{it}^k = f^k(x_{it}^k)$$

where the function $f^k \in \mathcal{F}$ is chosen to minimize the expectation of an empirical risk function like mean squared error

$$R[f^k] = \mathbb{E}_t[c(x_{it}^k, y_{it}^k, f^k(x_{it}^k))] = \mathbb{E}_t[(y_{it}^k - f^k(x_{it}^k))^2],$$

and the set of functions \mathcal{F} is chosen with the help of Vapnik-Chervonenkis (VC) theory to avoid overfitting.

A fundamental problem, however, is that the agent never observes $y_{it}^k = \sigma_{it}^{k*}$. Thus, the agent cannot construct the risk function, so choosing f^k according to this criterion is not a well-posed

problem.

4.3 A Feasible and General Solution

While the agent cannot solve her problem under the most general assumptions, she can solve her problem under more general assumptions than those required for DeGroot updating to be a solution.¹⁵ Note that the agent's problem can be posed as inferring the value of $\pi_{ijt}^k \equiv \sigma_{it}^{k*} - \sigma_{jt}^k$ so that she can properly adjust each socially observed signal σ_{jt}^k . To think about the agent's problem in terms of inference under uncertainty about π_{ijt}^k , first denote the sampling process for data W_{it}^* by Γ_i^* , which results in the sampling distribution $\sigma_{it}^{k*} \sim G_i^{k*}$, with mean μ_i^{k*} . Denote the sampling process for data W_{it} by Γ_i , where the sampling distribution $\sigma_{it}^k \sim G_i^k$ has mean μ_i^k . Define analogous sampling processes and sampling distributions for the signals from each agent j .

There are five factors that drive the distribution of π_{ijt}^k :

I (Random Sampling Error): $\Gamma_i^* = \Gamma_j$

II (Biased Sampling Process): $\Gamma_i^* \neq \Gamma_j$

III (Different Models): $\varphi_i^k \neq \varphi_j^k$

IV (Social Influence): j 's model of social learning and/or her network

V (Strategic Reporting): Part of what determines the function φ_j^k is strategic reporting

Assuming that only factors (I)-(III) drive the distribution of $\sigma_{it}^{k*} - \sigma_{jt}^k$, we have

$$\mathbb{E}[\lambda_{it}^{k*} - \lambda_{jt}^k] = \mathbb{E}[\sigma_{it}^{k*} - \sigma_{jt}^k] = \mu_i^{k*} - \mu_j^k,$$

so that linear opinion pooling of signals adjusted using

$$s_{jt}^k = \sigma_{jt}^k + (\lambda_{it}^{k*} - \lambda_{jt}^k) \tag{11}$$

solves the agent's problem.¹⁶ Formally stated,

Proposition 4 (Biased Social Signals with Known Sources of Bias). *Maintain assumptions A1-A3.*

Now suppose that the agent receives biased signals in the sense that $\mathbb{E}[\sigma_{jt}^k] \neq \mu_{it}^{k}$, but that:*

(A4c) *Only factors (II) and (III) contribute to the fact that $\mathbb{E}[\sigma_{jt}^k] \neq \mu_{it}^{k*}$.*

¹⁵This moves beyond the cases studied in Aliprantis (2017).

¹⁶DeGroot (1974) studies social learning under a version of (I). Manski (2004) studies social learning under a version of (II) focusing more attention on the construction of the interval $\Lambda_i^{k*} \equiv \mathbb{E}[\Lambda_{it}^{k*}]$. Manski (2004) models communication differently than (A1): assuming that outcomes are directly observed rules out (III) and makes the agent's problem more specifically about how to shrink the interval Λ_i^{k*} after observing potentially non-iid data. This paper adopts a more stylized version of (II) in order to study social learning under (I)+(II)+(III).

