



PIER

PENN INSTITUTE *for* ECONOMIC RESEARCH
UNIVERSITY *of* PENNSYLVANIA

The Ronald O. Perelman Center for
Political Science and Economics (PCPSE)
133 South 36th Street
Philadelphia, PA 19104-6297

pier@econ.upenn.edu

<http://economics.sas.upenn.edu/pier>

PIER Working Paper

19-001

The Wisdom of a Confused Crowd: Model-Based Inference

GEORGE J. MAILATH
University of Pennsylvania

LARRY SAMUELSON
Yale University

January 16, 2019

<https://ssrn.com/abstract=3317592>

The Wisdom of a Confused Crowd: Model-Based Inference*

George J. Mailath[†] Larry Samuelson[‡]

January 16, 2019

Abstract

“Crowds” are often regarded as “wiser” than individuals, and prediction markets are often regarded as effective methods for harnessing this wisdom. If the agents in prediction markets are Bayesians who share a common model and prior belief, then the no-trade theorem implies that we should see no trade in the market. But if the agents in the market are *not* Bayesians who share a common model and prior belief, then it is no longer obvious that the market outcome aggregates or conveys information. In this paper, we examine a stylized prediction market comprised of Bayesian agents whose inferences are based on different models of the underlying environment. We explore a basic tension—the differences in models that give rise to the possibility of trade generally preclude the possibility of perfect information aggregation.

*We thank many seminar audiences for helpful comments and discussions. We thank the National Science Foundation (SES-1459158 and SES-1559369) for financial support.

[†]Department of Economics, University of Pennsylvania, and Research School of Economics, Australian National University; gmailath@econ.upenn.edu

[‡]Department of Economics, Yale University, New Haven, CT 06525, larry.samuelson@yale.edu

Contents

1	Introduction	1
2	The Setting	3
2.1	The Environment	3
2.2	Model-Based Reasoning	3
2.3	Beliefs	6
2.3.1	Prior and Full-Information Beliefs	6
2.3.2	Interim Beliefs	7
2.3.3	Why Don't Agents Choose the Right Model?	9
3	Learning from Others	10
3.1	Updating	11
3.2	How Revealing are Beliefs?	17
3.3	Properties of the Belief Updating Process	21
3.4	Common Knowledge	23
4	When Are Crowds Confused?	25
4.1	Do Crowds Agree?	25
4.2	Are Crowds Wise?	26
5	What Makes A Crowd Wiser?	29
5.1	Correlated Model Predictions	30
5.2	Models with a Common Component	32
5.3	Crowds with Enough (Dispersed) Information	34
6	Related Literature and Discussion	37
A	Appendices	39
A.1	What's Wrong with Different Priors?	39
A.2	An Example with Pooling	43
A.3	An Example with Infinite Iterations	44
A.4	Proposition 1.5 and Oracles	46
A.5	Common Knowledge when Models are Infinite	47
A.6	An Example Illustrating Redundancy and Correlation	48
A.7	Proof of Lemma 2	49
A.8	Proof of Lemma 3	51
	References	52

The Wisdom of a Confused Crowd: Model-Based Inference

“For the many, of whom each individual is but an ordinary person, when they meet together may very likely be better than the few good, if regarded not individually but collectively, For each individual among the many has a share of excellence and practical wisdom.”

Aristotle, *Politics*, part 3.¹

1 Introduction

The idea that groups of people may make better decisions than individuals is an old one. More recently, interest in the “wisdom of the crowd” was catalyzed by Francis Galton (Galton, 1907), who observed a contest calling on participants to guess the dressed weight of a live ox. Examining the nearly 800 entries, Galton wrote that “the middlemost estimate expresses the vox populi” and reported this value (the median) as 1207 pounds, within about 0.8% of the actual weight of 1198 pounds.

The fairgoers had the advantage of having Galton himself to aggregate their information, turning the exercise into an arguably straightforward illustration of the law of large numbers. More recently, interest in the wisdom of the crowd has focussed on the ability of *markets* to aggregate information (e.g., Surowiecki, 2004), reflected in the efficient market hypothesis of finance (Fama, 1970, 1991), auction-based models of rational expectations (Wilson, 1977, Milgrom, 1979), and the design of prediction markets tailored specifically to aggregate information (Wolfers and Zitzewitz, 2004, and Berg, Forsythe, Nelson, and Rietz, 2008). Here, the market price takes Galton’s place as information aggregator.

The seemingly compelling intuition that prediction markets should aggregate information runs squarely into the no-trade theorem (Milgrom and Stokey, 1982): If the participants are Bayesians who share a common model and common prior belief, then there should be no trade in a prediction market.² Nonetheless, trade does occur (Wolfers and Zitzewitz, 2004), and hence the participants are *not* Bayesians who share a common model and prior belief. *Can we then expect information aggregation?*

¹We thank Rann Smorodinsky for calling this discussion to our attention.

²Reflecting this difficulty, Ottaviani and Sørensen (2010) examine bettors who derive a recreational utility from betting in a parimutuel market. Koessler, Noussair, and Ziegelmeyer (2008) examine bettors who are *required* to bet a fixed amount in a parimutuel market.

Given the implausibility of the common prior assumption, a natural response to the observations of the previous paragraph is to allow agents to hold different prior beliefs.³ Our response is different: We view different priors as a symptom of more basic underlying differences, namely different models. The important advantage of working directly with different models is that we can then reasonably insist that agents have a common prior on the *common* elements of their models. This restores much of the discipline whose absence typically pushes research away from models with heterogeneous priors. Appendix A.1 illustrates the lack of discipline that arises with heterogeneous priors.

We find a basic tension. In Section 4, we establish a pair of negative results. Unless the differences in agents' models are trivial, market interactions will not lead agents to common beliefs. More problematically, any conventional aggregate of the agents' beliefs will often be off the mark, in the sense that the correct-model belief will lie outside the convex hull of the agents' beliefs. In general, crowds are not wiser than their constituents.

And yet at times prediction markets seem to work. Section 5 shows that if the agents' models have enough in common, then interacting will lead agents to similar beliefs, even if their models also exhibit some bizarre idiosyncracies. Perhaps more importantly, if the agents collectively have sufficient information, then the average belief in the crowd will be close to the correct-model belief, even if their information is widely dispersed. The key to the wisdom of the crowd is *not* that its members have common information or a common model, but that the different models the agents use imply a sufficiently common interpretation of whatever information they have.

While our original motivation came from prediction markets, our interest is *not* in modeling prediction markets per se. We believe that the principles uncovered in the course of this analysis are applicable to a broad range of belief aggregation problems, and so deliberately work with a stylized model of belief exchange that abstracts from market microstructure details.

³Morris (1994) is a good point of entry for work in this area. There are many variations on this idea. For example, Geanakoplos (1989) and Brandenburger, Dekel, and Geanakoplos (1992) explore the extent to which trade can arise between agents who have common priors but nonpartitional information structures, and explore the extent to which reasoning based on different priors is equivalent to reasoning based on common priors but nonpartitional information structures.

2 The Setting

2.1 The Environment

The logical structure of our model is best revealed by working with a general model: States of the world are given by a complete, separable metric space Ω , and the associated collection of events is the Borel σ -algebra, denoted by \mathcal{F} . However, the model is more readily interpreted, and the statement of some of the results is significantly simplified, by considering the *product case*: $\Omega = X^N$, where $X \subseteq \mathbb{R}$ and $N \subseteq \mathbb{N}$ is possibly infinite, and where \mathcal{F} is the usual σ -algebra.⁴

We present the model for the general case, but freely appeal to the product case when developing intuition or presenting results. For illustration and intuition, we often further assume that X is finite (typically $\{0, 1\}$) and N is finite.

Nature draws a state ω from Ω according to the probability measure ρ on Ω . Agents form beliefs about the occurrence of a measurable event $F \in \mathcal{F}$. It is convenient to describe this event in terms of its indicator function $f : \Omega \rightarrow \{0, 1\}$, where $f(\omega) = 1$ if and only if $\omega \in F$.

2.2 Model-Based Reasoning

It is standard in economic analyses to equip agent i with the state space Ω , prior belief ρ , and description f of an event, at which point the agent would use Bayes' rule to process her information. We refer to such a reasoner as an *agent- i oracle* or simply an *agent oracle*.

In contrast, we are concerned with *model-based reasoners*. A model-based reasoner is a faithful adherent of Savage's (1972) *Foundations of Statistics*. Savage explains that it is a hopeless undertaking to work with a state space that specifies "[t]he exact and entire past, present, and future history of the universe, understood in any sense, however wide" (Savage, 1972, p. 9).⁵ Savage argues on the same page that "the use of modest little worlds, tailored to particular contexts, is often a simplification, the advantage of which is justified". A model-based reasoner's "modest little world" will partition the state space into equivalence classes that he or she believes

⁴The usual σ -algebra is the σ -algebra generated by the open sets of the product topology, which is the one generated by sets of the form $\{\omega \in \Omega : \omega_k = \tilde{\omega}_k \forall k \in K\}$ for $\tilde{\omega}_k \in X$, $k \in K$, and some finite set $K \subseteq N$.

⁵Savage (1972, p. 16) describes this logical extreme of "look before you leap" as "utterly ridiculous".

capture relevant information about F while ignoring irrelevant information. These equivalence classes then become “states” in the reasoner’s model.

We capture this reliance on models by assuming that agent i explains the connection between the states in Ω and the occurrence of the event F by a function (her *theory*) $f^i : \Omega \rightarrow [0, 1]$ that is measurable with respect to a sub- σ -algebra \mathcal{M}^i (her *model*) of \mathcal{F} . If $F \in \mathcal{M}^i$, then agent i ’s theory is the indicator function f , and i correctly understands the events determining F . We are most interested in cases in which agent i simplifies her problem by working with a strict subset $\mathcal{M}^i \subsetneq \mathcal{F}$ for which $F \notin \mathcal{M}^i$. The assumption that i ’s theory f^i is measurable with respect to her model \mathcal{M}^i means that i believes that only the events in \mathcal{M}^i are relevant to the determination of F . Agent i realizes that such a model cannot be expected to perfectly explain the event F , reflected in the fact that f^i maps into $[0, 1]$ (rather than $\{0, 1\}$), giving the probability that the event F has occurred.

In the product case of $\Omega = X^N$, agent i ’s assessment of the event F depends only on the realizations in the dimensions contained in the subset $M^i \subseteq N$. Hence, \mathcal{M}^i is the σ -algebra generated by sets of the form $\{\omega_{M^i}\} \times X^{N \setminus M^i} =: \{\omega_{M^i}\} \times X^{-M^i}$, and we also use “model” to refer to M^i . In terms of Savage’s procedure for creating a “modest little world”, the equivalence classes for agent i are of the form $\{\omega_{M^i}\} \times X^{-M^i}$. Since agent i is concerned only with ω_{M^i} , we can use $f^i(\omega_{M^i})$ to denote the (constant) value of f^i on the set $\{\omega_{M^i}\} \times X^{-M^i}$.

For example, suppose a collection of agents is called upon to predict the price movements of a financial asset, so that the event F corresponds to an increase in the price of that asset. Even upon restricting attention to professionals, we encounter a variety of approaches. A fundamentalist will typically seek information on the cash reserves, debt load, volume of sales, profit margin (and so on) of the underlying firm; these are the variables that would appear in her model M^i . A chartist will inquire about (and include in his M^i variables corresponding to) recent sales volumes, price trends, reversals in price movements, the existence of apparent price ceilings, and so on. An efficient marketer will respond by asking for a coin to flip. And even among professionals, there are forecasters whose models focus on astrological data. The fundamentalist is likely to exclude much of the asset-price history from her model, while the chartist may neglect various aspects of the firm’s current financial position. Both will typically exclude information about zodiac signs. All of the agents are likely to miss factors whose relevance has not yet been imagined, as well as factors they are convinced are irrelevant, while possibly including irrelevant factors.

Remark 1 (unimagined or unappreciated events?) We are agnostic about choosing between two formally-equivalent interpretations of our model of model-based reasoning. We can think of agent i as recognizing all of the information contained in a state $\omega \in \Omega$, but adopting a simple theory f^i that depends on only some of this information. The fundamentalist may recognize the events specifying the signs of the zodiac, but will ignore them when reasoning about stock prices. Alternatively, we can view agent i as recognizing the events in \mathcal{M}^i , but as having no idea what other events appear in other agents' models. \blacklozenge

The only potential glitch in agent i 's reasoning arises from her reliance on the model \mathcal{M}^i . In particular, conditional on using such a model, her theory f^i must be consistent with F 's indicator function f . In the product case with finite N , this is the requirement that for every positive probability ω_{M^i} , we have

$$f^i(\omega_{M^i}) = \sum_{\omega \in \Omega} f(\omega) \rho(\omega | \omega_{M^i}).$$

Intuitively, the probability agent i attaches to the event F upon observing ω_{M^i} matches the probability attached to the event F by the measure ρ conditional on ω_{M^i} .

The general case requires nothing more conceptually, but does involve more abstract notation. We require

$$f^i(\omega) = \mathbb{E}[f | \mathcal{M}^i](\omega). \tag{1}$$

This is again the statement that the probability agent i attaches to event F upon having observed any event in \mathcal{M}^i (which are the only events i thinks are relevant) is the probability attached to the event F by the measure ρ , conditional on that event.⁶

A motivation for (1) is that agent i builds her view of the world from her model \mathcal{M}^i and her access to a historical record of an unlimited number of independent draws from the prior distribution ρ . Focusing for convenience on the finite product case, the agent calculates the frequency of the realizations of ω_{M^i} in these observations. While the agent could infer the probability $\rho(\omega)$ for any state ω , all the agent requires to calculate the implication of her theory f^i are the probabilities of the realizations of ω_{M^i} (i.e., of sets of the form $\{\omega_{M^i}\} \times X^{-M^i}$). Agent i can also observe which realizations of ω_{M^i} in the record of past draws correspond to the occurrence of the event

⁶Because Ω is complete, separable, and metric (i.e., Polish), we can assume that conditional beliefs exist for *all* ω (Stroock, 2011, Theorem 9.2.1).

F . Hence, for each value of ω_{M^i} , she can calculate the probability $f^i(\omega_{M^i})$ given by (1).

We do not expect agent i to literally have access to an infinite number of draws. In practice, the agent estimates f^i on the basis of a finite number of observations. Hence, even while motivating (1) in terms of an infinite record, we still feel free to appeal to the case of finite data when developing intuition for our analysis. Our goal in imposing (1) is to isolate the implications of agent i 's model-based inference from the estimation problems that invariably arise with finite sets of data, much as econometricians prefer to separate questions of estimation and identification, thus removing every obstacle from information processing other than agent i 's reliance on the model \mathcal{M}^i . In particular, (1) imposes a basic consistency requirement across agents.

Remark 2 (choosing models) People go to great lengths to advocate for their models. Einstein is reputed to have argued that “God does not play dice with the universe”, and Dirac to have argued that “God used beautiful mathematics in creating the world”. Both are (in our view) examples of advocating particular (types of) model. We do not examine the process by which agent i comes to focus on \mathcal{M}^i . One point worth emphasizing, however, is that we should *not* expect the model selection process to eliminate differences in models, since different agents may follow different model selection processes.

Section 2.3.3 discusses our assumption that agents do not infer an appropriate model $\mathcal{M}^i \subseteq \mathcal{F}$ from the data. \blacklozenge

Our agents are certain about their models. For the product case, this can be viewed as a form of correlation neglect. Agent i believes that conditional on ω_{M^i} , the event F and ω_{-M^i} are uncorrelated. Agent i 's model can be viewed as a simple Bayesian network (see Pearl (2009) for an exposition, and Spiegler (2016), who also assumes agents have access to infinite data, for an application to misspecified models). That work focuses on agents interpreting correlations as causation, while our focus is on aggregation.

