

Differences of Opinion

Dionissi Aliprantis

Federal Reserve Bank of Cleveland (US)

DIONISSI.ALIPRANTIS@CLEV.FRB.ORG

Abstract

This paper considers the resolution of ambiguity according to the scientific ideal of direct observation when there is a practical necessity for social learning. An agent faces ambiguity when she directly observes low-quality data yielding set-identified signals. I suppose the agent’s objective is to choose the single belief replicating what would occur with high-quality data yielding point-identified signals. I allow the agent to solve this missing data problem using signals 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; social learning; partial identification; causal inference; DeGroot learning rule; bounded confidence.

1. Introduction

We all hold beliefs based on limited personal experience. This is often due to logistical, and not philosophical, limitations. The scientific ideal of “**seeing for one’s self**” is subject to time and resource constraints that make it infeasible to personally verify all claims. How do we form beliefs based on evidence beyond our personal experience?

This paper studies scenarios of partial identification in which personal experience offers no guidance for belief formation beyond a range of possibilities. Consider the example of a high school guidance counselor advising minority students on whether to attend a selective or non-selective college. What is the probability that an advisee will graduate from the selective college? The counselor would face partial identification if the high school had not tracked the experiences of recent graduates, or had sent few students to selective colleges.

When facing partial identification, the counselor could provide his students with a range of probabilities. Alternatively, the counselor could provide a single probability based on information beyond his directly-observed data. The choice of a single probability would use the counselor’s judgment to combine his own experience; his discussions with others like counselors or teachers; and the conflicting estimates in the literature ([Arcidiacono and Lovenheim, 2016](#); [Alon and Tienda, 2005](#)). This paper models the counselor’s choice of a single probability.

The general setting begins with an agent who must form beliefs about a set of propositions. The agent can use a model to translate data into signals about the truth of each proposition. Under frequentist inference she may form beliefs as the mean of her signals observed over discrete time.¹

There are many situations in which the available data might only allow the agent to partially identify a signal. An obvious scenario pertains to causal propositions when one cannot easily ob-

1. For independent and identically distributed (iid) signals, the Law of Large Numbers ensures such beliefs will converge to the mean of the signal distribution.

serve the Data Generating Process (DGP) under controlled intervention. This situation is widespread in economics, with many important counterfactual outcomes waiting to be definitively quantified.² Beliefs of an agent observing iid partially-identified signals will converge to a set when formed by averaging signals observed over discrete time (Artstein and Vitale, 1975).

When the agent has a set of possible beliefs, or faces ambiguity, prominent decision rules instruct her to choose the single belief generating an extreme utility (Gilboa and Marinacci, 2013). For 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).

This paper separates belief formation from preferences: When choosing one belief, the objective is to accurately represent the DGP. While the scientific ideal is to attain this objective based on direct observation, no single belief cannot satisfy this ideal when directly-observed data are only capable of partial identification. However, a single belief can approximate the scientific ideal if data yielding point-identification can be inferred from second-hand observations.

I specify the agent’s problem as an attempt to replicate the beliefs she would have formed had she directly observed data yielding point-identified signals. The agent’s problem can be viewed as a missing data problem to be solved with signals observed through her social network. I assume that communication is imperfect, so that socially-observed signals are communicated alone, without the data or model used in their construction.

I first show that if the agent uses linear opinion pooling of signals, a common method for combining forecasts and estimates, she will follow the canonical DeGroot (1974) learning rule under a special case of observed data. I then show that such DeGroot updating solves the agent’s problem under additional assumptions on the homogeneity of data and models in the agent’s network.

Two issues argue for pushing beyond the assumptions necessary for DeGroot updating to solve the agent’s problem. The first is that the assumptions justifying DeGroot updating are strong. For example, individuals can be justified in using different models to interpret the same data (Al-Najjar, 2009), and the agent might observe new data over time (Jadbabaie et al., 2012).

Second, while DeGroot updating is the benchmark for non-Bayesian learning on social networks, a combination of theory and evidence motivates the desideratum of an alternative capable of generating polarization (Golub and Sadler, 2016). DeGroot learning and many of its generalizations converge to a degenerate distribution for connected networks (Jackson, 2008; Dandekar et al., 2013).³ However, an empirical analogue of a connected network - individuals exposed to sources of information contradicting their beliefs - is often observed together with persistent disagreement. Examples include scientific opinions when journals publish opposing research and public opinion when individuals are exposed to diverse news sources (Gentzkow and Shapiro, 2011).⁴ The emergence of “fake news” highlights this limitation of DeGroot updating.