Then linear opinion pooling where the agent constructs unobserved high-quality signals with her model as

$$\widehat{\sigma}_{it}^k = \theta_{it}^k \sigma_{it}^k + (1 - \theta_{it}^k) \sigma_{Jt}^k \quad \text{where} \quad (12)$$

$$\sigma_{Jt}^k = \sum_{j \in \mathcal{J}^k} w_{jt}^k s_{jt}^k \quad \text{with } w_{jt}^k \geq 0 \quad \forall j \in \mathcal{J}^k, \quad \sum_{j \in \mathcal{J}^k} w_{jt}^k = 1 \quad (13)$$

$$s_{jt}^k = \sigma_{jt}^k + \left(\lambda_{it}^{k*} - \lambda_{jt}^k \right) \quad (14)$$

solves the agent's problem.

Proof. Again repeatedly applying the linearity of the expectations operator, A3 and A4c imply that

$$\begin{aligned} \lim_{t \rightarrow \infty} \lambda_{it+1}^k &= \mathbb{E}[\widehat{\sigma}_{it}^k] = \mathbb{E}[\overline{\theta}_i^k \sigma_{it}^k + (1 - \overline{\theta}_i^k) \sigma_{Jt}^k] &&= \overline{\theta}_i^k \mathbb{E}[\sigma_{it}^k] + (1 - \overline{\theta}_i^k) \mathbb{E}[\sigma_{Jt}^k] \\ &= \overline{\theta}_i^k \mathbb{E}[\sigma_{it}^k] + (1 - \overline{\theta}_i^k) \mathbb{E}\left[\sum_{j \in \mathcal{J}^k} w_{jt}^k s_{jt}^k \right] &&= \overline{\theta}_i^k \mathbb{E}[\sigma_{it}^k] + (1 - \overline{\theta}_i^k) \sum_{j \in \mathcal{J}^k} w_{jt}^k \mathbb{E}[s_{jt}^k] \\ &= \overline{\theta}_i^k \mu_{it}^k + (1 - \overline{\theta}_i^k) \sum_{j \in \mathcal{J}^k} w_{jt}^k \left(\mathbb{E}[\sigma_{jt}^k] - \lambda_{jt}^k + \lambda_{it}^{k*} \right) && (15) \\ &= \mu_i^{k*}. \end{aligned}$$

□

5 Empirical Implementation and Opinion Dynamics

5.1 Empirical Implementation

Not observing the right hand side of Equation 11, the agent might assume she can replace the unobserved quantity λ_{it}^{k*} with the observed quantity λ_{it}^k to translate signals as:

$$\text{(A5): } s_{jt}^k = \sigma_{jt}^k + \left(\lambda_{it}^k - \lambda_{jt}^k \right).$$

The agent knows that A5 does not necessarily solve her problem if factors (IV) and (V) help to drive the distribution of π_{ijt}^k . Thus, she would like to lean on A5 as little as possible, or to use signals from the senders for whom A5 is the most credible.

Define Δ_{ijt} as the credibility that the agent i deems to A5 as applied to sender j 's signal, and $\Delta_{it}^k = f(\{\Delta_{ijt}\}_{j \in \mathcal{J}^k}) \in [0, 1]$ as the total credibility agent i deems to A5 as applied to the socially-available information on proposition k . Then if we define $s_{jt}^k = g^k(\mathcal{I}_{it}, \mathcal{I}_{Jt})$ by A5 and denote the weights attached to each of these interpreted signals as the sender's relative credibility

$w_{jt}^k = \frac{\Delta_{ijt}}{\sum_{j=1}^{J^k} \Delta_{ijt}}$, we can define

$$s_{Jt}^k = \sum_{j=1}^{J^k} w_{jt}^k s_{jt}^k$$

(assuming $\Delta_{ijt} \neq 0$ for some $j \in \mathcal{J}^k$), and the agent could infer signals using $\widehat{\sigma}_{it}^k = f^k(\mathcal{I}_{it}, \mathcal{I}_{Jt})$ defined as

$$\widehat{\sigma}_{it}^k = \theta_{it}^k \sigma_{it}^k + (1 - \theta_{it}^k) \left[\Delta_{it}^k s_{Jt}^k + (1 - \Delta_{it}^k) \sigma_{it}^k \right]. \quad (f^k)$$

Note that A5 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).