2.3 Beliefs

2.3.1 Prior and Full-Information Beliefs

A model-based reasoner with no information about the state attaches to the event F the probability

$$\mathbb{E}[f^i(\omega)] = \mathbb{E}[\mathbb{E}[f|\mathcal{M}^i]] = \mathbb{E}[f(\omega)],$$

where the first equality is from (1) and the second is the first of many applications of the law of iterated expectations. This indicates that agent i 's prior belief matches that of an agent- i oracle.

If agent i observed all of the information she deemed relevant, i.e., if agent i observed the realization of the events in \mathcal{M}^i , then she would regard herself as having full information,⁷ and would attach to the event F the probability

$$f^i(\omega),$$

whose value is given by (1). We refer to this as a *full information belief*. Agent i *ignores* any empirical evidence of events outside of \mathcal{M}^i , but she *correctly* uses empirical frequencies in assessing the implications of the information she *does* think relevant, namely events in \mathcal{M}^i . It follows immediately from (1) that the full-information beliefs of a model-based reasoner agree with those of an agent oracle.

2.3.2 Interim Beliefs

We allow agent i to observe information about some (but perhaps not all) of the events in \mathcal{M}^i , and to possibly have information about some events not in \mathcal{M}^i . Let $\mathcal{I}^i \subseteq \mathcal{F}$ denote the information (σ -algebra) agent i observes.

In the product case, agent i observes information about some (but perhaps not all) of the dimensions of M^i , and possibly some additional dimensions; we denote by I^i the dimensions i observes.⁸

Given her information at ω , agent i assigns to F an interim probability, which we denote by $\beta^i(\omega)$, given by

$$\beta^i(\omega) = \mathbb{E}[f^i \mid \mathcal{I}^i](\omega).$$

We can again think of agent i as observing her information and then consulting the record of past observations. Agent i considers those realizations consistent with the events she observes and takes the expectation of her full-information beliefs $f^i(\omega)$ on this set of realizations.

Agent i 's updating takes place in two steps. To fix ideas, consider the product case. The agent first estimates the full-information beliefs $f^i(\omega_{M^i})$. Then, upon observing a realization ω_{I^i} of the dimensions in I^i , the agent restricts her attention to the set of states matching this realization and

⁷Recall the agent understands her model may be incomplete, and so full information refers to the information needed for her model.

⁸In this case, $\mathcal{I}^i = \mathcal{G}^{I^i}$, where for any subset $K \subseteq N$, we denote by \mathcal{G}^K the σ -algebra generated by the cylinder sets $\{\omega_K\} \times X^{-K}$.

$\overbrace{\hspace{2cm}}^{\Omega}$					ω_1 $f^i(\omega_1)$	
ω_1	ω_2	$\rho(\omega)$	$f(\omega)$			
0	0	1/4	0	}	0	0
0	1	1/4	0	}		
1	0	1/4	0	}	1	1/2
1	1	1/4	1	}		

Figure 1: The structure for Example 1. The first four columns present the environment, (Ω, ρ, f) , and the last three columns present agent i 's model, which has only two states, and full information beliefs.

takes the expected value of $f^i(\omega_{M^i})$ over this set. This two-step feature is the essence of model-based reasoning. In contrast, agent i 's oracle substitutes the true indicator function f for the first step and, after observing the information $\omega_{\mathcal{I}^i}$, takes the expected value of $f(\omega)$ over this set.

If $\mathcal{I}^i \subseteq \mathcal{M}^i$, so that agent i observes only information she deems relevant (which for the product case, is equivalent to $I^i \subseteq M^i$), then the two updating procedures are equivalent. In this case, we have

$$\begin{aligned}
 \beta^i(\omega) &= \mathbb{E}[f^i \mid \mathcal{I}^i](\omega) \\
 &= \mathbb{E}[\mathbb{E}[f \mid \mathcal{M}^i] \mid \mathcal{I}^i](\omega) \\
 &= \mathbb{E}[f \mid \mathcal{I}^i](\omega),
 \end{aligned}$$

where the first equality repeats our definition of the interim belief β^i , the next line follows by inserting the definition of the full-information belief from (1), and the final line follows from the law of iterated expectations. This equivalence between model-based and oracular reasoning breaks down if \mathcal{I}^i is not a subset of \mathcal{M}^i . Formally, we can no longer apply the law of iterated expectations that capped the previous argument. A product case example illustrates.

Example 1 Suppose $\Omega = \{0, 1\}^2$, with each state equally likely. The event F consists of the state $(1, 1)$. We summarize this information in the left array in Figure 1.

Let $M^i = \{1\}$ and $I^i = \{2\}$. Agent i views dimension 1 as the only relevant dimension, and $\omega_{M^i} = \omega_1$ takes on two values, 0 and 1. Agent i 's full-information beliefs $f^i(\omega_1)$ are given by

$$f^i(0) = 0 \quad \text{and} \quad f^i(1) = 1/2.$$

We summarize agent i 's model and full-information beliefs in the right array in Figure 1.

Agent i observes dimension 2, or $I^i = \{2\}$. An agent- i oracle who observed ω_2 would form the posteriors $\rho(\omega|\omega_2)$ given by (economizing on notation by shortening $\rho((0, 0) | 0)$ to $\rho(0, 0 | 0)$)

$$\begin{array}{ccc} \rho(0, 0 | 0) = 1/2, & & \rho(1, 0 | 0) = 1/2, \\ \rho(0, 1 | 1) = 1/2, & \text{and} & \rho(1, 1 | 1) = 1/2, \end{array}$$

with all other conditional probabilities equaling 0, and hence would have beliefs

$$\begin{array}{l} \mathbb{E}[f(\omega_1, \omega_2) | 0] = \rho(0, 0 | 0)f(0, 0) + \rho(1, 0 | 0)f(1, 0) = 0 \\ \text{and } \mathbb{E}[f(\omega_1, \omega_2) | 1] = \rho(0, 1 | 1)f(0, 1) + \rho(1, 1 | 1)f(1, 1) = 1/2. \end{array}$$

The model-based reasoner forms an identical posterior over states, but then takes the expectation of her full-information belief to obtain

$$\begin{array}{l} \beta^i(0) = \mathbb{E}[f^i(\omega_1) | 0] = \rho(0, 0 | 0)f^i(0) + \rho(1, 0 | 0)f^i(1) = 1/4 \\ \text{and } \beta^i(1) = \mathbb{E}[f^i(\omega_1) | 1] = \rho(0, 1 | 1)f^i(0) + \rho(1, 1 | 1)f^i(1) = 1/4. \quad \blacklozenge \end{array}$$

2.3.3 Why Don't Agents Choose the Right Model?

In a setting as simple as Example 1, how can the agent fail to simply calculate the empirical frequency of F given either $\omega_2 = 0$ or $\omega_2 = 1$? More generally, given an unlimited record of previous draws from the distribution ρ , why not simply use the historical record to identify the correct model?

In practice, agents are confronted with finite data and a state space of potentially infinite complexity, and an agent will never encounter data that unambiguously contradicts whatever model she holds. Instead, anomalous observations can always be explained away by unobserved factors. Indeed, for every event and every set of data, there will an infinite collection of models that explain the data perfectly, making it impossible to use the data to find the "right" model. And for every event, there will be an infinite list of variables about which the agent could collect information, making it impossible to be a pure empiricist. Al-Najjar (2009) and Gilboa and Samuelson

(2012) elaborate on the futility of interpreting data without models.⁹ If she is to make any meaningful use of the data, agent i must then appeal to *some* model \mathcal{M}^i . In order to focus on how and whether agents can “learn” from each other, we have then made the extreme assumption (as does [Spiegler, 2016](#)) that the restriction captured in (1) holds with equality rather than approximately, so that the agent makes perfect use of whatever model she has.

[Giacomini, Skreta, and Turén \(2007\)](#) examine the behavior of 75 professional forecasters. The object of each participant was to predict the US inflation rate, for each of the years 2007–2014. Forecasting typically began at the beginning of July of the preceding year (with slightly later initial forecasts for 2007 and 2008), with individual forecasters updating their predictions at any time until the end of the year in question. [Giacomini, Skreta, and Turén \(2007\)](#) argue that the forecasters in their sample appear to be Bayesians (albeit much more so in non-crisis years), but with different models that lead them to different forecasts. In response to this disagreement, the agents persevere in their belief in their models (again, more so in non-crisis years) and in their disagreement. Such agents would find themselves well at home in our setting.

3 Learning from Others

We now examine how agents update their beliefs in response to information about others’ beliefs. There are K agents. Each agent i has a model \mathcal{M}^i exhibiting the properties outlined in Section 2.2, and has access to information \mathcal{I}^i . We refer to such a collection as a *crowd*.

In Section 2.3.2, we saw that model-based and oracular updating are equivalent when $\mathcal{I}^i \subseteq \mathcal{M}^i$. To focus on the implications for model-based inference that arise from the interaction between agents, we henceforth assume $\mathcal{I}^i \subseteq \mathcal{M}^i$ for each agent i .

Remark 3 (different models or different events?) We interpret our analysis as that of agents forming beliefs about a single event F , but with differ-

⁹Other papers examining interactions between agents who persist in holding different models include [Acemoglu, Chernozhukov, and Yildiz \(2016\)](#) and [Eyster and Piccione \(2013\)](#). [Hong, Stein, and Yu \(2007\)](#) examine a model in which agents restrict attention to a class of models simpler than that of an (in our terms) agent oracle, but update their beliefs about which model in the simple class is appropriate. The extension of our model to such a setting would involve agents who restrict attention to *different* classes of models, or follow different model-updating rules. The difficulties agents face in learning from other agents would only be exacerbated in such a setting.

ent models. The challenge is then to examine how agents infer information relevant to their own models from other agents who have different models. Returning to our example, the fundamentalist may recognize that there is information to be gleaned about fundamentals from another agent who is primarily concerned with charts.

Much of the analysis of this section could be recast as one in which every agent is an oracle, but the agents are forming beliefs about different events. The challenge is then to examine how agents infer information about their own events from other agents who are concerned with other events. One fundamentalist may be concerned with industrial stocks, while recognizing that there is useful information to be gleaned from the beliefs of an analyst who specializes in agricultural futures. For the purposes of much of this section, it is a taste question which interpretation is most congenial. However, our positive results on the various forms of agreement are most consistent with our preferred interpretation.



3.1 Updating

We are interested in how agents update their beliefs in response to information about others' beliefs. There are three intertwined questions here. First, how does agent i update her beliefs upon receiving information from agent j , even while being convinced that j is confused in her model of the world and hence that some of j 's information is irrelevant? Second, what if agent i does not know player j 's model \mathcal{M}^j , and hence does not know the extent of j 's confusion? Third, how does player j transmit information to player i ? For example, is the price that agent i observes in a prediction market enough for i to identify j 's beliefs?

We focus on the first question—how agent i updates her beliefs upon receiving information from agent j . Agent i may find j 's information relevant for two reasons. First, j may observe an event (or a dimension in the product case) that appears in i 's model but i does not observe. A fundamentalist may be convinced that the outcome of a firm's recent drug trial is important, but may not be privy to that outcome, and so may glean inferences from the beliefs of an insider. Second, there may be correlations between the events, so that knowledge of an event in j 's model but not in i 's model may be useful to i . We allow for both.

We take an agnostic position on the second question. Our model is consistent with two formally-equivalent assumptions of what agents know

about each others' models. On the one hand, each agent may know the model of each other agent. Agent i will in general think that agent j is confused in her choice of model, but can infer from j 's beliefs information that i considers relevant. On the other hand, agents may know nothing about others' models. In this case, we assume that agent i can again appeal to the historical record to infer, from any belief announced by j , information that is relevant to i 's model.

Turning to the third question, the transmission of information between agents depends on potentially fine institutional details. We might be interested in a centralized market, a decentralized market, or some other interaction. The agents in a market may be able to observe the entire order book of a market, or only the marginal offers, or only successful trades, or no offers at all. We abstract from these and a host of other details by studying a process of flawless belief exchange. We follow [Geanakoplos and Polemarchakis \(1982\)](#) in examining the following protocol:

- (a) First, each agent i observes their information in \mathcal{I}^i (i.e., which events in \mathcal{I}^i realized or, in the product case, the realization ω_{I^i}) and forms her interim belief.
- (b) The agents (by assumption truthfully) simultaneously announce their interim beliefs, and revise their beliefs in response to these announcements.
- (c) The agents announce their revised beliefs, and then again revise, and again announce, and so on.

Formally, we assume this process continues indefinitely; we will say that the process terminates if a stage is reached at which beliefs are not subsequently revised.¹⁰

To make this process precise, fix ω and suppose that agent i has observed her information while the other agents have announced the vector $b^{-i} = (b^1, \dots, b^{i-1}, b^{i+1}, \dots, b^K)$, where the j^{th} -element of the vector b^{-i} corresponds to agent j 's announced posterior probability b^j of the event

¹⁰[Geanakoplos and Polemarchakis \(1982\)](#) assume the agents have finite information partitions, ensuring that the belief revision process terminates in a finite number of steps. [Sethi and Yildiz \(2012\)](#) apply [Geanakoplos and Polemarchakis \(1982\)](#) to a model of deliberation with different priors. While the signals are normally distributed (and so agents have infinite information partitions), there is no belief updating in their model after the first round because the first round signals are a sufficient statistic of the relevant private information.

F (determined by j 's observation of his information). Letting \mathcal{B}^{-i} be the σ -algebra on Ω generated by the announcement b^{-i} , agent i forms her *model-based* belief about the event F as

$$\beta^i(\omega, b^{-i}) = \mathbb{E}[f^i | \mathcal{I}^i, \mathcal{B}^{-i}](\omega). \quad (2)$$

For the product case, this can be written as (recalling footnote 8)

$$\beta^i(\omega_{I^i}, b^{-i}) = \mathbb{E}[f^i | \mathcal{G}^{I^i}, \mathcal{B}^{-i}](\omega). \quad (3)$$

We see here again the nature of model-based updating. An agent- i oracle, having collected the information (ω_{I^i}, b^{-i}) , would form the belief

$$\mathbb{E}[f | \mathcal{I}^i, \mathcal{B}^{-i}](\omega).$$

Intuitively, we can think of the agent- i oracle as collecting from the empirical record all those observations characterized by (ω_{I^i}, b^{-i}) , and taking the frequency of the event F in these observations as the posterior probability of F . In calculating $\beta^i(\omega_{I^i}, b^{-i})$, a model-based agent i similarly begins by collecting all of the information available, consisting of (ω_{I^i}, b^{-i}) , but then first uses this information to revise her distribution over the elements ω_{M^i} , and then given this revised distribution, takes an expectation of her full-information beliefs, yielding (3). Example 2 below shows how this leads to differences in model-based and oracular beliefs.