2. In microeconomics alone it has proven difficult to ascertain outcomes under controlled interventions to neighborhood characteristics (Ludwig et al., 2008; Aliprantis, 2017), teacher characteristics (Rothstein, 2010; Kinsler, 2012), educational attainment (Angrist and Krueger, 1991; Aliprantis, 2012), minimum wages (Card and Krueger, 1994; Neumark and Wascher, 2000), unemployment benefits (Hagedorn et al., 2013; Farber and Valletta, 2015), income taxes (Manski, 2014), and right-to-carry laws (Manski and Pepper, 2015; Durlauf et al., 2016).

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

4. For example, there 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.

I show that linear opinion pooling can still solve the agent’s problem after weakening the assumptions justifying DeGroot updating. In contrast to DeGroot updating, though, this solution requires 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.

The agent might nevertheless choose a model for interpreting signals, just as methods for causal inference attempt to overcome the fundamental problem of evaluation. I study how the agent might use the model implied by a “reasonable” heuristic. The agent first interprets signals according to the model. The agent then uses the relative entropy of disagreement over all propositions to assess the credibility of applying the heuristic to each sender. The agent then combines interpreted signals, giving more weight to the interpreted signals from senders deemed most credible.

Although the updating rule tends to reach a 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. A key mechanism is generated by the use of relative entropy to assess the credibility of interpreted signals. 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.

Polarization is possible because in contrast to updating in DeGroot or bounded confidence models, the agent can update her beliefs away from a signal if it comes from a sender with whom she tends to disagree. 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 one way the agent might try to resolve the ambiguity she faces, using the signals she receives from individuals in her social network to form her beliefs. In the full paper I also show why finding a model of social learning to solve the agent’s problem is an ill-posed problem, and describe the implications of a heuristic the agent might use to specify a model of social learning. I further investigate the implications of this heuristic in greater detail, studying belief dynamics under one specification of the updating rule for several parameterizations under various network and proposition structures. Section 4 concludes.

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} .

5. This greatly simplifies the analysis. See Paris and Vencovská (1990) and Wilmers (2010) for implications of propositional calculus when considering propositions formed as compound propositions.

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

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

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

Now consider a setting in which the agent's directly-observed data W_{it} only allows her to set identify the true iid 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

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

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 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 $\theta_{it}^k = 1$.

The canonical example of the proposition $p^1 = \text{‘‘A given coin will land Heads.’’}$ helps to illustrate the difference between these settings. Suppose that high-quality data maps into signals generated by iid draws from a binomial distribution with probability 0.5 where $\sigma = 1$ if the coin lands Heads and $\sigma = 0$ if the coin lands Tails. In the case of high-quality data where $\theta_{it}^1 = 1$ for all t , $\sigma_{it}^k = \sigma_{it}^{k*}$, and so λ_{it+1}^{k*} can be calculated from (1) as the relative frequency of Heads, and will converge to 0.5 as $t \rightarrow \infty$.

In contrast, an agent with low-quality data mapping into signals represented by $\theta_{it}^1 = 0.2$ for all t will be subject to ambiguity in addition to risk.⁷ If the observed signal is Heads, then the agent can bound the true signal to be within $[0.2, 1]$. If the observed signal is Tails, then the agent bounds the true signal to be in $[0, 0.8]$. Thus as $t \rightarrow \infty$, the agent will infer that the mean of the true signals is $\mu_i^{k*} \in \Lambda_i^{k*} = [0.1, 0.9]$.⁸

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

8. Confidence 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 (2008).

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 ($\theta_{it}^k < 1$). For causal propositions, the difficulty of achieving identification is an obvious interpretation of signals having low quality. Examples abound of counterfactual outcomes that are difficult to quantify in microeconomics, macroeconomics, and finance because one cannot easily observe the Data Generating Process (DGP) under controlled intervention.⁹

Non-causal propositions can also have low-quality signals for reasons like survey non-response (Manski, 2015). 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 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 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).

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

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

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

9. See Footnote 2 for some examples from microeconomics.

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 Bayesian updating, or Bayesian social learning, 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

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

[Acemoglu et al. \(2016\)](#) show in a related setting that strong restrictions 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](#)).¹⁰ Weakening these assumptions is a key motivation of the literature on non-Bayesian social learning ([Molavi et al., 2015](#)).

Correctly specifying the likelihood function is the same as specifying

$$f(\varphi_{it}^k(W_{it}^*), \{\varphi_j^k(W_{jt})\}_{j \in \mathcal{J}^k} | T(p^k)),$$

which 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 Λ_{jt}^{k*} , θ_{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 relative to directly observing the data when information is obtained socially. Among other reasons, this assumption is positively appealing because there is a well-documented tendency for researchers and statistical agencies 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, 2015](#)).