The final hurdle to empirically implementing the agent's model of social learning is empirically determining the credibility the agent gives to A5 as a means of adjusting a signal from sender j , Δ_{ijt} . 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.¹⁷

(A6) $\lambda_{it}^{k*} - \lambda_{jt}^k = \lambda_{it}^{k'*} - \lambda_{jt}^{k'}$ for all $t \in \mathbb{N}$, $j \in \mathcal{J}$, and $p^k, p^{k'} \in \mathcal{K}$

Under A6, the agent could use the distribution of $\delta_{ijt}^k \equiv \lambda_{it}^k - \lambda_{jt}^k$ over the proposition space as a means of assigning credibility to A5 applied to sender 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 agent j if the distribution of δ_{ijt}^k were a degenerate distribution, and would be least informed were δ_{ijt}^k to follow a uniform distribution. In the Figure, $f(\delta_{i \max t}^k) = U[-1, 1]$ has the maximum entropy (representing the least informative sender), senders $j = 1, 2, 3$ have high entropy (representing low information), senders $j = 4, 5, 6$ have medium entropy (representing moderate information), and senders $j = 7, 8, 9$ have low entropy (representing high information).

¹⁷We 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 A6 could help to explain why so much attention is paid to seemingly unimportant propositions in the US' "culture war."

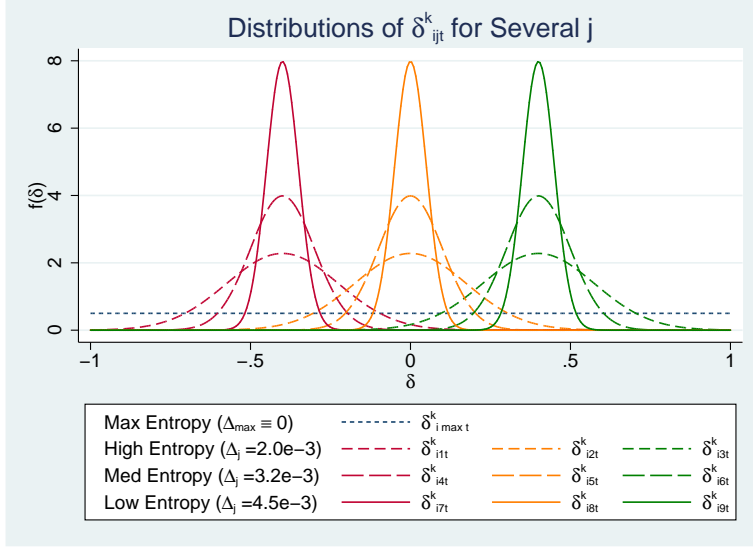


Figure 2: Entropy of Sender j

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 how uncertain she is about her disagreement with a sender.

Assuming that the distribution of δ_{ijt}^k across propositions $k \in \{1, \dots, K\}$ has a probability density function Q_{ijt} , the informational content of sender j ’s signal (or certainty about δ_{ijt}^k) can be defined as the Kullback-Leibler divergence from the uniform distribution over $[-1, 1]$:

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

which is a measure of the difference in entropy of Q_{ijt} relative to the maximum entropy (uncertainty) distribution.¹⁸ I assume that the credibility the agent assigns to interpreted signals from sender j is $\Delta_{ijt} \equiv \rho(D_{KL}(Q_{ijt} : U))$.

5.2 Simulations

I now show simulations illustrating some of the belief dynamics that f^k 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

¹⁸In another setting, Zanardo (2017) shows that Kullback-Leibler relative entropy satisfies desirable axioms as a measure of disagreement.

settings.

In both numerical experiments, I consider a network of $J + 1 = 300$ individuals learning about $K = 30$ propositions, with $\theta_{it}^k = 0.1$ for all k, i , and t . Each agent assesses the credibility of A5 applied to sender j 's signal according to

$$\Delta_{ijt} \equiv \rho(D_{KL}(Q_{ijt} : U)) = [\gamma_1 D_{KL}(Q_{ijt} : U)]^{\gamma_2}$$

where $(\gamma_1, \gamma_2) = (100, 8) \in [0, \infty) \times [0, \infty)$ can be thought of as distrust parameters, and the credibility given to the sum of social information in the network $\Delta_{it}^k = 1$ for all k, i , and t . I consider directly-observed data most likely to result in agreement, or convergence to a degenerate distribution. I assume all individuals directly-observe data generating the same signals at all times for all propositions:

$$\sigma_{it}^k = 0.5 \quad \forall \quad t = 2, 3, 4, \dots; \quad k = 1, \dots, K; \quad \text{and} \quad i = 1, \dots, J + 1.$$