Denote the first posterior announced by agent i by b_0^i , the second posterior by b_1^i , and so on; the vector of announced posteriors is similarly denoted by $b_0 = (b_0^i, b_0^{-i})$, $b_1 = (b_1^i, b_1^{-i})$, and so on. The beliefs we have examined to this point are

$$b_0^i = \beta^i(\omega) \text{ and } b_1^i = \beta^i(\omega, b_0^{-i}).$$

Let \mathcal{B}_1^{-i} denote the σ -algebra induced by the announcements (b_0^{-i}, b_1^{-i}) . Then given the beliefs $b_0^i = \beta^i(\omega)$ and $b_1^i = \beta^i(\omega, b_0^{-i})$, an announcement by the remaining agents of their updated posteriors $b_1^j = \beta^j(\omega, b_0^{-j})$ results in agent i updating her beliefs to

$$b_2^i = \beta^i(\omega, b_0^{-i}, b_1^{-i}) = \mathbb{E}[f^i | \mathcal{I}^i, \mathcal{B}_1^{-i}](\omega).$$

Letting \mathcal{B}_n denote the σ -algebra induced by (b_0, \dots, b_n) , we have, for all n ,

$$b_{n+1}^i = \beta^i(\omega, b_0^{-i}, \dots, b_n^{-i}) = \mathbb{E}[f^i | \mathcal{I}^i, \mathcal{B}_n^{-i}](\omega) = \mathbb{E}[f^i | \mathcal{I}^i, \mathcal{B}_n](\omega),$$

where the last equality follows from $\sigma(\mathcal{I}^i, \mathcal{B}_n^{-i}) = \sigma(\mathcal{I}^i, \mathcal{B}_n)$ (and $\sigma(\mathcal{A}, \mathcal{B})$ is the σ -algebra generated by the σ -algebras \mathcal{A} and \mathcal{B}). Let $\mathbf{b} := (b_0, b_1, \dots)$

denote the infinite collection of announcements and \mathcal{B}_∞ the σ -algebra induced by \mathbf{b} . It will also be useful to keep track of the beliefs of the *public oracle*,

$$\mathbb{E}[f|\mathcal{B}_n] \text{ and } \mathbb{E}[f|\mathcal{B}_\infty].$$

Intuitively, a public oracle is an agent whose theory is given by f , model by \mathcal{F} , and who observes the announcements of all players, but no other information.

Since each agent and the public oracle follow Bayesian updating on the sequence of increasingly informative announcements (filtrations) (\mathcal{B}_n^{-i}) and (\mathcal{B}_n) , the resulting sequence of updates are martingales and so converge (with probability one under ρ) to limits which are measurable with respect to the limit σ -algebras. Summarizing this discussion, we have the following.

Lemma 1 *The updated beliefs*

$$(\mathbb{E}[f^i|\mathcal{I}^i, \mathcal{B}_n^{-i}])_{n=1}^\infty \text{ and } (\mathbb{E}[f|\mathcal{B}_n])_{n=1}^\infty$$

are martingales, with ρ -almost-sure limits

$$\mathbb{E}[f^i|\mathcal{I}^i, \mathcal{B}_\infty^{-i}] \text{ and } \mathbb{E}[f|\mathcal{B}_\infty].$$

Example 2 Suppose the state space is given by $\Omega = \{0, 1\}^4$. The pair (ω_1, ω_2) is drawn from the distribution $\Pr\{(\omega_1, \omega_2) = (0, 0)\} = \Pr\{(\omega_1, \omega_2) = (1, 1)\} = \frac{3}{8}$, $\Pr\{(\omega_1, \omega_2) = (0, 1)\} = \Pr\{(\omega_1, \omega_2) = (1, 0)\} = \frac{1}{8}$, and the pair (ω_3, ω_4) is independently drawn from a distribution with an identical correlation structure. The event is

$$F = \left\{ \omega : \sum_{k=1}^4 \omega_k \geq 2 \right\}.$$

There are two agents. Agent 1's model and information are given by

$$M^1 = \{1, 2, 3, 4\} \quad \text{and} \quad I^1 = \{2, 3\}$$

while agent 2's are given by

$$M^2 = \{3, 4\} \quad \text{and} \quad I^2 = \{4\}.$$

This information is summarized in Figure 2, together with the interim beliefs $\beta^1(\omega_{I^1})$ and $\beta^2(\omega_{I^2})$.

Now we turn to updating in response to others' beliefs. First, consider agent 1. Agent 2 observes only one piece of information, namely ω_4 , and agent 2's belief b^2 reveals the value of ω_4 . Agent 1's model-based belief,

State $(\omega_1, \omega_2, \omega_3, \omega_4)$	Prior ρ	$f(\omega)$	2's theory $f^2(\omega_{M^2})$	Interim beliefs $\beta^1(\omega_{I^1})$ $\beta^2(\omega_{I^2})$		First-round updates $\beta^1(\omega_{I^1}, b_0^2)$ $\beta^2(\omega_{I^2}, b_0^1)$		Second round $\beta^2(\omega_{I^2}, b_0^1, b_1^1, b_0^2)$
(0, 0, 0, 0)	9/64	0	3/8	1/16	14/32	0	3/8	3/8
(0, 0, 0, 1)	3/64	0	5/8	1/16	29/32	1/4	5/8	5/8
(0, 0, 1, 0)	3/64	0	5/8	13/16	14/32	1/4	14/32	5/8
(0, 1, 0, 0)	3/64	0	3/8	13/16	14/32	3/4	14/32	3/8
(1, 0, 0, 0)	3/64	0	3/8	1/16	14/32	0	3/8	3/8
(0, 0, 1, 1)	9/64	1	1	13/16	29/32	1	29/32	29/32
(0, 1, 0, 1)	1/64	1	5/8	13/16	29/32	1	29/32	29/32
(1, 0, 0, 1)	1/64	1	5/8	1/16	29/32	1/4	5/8	5/8
(0, 1, 1, 0)	1/64	1	5/8	1	14/32	1	5/8	5/8
(1, 0, 1, 0)	1/64	1	5/8	13/16	14/32	1/4	14/32	5/8
(1, 1, 0, 0)	9/64	1	3/8	13/16	14/32	3/4	14/32	3/8
(1, 1, 1, 0)	3/64	1	5/8	1	14/32	1	5/8	5/8
(1, 1, 0, 1)	3/64	1	5/8	13/16	29/32	1	29/32	29/32
(1, 0, 1, 1)	3/64	1	1	13/16	29/32	1	29/32	29/32
(0, 1, 1, 1)	3/64	1	1	1	29/32	1	1	1
(1, 1, 1, 1)	9/64	1	1	1	29/32	1	1	1

$$M^1 = \{1, 2, 3, 4\}, \quad M^2 = \{3, 4\},$$

$$I^1 = \{2, 3\}, \quad I^2 = \{4\}.$$

Figure 2: The beliefs for Example 2. Agent 1's theory agrees with the indicator f and so is not listed separately.

given by (2), is then identical to the interim belief agent 1 would have if 1 observed $\{\omega_2, \omega_3, \omega_4\}$, and is identical to the belief of an agent-1 oracle. We report these beliefs in Figure 2, in the column labeled $\beta^1(\omega_{I^1}, b_0^2)$.

We now turn to agent 2. Suppose, first, $b_0^1 = 1$. Agent 2 observes ω_4 , and infers that agent 1 has observed $\omega_3 = 1$, leading agent 2's posteriors $\rho(\omega_{M^2} | \omega_{I^2}, b_0^1)$ to take on the values

$$\rho(1, 0 | 0, 1) = \rho(1, 1 | 1, 1) = 1.$$

Any information about ω_2 in agent 1's belief agent 2 considers irrelevant. Agent 2's updated beliefs $\beta^2(\omega_{I^2}, b_0^1)$ about the event F are then given by

$$\beta^2(0, 1) = f^2(1, 0) = 5/8$$

and $\beta^2(1, 1) = f^2(1, 1) = 1.$

We see here the difference between model-based and oracular updating. An agent-2 oracle who observed $\omega_4 = 0$ and $b_0^1 = 1$ would infer that the state is $(0, 1, 1, 0)$ with probability $1/2$ and $(1, 1, 1, 0)$ with probability $1/2$, thus limiting attention to the relevant lines of Figure 2 (just as an oracular reasoner observing $\omega_2 = 0$ limits attention to the first and third lines of Figure 1 in Example 1). Both states give rise to the event F , and so the agent-2 oracle would attach posterior probability 1 to the event. In contrast, the model-based updater who has observed $\omega_4 = 0$ and $b_0^1 = 1$ draws the inference that $(\omega_3, \omega_4) = (1, 0)$. The agent then calculates her full information probability of F , given only $(\omega_3, \omega_4) = (1, 0)$, or equivalently calculates the empirical frequency of the event F in all observations in the record in which $(\omega_3, \omega_4) = (1, 0)$, which is $5/8$. Notice that these observations include $(0, 0, 1, 0)$, $(0, 1, 1, 0)$, $(1, 0, 1, 0)$, and $(1, 1, 1, 0)$. The agent-2 oracle ignores the first and third of these, on the grounds that agent 1's announcement of $b_0^1 = 1$ precludes these two states. However, the model-based reasoner views ω_1 and ω_2 as irrelevant—indeed, may not even recognize the existence of these dimensions—and hence ignores this information, making use only of the information that $(\omega_3, \omega_4) = (1, 0)$.

The case of $b_0^1 = 1/16$ is similar.

Finally, suppose $b_0^1 = \frac{13}{16}$. Unlike the previous two cases, this observation does not unambiguously identify player 1's observation, pooling $(0, 1)$ and $(1, 0)$. Instead, player 2's updated distribution $\rho^2(\omega_{M^2} | \omega_{I^2}, b_0^1)$ takes on the values

$$\begin{aligned} \rho(0, 0 | 0, 13/16) &= 3/4, \\ \rho(1, 0 | 0, 13/16) &= 1/4, \\ \rho(0, 1 | 1, 13/16) &= 1/4, \\ \text{and } \rho(1, 1 | 1, 13/16) &= 3/4. \end{aligned}$$

Agent 2's updated beliefs $\beta^2(\omega_{I^2}, b_0^1)$ about the event F are then given by

$$\beta^2(0, 13/16) = \rho(0, 0 | 0, 13/16)f^2(0, 0) + \rho(1, 0 | 0, 13/16)f^2(1, 0) = 14/32$$

and

$$\beta^2(1, 13/16) = \rho(0, 1 | 0, 13/16)f^2(0, 1) + \rho(1, 1 | 0, 13/16)f^2(1, 1) = 29/32.$$

Again, these beliefs differ from those of an agent-2 oracle, who attaches probabilities $5/8$ (after observing $(\omega_4, b_0^1) = (0, 13/16)$) and 1 (after observing $(\omega_4, b_0^1) = (1, 13/16)$) to event F . The results of agent 2's updating

are reported in the column $\beta^2(\omega_{I^2}, b_0^1)$. This concludes the first round of updating.

The next step calls for the agents to announce their updated beliefs to one another. Agent 1 learns nothing new from this new announcement. Agent 2's original announcement revealed all of 2's information to 1, namely the value of ω_4 , and so agent 1 draws no further inferences (and the table contains no further column for agent 1).

Agent 2 does update in response to agent 1's announcement, giving rise to the column $\beta_2(\omega_{I^2}, b_0^1, b_1^1, b_0^2)$. First, suppose agent 1 announces the belief $1/16$ on the first round. This announcement reveals to agent 2 that $\omega_3 = 0$ (and also that $\omega_2 = 0$, though 2 considers this information irrelevant), and there is nothing more for 2 to learn from 1's subsequent announcement of either 0 or $1/4$. Agent 2's beliefs are unchanged in this case. A similar argument applies if agent 1 announces a belief of 1.

Suppose that 1's initial announcement was $13/16$, and 2's observation is $\omega_4 = 1$ (and hence 2's report was $29/32$). Agent 1's updated belief is always 1 in this case, and hence there is no new information for agent 2 to process on the second round. In this case, 2's beliefs remain unchanged. Suppose, however, that 2's initial observation was $\omega_4 = 0$ (and hence 2's report was $14/32$). Now suppose 2 observes that 1 has revised her belief to $1/4$. This reveals to 2 that $\omega_3 = 1$. (It also potentially reveals that $\omega_2 = 0$, but 2 considers this information irrelevant and ignores it.) Agent 2 then notes that when $(\omega_3, \omega_4) = (1, 0)$, the full-information belief of the event F is $5/8$, and this becomes 2's new belief. Analogously, suppose that 2's initial observation was $\omega_4 = 0$ (and hence 2's report was $14/32$). Now 2 observes that 1 has revised her belief to $3/4$. This reveals to 2 that $\omega_3 = 0$. Agent 2 then notes that when $(\omega_3, \omega_4) = (0, 1)$, the full information belief of the event F is $3/8$, and this becomes 2's new belief. We report these beliefs in column $\beta_2(\omega_{I^2}, b_0^1, b_1^1, b_0^2)$.

It is straightforward to check that subsequent rounds of announcements have no further effect on beliefs. \blacklozenge

3.2 How Revealing are Beliefs?

We might go further in our quest to give our agents the best chance at aggregating information, by simply having them announce their *information* rather than their beliefs to one another. However, while we are comfortable in abstracting from the details of market microstructure by using the exchange of beliefs as a convenient proxy for the workings of a market, we are

not comfortable simply assuming the market will reveal *all* of the agents' information.

This difference matters. As illustrated by Geanakoplos and Polemarchakis (1982, Proposition 3) in the product case, an agent oracle need not hold the same beliefs as someone who can observe the information contained in $\cup_{k=1}^K I^k$.¹¹ Instead, some player k 's belief announcements may pool together some of the information contained in I^k .

Constructing such product case examples is straightforward, and does not exploit misspecification in the agents' models. Suppose each of K oracular agents observes an independent, equiprobable draw from $\{0, 1\}$, and that the event F occurs if and only if these draws agree. Then the agents' beliefs are uninformative, and the agents will not revise their beliefs no matter how many announcements they make, even though a single announcement of their *information* would be instantly informative.

One might object that this example is special, both because it would take only a single agent whose information set includes more than one dimension to ensure that some information is transmitted, and because the probability of the event F is small if K is large. Appendix A.2 describes an example that addresses these concerns.

One might then counter that the pooling encountered in these examples is nongeneric (Geanakoplos and Polemarchakis, 1982, Proposition 4). Indeed, one might argue that for a generic specification of prior beliefs, each agent's *first* announcement reveals that agent's information, and hence we need not worry about the difference between beliefs and information, and certainly need not worry about multiple rounds of announced beliefs.

We first note that if the state space is a (multi-dimensional) continuum with agents receiving continuously distributed signals, and if an agent observed several signals, then a *one*-dimensional announcement will typically (and generically) not reveal all the agent's information. We find it convenient in the examples to strip away complications by working with discrete signals, but are then unwilling to appeal to genericity arguments. Second, even within a discrete framework, the space of prior beliefs may not be the appropriate space to seek genericity. For example, the factors determining which state has occurred may be summarized by a tree, with random moves at decision nodes and terminal nodes corresponding to states. We would then apply genericity arguments to the mixtures appearing in the tree. If this tree has a nontrivial structure, then generic specifications of the proba-

¹¹The limit beliefs held by an agent oracle are Geanakoplos and Polemarchakis's (1982) indirect communication equilibrium beliefs.

bilities appearing in the tree will induce probability distributions over states that appear nongeneric, but that we nonetheless view as robust.

We believe that the repeated announcement of beliefs gives us information transmission similar to that allowed by (for example) the common knowledge that agents are willing to trade, sufficiently so that we are willing to avoid modeling the fine details of the market interaction by working directly with sequences of belief announcements. However, we are not convinced that market or other interactions will necessarily reveal every detail of every agent's information, and so would be skeptical of a model that precluded pooling.

We have noted that pooling can prevent the complete learning of agents' information, even with correctly specified models. Our next example illustrates a phenomenon that can only arise with agents having different models: increasing the information of one agent (even when another agent thinks the information is valuable) can result in a deterioration of inferences.

Example 3 We jump immediately to the tabular presentation of this product case example, which includes all the relevant information, presented in Figure 3. In contrast to the presentation of our earlier examples, we replace the column specifying the indicator function, f , with f^* , its expected value conditional on all the agents' model dimensions, i.e., $f^*(\omega) := \mathbb{E}[f(\omega) \mid \omega_1, \omega_2, \omega_3]$. In Example 2, $N = \cup_i M^i$, which is to say that the dimensions contained in $\{0, 1\}^{\cup_i M^i}$ suffice to determine the value of f . In the current example, there are additional dimensions in the state space that we have not presented. These dimensions lie outside all agents' models, and play a role in the analysis only to the extent that they shape the values of f^* and so we omit them from the table.