10. [Benoît 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 $p^{k''}$.

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

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

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

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

This restriction reframes the choice of f^k as the choice of σ_{Jt}^k .¹¹ 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.

3.2 Some Solutions to the Agent's Problem

When faced with problems like the agent's problem, a popular set \mathcal{F} is linear opinion pooling (Ranjan and Gneiting, 2010). It turns out that using repeated linear opinion pooling to solve the agent's problem 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*

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

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

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

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

11. Assuming 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 .

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 Signals) *Assume again, as we did in the case of private learning, that*

(A2) Averaging Signals: $\beta_t = 1/t$, 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) Private signals are iid with $\mathbb{E}[\sigma_{it}^{k*}] \equiv \mu_i^{k*} = \mu_i^k \equiv \mathbb{E}[\sigma_i^k]$, and

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

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 . 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

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

■

We can imagine scenarios in which the agent observes data and signals in each period, but this additional information is potentially 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) *Now suppose that the agent receives biased signals in the sense that $\mathbb{E}[\sigma_{jt}^k] \neq \mu_j^{k*}$, 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

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

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

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

solves the agent's problem.

Proof By A2 we know that $\lim_{t \rightarrow \infty} \lambda_{it+1}^k = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{n=1}^t \widehat{\sigma}_{in}^k$. If the signals are iid, then since the sum of iid random variables is itself an iid random variable, by the law of large numbers we know that $\lim_{t \rightarrow \infty} \lambda_{it+1}^k = \mathbb{E}[\widehat{\sigma}_{it}^k]$. After repeatedly applying the linearity of the expectations operator, A3 and A4a imply that

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

■

4. 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-identified signals.

Acknowledgment

Many of the ideas in this paper originated from discussions with Alon Bergman and Gregorio Caetano. I also thank Michalis Haliassos, Charles Manski, Jan-Peter Siedlarek, Alireza Tahbaz-Salehi, and seminar participants at the Cleveland Fed and Goethe University Frankfurt 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.

References

- D. Acemoglu and A. Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.
- D. Acemoglu, V. Chernozhukov, and M. Wernz. Fragility of asymptotic agreement under Bayesian learning. *Theoretical Economics*, 11:187–225, 2016.
- N. I. Al-Najjar. Decision makers as statisticians: Diversity, ambiguity, and learning. *Econometrica*, 77(5):1371–1401, 2009.
- D. Aliprantis. Redshirting, compulsory schooling laws, and educational attainment. *Journal of Educational and Behavioral Statistics*, 37(2):316–338, 2012.
- D. Aliprantis. Assessing the evidence on neighborhood effects from Moving to Opportunity. *Empirical Economics*, 2017. doi:[10.1007/s00181-016-1186-1](https://doi.org/10.1007/s00181-016-1186-1).
- S. Alon and M. Tienda. Assessing the “mismatch” hypothesis: Differences in college graduation rates by institutional selectivity. *Sociology of Education*, 78(4):294–315, 2005.
- J. Andreoni and T. Mylovannov. Diverging opinions. *American Economic Journal: Microeconomics*, 4(1):209–232, 2012.
- J. D. Angrist and A. B. Krueger. Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics*, 106(4):979–1014, 1991.
- P. Arcidiacono and M. Lovenheim. Affirmative action and the quality-fit trade-off. *Journal of Economic Literature*, 54(1):3–51, 2016.
- Z. Artstein and R. A. Vitale. A strong law of large numbers for random compact sets. *The Annals of Probability*, 3(5):879–882, 1975.
- J.-P. Benoît and J. Dubra. A theory of rational attitude polarization. *Mimeo., London Business School*, 2015.
- C. R. Bollinger and M. van Hasselt. A Bayesian analysis of binary misclassification: Inference in partially identified models. *Mimeo., University of Kentucky*, 2008.
- D. Card and A. B. Krueger. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *The American Economic Review*, 84(4):772–793, 1994.
- V. Chernozhukov, H. Hong, and E. Tamer. Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284, 2007.
- F. P. Coolen, M. C. Troffaes, and T. Augustin. Imprecise probability. In M. Lovric, editor, *International Encyclopedia of Statistical Science*. Springer, 2011.
- P. Dandekar, A. Goel, and D. T. Lee. Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*, 110(15):5791–5796, 2013. doi:[10.1073/pnas.1217220110](https://doi.org/10.1073/pnas.1217220110).