What varies across Experiments 1 and 2 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 λ_{i1}^k are generated as follows:

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

with

$$\lambda_{i1}^k = \min \left\{ \max \left\{ \bar{\lambda}_i^k, 0 \right\}, 1 \right\} \quad \forall k = 1, \dots, K.$$

Figure 3a shows the initial distribution of beliefs by clusters for proposition p^1 .

Figure 3b shows that as individuals in the network update using f^k , 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 tends to agree with those in a widely-distributed cluster (unbiased but imprecise), but tends to disagree with those in a tightly-distributed cluster (biased but precise), that agent will rely more on interpreted signals from the disagreeing cluster, and this can cause her to overcompensate when they provide her with unbiased signals.

This is what happens to agents in cluster \mathcal{C}_1 in Experiment 1 (Figure 3c). 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 σ_{it}^{1*} , so the interpreted signal s_{jt}^1 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).

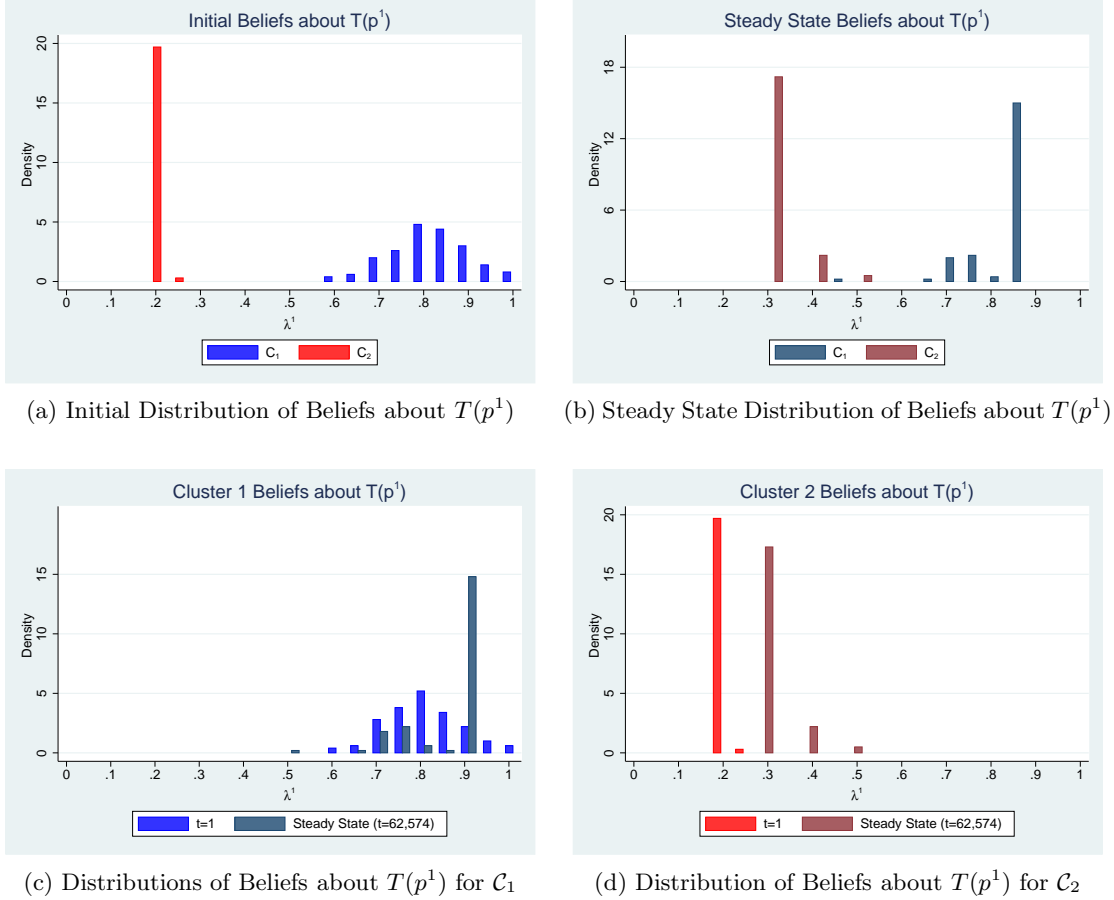


Figure 3: Beliefs in Experiment 1

Steady state belief distributions are those for which $\max_{k,i} \|\lambda_{it}^k - \lambda_{it+1}^k\| < 1e-6$.

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.

6 Conclusion

This paper presented a positive theory of belief formation. I proposed one way that an agent might choose a single subjective probability from a set of possible probabilities. When the agent faces ambiguity because her directly-observed data only allow her to partially identify a signal about the truth of a proposition, she might seek to learn from individuals in her social network. Assuming that communication is imperfect, so that individuals can only communicate a point estimate of their signals and beliefs, the agent must determine how to combine the signals she observes.

I showed that when signals are unbiased, linear opinion pooling of signals generates DeGroot updating, and is able to replicate classical inference with high-quality data yielding point-valued signals. When individuals in the agent's network have different models or access to different quality data, then their signals will not be unbiased. In this case, the agent might still form beliefs by