Since ω_1 is independent of (ω_2, ω_3) , agent 2 learns nothing from agent 1 and does no updating. Agent 1's learns the realization of ω_2 from agent 2, and so does one round of updating. In four of the states, agent 1 learns the probability of F , namely 0. Agent 1 overestimates the value of F in two of the remaining four states and underestimates it in the remaining two states.

Now suppose we give agent 2 more information, as displayed in Figure 4. Agent 2 again does not update, while agent 1 does one round of updating. As a result of the additional information, agent 2 now pools her states. Agent 1 does *not* estimate the probability of F correctly in any state. \blacklozenge

State $(\omega_1, \omega_2, \omega_3)$	Prior ρ	$f^*(\omega)$	Interim beliefs		First-round update
			$\beta^1(\omega_{I^1})$	$\beta^2(\omega_{I^2})$	$\beta^1(\omega_{I^1}, b_0^2)$
(0, 0, 0)	1/10	0	$(2x + y)/5$	$(x + y)/4$	0
(1, 0, 0)	1/10	x	$(x + y)/5$	$(x + y)/4$	$(x + y)/2$
(0, 0, 1)	1/10	0	$(2x + y)/5$	$(x + y)/4$	0
(1, 0, 1)	1/10	y	$(x + y)/5$	$(x + y)/4$	$(x + y)/2$
(0, 1, 0)	2/10	x	$(2x + y)/5$	$(2x + y)/6$	$(2x + y)/2$
(1, 1, 0)	2/10	0	$(x + y)/5$	$(2x + y)/6$	0
(0, 1, 1)	1/10	y	$(2x + y)/5$	$(2x + y)/6$	$(2x + y)/2$
(1, 1, 1)	1/10	0	$(x + y)/5$	$(2x + y)/6$	0

$$\Omega = \{0, 1\}^3, \quad M^1 = \{1, 2, 3\}, \quad M^2 = \{2, 3\}, \\ I^1 = \{1\}, \quad I^2 = \{2\}.$$

Figure 3: The beliefs for Example 3.

State $(\omega_1, \omega_2, \omega_3)$	Prior ρ	$f^*(\omega)$	Interim beliefs		First-round update
			$\beta^1(\omega_{I^1})$	$\beta^2(\omega_{I^2})$	$\beta^1(\omega_{I^1}, b_0^2)$
(0, 0, 0)	1/10	0	$(2x + y)/5$	$x/2$	$2x/3$
(1, 0, 0)	1/10	x	$(x + y)/5$	$x/2$	$x/3$
(0, 0, 1)	1/10	0	$(2x + y)/5$	$y/2$	$y/2$
(1, 0, 1)	1/10	y	$(x + y)/5$	$y/2$	$y/2$
(0, 1, 0)	2/10	x	$(2x + y)/5$	$x/2$	$2x/3$
(1, 1, 0)	2/10	0	$(x + y)/5$	$x/2$	$x/3$
(0, 1, 1)	1/10	y	$(2x + y)/5$	$y/2$	$y/2$
(1, 1, 1)	1/10	0	$(x + y)/5$	$y/2$	$y/2$

$$\Omega = \{0, 1\}^3, \quad M^1 = \{1, 2, 3\}, \quad M^2 = \{2, 3\}, \\ I^1 = \{1\}, \quad I^2 = \{2, 3\}.$$

Figure 4: The result of giving agent 2 in Figure 3 increased information.

3.3 Properties of the Belief Updating Process

The following proposition gathers some information about the belief-updating process. Recall that throughout, we maintain the assumption that $\mathcal{I}^i \subseteq \mathcal{M}^i$ for all i , and that $\mathbf{b} = (b_0, b_1, \dots)$ denotes the complete sequence of publicly announced beliefs with associated σ -algebra \mathcal{B}_∞ . We introduce the *omniscient* oracle, who (in addition to having the model \mathcal{F} , or N in the product case) knows the realization of the state.

Proposition 1

1. If \mathcal{M}^j is finite for all $j \in \{1, \dots, K\}$, then \mathbf{b} is eventually constant, i.e., the updating process terminates. Upon termination, posterior beliefs need not be equal.
2. If \mathcal{M}^j is infinite for at least two $j \in \{1, \dots, K\}$, then while the updating process may not terminate, the beliefs do converge to limits that need not be equal.
3. Once an agent's belief equals 0 or 1, that agent's belief agrees with those of the omniscient oracle, and so are never subsequently revised.¹²
4. Two agents cannot simultaneously assign a belief of 0 and 1 to the event F .
5. Agent i 's private information is only pooled in the limit if it doesn't make any difference to agent i . Formally,

$$\mathbb{E}[f^i \mid \mathcal{I}^i, \mathcal{B}_\infty] = \mathbb{E}[f^i \mid \mathcal{B}_\infty].$$

6. If f is measurable with respect to \mathcal{M}^i , then $f^i = f$ and agent i 's limit belief equals the agent- i oracular and public oracular belief, that is,

$$\begin{aligned} \mathbb{E}[f \mid \mathcal{I}^i, \mathcal{B}_\infty] &= \mathbb{E}[f^i \mid \mathcal{I}^i, \mathcal{B}_\infty] \\ &= \mathbb{E}[f^i \mid \mathcal{B}_\infty] \\ &= \mathbb{E}[f \mid \mathcal{B}_\infty]. \end{aligned}$$

¹²So model-based reasoners cannot match the common description of being “often wrong but never in doubt.”

7. If $\cup_{j=i}^K \mathcal{I}^j \subseteq \mathcal{M}^i$, then agent i 's limit belief equals the agent- i oracular and public oracular belief,

$$\begin{aligned} \mathbb{E}[f \mid \mathcal{I}^i, \mathcal{B}_\infty] &= \mathbb{E}[f^i \mid \mathcal{I}^i, \mathcal{B}_\infty] \\ &= \mathbb{E}[f^i \mid \mathcal{B}_\infty] \\ &= \mathbb{E}[f \mid \mathcal{B}_\infty]. \end{aligned}$$

Proof.

1. At each round n of the updating process, agent i 's belief about the event F is the expectation of i 's full-information belief conditioning on the σ -algebra reflecting information revealed by the collective announcements of the agents and i 's information, $\sigma(\mathcal{I}^i, \mathcal{B}_n)$. The sequence $(\sigma(\mathcal{I}^i, \mathcal{B}_n))_{n=0}^\infty$ is a filtration, with each σ -algebra being coarser than $\sigma(\mathcal{I}^1, \dots, \mathcal{I}^K)$.

If all \mathcal{M}^j are finite, then each $\sigma(\mathcal{I}^i, \mathcal{B}_n)$ and $\sigma(\mathcal{I}^1, \dots, \mathcal{I}^K)$ are generated by finite partitions, and so the filtration must eventually be constant, ensuring that the updating process terminates. Example 2 shows that the limit beliefs need not agree.

2. The convergence of agent i 's beliefs is an implication of the observation that beliefs are a martingale (Lemma 1). Appendix A.3 describes a product case example with infinite M^1 and M^2 in which updating proceeds for an infinite number of rounds, with beliefs converging to limits that (with positive probability) are not equal.
3. A belief b_n^i for agent i can equal an extreme value (0 or 1) at some round n if and only if the full-information belief $f^i(\omega)$ takes the same extreme value on a full ρ -measure event in $\sigma(\mathcal{I}^i, \mathcal{B}_n)$, which implies the omniscient oracle has the same beliefs on a full ρ -measure event in $\sigma(\mathcal{I}^i, \mathcal{B}_n)$, and so on every subsequent subevent in the sequence.
4. This is immediate since the omniscient oracle cannot have two distinct beliefs.
5. Proof is by contradiction. Suppose that

$$\mathbb{E}[f^i \mid \mathcal{I}^i, \mathcal{B}_\infty] \neq \mathbb{E}[f^i \mid \mathcal{B}_\infty].$$

Then, \mathcal{B}_∞ must pool together some states that agent i does not pool

together, and on which f^i is not constant.¹³ But if this were the case, then there would be an announcement from agent i not contained in \mathcal{B}_∞ , a contradiction.¹⁴

6. Immediately follows from the definitions and item 5.
7. We verify the first equality. Since $\sigma(\mathcal{I}^i, \mathcal{B}_\infty) \subseteq \sigma(\cup_j \mathcal{I}^j)$, if $\cup_j \mathcal{I}^j \subseteq \mathcal{M}^i$, then $\sigma(\mathcal{I}^i, \mathcal{B}_\infty) \subseteq \sigma(\mathcal{M}^i)$, and so

$$\begin{aligned} \mathbb{E}[f^i | \mathcal{I}^i, \mathcal{B}_\infty] &= \mathbb{E}[\mathbb{E}[f | \mathcal{M}^i] | \mathcal{I}^i, \mathcal{B}_\infty] \\ &= \mathbb{E}[f | \mathcal{I}^i, \mathcal{B}_\infty], \end{aligned}$$

where the first equality is just (1) and the second applies the law of iterated expectations.

The second equality is just item 5 above, while the third equality is established by an identical argument to that which verified the first equality.

■

Agent 2's limit model-based beliefs in Example 2 are equal to those of an omniscient oracle for the states reported in the final two lines of Figure 2, in keeping with Proposition 1.3. In every other case, agent 2's limit model-based beliefs differ from her agent oracular beliefs (which equal the column $\beta^1(\omega_{I^1}, b_0^2)$ in Figure 2, since agent 1 is an agent oracle).

3.4 Common Knowledge

We now explore the sense in which, once beliefs in the belief revision process have converged, the resulting beliefs, though different, are common knowledge. Here, we find it most natural to adopt the interpretation that the agents understand each others' models.

Suppose each agent's model \mathcal{M}^i is finite. We can think of agent i 's model as described by a finite partition of Ω and, since $\mathcal{I}^i \subseteq \mathcal{M}^i$, agent i 's information as a coarser partition of Ω . The announcement of a belief b^i implies that the event that led agent i to having that belief is common knowledge, and so all agent's information partitions are refined. After round

¹³More precisely, there exist two positive probability events E and E' in $\sigma(\mathcal{I}^i, \mathcal{B}_\infty)$ not separated by \mathcal{B}_∞ (i.e., for all events $B \in \mathcal{B}_\infty$, we have either $E, E' \subseteq B$ or $(E \cup E') \cap B = \emptyset$) for which $\mathbb{E}[f^i | E] \neq \mathbb{E}[f^i | E']$.

¹⁴Appendix A.4 shows that Proposition 1.5 does not hold for agent oracles.

n announcements, all agents have new partitions, and the intersection of the events leading to the round n announcements is common knowledge (though beliefs conditional on the intersection need not be common knowledge).

We say that a vector of beliefs (b^1, \dots, b^K) is *common knowledge at state* ω if these beliefs prevail at every state in that element of the meet of the agents' partitions containing ω , and their announcement does not lead to further revision of the partitions.

Intuitively, if the true state was not contained in a common knowledge event containing the final posteriors to be announced, then there would be further revision. This leads to:

Proposition 2 *If \mathcal{M}^i is finite for all i , then once the updating process terminates, the resulting beliefs are common knowledge.*

Proof. Each agent's interim belief, and each subsequent announcement by that agent, must be measurable with respect to the agent's partition. Each announcement thus gives rise to a common knowledge event. Moreover, for each player, these common knowledge events are descending, and hence form a sequence that is eventually constant. By Proposition 1.5, the limit beliefs are constant on this limit set, and so their announcement does not change agents' partitions. Moreover, since the \mathcal{M}^i are finite, all players know the finite time by which the updating process terminates, and so at that time the beliefs are common knowledge. ■

Remark 4 (common knowledge with infinite models) When the models are infinite, as in Example A.3, the belief revision process may continue without end. At no stage during the belief-revision process in Appendix A.3 are the beliefs common knowledge. Despite this difficulty, there is an appropriate notion of common knowledge when the models are infinite (due to Brandenburger and Dekel, 1987) that we apply in Appendix A.5 and show that the limit beliefs are again common knowledge. ◆

The common knowledge of limit beliefs implies an agreement theorem.

Proposition 3 *If all agents have the same (finite) model \mathcal{M} , then all agents have the same limit beliefs, for all possible information structures.*

Proof. In each round, all agents are updating their beliefs on the same partition \mathcal{M} , and since beliefs are common knowledge, they must agree (Aumann, 1976). ■

We restrict attention in the statement of this result to finite models in order to apply Aumann’s theorem. This result is immediately intuitive, since the coincidence of agents’ models removes the friction that leads agents to different beliefs. Technical complications, arising out of the potential need to condition on zero-probability events, appear in extending the result to infinite models. Green (2012) presents an agreeing-to-disagree result for infinite models that would allow such an extension. Section 4.1 shows that there is little hope short of common models for agents’ beliefs to coincide.

4 When Are Crowds Confused?

This section establishes a pair of negative results. We focus throughout this section on the product case. More general results hold that are similar in spirit, but considerably more cumbersome to state and interpret.

First, if realizations are drawn independently across dimensions and agents’ models and information exhibit nontrivial differences, then the process of belief exchange will not lead agents to the same beliefs. This first result is expected. Crowds are expected to be wise not because they lead people to the same belief, but because they allow diverse beliefs to be aggregated in an informative way. Our second result is that in general, the process of belief exchange will not ensure that some aggregate measure of their beliefs will be close to the belief of an agent who held the correct model.

4.1 Do Crowds Agree?

In one sense, it is obvious (and our earlier examples confirm) that when agents have substantively different models, their limit beliefs may not agree. In this section, we show that this failure is unavoidable—people who have different models are bound to disagree.

To ease notation, we consider the product case and suppose there are two agents. We say that limit beliefs *necessarily agree* if, for all $\omega \in \Omega$, the limit beliefs of agents 1 and 2 are equal, i.e.,

$$\mathbb{E}[f^1 | \mathcal{G}^{I^1}, \mathcal{B}_\infty] = \mathbb{E}[f^2 | \mathcal{G}^{I^2}, \mathcal{B}_\infty]. \quad (4)$$

The left side is agent 1’s model-based belief, giving 1’s observation of ω_{I^1} and the announced sequence of beliefs, and the right side is agent 2’s corresponding belief.

We say that the subset \tilde{I}^i is *redundant* in agent i ’s model if f^i is measurable with respect to $\mathcal{G}^{M^i \setminus \tilde{I}^i}$.

Proposition 4 *Suppose in the product case, the values ω_k are drawn independently, and $K = 2$. If for some $i \in \{1, 2\}$, $I^i \setminus M^j$ is not redundant for agent i , then the limit beliefs of agents 1 and 2 do not necessarily agree.*

Proof Suppose $I^1 \setminus M^2$ is not redundant for agent 1. Since the dimensions are independent, the announcements of agent 2 only convey information about the dimensions in M^2 , and since agent 1's theory f^1 is not measurable with respect to $\mathcal{G}^{M^1 \setminus (I^1 \setminus M^2)}$, agent 1's limit beliefs will vary over $\mathcal{G}^{I^1 \setminus M^2}$. Agent 2's limit beliefs, however, must ignore any information about dimensions outside M^2 (again because of independence), and so limit beliefs cannot necessarily agree. ■

This result indicates that agents' limit beliefs always agree only if agents have *no* useful information about the sources of disagreement between their models.

The argument in Proposition 4 leaves open the possibility that if there is correlation between these values, then it may well be that agent 1 observes information that is useful to agent 2, not because it appears in 2's model but because it is correlated with the values of other dimensions in 2's model (that 2 does not observe), all while beliefs necessarily agree. The example in Appendix A.6 illustrates that this can indeed occur. Moreover, Section 5 shows that strong correlation implies limit beliefs will be close,

4.2 Are Crowds Wise?

The point behind the “wisdom of the crowd” is not that the crowd makes individuals wise, but that the crowd itself is wise. We can thus reasonably assert that information is aggregated, even though various agents disagree, as long as the crowd forms beliefs that are “correct on average.”