- M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121, 1974.
- S. N. Durlauf, S. Navarro, and D. A. Rivers. Model uncertainty and the effect of shall-issue right-to-carry laws on crime. *European Economic Review*, 81:32–67, 2016.
- H. S. Farber and R. G. Valletta. Do extended unemployment benefits lengthen unemployment spells?: Evidence from recent cycles in the U.S. labor market. *Journal of Human Resources*, 50(4):873–909, 2015.
- M. Gentzkow and J. M. Shapiro. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839, 2011. doi:[10.1093/qje/qjr044](https://doi.org/10.1093/qje/qjr044).
- I. Gilboa and M. Marinacci. Ambiguity and the Bayesian paradigm. In D. Acemoglu, M. Arelano, and E. Dekel, editors, *Advances in Economics and Econometrics: Tenth World Congress: Economic Theory*, volume I, chapter 7, pages 179–242. Cambridge University Press, 2013.
- I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- B. Golub and M. O. Jackson. How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3):1287–1338, 2012. doi:[10.1093/qje/qjs021](https://doi.org/10.1093/qje/qjs021).
- B. Golub and E. Sadler. Learning in social networks. In Y. Bramoullé, A. Galeotti, and B. Rogers, editors, *The Oxford Handbook of the Economics of Networks*. Oxford University Press, 2016.
- M. Hagedorn, F. Karahan, I. Manovskii, and K. Mitman. Unemployment benefits and unemployment in the great recession: The role of macro effects. Technical report, 2013.
- G. W. Imbens and C. F. Manski. Confidence intervals for partially identified parameters. *Econometrica*, 72(6):1845–1857, 2004.
- M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, 2008.
- A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi. Non-Bayesian social learning. *Games and Economic Behavior*, 76(1):210 – 225, 2012.
- J. Kinsler. Assessing Rothstein’s critique of teacher value-added models. *Quantitative Economics*, 3(2):333–362, 2012.
- J. Ludwig, J. B. Liebman, J. R. Kling, G. J. Duncan, L. F. Katz, R. C. Kessler, and L. Sanbonmatsu. What can we learn about neighborhood effects from the Moving to Opportunity experiment? *American Journal of Sociology*, 114(1):144–188, 2008.
- C. F. Manski. *Identification for Prediction and Decision*. Harvard University Press, 2007.
- C. F. Manski. Choosing treatment policies under ambiguity. *Annual Review of Economics*, 3:25–49, 2011.
- C. F. Manski. Identification of income-leisure preferences and evaluation of income tax policy. *Quantitative Economics*, 5(1):145–174, 2014.

- C. F. Manski. Communicating uncertainty in official economic statistics: An appraisal fifty years after Morgenstern. *Journal of Economic Literature*, 53(3):631–653, 2015.
- C. F. Manski and J. V. Pepper. How do right-to-carry laws affect crime rates? coping with ambiguity using bounded-variation assumptions. *NBER Working Paper 21701*, 2015.
- P. Molavi, A. Tahbaz-Salehi, and A. Jadbabaie. Foundations of non-Bayesian social learning. *Mimeo., Columbia University*, 2015.
- H. R. Moon and F. Schorfheide. Bayesian and frequentist inference in partially identified models. *Econometrica*, 80(2):755–782, 2012.
- M. Mueller-Frank. Reaching consensus in social networks. *IESE Working Paper 1116-E*, 2015.
- D. Neumark and W. Wascher. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania: Comment. *The American Economic Review*, 90(5):1362–1396, 2000.
- J. B. Paris and A. Vencovská. A note on the inevitability of maximum entropy. *International Journal of Approximate Reasoning*, 4(3):183–223, 1990.
- R. Ranjan and T. Gneiting. Combining probability forecasts. *Journal of the Royal Statistical Society. Series B*, 72(1):71–91, 2010.
- J. P. Romano and A. M. Shaikh. Inference for the identified set in partially identified econometric models. *Econometrica*, 78(1):169–211, 2010.
- J. Rothstein. Teacher quality in educational production: Tracking, decay, and student achievement. *The Quarterly Journal of Economics*, 125(1):175–214, 2010.
- J. Stoye. More on confidence intervals for partially identified parameters. *Econometrica*, 77(4):1299–1315, 2009.
- E. Tamer. Partial identification in econometrics. *Annual Review of Economics*, 2(1):167–195, 2010.
- G. Wilmers. The social entropy process: Axiomatising the aggregation of probabilistic beliefs. *Mimeo., Manchester Institute for Mathematical Sciences*, 2010.