linear opinion pooling on interpreted signals. I considered how the agent might engage in social learning by interpreting social signals, and how this approach to belief formation can lead to a non-constricting rule of thumb capable of generating polarization and clustered, permanent disagreement on a connected network where everyone observes the same data and processes that data with the same model.

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), and 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)). 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 social information to the empirical studies on belief revision surveyed in Manski (2017) would help to discipline all of this theoretical work.

References

- Acemoglu, D., V. Chernozhukov, and M. Yildiz (2016). Fragility of asymptotic agreement under Bayesian learning. *Theoretical Economics* 11, 187–225.
- Acemoglu, D. and A. Ozdaglar (2011). Opinion dynamics and learning in social networks. *Dynamic Games and Applications* 1(1), 3–49.
- 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. (2017, 10–14 Jul). Differences of opinion. In A. Antonucci, G. Corani, I. Couso, and S. Destercke (Eds.), *Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications*, Volume 62 of *Proceedings of Machine Learning Research*, pp. 1–12.
- Aliprantis, D. and F. G.-C. Richter (2016). Evidence of neighborhood effects from MTO: LATEs of neighborhood quality. *Federal Reserve Bank of Cleveland WP 12-08R*.
- Andreoni, J. and T. Mylovonov (2012). Diverging opinions. *American Economic Journal: Microeconomics* 4(1), 209–232.
- Artstein, Z. and R. A. Vitale (1975). A strong law of large numbers for random compact sets. *The Annals of Probability* 3(5), 879–882.

- Banerjee, A. and D. Fudenberg (2004). Word-of-mouth learning. *Games and Economic Behavior* 46, 1–22.
- Benoît, J.-P. and J. Dubra (2015). A theory of rational attitude polarization. *Mimeo., London Business School*.
- 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.
- Bollinger, C. R. and M. van Hasselt (2017). A Bayesian analysis of binary misclassification. *Economics Letters* 156, 68 – 73.
- Bradley, R. (2017, Jul). Learning from others: conditioning versus averaging. *Theory and Decision*.
- Burszтын, 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.
- Chernozhukov, V., H. Hong, and E. Tamer (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica* 75(5), 1243–1284.
- 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.
- Ellison, G. and D. Fudenberg (1995). Word-of-mouth communication and social learning. *The Quarterly Journal of Economics* 110(1), 93–125.