We have introduced agent oracles, the public oracle and the omniscient oracle. We now introduce the *universal oracle*, who has access to all of the agents' information and hence has beliefs $\mathbb{E}[f \mid \mathcal{I}^1, \dots, \mathcal{I}^K]$.

All oracular beliefs are based on the true indicator function f . The difference between the different oracles is the information on which the oracles condition. In order of increasing information, the public oracle has the least information (namely, \mathcal{B}_∞), followed by an agent's oracle (who has both \mathcal{B}_∞ and that agent's information \mathcal{I}^i), then the universal oracle, and finally the omniscient oracle.

The least demanding standard for beliefs being correct on average is that the universal oracular belief lies in the convex hull of the agents' updated

State (ω_1, ω_2)	Prior ρ	$f^*(\omega)$	Theories $f^1(\omega_{M^1})$ $f^2(\omega_{M^2})$	
(0, 0)	1/4	7/8	1/2	9/16
(0, 1)	1/4	1/8	1/2	1/2
(1, 0)	1/4	2/8	9/16	9/16
(1, 1)	1/4	7/8	9/16	1/2

$$X = \{0, 1\}, \quad M^1 = \{1\}, \quad M^2 = \{2\},$$

$$I^1 = \{1\}, \quad I^2 = \{2\}.$$

Figure 5: The universal oracular beliefs are not in the convex hull of agent beliefs.

beliefs. Unfortunately, even this mild requirement is not guaranteed.

Example 4 We examine a product case in which $M^1 \cup M^2 = \Omega = I^1 \cup I^2 = \Omega$, so every dimension appears in the model of at least one agent and is also observed by at least one agent. This presents conditions most favorable to information aggregation. Consider the environment in Figure 5. Both agents observe the information they deem relevant, neither updates, and their beliefs are given by their theories. In every state, the universal oracular belief (given by $f^*(\omega)$) lies outside the convex hull of the agents' limit beliefs. \blacklozenge

Our next proposition shows it is a more general result that the universal oracular belief lies outside the convex hull of the model-based beliefs. We say a model $\tilde{\mathcal{G}}$ is *sufficient* if f is measurable with respect to the completion of $\tilde{\mathcal{G}}$, i.e., if the events in $\tilde{\mathcal{G}}$ suffice to determine whether the event F has occurred.¹⁵ In the product case, a subset $\tilde{N} \subset N$ is sufficient if the dimensions in \tilde{N} suffice to determine whether F has occurred. Obviously N is always a sufficient set. If there are smaller sufficient sets, then there will be many smaller sufficient sets. There is always at least one minimal sufficient set, and there may be multiple minimal sufficient sets (e.g., if the realizations on some dimensions are perfectly correlated).

Proposition 5 *In the product case, let $\cup_{k=1}^K M^k = \cup_{k=1}^K I^k = \tilde{N}$ for some minimal sufficient set \tilde{N} , with the M^k mutually disjoint and $M^k \subsetneq \tilde{N}$ for*

¹⁵The completion of a σ -algebra \mathcal{G} is the σ -algebra generated by \mathcal{G} and all the zero measure subsets.

each k . Then there exist states for which the universal oracular belief lies outside the convex hull of the model-based beliefs.

Proof Suppose first that beliefs reveal the agent’s information, i.e., $\mathcal{B}_\infty = \sigma(\mathcal{G}^{I^1}, \dots, \mathcal{G}^{I^K})$. Then the agents’ limit model-based beliefs will be their full-information beliefs. However, because f is not measurable with respect to the completion of any \mathcal{G}^{M^j} , for every agent there is a state at which her beliefs do not equal 0 or 1.

We now argue that there is a state at which no agent’s beliefs equal 0 or 1. Suppose not. Then (up to sets of ρ -measure zero), from which (along with the fact that f is not constant) it cannot be the case that for every state, there is at least one agent whose belief is either 0 or 1.¹⁶ Because \tilde{N} is sufficient, the universal oracular belief will always be 0 or 1, and hence must sometimes lie outside the convex hull of the agents beliefs. If beliefs are not revealing, then the agents have less information, and so again there cannot be an agent whose beliefs are always either 0 or 1. ■

The principle behind this result is reminiscent of Roux and Sobel’s (2015) argument that groups will typically have more precise information than any individual in the group, and hence will react more strongly to that information.¹⁷ Notice that our result holds in the absence of pooling announcements, and so is not simply a statement that the universal oracle has more information than does any single agent—in the limit they will often have identical information. However, the universal oracle has a more encompassing model than any of the individuals in the crowd, and hence makes use of more information, leading to more extreme beliefs.

¹⁶Recall that a full-information belief $f^i(\omega_{M^i})$ can equal 1 only if $f(\omega_{M^i}, \omega_{-M^i}) = 1$ for almost all ω_{-M^i} . We couple this with the product structure of the agent’s disjoint models to conclude that if there is in every state an agent whose full-information belief is either 0 or 1, then these beliefs must either be *almost all* 0 or *almost all* 1, which is to say that f must be almost surely constant.

¹⁷Roux and Sobel (2015) examine Bayesian agents with common preferences and holding a common view of a monotonic decision problem (so that only the magnitude, and not the direction, of the optimal decision is in question). They identify conditions under which the optimal action of the group is more variable than the distribution of actions taken by the members of the group. The underlying idea is that the group has more precise information than any individual, and hence will respond to this information more vigorously than will an individual. As a result, in those circumstances in which an extreme decision is optimal, individuals will move only slightly toward that extreme, even if they have extreme signals, because their individual signals are noisy. The group, facing a collection of extreme signals, essentially has a very precise extreme signal, and hence will move far in that direction.

State $(\omega_1, \omega_2, \omega_3)$	Prior ρ	$f(\omega)$	Interim beliefs	
			$\beta^1(\omega_{I^1})$	$\beta^2(\omega_{I^2})$
(0, 0, 0)	1/8	1	1	1/2
(0, 0, 1)	1/8	1	1	1/2
(0, 1, 0)	1/8	1	1/2	1
(0, 1, 1)	1/8	0	1/2	0
(1, 0, 0)	1/8	0	0	1/2
(1, 0, 1)	1/8	0	0	1/2
(1, 1, 0)	1/8	1	1/2	1
(1, 1, 1)	1/8	0	1/2	0

$$\Omega = \{0, 1\}^3, \quad M^1 = \{1, 2\}, \quad M^2 = \{2, 3\}, \\ I^1 = \{1, 2\}, \quad I^2 = \{2, 3\}.$$

Figure 6: The beliefs for Example 5. The updating process terminates with the interim beliefs.

Proposition 5 shows that if the agents have distinct views of the world, in the sense that their models do not intersect, then the universal oracular belief will invariably sometimes lie outside the convex hull of the agents' beliefs. Example 5 shows that if the agents have more similar models, in the sense of sharing some common dimensions, then it is possible (though not guaranteed) that the convex hull of the agents' beliefs may contain the universal oracular belief. This provides the first hint of a theme we develop in the next section: crowds can effectively aggregate information if the agents' models have enough in common, even though information may be widely dispersed and the agents' models may yet exhibit significant (even pervasive) differences.

Example 5 Consider the crowd presented in Figure 6. The models of agents 1 and 2 share one dimension in common. In each state ω , one agent has a belief matching the universal oracular belief $f(\omega)$, and hence the latter lies in the convex hull of the agents' beliefs. \blacklozenge

5 What Makes A Crowd Wiser?

Section 4 makes it clear that for an arbitrarily fixed crowd, we cannot in general expect limiting model-based beliefs to reflect universal oracular beliefs. Nonetheless, Surowiecki (2004) argues that appropriately configured

crowds are wise, emphasizing that crowds should have diverse points of view, independent reasoning, and a decentralized structure.

This section presents three variations on the result that crowds will effectively aggregate information if their members have a sufficiently common understanding. The agents need not have similar information, and indeed each individual agent may have very little information. The market will aggregate their information, as long as their models by which they interpret this information are not too different. We develop these results for the general case, with the exception that the attendant notational complexities force us to seek refuge in the product case when presenting Proposition 8.

5.1 Correlated Model Predictions

If the realizations in the agents' different models are sufficiently correlated across models, then their limit beliefs will be close. We view this correlation as an indication that the agents' models are, for practical purposes, nearly the same. The extreme case involves agents whose models are disjoint, but whose realizations are perfectly correlated, so that the agents effectively have the same model described in different languages.

Proposition 6 *Fix agent i 's theory f^i . For any $\varepsilon > 0$, there is an $\eta < 1$ such that if the coefficient of correlation between agent i 's theory f^i and agent j 's theory f^j is at least $1 - \eta$, then agents i 's and j 's limiting beliefs are within ε of one another with probability $1 - \varepsilon$.*

This proposition imposes the correlation requirement on f^i and f^j , rather than, in the product case, imposing a (stronger) requirement on the correlation between ω_{M^i} and ω_{M^j} , because correlation is relevant only for those dimensions that play a role in affecting beliefs about F .

The proof first shows that if two agents' theories are perfectly correlated ex ante, then their updated beliefs must be identical. In this case, the agents effectively have identical models with different descriptions. We then show that if two agents' theories are *close* ex ante, then their limit beliefs must, with high probability, be close. The delicateness in establishing this seemingly intuitive result arises in showing that it holds *irrespective* of the nature of the updating. The following lemma (proved in Appendix A.7), which will also be useful later, provides a key technical result.

Lemma 2 *Suppose $(f_n)_n$ is a sequence of \mathcal{F} -measurable functions converging almost surely to the \mathcal{F} -measurable function f^\dagger . For all $\delta > 0$, there exists*

a set Ω_δ with $\rho(\Omega_\delta) > 1 - \delta$ and an integer N_δ such that for all $n > N_\delta$ and for all σ -algebras $\mathcal{H} \subseteq \mathcal{F}$,

$$\left| \mathbb{E}[f^\dagger | \mathcal{H}](\omega) - \mathbb{E}[f_n | \mathcal{H}](\omega) \right| < \delta \quad \forall \omega \in \Omega_\delta.$$

Proof of Proposition 6. Suppose first that the coefficient of correlation between f^i and f^j equals 1. Then $f^j - \mathbb{E}f = \alpha(f^i - \mathbb{E}f)$ ρ -almost surely for some constant $\alpha > 0$. Suppose f^i is not constant (if it were, the result is trivial), so that for some $x > 0$, $\mathbb{E}f + x$ is in the support of f^i .

We now argue that $\alpha = 1$. En route to a contradiction, suppose $\alpha > 1$ (a similar argument rules out $\alpha < 1$). Fix $\varepsilon > 0$ so that $\alpha(x - \varepsilon) > x$ and set $B(x) := \{\omega : x - \varepsilon \leq f^i(\omega) - \mathbb{E}f \leq x\}$. We may assume $\rho(B(x)) > 0$ (if not, marginally increasing the value of x yields a positive measure set). Then, for $y = \alpha x$ and $B'(y) := \{\omega : y - \alpha\varepsilon \leq f^j(\omega) - \mathbb{E}f \leq y\}$, we have $\rho(B(x) \Delta B'(y)) = 0$. From (1), we then have

$$\begin{aligned} x\rho(B(x)) &\geq \int_{B(x)} (f^i(\omega) - \mathbb{E}f) d\rho \\ &= \int_{B(x)} (f(\omega) - \mathbb{E}f) d\rho \\ &= \int_{B'(y)} (f(\omega) - \mathbb{E}f) d\rho \\ &= \int_{B'(y)} (f^j(\omega) - \mathbb{E}f) d\rho \geq (y - \alpha\varepsilon)\rho(B'(y)), \end{aligned}$$

and so $x \geq \alpha(x - \varepsilon)$, a contradiction.

From Proposition 1.5, we have that agent i 's limiting belief $\mathbb{E}[f^i | \mathcal{I}^i, \mathcal{B}_\infty]$ equals $\mathbb{E}[f^i | \mathcal{B}_\infty]$ which equals $\mathbb{E}[f^j | \mathcal{B}_\infty]$, and so agent i and j 's limiting beliefs agree on any sequence of announced posteriors.

Turning to the approximation, it is enough to prove that we can make $\mathbb{E}|f^i - f^j|$ arbitrarily small by choosing η sufficiently small. We prove the latter by contradiction. If not, then there exists $\varepsilon > 0$ such that for all $n > 0$ there exists f_n^j such that the correlation between f^i and f_n^j is at least $1 - 1/n$ and yet $\mathbb{E}|f^i - f_n^j| > \varepsilon$.

Define $X := f^i - \mathbb{E}f$ and $Y_n := f_n^j - \mathbb{E}f$. Then,

$$\begin{aligned} \mathbb{E}[Y_n \mathbb{E}(X^2) - X \mathbb{E}(XY_n)]^2 &= \mathbb{E}(X^2) [\mathbb{E}(X^2) \mathbb{E}(Y_n^2) - \mathbb{E}(XY_n)^2] \\ &\leq \mathbb{E}(X^2) [\mathbb{E}(X^2) \mathbb{E}(Y_n^2) - (1 - 1/n)^2 \mathbb{E}X^2 \mathbb{E}Y_n^2] \\ &= (\mathbb{E}X^2)^2 \mathbb{E}Y_n^2 [1 - 1 + 2/n - 1/n^2], \end{aligned}$$

and so $Y_n \mathbb{E}(X^2) - X \mathbb{E}(XY_n)$ converges in mean square to 0 (since $\mathbb{E}Y_n^2$ is bounded above by $\frac{1}{4}$). If (Y_n) (or any subsequence) has a limit in mean square (and so a limit in mean), then that limit must equal X (for the reasons above). We will show that every subsequence has a convergent subsubsequence, which implies that the original sequence converges to X .

We use n to index an arbitrary subsequence and let $\alpha_n := \mathbb{E}XY_n / \mathbb{E}(X^2)$, so that $Y_n - \alpha_n X$ converges to 0 in mean square. We claim that (α_n) has a convergent subsequence. For, if not, then $|\alpha_n| \rightarrow \infty$, which implies $\mathbb{E}Y_n^2 \rightarrow \infty$, which is impossible.

Suppose (α_{n_k}) converges to some α . Then,

$$0 \leq \mathbb{E}(Y_{n_k} - \alpha X)^2 \leq \mathbb{E}(Y_{n_k} - \alpha_{n_k} X)^2 + \mathbb{E}(\alpha - \alpha_{n_k})^2 X^2 \rightarrow 0,$$

and so Y_n converges in mean square to αX , and so $\alpha = 1$.

It remains to argue that for n sufficiently large, with probability at least $1 - \varepsilon$,

$$|\mathbb{E}[Y_n | \mathcal{I}^j, \mathcal{B}_\infty] - \mathbb{E}[X | \mathcal{I}^i, \mathcal{B}_\infty]| < \varepsilon.$$

By Proposition 1.5, this inequality can be rewritten as

$$|\mathbb{E}[Y_n | \mathcal{B}_\infty] - \mathbb{E}[X | \mathcal{B}_\infty]| < \varepsilon.$$

Since every subsequence of $(Y_n)_n$ has a sub-subsequence almost surely converging to X , the desired result is implied by Lemma 2. \blacksquare

5.2 Models with a Common Component

The following result shows that if the agents' models share a large enough common component, then their beliefs cannot be too different from one another. We maintain our assumption that for each agent i , we have $\mathcal{I}_n^i \subseteq \mathcal{M}_n^i$.

Proposition 7 *Consider a sequence of crowds with models $(\mathcal{M}_n^1, \dots, \mathcal{M}_n^K)_{n=1}^\infty$, where each $(\mathcal{M}_n^i)_n$ is a filtration. Suppose that for all i and j , $\sigma(\mathcal{M}_1^i, \mathcal{M}_2^i, \dots) = \sigma(\mathcal{M}_1^j, \mathcal{M}_2^j, \dots)$. Then, for all (ρ, f) and for every $\delta > 0$, there exists N_δ such that for any accompanying sequence of information $(\mathcal{I}_n^1, \dots, \mathcal{I}_n^K)_{n=1}^\infty$, for all $n > N_\delta$, with probability at least $1 - \delta$, the limit beliefs of any agent i is within δ of the public oracular belief, and so with probability at least $1 - 2\delta$, within 2δ of the limit beliefs of any other agent j .*

The second sentence of the proposition ensures that as n grows, if agents' models also grow, the agents' models have more and more in common (and if the models are constant, then they agree). Notice that we *do* allow $\sigma(\mathcal{M}_1^i, \mathcal{M}_2^i, \dots) \subsetneq \mathcal{F}$, so that no agent i ever acquires a complete understanding of the event F , and at the extreme the hypotheses of the proposition are also consistent with all agents having a common fixed model.

In the product case, the proposition requires that as the agents' models grow, they have more dimensions in common. A simple but more demanding assumption in the product case is that any dimension that appears in an agent's model eventually appears in every agent's model. Notice that this requirement is consistent with the agents having an arbitrarily small, even zero, proportion of their models in common, for every term in the sequence of crowds.¹⁸

As agents' models share an increasing common component, the agents come to share common beliefs, and these beliefs will be close to the public oracular belief. If the agents have access to little information, their beliefs will be rather uninformative, while agents with access to ample information will have beliefs close to those of an omniscient oracle.

The idea behind the proof is that as the models become increasingly sophisticated, differences in the explanatory power of the excluded events disappear.

Proof. Fix a value $\delta > 0$. Define

$$\mathcal{M}_\infty := \sigma(\mathcal{M}_1^i, \mathcal{M}_2^i, \dots)$$

(which is, by assumption, independent of i) and set $\hat{f} := \mathbb{E}[f \mid \mathcal{M}_\infty]$.

Agent i 's theory under her n^{th} model is given by

$$f_n^i = \mathbb{E}[f \mid \mathcal{M}_n^i] = \mathbb{E}[\hat{f} \mid \mathcal{M}_n^i]$$

(where the second equality follows from \mathcal{M}_n^i being coarser than \mathcal{M}_∞ and the law of iterated expectations). Since $(\mathcal{M}_n^i)_n$ is a filtration, with limit σ -algebra \mathcal{M}_∞ ,

$$f_n^i \rightarrow \hat{f} \quad \rho\text{-a.s.}$$

Our goal is to show that with probability at least $1 - \delta$, we have

$$\left| \mathbb{E}[f_n^i \mid \mathcal{G}^{I_n^i}, \mathcal{B}_\infty(n)] - \mathbb{E}[f \mid \mathcal{B}_\infty(n)] \right| < \delta,$$

¹⁸For example, in the product case, agent 1's n th model may be $\{1, 2, 3, \dots, n\} \cup \{1, 3, 5, 7, \dots\}$ and agent 2's n th model may be $\{1, 2, 3, \dots, n\} \cup \{2, 4, 6, 8, \dots\}$.

where $\mathcal{B}_\infty(n)$ is the σ -algebra induced by the sequence of publicly announced beliefs for the n th crowd.

By Lemma 2, with probability at least $1 - \delta$, we have

$$\left| \mathbb{E}[f_n^i \mid \mathcal{B}_\infty(n)] - \mathbb{E}[\hat{f} \mid \mathcal{B}_\infty(n)] \right| < \delta.$$

By Proposition 1.5,

$$\mathbb{E}[f_n^i \mid \mathcal{I}_n^i, \mathcal{B}_\infty(n)] = \mathbb{E}[f_n^i \mid \mathcal{B}_\infty(n)],$$

and so with probability at least $1 - \delta$, we have

$$\left| \mathbb{E}[f_n^i \mid \mathcal{I}_n^i, \mathcal{B}_\infty(n)] - \mathbb{E}[\hat{f} \mid \mathcal{B}_\infty(n)] \right| < \delta.$$

Finally, since $\mathcal{B}_\infty(n)$ is coarser than \mathcal{M}_∞ (since $\cup_j \mathcal{I}_n^j \subseteq \cup_j \mathcal{M}_n^j \subseteq \mathcal{M}_\infty$) and $\hat{f} := \mathbb{E}[f \mid \mathcal{M}_\infty]$, we have that with probability at least $1 - \delta$,

$$\left| \mathbb{E}[f_n^i \mid \mathcal{I}_n^i, \mathcal{B}_\infty(n)] - \mathbb{E}[f \mid \mathcal{B}_\infty(n)] \right| < \delta.$$

■

Proposition 7 gives sufficient conditions for agents' limiting beliefs to converge (as their models become more similar) to the public oracular belief. This establishes a relationship between the beliefs of different agents, but says little about how this belief relates to the event F . Strengthening the result to involve the universal or omniscient oracular belief would require placing conditions on the agents' information sets as well as models. We turn to one such result.

5.3 Crowds with Enough (Dispersed) Information

It is more in keeping with the wisdom-of-the-crowd spirit to seek conditions under which some measure of the central belief converges to the omniscient belief as the crowd grows large. This section presents such a result simplifying the exposition by restricting attention to the product case.

It is clear that this will require some conditions. Section 3.2 presents examples in which agents' announcements convey no information, despite the fact that all agents are agent oracles. We cannot expect model-based reasoners to come closer to omniscient beliefs than do agent oracles. Our goal is to identify conditions under which crowds of model-based reasoners will aggregate information, given that a crowd of agent oracular reasoners would do so.

Toward this end, we say that f is *discernible* if, for any sets $I^1, I^2 \subseteq N$, we have

$$\mathbb{E}[f|\mathcal{G}^{I^1 \cup I^2}] = \mathbb{E}[f|\sigma(\mathbb{E}[f|\mathcal{G}^{I^1}], \mathbb{E}[f|\mathcal{G}^{I^2}])],$$

where $\sigma(\mathbb{E}[f|\mathcal{G}^{I^1}], \mathbb{E}[f|\mathcal{G}^{I^2}])$ is the σ -algebra induced by the announcements of the beliefs $\mathbb{E}[f|\mathcal{G}^{I^1}]$ and $\mathbb{E}[f|\mathcal{G}^{I^2}]$. Discernibility requires that the announcements $\mathbb{E}[f|\mathcal{G}^{I^1}]$ and $\mathbb{E}[f|\mathcal{G}^{I^2}]$ allow an agent oracle to infer the information contained in $I^1 \cup I^2$ that is relevant for determining f (but may not allow the agent oracle to identify $\omega_{I^1 \cup I^2}$, and hence allows the pooling of information that is irrelevant for determining F). For any f , discernibility will approximately hold if I_1 and I_2 are sufficiently large. Discernibility fails in the pooling examples of Section 3.2.

The proof of the following Lemma is in Appendix A.8.

Lemma 3 *For all $\varepsilon > 0$, there exists a finite set $K_\varepsilon \subseteq \mathbb{N}$ such for all σ -algebras \mathcal{H} ,*

$$\rho \{ |\mathbb{E}[f|\mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) - f(\omega)| < \varepsilon \} \geq 1 - \varepsilon.$$

Remark 5 (interpreting Lemma 3) Lemma 3 implies that if an agent's model includes the dimensions K_ε , not only is the agent's theory close to the true description, but adding further dimensions to the model cannot change the agent's theory significantly. The first implication is relatively straightforward, while establishing the second requires somewhat more care. Notice that this lemma allows a variation on Proposition 7. Any sequence of crowds whose agents' models eventually include K_ε must eventually have beliefs that are close to those of a public oracle (and hence each other). Lemma 3 plays the role of Lemma 2 in the proof of Proposition 7, with attendant straightforward changes. \blacklozenge

We assume that as the crowd grows, it becomes increasingly likely that the models (though not necessarily the information) of most agents become sufficiently sophisticated. Let λ be a full support measure on the space of all finite models. Let $\Gamma \in \mathbb{N}$. For each finite model $M \subseteq \mathbb{N}$, let μ_M be a probability distribution with full support over the subsets of M with at most $\min\{|M|, \Gamma\}$ elements.

Proposition 8 *Suppose f is discernible. Consider a crowd of n agents, with each agent i 's model M^i independently drawn according to λ and information set I^i then drawn independently according to μ_{M^i} . For all $\varepsilon > 0$,*

[8.1] there exists an N_ε such that for all finite models M containing K_ε and all $n > N_\varepsilon$, the probability that every agent with the model M has a belief within ε of a hypothetical agent with the same model M and observing data $I = K_\varepsilon$ is at least $1 - \varepsilon$, and

[8.2] there exists N_ε^* such that for all $n > N_\varepsilon^*$, the probability the average of beliefs is within $2\varepsilon + (1 - \eta)$ of the omniscient belief is at least $1 - \varepsilon$, where η is the probability under λ that a randomly drawn model does not include K_ε .

Proof. We first note that discernibility extends to finite numbers of sets. We have

$$\begin{aligned} \mathbb{E}[f|\mathcal{G}^{I^1 \cup I^2 \cup I^3}] &= \mathbb{E}[f|\sigma(\mathbb{E}[f|\mathcal{G}^{I^1 \cup I^2}], \mathbb{E}[f|\mathcal{G}^{I^3}])] \\ &= \mathbb{E}[f|\sigma(\mathbb{E}[f|\mathcal{G}^{I^1}], \mathbb{E}[f|\mathcal{G}^{I^2}], \mathbb{E}[f|\mathcal{G}^{I^3}])], \end{aligned}$$

where the first equality is the statement of discernibility, and the second again applies discernibility (since $\sigma(\mathbb{E}[f|\mathcal{G}^{I^1 \cup I^2}], \mathbb{E}[f|\mathcal{G}^{I^3}]) = \sigma(\mathbb{E}[f|\mathcal{G}^{I^1}], \mathbb{E}[f|\mathcal{G}^{I^2}], \mathbb{E}[f|\mathcal{G}^{I^3}])$).

For [8.1], fix a model M containing K_ε , and denote by \mathcal{K} a collection of subsets of \mathbb{N} satisfying $K_\varepsilon \subseteq \cup \mathcal{K}$ such that no set in \mathcal{K} has more than Γ elements. We can choose N_ε sufficiently large that for $N > N_\varepsilon$ with probability at least $1 - \varepsilon$, for every set $I^i \in \mathcal{K}$, there is an agent in the crowd whose model and information set consist precisely of that set. This agent's first-round belief will be $\mathbb{E}[f|\mathcal{G}^{I^i}]$. Applying discernibility, on the second and each subsequent round, every agent in the crowd whose model contains K_ε will have a belief $\mathbb{E}[f|\mathcal{G}^{K_\varepsilon}, \mathcal{H}]$ for some σ -algebra \mathcal{H} . Lemma 3 then implies the result. For [8.2], we note that we can choose $N_\varepsilon^* \geq N_\varepsilon$ so that for all $n > N_\varepsilon^*$, with probability at least $1 - \varepsilon$ the proportion of agents whose models include K_ε will be at least $\eta - \varepsilon$. The difference between the average belief and the belief of an omniscient oracle is then at most

$$(\eta - \varepsilon)\varepsilon + (1 - \eta + \varepsilon) \leq 2\varepsilon + 1 - \eta.$$

■

The substantive work in Proposition 8 is done by Lemma 3, which guarantees that any agent with a model containing K_ε and full information has beliefs that must be close to the omniscient oracle with high probability. It is not surprising that if we give an agent enough information, then their belief will be close to the omniscient oracular belief. The somewhat more delicate part of Lemma 3 lies in showing that nothing else that the agent could possibly observe can drive the agent outside the ε margin of error.

Proposition 8.1 establishes that as the crowd grows large, agents whose models include K_ε will have beliefs close to those of an omniscient oracle, no matter what they observe. Here, the exchange of information between agents obviates the need for each agent to have copious information. It may be that each individual agent has very little information—our requirement is only that some agents have at least Γ dimensions in their information, where Γ may be quite small.

It is clear that model-based reasoners have no hope of reaching omniscient beliefs if the announcements of oracles pool information. The discernibility requirement in Proposition 8 addresses pooling in a particularly brutal way, ensuring that for any information set I , the first announcement by an agent i with $M^i = I^i = I$ reveals all the relevant information contained in I . It would suffice that *limiting* beliefs reveal such information. We could accordingly introduce less demanding versions of discernibility, at the costs of greater complexity and pushing the assumption further away from the fundamentals of the problem. We view discernibility as being more demanding when applied to small information sets. We could work with versions of discernibility that apply only to larger sets, but then would need to place stronger requirements on the presence of agents with larger information sets.

Proposition 8.2 establishes that if *enough* agents have large enough (i.e., containing K_ε) models, then it is very likely that the average belief in the crowd will be very close to that of an omniscient oracle. This may appear to be nothing more than the statement that if enough people get it right, then the average will be about right. Again, the more delicate part of the argument is handled by Lemma 3, ensuring that the beliefs of those who would otherwise “get it right” are not disrupted by the presence of some agents with bizarre models. The average belief is then driven toward the omniscient belief, not by having those who get it right convincing or converting those who are confused, but by having the former swamp the latter. Notice, however, that for this to happen there must be sufficiently many agents with sufficiently large and common models. There is no similar requirement on the commonality of information. Prediction markets can indeed effectively aggregate dispersed information, *if* the agents have a sufficiently common understanding of the meaning of that information.

6 Related Literature and Discussion

The wisdom of the crowd has attracted considerable attention (the introduction mentions some points of entry into the literature). Arieli, Babichenko,

and Smorodinsky (2018) examine a model in which the members of a crowd receive a signal, update their beliefs, and then report their beliefs. The question is when an observer can infer the identity of the underlying state, despite knowing nothing about the agents' signal structures. The (rough) answer is that even if the crowd is arbitrarily large, no inferences can be drawn unless a signal drives a posterior belief to either 0 or 1. The flavor of this result is reminiscent of our observation (Proposition 1) that beliefs of 0 or 1 must match those of an omniscient oracle and that (from Appendix A.1) when beliefs are interior, Bayes rule places very little discipline on models in the absence of a common prior. The spirit of this exercise differs from our attempt to give the agents as good a chance as possible of aggregating information by effectively assuming that they understand each others' models.

There is a growing literature on misspecified models. The most closely related work is Spiegler (2016, 2018). Bohren (2016) and Bohren and Hauser (2018) study a model in which agents have difficulties extracting information from the actions of previous agents because they are unsure as to how much of the information contained in previous actions is new and how much is redundant. In contrast, our agents never have difficulties extracting information from others, but have different views as to which of this information is relevant.

Economists are sometimes criticized for pursuing ideas that work in theory but not in practice. In contrast, prediction markets work in practice but not theory. Wolfers and Zitzewitz (2004) describe the growing use of prediction markets and describe their effectiveness.¹⁹ On the one hand, our positive results hold only under strong assumptions. Proposition 4 indicates that under precisely the conditions in which prediction markets are thought to be valuable, namely dispersed information, we cannot expect agents to come to agreement. Proposition 5 indicates that we cannot in general expect the more reasonable goal of extracting useful information from the collection of market beliefs.

On the other hand, Propositions 6, 7 and 8 point to some circumstances under which we might reasonably expect prediction markets to work well. The basic requirement of these positive results is that the agents' models are sufficiently similar. Information may well be seemingly hopelessly dispersed, and yet will be effectively aggregated by the market, if enough agents have

¹⁹In a similar vein, various versions of double auctions with small numbers of agents do a remarkable job of producing competitive prices (Friedman and Rust (1993)), despite the lack of an underlying theory.

enough of a common view as to the meaning of such information.