- 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.
- Gajdos, T. and J.-C. Vergnaud (2013). Decisions with conflicting and imprecise information. *Social Choice and Welfare* 41(2), 427–452.
- Gentzkow, M. and J. M. Shapiro (2011). Ideological segregation online and offline. *The Quarterly Journal of Economics* 126(4), 1799–1839.
- Gilboa, I. and M. Marinacci (2013). Ambiguity and the Bayesian paradigm. In D. Acemoglu, M. Arellano, and E. Dekel (Eds.), *Advances in Economics and Econometrics: Tenth World Congress: Economic Theory*, Volume I, Chapter 7, pp. 179–242. Cambridge University Press.
- Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2), 141–153.
- Golub, B. and M. O. Jackson (2012). How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics* 127(3), 1287–1338.
- 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).
- Imbens, G. W. and C. F. Manski (2004). Confidence intervals for partially identified parameters. *Econometrica* 72(6), 1845–1857.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton: Princeton University Press.
- Jadbabaie, A., P. Molavi, A. Sandroni, and A. Tahbaz-Salehi (2012). Non-Bayesian social learning. *Games and Economic Behavior* 76(1), 210 – 225.
- 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]*.
- 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. (2004). Social learning from private experiences: The dynamics of the selection problem. *The Review of Economic Studies* 71(2), 443–458.
- 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. (2015). Communicating uncertainty in official economic statistics: An appraisal fifty years after Morgenstern. *Journal of Economic Literature* 53(3), 631–653.
- Manski, C. F. (2016). Credible interval estimates for official statistics with survey nonresponse. *Journal of Econometrics* 191(2), 293–301.
- Manski, C. F. (2017). Survey measurement of probabilistic macroeconomic expectations: Progress and promise. *NBER WP 23418*.
- Molavi, P., A. Tahbaz-Salehi, and A. Jadbabaie (2017). Foundations of non-Bayesian social learning. *Mimeo., Northwestern University*.
- Moon, H. R. and F. Schorfheide (2012). Bayesian and frequentist inference in partially identified models. *Econometrica* 80(2), 755–782.
- 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.
- Romano, J. P. and A. M. Shaikh (2010). Inference for the identified set in partially identified econometric models. *Econometrica* 78(1), 169–211.
- Schölkopf, B. and A. J. Smola (2002). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press.

- 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.
- Smithson, M. (2013). Conflict and ambiguity: Preliminary models and empirical tests. In *Proceedings of the 8th International Symposium on Imprecise Probability: Theories and Applications (Compiègne)*.
- 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.
- Stoye, J. (2009). More on confidence intervals for partially identified parameters. *Econometrica* 77(4), 1299–1315.
- Tamer, E. (2010). Partial identification in econometrics. *Annual Review of Economics* 2(1), 167–195.
- 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 Appendix: Notation

The Agent's Problem

p^k	Proposition $k \in \{1, \dots, K\}$
$T(p^k)$	Truth-value of p^k , taking values in $\{0, 1\}$
λ_{it}^k	Agent i 's beliefs at time t about $T(p^k)$, taking values in $[0, 1]$
φ_i^k	Agent i 's model for transforming data into signals about $T(p^k)$
W_{it}^*	Data directly-observed by agent i yielding point-valued signals about $T(p^k)$ via φ_i^k
σ_{it}^{k*}	A point-valued signal
μ_i^{k*}	Agent i 's ideal belief when $\mu_i^{k*} = \mathbb{E}_t[\sigma_{it}^{k*}]$
W_{it}	Data yielding set-valued signals about $T(p^k)$ via φ_i^k
Λ_{it}^{k*}	Set of possible beliefs (imprecise probability)
θ_{it}^k	Parameter that, along with σ_{it}^k , defines the correspondence $\varphi_i^k : W_{it} \Rightarrow [\underline{\sigma}_{it}^{k*}, \overline{\sigma}_{it}^{k*}]$

The Agent's Model of Social Learning

\mathcal{J}	Network of J individuals sending signals to agent i
$\mathcal{J}^k \subseteq \mathcal{J}$	Set of J^k individuals sending signals to agent i about proposition p^k
\mathcal{I}_{it}	Agent i 's directly-observed information set
\mathcal{I}_{Jt}	Agent i 's socially-observed information set
φ_j^k	Individual j 's model for transforming data into signals about $T(p^k)$
W_{jt}	Individual j 's directly-observed data
σ_{jt}^k	Individual j 's point-valued signal
f^k	Agent i 's model of social learning, or for interpreting signals from senders in her network
s_{jt}^k	Agent i 's interpreted signal from sender j regarding proposition p^k
δ_{ijt}^k	Disagreement between agent i and sender j about proposition k , or $\lambda_{it}^k - \lambda_{jt}^k$
Δ_{ijt}	Credibility agent i assesses to A5 when applied to sender j 's signals
Δ_{it}^k	Total credibility of interpreted signals on proposition p^k
w_{jt}	Weight given to interpreted signals from sender j
s_{Jt}^k	Inferred signal from all socially-observed signals
$\widehat{\sigma}_{it}^k$	Inferred signal from combining directly-observed data and socially-observed signals

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 $p_1 \in \{p_0, p_0 + \theta\}$.¹⁹ 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:

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

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

$$\begin{aligned} D = 1 &\iff U_1 \geq U_0 \iff p_1w - c \geq p_0c \\ &\iff \pi(p_0 + \theta)w + (1 - \pi)p_0w - c \geq \pi p_0w + (1 - \pi)p_0w \\ &\iff \pi\theta w \geq c & (16) \\ &\iff \pi \geq \frac{c}{\theta w}. & (17) \end{aligned}$$

Equation 16 shows that this comparison in Equation 17 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

$$D = 1 \iff \underline{U}_1 \geq \underline{U}_0 \tag{18}$$

¹⁹Analogous to the binary state of the world studied in the text, 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 18 can be stated as

$$\begin{aligned}D = 1 &\iff \underline{\pi}(p_0 + \theta)w + (1 - \underline{\pi})p_0w - c \geq \underline{\pi}p_0w + (1 - \underline{\pi})p_0w \\ &\iff \underline{\pi}\theta w \geq c\end{aligned}\tag{19}$$

$$\iff \underline{\pi} \geq \frac{c}{\theta w}.\tag{20}$$

Note the similarity between Equations 19 and 20 and Equations 16 and 17. 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

$$\begin{aligned}D = 1 &\iff \bar{R}_1 \leq \bar{R}_0 \\ &\iff c - \underline{\pi}\theta w \leq \bar{\pi}\theta w - c \\ &\iff (\underline{\pi} + \bar{\pi})\theta w \geq 2c\end{aligned}\tag{21}$$

$$\iff \frac{\underline{\pi} + \bar{\pi}}{2} \geq \frac{c}{\theta w}.\tag{22}$$

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 λ_{i1}^k are generated as follows:

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

with

$$\lambda_{i1}^k = \min \left\{ \max \left\{ \bar{\lambda}_i^k, 0 \right\}, 1 \right\} \quad \forall k = 1, \dots, K.$$

Figures 4a and 5a shows the initial distribution of beliefs by clusters for propositions p^1 and p^2 , respectively. Figures 4 and 5 show how beliefs evolve by cluster.

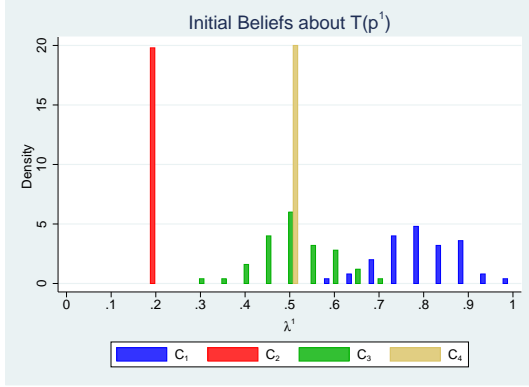
In Experiment 3, I also assume there are four clusters of individuals, \mathcal{C}_1 , \mathcal{C}_2 , \mathcal{C}_3 , and \mathcal{C}_4 with $\text{card}(\mathcal{C}_1) = 50$, $\text{card}(\mathcal{C}_2) = 100$, $\text{card}(\mathcal{C}_3) = 50$, and $\text{card}(\mathcal{C}_4) = 100$. Initial beliefs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

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


with

$$\lambda_{i1}^k = \min \left\{ \max \left\{ \bar{\lambda}_i^k, 0 \right\}, 1 \right\} \quad \forall k = 1, \dots, K.$$


Figures 6a and 7a shows the initial distribution of beliefs by clusters for propositions p^1 and p^2 , respectively. Figures 6 and 7 show how beliefs evolve by cluster.