Proposition 7 indicates that crowds can be effective in aggregating information, no matter how dispersed the information, if the participants have a sufficiently common understanding of what information is important. Galton's (1907) ox tale is one case where it seems reasonable to posit that the participants, presumably having had significant experience, would have similar understandings of oxen, even if their brief impressions lead them to different initial estimates of the ox standing before them. Similarly, the sales force of a firm may have a reasonably common understanding of which factors portend brisk sales, even if they extract different information from their experiences with their idiosyncratic client lists.

Predicting one-off political events appears to be qualitatively different. Nonetheless, Proposition 6 points to the forces that could lead to effective prediction markets. Suppose F is the event that some candidate, Mr. Smith, goes to Washington. Agent 1 is an economist who believes that Mr. Smith's election is determined by factors such as the current rates of employment, inflation, and economic growth. Agent 2 is a political scientist who believes the election hinges on social factors such as the electorate's belief that the country is "on the right track" and that "politicians are sensitive to my problems." The models of the two agents might well be disjoint, but as long as there is sufficient correlation between the relevant realizations, then the agents will come to similar views.

An appealing intuition is that large crowds are more effective in aggregating information than smaller crowds. Proposition 8 makes this intuition precise. The important implication here is that the advantage of a large crowd arises not because agents who have effective models "correct" the reasoning of those who do not, but because the former swamp the latter.

A Appendices

A.1 What's Wrong with Different Priors?

Perhaps the most common response to the no-trade theorem is to allow agents to hold different prior beliefs. Could it be that our analysis of model-based reasoning is simply a repackaged version of allowing agents to hold different priors?

The starkest difference is that models with different prior beliefs impose virtually no discipline on the relationship of the beliefs of different agents, and hence on the "collective" beliefs of the agents. In contrast, model-based reasoning insures that agents' beliefs about the events they deem relevant

are firmly anchored to the data. This imposes restrictions on the beliefs of individual agents as well as restrictions on how the beliefs of various agents can differ.

It is a common characterization of Bayesian updating that at least eventually “the data swamps the prior.” This suggests that the discordance allowed by differing priors should be only temporary, with the data eventually imposing as much discipline on a crowd with different priors as it does on crowd of model-based reasoners.²⁰ To investigate this, we consider the product case and examine a sequence of crowds that receive increasing amounts of information. In order to focus clearly in the discipline imposed on beliefs by this information, we assume the agents have common information. In particular, let $(I_n)_{n=0}^\infty$ be an increasing sequence of subsets of \mathbb{N} . We consider a sequence of crowds, with every agent’s information set I_n^i in crowd n given by I_n .

We begin with a model of different priors, holding fixed the other aspects of agents’ models. Suppose each agent has the correct state space and description f , but we place no restrictions on the priors ρ^i , and in particular no restrictions on how these priors may differ across agents.

Given the sequence of crowds, let $(\beta_n^i)_{i=1, n=0}^{K, \infty}$ be the sequence of induced limiting beliefs, for each agent in each crowd, about the event F . We now argue that once we allow priors to differ, there are few restrictions placed on the sequence of limit posteriors $(\beta_n^i)_{i=1, n=1}^{K, \infty}$, even though the agents are oracles in that their theories match the description f .

Of course, the agents’ limit posteriors are not completely arbitrary, as the mere fact that they are derived from Bayes’ rule imposes some restrictions. Say that the sequence $(\beta_n^i)_{i=1, n=0}^{K, \infty}$ has the *martingale property* if, for any agent i and ω_{I_n} , there exists $\omega_{I_{n+1}}$ consistent with ω_{I_n} with

$$\beta_{n+1}^i(\omega_{I_{n+1}}) < \beta_n^i(\omega_{I_n}), \quad (5)$$

if and only if there also exists $\omega'_{I_{n+1}}$ consistent with ω_{I_n} with

$$\beta_{n+1}^i(\omega'_{I_{n+1}}) > \beta_n^i(\omega_{I_n}). \quad (6)$$

Intuitively, an agent can receive encouraging news if and only if it is also possible for the agent to receive discouraging news. Note that this implies that zero and unitary beliefs are absorbing.

We also impose minimal consistency with f . The consistency requirement is the following, where the antecedents should be interpreted as the

²⁰This implicitly assumes the truth is in the support of the prior (technically, that the true distribution is absolutely continuous with respect to the prior).

joint hypothesis that the limit exists and has the indicated sign, and $[\omega_{I_n}]$ is the cylinder set given by $\{\omega_{I_n}, \omega_{-I_n}\}$,

$$\lim_n \beta_n^i(\omega_{I_n}) > 0 \implies \exists \omega \in [\omega_{\cup_n I_n}] \text{ s.t. } f(\omega) = 1 \quad (7)$$

$$\text{and } \lim_n \beta_n^i(\omega_{I_n}) < 1 \implies \exists \omega \in [\omega_{\cup_n I_n}] \text{ s.t. } f(\omega) = 0. \quad (8)$$

Requirements (7) and (8) are the only ones that connect the event F with agent beliefs. Without them, there is nothing precluding an agent from, for example, assigning positive probability to F on the basis of some information $\omega_{\cup_n I_n}$ when F is inconsistent with that information. If that were to happen, there is clearly no hope for $\beta_n^i(\omega_{I_n}) = \mathbb{E}_{\rho^i}[f \mid \omega_{I_n}]$.

Proposition 9 *Consider a sequence of crowds indexed by $n = 0, \dots$, with each agent i 's information set I_n^i in crowd n given by I_n , where the sequence $(I_n)_{n=0}^\infty$ is increasing. Suppose the sequences $(\beta_n^i)_{i=1, n=0}^{K, \infty}$ satisfy the martingale property and (7) and (8). Then there exists a vector of prior beliefs (ρ^1, \dots, ρ^K) generating the limiting posterior beliefs $(\beta_n^i)_{i=1, n=0}^{K, \infty}$, i.e., $\beta_n^i(\omega_{I_n}) = \mathbb{E}_{\rho^i}[f \mid \omega_{I_n}]$.*

Before proving this result, we make three observations. First, if $\cup_{n=0}^\infty I_n = \Omega$, then since beliefs are a martingale, $\beta_n^i \rightarrow f$ ρ^i -almost surely. For states that positive probability under ρ^i and ρ , the data then swamps the prior—agent i attaches probability one to the event that her beliefs about F converge to those of an omniscient oracle. However, the convergence in the previous observation is pointwise, not uniform. That is, for any finite sequence (β_n^i) satisfying the martingale property given in (5)–(6), there is a prior rationalizing (β_n^i) . Notice that there need be no connection between such a sequence and the event F . Hence, Bayesian updating from different priors places no restrictions on finite sequences of agents' beliefs, no matter how long. Moreover, if $\cup_{n=0}^\infty I_n \subsetneq \Omega$, the beliefs over states conditional on $\cup_{n=0}^\infty I_n$ are essentially arbitrary, needing only to satisfy the property that the conditional probability of F equals the limit of β_n^i . Hence, unless we are dealing with a case in which the agents will eventually resolve every vestige of uncertainty, updating places few restrictions on beliefs. If agents with different priors are also sufficiently romantic as to think the world will always contain some mystery, then we cannot expect their beliefs to be coherent.

Proof. We fix an agent i and construct the prior belief ρ^i , proceeding by induction. Note that β_0^i is the agent's prior probability of F . If this prior is either 0 or 1, then so must be all subsequent updates, and then any prior

belief with support contained either on the event F^c or on the event F (respectively, with the requisite set nonempty, by the martingale property) suffices.

Suppose $\beta_0^i \in (0, 1)$. By assumption, the measure β_1^i attaches conditional probabilities to a collection of cylinder sets of the form $[\omega_{I_1}]$, with some of these values larger than β_0^i and some smaller. Assign probabilities $\rho^i([\omega_{I_1}])$ to these sets so that the average of the conditional probabilities is β_0^i . Continuing in this fashion, we attach a probability to every cylinder set $[\omega_{I_n}]$. It follows from Kolmogorov's theorem (Billingsley, 2012, p. 517) that this measure extends to a probability measure ψ^i over $X^{\cup_{n=0}^{\infty} I_n}$. By construction, (β_n^i) is a martingale with respect to ψ^i , and so converges almost surely $[\psi^i]$ to some β_∞^i (which is measurable with respect to $\cup_{n=0}^{\infty} I_n$).

Suppose f is measurable with respect to $\cup_{n=0}^{\infty} I_n$. Then (7) and (8) imply that $\beta_\infty^i = f$ almost surely: If $\cup_{n=0}^{\infty} I_n = N$, set $\rho^i = \psi^i$ and we have $\beta_n^i(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$. If $\cup_{n=0}^{\infty} I_n$ is a strict subset of N , then let ρ^i be any probability measure whose marginal on $X^{\cup_n I_n}$ agrees with ψ^i and we again have $\beta_n^i(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$.

Suppose f is not measurable with respect to $\cup_{n=0}^{\infty} I_n$. This implies that $\cup_{n=0}^{\infty} I_n$ is a strict subset of N . Requirements (7) and (8) imply that we can choose $\rho^i \in \Delta(\Omega)$ so that its marginal on $X^{\cup_n I_n}$ agrees with ψ^i and $\beta_\infty^i(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$. This then implies $\beta_n^i(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$. ■

We now contrast this result with a crowd of different models. We again consider a sequence of crowds that receive increasing amounts of information (I_n) and assume the agents have common information. We maintain our running assumption that agents observe information contained in their models.

Proposition 1 immediately implies the following.

Remark 6 (comparison of model-based updating and different priors)

Consider a sequence $n = 1, \dots$, of crowds, with agent i 's model given by M^i , and each agent i 's information set I_n^i in crowd n given by I_n . For each n and each agent i , $I_n \subseteq M^i$. Then every agent's limit belief equals the public oracular belief. ◆

Model-based updating thus places considerably more structure on agents' beliefs. Even when removing all other obstacles to disagreement, including making information common, agents with different priors face virtually unlimited possibilities for disagreement. In contrast to the case of different

State ($\omega_1, \omega_2, \omega_3$)	Prior ρ	$f^*(\omega)$	Interim beliefs					
			$\beta^1(\omega_{I^1})$	$\beta^2(\omega_{I^2})$	$\beta^3(\omega_{I^3})$	$\beta^4(\omega_{I^4})$	$\beta^5(\omega_{I^5})$	$\beta^6(\omega_{I^6})$
(0, 0, 0)	1/8	α	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(0, 0, 1)	1/8	β	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(0, 1, 0)	1/8	β	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(0, 1, 1)	1/8	α	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(1, 0, 0)	1/8	β	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(1, 0, 1)	1/8	α	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(1, 1, 0)	1/8	α	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$
(1, 1, 1)	1/8	β	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$	$(\alpha + \beta)/2$

$$M^1 = M^2 = M^3 = M^4 = M^5 = M^6 = \{1, 2, 3\}$$

$$I^1 = \{1\}, \quad I^2 = \{2\}, \quad I^3 = \{3\}, \quad I^4 = \{1, 1\}, \quad I^5 = \{1, 3\}, \quad I^6 = \{2, 3\}.$$

Figure 7: The beliefs for Appendix A.2. The function f^* is the expected value of f conditional on all agents' model dimensions.

priors, the only sources of disagreement among agents with different models arise out of the different ways agents interpret information they think irrelevant.

A.2 An Example with Pooling

The model and information are presented in Figure 7. See Example 3 for the explanation of f^* , the expected value of f conditional on all the agents' model dimensions, i.e., $f^*(\omega) := \mathbb{E}[f(\omega) \mid \omega_1, \omega_2, \omega_3]$.

The six agents have a common model $M = \{1, 2, 3\}$. Their beliefs would remain unchanged if we made them agent oracles. For every possible information set $I \subsetneq M$, there exists an agent who observes the information contained in I .²¹ Nonetheless, no information is revealed in this example, no matter how often agents exchange beliefs.

The distribution of realizations in Figure 7 is independent across dimensions, but the construction of such an example does not depend on such independence. Each of the interim beliefs in the first three columns of beliefs

²¹It is immediate that information will be conveyed if there is an agent who observes every dimension the agents think relevant, i.e., $I = M$.

in Figure 7 reflects of calculation of the form

$$\frac{\rho(\omega^k)f(\omega^k) + \rho(\omega^\ell)f(\omega^\ell) + \rho(\omega^m)f(\omega^m) + \rho(\omega^n)f(\omega^n)}{\rho(\omega^k) + \rho(\omega^\ell) + \rho(\omega^m) + \rho(\omega^n)} = \frac{\alpha + \beta}{2},$$

while each belief in the last three columns reflects a calculation of the form

$$\frac{\rho(\omega^k)f(\omega^k) + \rho(\omega^\ell)f(\omega^\ell)}{\rho(\omega^k) + \rho(\omega^\ell)} = \frac{\alpha + \beta}{2}.$$

We can rearrange these as

$$\begin{aligned} \rho(\omega^k)f(\omega^k) + \rho(\omega^\ell)f(\omega^\ell) + \rho(\omega^m)f(\omega^m) + \rho(\omega^n)f(\omega^n) \\ = [\rho(\omega^k) + \rho(\omega^\ell) + \rho(\omega^m) + \rho(\omega^n)] \frac{\alpha + \beta}{2} \end{aligned}$$

and

$$\rho(\omega^k)f(\omega^k) + \rho(\omega^\ell)f(\omega^\ell) = [\rho(\omega^k) + \rho(\omega^\ell)] \frac{\alpha + \beta}{2}.$$

Replacing (ρ, f) in Figure 7 with any alternative (ρ', f') satisfying

$$\rho(\omega^k) \left[f(\omega^k) - \frac{\alpha + \beta}{2} \right] = \rho'(\omega^k) \left[f'(\omega^k) - \frac{\alpha + \beta}{2} \right]$$

for all states ω^k gives an alternative formulation, potentially including correlation across dimensions, generating identical interim beliefs. Indeed, this gives a recipe for constructing many such examples.

A.3 An Example with Infinite Iterations

Let $N = \mathbb{N}$ and $\Omega = \{0, 1\}^\infty$. There are two agents, with $M^1 = \mathbb{N} \setminus \{1\}$ and $M^2 = \mathbb{N} \setminus \{2\}$. The data generating process ρ independently chooses each dimension to be 0 or 1 with probability 1/2. Agents 1 and 2 observe

$$I^1 = \{1, 3, 4, 6, 8, 10, \dots\} \text{ and } I^2 = \{2, 3, 5, 7, 9, 11, \dots\}.$$

We first define two events, G and H , which are constituents of the event F .

The event G occurs if and only if $(\omega_1, \omega_2) = (1, 0)$.

The event H occurs if at least one of the following statements holds:

$$\begin{aligned}
&\omega_3 = \omega_4 = \omega_5, \\
&(\omega_3 + \omega_5)_{\text{mod } 2} = \omega_6 = (\omega_8 + \omega_9)_{\text{mod } 2} = (\omega_{10} + \omega_{11})_{\text{mod } 2}, \\
&(\omega_3 + \omega_4)_{\text{mod } 2} = \omega_7 = (\omega_8 + \omega_9)_{\text{mod } 2} = (\omega_{10} + \omega_{11})_{\text{mod } 2}, \\
&(\omega_3 + \omega_7)_{\text{mod } 2} = \omega_8 = (\omega_{10} + \omega_{11})_{\text{mod } 2} = (\omega_{12} + \omega_{13})_{\text{mod } 2} \\
&\quad = (\omega_{14} + \omega_{15})_{\text{mod } 2}, \\
&(\omega_3 + \omega_6)_{\text{mod } 2} = \omega_9 = (\omega_{10} + \omega_{11})_{\text{mod } 2} = (\omega_{12} + \omega_{13})_{\text{mod } 2} \\
&\quad = (\omega_{14} + \omega_{15})_{\text{mod } 2}, \\
&(\omega_3 + \omega_9)_{\text{mod } 2} = \omega_{10} = (\omega_{12} + \omega_{13})_{\text{mod } 2} = (\omega_{14} + \omega_{15})_{\text{mod } 2} \\
&\quad = (\omega_{16} + \omega_{17})_{\text{mod } 2} = (\omega_{18} + \omega_{19})_{\text{mod } 2}, \\
&(\omega_3 + \omega_8)_{\text{mod } 2} = \omega_{11} = (\omega_{12} + \omega_{13})_{\text{mod } 2} = (\omega_{14} + \omega_{15})_{\text{mod } 2} \\
&\quad = (\omega_{16} + \omega_{17})_{\text{mod } 2} = (\omega_{18} + \omega_{19})_{\text{mod } 2}, \\
&\quad \vdots
\end{aligned}$$

The probability of event H lies between $1/4$ (the probability that $\omega_3 = \omega_4 = \omega_5$) and $3/4$ (the sum of the probabilities of each of the statements on the list).