(a) Initial Distribution of Beliefs about $T(p^1)$




(b) Steady State Distribution of Beliefs about $T(p^1)$




(c) Distributions of Beliefs about $T(p^1)$ for C_1



(d) Distribution of Beliefs about $T(p^1)$ for C_2



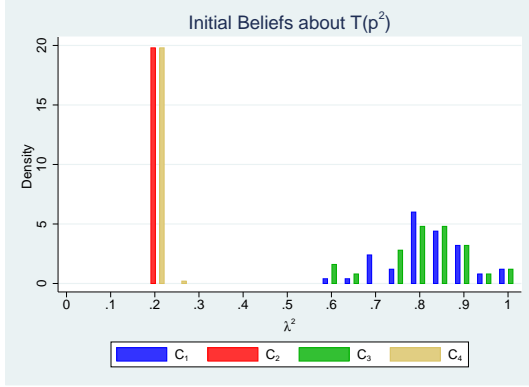
(e) Distributions of Beliefs about $T(p^1)$ for C_3




(f) Distribution of Beliefs about $T(p^1)$ for C_4

Figure 4: Beliefs in Experiment 2


Steady state belief distributions are those for which $\max_{k,i} \|\lambda_{it}^k - \lambda_{it+1}^k\| < 1e-6$.




(a) Initial Distribution of Beliefs about $T(p^1)$



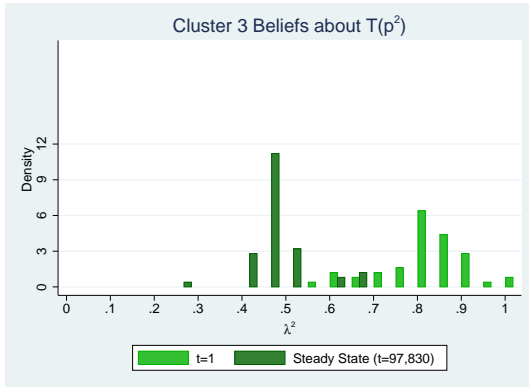
(b) Steady State Distribution of Beliefs about $T(p^1)$




(c) Distributions of Beliefs about $T(p^1)$ for C_1



(d) Distribution of Beliefs about $T(p^1)$ for C_2




(e) Distributions of Beliefs about $T(p^1)$ for C_3




(f) Distribution of Beliefs about $T(p^1)$ for C_4

Figure 5: Beliefs in Experiment 2


Steady state belief distributions are those for which $\max_{k,i} \|\lambda_{it}^k - \lambda_{it+1}^k\| < 1e-6$.




(a) Initial Distribution of Beliefs about $T(p^1)$




(b) Steady State Distribution of Beliefs about $T(p^1)$




(c) Distributions of Beliefs about $T(p^1)$ for C_1



(d) Distribution of Beliefs about $T(p^1)$ for C_2




(e) Distributions of Beliefs about $T(p^1)$ for C_3




(f) Distribution of Beliefs about $T(p^1)$ for C_4

Figure 6: Beliefs in Experiment 3


Steady state belief distributions are those for which $\max_{k,i} \|\lambda_{it}^k - \lambda_{it+1}^k\| < 1e-6$.




(a) Initial Distribution of Beliefs about $T(p^1)$




(b) Steady State Distribution of Beliefs about $T(p^1)$




(c) Distributions of Beliefs about $T(p^1)$ for C_1



(d) Distribution of Beliefs about $T(p^1)$ for C_2



(e) Distributions of Beliefs about $T(p^1)$ for C_3



(f) Distribution of Beliefs about $T(p^1)$ for C_4

Figure 7: Beliefs in Experiment 3

Steady state belief distributions are those for which $\max_{k,i} \|\lambda_{it}^k - \lambda_{it+1}^k\| < 1e-6$.