Now consider beliefs about the event $F := G \cup H$.

Upon observing ω_{I_1} , agent 1's posterior belief about every statement in the definition of H other than the first is unchanged. However, 1 updates positively the posterior probability that H holds if $\omega_3 = \omega_4$, and updates negatively if this equality fails. Agent 1's first announcement of the probability of F thus reveals the realization of ω_4 to agent 2, but reveals no additional information. Similarly, agent 2's first announcement of the probability of F reveals the realization of ω_5 (but no additional information) to agent 1.

The first round of announcements may reveal that the event H occurs, but with positive probability this is not the case. In the latter case, the agents now update their posteriors about the second and third statements in the definition of H (and no others), depending on their realizations of ω_6 and ω_7 , and their next announcements of the probability of F reveal these values. This in turn allows them to update their beliefs about the fourth and fifth statements (and no others), and so on.

With positive probability, the event H has indeed occurred, in which case the belief updating about the event H terminates after a finite number of iterations, with probability 1 attached to H . However, with positive

probability H has not occurred, in which case beliefs about H are revised forever.

We then have the following possibilities concerning the event $F = G \cup H$ (in all cases, after the initial exchange, subsequent exchanges of beliefs have no effect on the probability they attach to event G , and cause them to update the probability that H as described above):

- $(\omega_1, \omega_2) = (0, 1)$. Both agents attach interim probability 0 to event G , and each agent attaches the same probability to event F as they do to event H . Beliefs about H converge to a common limit.
- $(\omega_1, \omega_2) = (1, 0)$. Both agents attach interim probability 1/2 to the event that G has occurred. Beliefs about F converge to either 1/2 (if H has not occurred) or 1 (if H has occurred). In either case, beliefs converge to a common limit.
- $(\omega_1, \omega_2) = (0, 0)$. Agent 1 attaches interim probability 0 and agent 2 attaches interim probability 1/2 to event G . If H has occurred, the beliefs of both agents will eventually place probability 1 on event F . However, if H has not occurred, it will take an infinite number of exchanges for beliefs about event F to converge to 0 for agent 1 and 1/2 for agent 2.
- $(\omega_1, \omega_2) = (1, 1)$. This duplicates the previous case, with the roles of agents 1 and 2 reversed.

Remark 7 A simplification of this example shows that [Geanakoplos and Polemarchakis's \(1982\)](#) protocol on an infinite space with a common prior and model also need not terminate in a finite number of steps. Take the event to be H , the common model to be $N \setminus \{1, 2\}$, and let agent 1 observe $\{3, 4, 6, 8, \dots\}$, and agent 2 observe $\{3, 5, 7, 9, \dots\}$. ◆

A.4 Proposition 1.5 and Oracles

Proposition 1.5 does not extend to agent oracles: Figure 8 presents an example in which

$$\mathbb{E}[f \mid \mathcal{G}^{I^i}, \mathcal{B}_\infty] \neq \mathbb{E}[f \mid \mathcal{B}_\infty].$$

There is no nontrivial updating since both agents believe they know all they need to know. Moreover, the public oracle's beliefs coincide with that of agent 2, which differ from the last column.

State $(\omega_1, \omega_2, \omega_3)$	Prior ρ	$f(\omega)$	Interim beliefs		$\mathbb{E}f \mathcal{G}^{I^1}, \mathcal{B}_\infty]$
			$\beta^1(\omega_{I^1})$	$\beta^2(\omega_{I^2})$	
(0, 0, 0)	$a + b$	0	$b(x + y)/(2a + 2b)$	0	0
(0, 0, 1)	$b/2$	$2x$	$b(x + y)/(2a + 2b)$	$2(x + y)/3$	$(x + y)$
(0, 1, 0)	$b/2$	$2y$	$b(x + y)/(2a + 2b)$	$2(x + y)/3$	$(x + y)$
(1, 0, 0)	$a + b$	0	$b(x + y)/(2a + 2b)$	0	0
(0, 1, 1)	a	0	$b(x + y)/(2a + 2b)$	0	0
(1, 0, 1)	b	y	$b(x + y)/(2a + 2b)$	$2(x + y)/3$	$(x + y)/2$
(1, 1, 0)	b	x	$b(x + y)/(2a + 2b)$	$2(x + y)/3$	$(x + y)/2$
(1, 1, 1)	$a - b$	0	$b(x + y)/(2a + 2b)$	0	0

$$X = \{0, 1\}, M^1 = I^1 = \{1\} \text{ and } M^2 = I^2 = \{2, 3\}.$$

Figure 8: Beliefs for Appendix A.4.

A.5 Common Knowledge when Models are Infinite

Since we now must deal with conditioning on potentially zero probability events, we follow [Brandenburger and Dekel \(1987\)](#) in defining knowledge as probability one belief, and requiring conditional probabilities to be regular and proper.²² Recall that the state space has prior ρ , and suppose that each player's information is described by a σ -algebra \mathcal{G}^i . For each agent i , there is a mapping $\rho^i : \mathcal{F} \times \Omega \rightarrow [0, 1]$, where $\rho^i(\cdot | \omega)$ is a probability measure on \mathcal{F} for all $\omega \in \Omega$; for each $G \in \mathcal{F}$, $\rho^i(G | \cdot)$ is a version of $\rho(G | \mathcal{G}^i)$; and $\rho^i(G | \omega) = \chi_G(\omega)$ for all $G \in \mathcal{G}^i$ (in other words, ρ^i is a regular and proper conditional probability). These are the beliefs used to define what it means for agent i to know (assign probability 1 to) an event. By [Brandenburger and Dekel \(1987, Lemma 2.1\)](#), an event G is common knowledge at some ω (in the sense that every agent assigns probability one to the event, every agent assigns probability one to every agent assigning probability one to the event, and so on) if there is a set G' in the meet $\wedge \mathcal{G}^i$ such that $\omega \in G'$ and $\rho^i(\{\omega' \in G' : \omega' \notin G\} | \omega'') = 0$ for all $\omega'' \in \Omega$.²³ The last requirement is simply that G' is a subset of G , up to a zero measure set, under each agent's

²²[Bogachev \(2007, Corollary 10.4.10\)](#) ensures the existence of such conditional probabilities.

²³This is a sufficient condition for common knowledge. The characterization requires a little more ([Brandenburger and Dekel, 1987, Lemma 2.3 and Proposition 2.1](#)), which we do not need.

State (ω_1, ω_2)	Prior ρ	$f^*(\omega)$	Theories		Interim beliefs		First-round update
			$f^1(\omega_{M^1})$	$f^2(\omega_{M^2})$	$\beta^1(\omega_{I^1})$	$\beta^2(\omega_{I^2})$	$\beta^2(\omega_{I^2}, b_0^1)$
(0, 0)	a	x	$\frac{ax+by}{a+b}$	$\frac{ax+cz}{a+c}$	$\frac{ax+by}{a+b}$	$ax+by+cz+dw$	$\frac{a}{a+b} \frac{ax+cz}{a+c} + \frac{b}{a+b} \frac{by+dw}{b+d}$
(0, 1)	b	y	$\frac{ax+by}{a+b}$	$\frac{by+dw}{b+d}$	$\frac{ax+by}{a+b}$	$ax+by+cz+dw$	$\frac{a}{a+b} \frac{ax+cz}{a+c} + \frac{b}{a+b} \frac{by+dw}{b+d}$
(1, 0)	c	z	$\frac{cz+dw}{c+d}$	$\frac{ax+cz}{a+c}$	$\frac{cz+dw}{c+d}$	$ax+by+cz+dw$	$\frac{c}{c+d} \frac{ax+cz}{a+c} + \frac{d}{c+d} \frac{by+dw}{b+d}$
(1, 1)	d	w	$\frac{cz+dw}{c+d}$	$\frac{by+dw}{b+d}$	$\frac{cz+dw}{c+d}$	$ax+by+cz+dw$	$\frac{c}{c+d} \frac{ax+cz}{a+c} + \frac{d}{c+d} \frac{by+dw}{b+d}$

$$X = \{0, 1\}, \quad M^1 = \{1\}, \quad M^2 = \{2\}, \\ I^1 = \{1\}, \quad I^2 = \emptyset.$$

Figure 9: Agreement need not imply redundancy in the presence of correlation. The details for Appendix A.6.

beliefs ρ^j .

We will say that limit beliefs are common knowledge if they are common knowledge given the information provided to the agents by the entire infinite sequence of belief announcements.

Proposition 10 *Limit beliefs are common knowledge.*

Proof. Recall that \mathcal{B}_n denotes the round n σ -algebra generated by the announcements from the first n rounds. For each $\omega \in \Omega$, all events G satisfying $\omega \in G \in \mathcal{B}_n$ are common knowledge at ω . Recall also that $(\mathcal{B}_n)_n$ is a filtration with limit \mathcal{B}_∞ , so that the beliefs $\beta_{n+1}^i = \mathbb{E}[f \mid \mathcal{I}^i, \mathcal{B}_n]$ are a martingale and converge almost surely to $\mathbb{E}[f \mid \mathcal{I}^i, \mathcal{B}_\infty] =: \beta_\infty^i$. Moreover, $\beta_\infty^i = \int f d\rho_\infty^i$.

Fix b^i in the range of β_∞^i and let $A := (\beta_\infty^i)^{-1}(b^i)$. We now prove that for all $\omega \in A$ there is a subset A' in the meet $\bigwedge \sigma(\mathcal{I}^i, \mathcal{B}_\infty)$ containing ω . Fix $\omega \in A$, and define $A_n := \bigcap_j (\beta_n^j)^{-1}(b^j)$ where $b^j = \beta_n^j(\omega)$. Since $A_n \in \bigwedge \sigma(\mathcal{I}^i, \mathcal{B}_\infty)$, we have $\bigcap_n A_n \in \bigwedge \sigma(\mathcal{I}^i, \mathcal{B}_\infty)$. Suppose $\bigcap_n A_n \not\subseteq A$, so that there exists $\tilde{\omega} \in \bigcap_n A_n \setminus A$. But then $\beta_n^i(\omega) = \beta_n^i(\tilde{\omega})$ for all n , and since the beliefs converge,²⁴ $\beta_\infty^i(\tilde{\omega}) = b^i$, a contradiction. ■

A.6 An Example Illustrating Redundancy and Correlation

We start with the general specification given in Figure 9. Agent 1 observes every dimension of 1's model, and so never does any updating past the

²⁴The sentence previously footnoted implies we can assume beliefs converge on $\bigcap_n A_n$.

interim belief. Agent 2, who observes nothing, ceases updating after the first round. If the values of ω_1 and ω_2 are independently drawn, then it follows immediately from Proposition 4 that beliefs can necessarily agree only if ω_1 is redundant for agent 1.

We now seek values of the parameters for which ω_1 is not redundant for player 1, i.e.,

$$\frac{ax + by}{a + b} \neq \frac{cz + dw}{c + d} \quad (9)$$

and for which there is necessary agreement, i.e. (after simplification),

$$(a + c)by = ac(z - x) + b\frac{a + c}{b + d}(by + dw) \quad (10)$$

and

$$(b + d)cz = bd(y - w) + c\frac{b + d}{a + c}(ax + cz). \quad (11)$$

Setting $b = c = 0$ gives the case where the two dimensions are perfectly correlated (ω_2 is simply a relabeling of ω_1), and we trivially have necessary agreement without redundancy.

It is straightforward that there are many parameters with the desired characteristics. If we set $z = x$ and $y = w$, then *any* specification of a, b, c, d satisfies these equations, including values that also satisfy (9). In this case, ω_1 plays no role in the determination of F , and agent 1's observation of ω_1 is informative only to the extent that it is correlated with ω_2 . In addition, agent 2 receives no information of her own, and so must similarly rely on gleaning information from the correlation of ω_1 with ω_2 , leading the two agents to agree. In the case of independence, or $a = b = c = d$, agent 1 learns nothing about the state, and the two agents necessarily agree on the uninformative posterior of $1/2$.

When at least one of $z = x$ and $y = w$ fails, then ω_1 plays a role in determining the event F . There then there exist particular values of a, b, c, d satisfying the equations (10)–(11) for necessary agreement.

A.7 Proof of Lemma 2

Fix a value $\delta > 0$. Choose λ and ε such that

$$\begin{aligned} \lambda &> 1/\delta \\ \varepsilon + \varepsilon\lambda &< \delta. \end{aligned}$$

By Egorov's theorem, the convergence of f_n to f^\dagger is uniform on a set of large measure. In particular, there exists a value N_δ and a set Ω_δ of measure at least $1 - \varepsilon$ (which is at least $1 - \delta$) with the property that for all $n > N_\delta$, we have

$$|f_n(\omega) - f^\dagger(\omega)| < \varepsilon \quad \forall \omega \in \Omega_\delta.$$

We now argue that with probability at least $1 - \delta$, we have

$$\left| \mathbb{E}[f_n | \mathcal{H}] - \mathbb{E}[f^\dagger | \mathcal{H}] \right| < \delta.$$

Let

$$h(\omega) = \begin{cases} \varepsilon, & \omega \in \Omega_\delta, \\ 1, & \omega \notin \Omega_\delta. \end{cases}$$

Then,

$$|f_n(\omega) - f^\dagger(\omega)| \leq h^i(\omega),$$

and

$$\begin{aligned} \mathbb{E}[h | \mathcal{H}] &= \varepsilon \Pr(\Omega_\delta | \mathcal{H}) + \Pr(\Omega \setminus \Omega_\delta | \mathcal{H}) \\ &\leq \varepsilon + \Pr(\Omega \setminus \Omega_\delta | \mathcal{H}). \end{aligned}$$

Let $A := \{\Pr(\Omega \setminus \Omega_\delta | \mathcal{H}) > \lambda\varepsilon\}$. Then $A \in \mathcal{H}$ and so

$$\begin{aligned} \lambda\varepsilon \Pr A &< \int_A \mathbb{E}[\chi_{\Omega \setminus \Omega_\delta} | \mathcal{H}] d\rho \\ &= \int_A \chi_{\Omega \setminus \Omega_\delta} d\rho \\ &\leq \varepsilon, \end{aligned}$$

and so

$$\Pr[\Pr(\chi_{\Omega \setminus \Omega_\delta} | \mathcal{H}) > \lambda\varepsilon] \leq 1/\lambda,$$

and hence we have

$$\Pr[\varepsilon + \Pr(\Omega \setminus \Omega_\delta | \mathcal{H}) < (1 + \lambda)\varepsilon] > 1 - 1/\lambda.$$

Invoking our conditions on λ and ε yields

$$\Pr[\mathbb{E}[h | \mathcal{H}] < \delta] > 1 - \delta,$$

and since

$$\left| \mathbb{E}[f_n | \mathcal{H}] - \mathbb{E}[f^\dagger | \mathcal{H}] \right| \leq \mathbb{E}[|f_n - f^\dagger| | \mathcal{H}],$$

we have the desired inequality.

A.8 Proof of Lemma 3

Recall that \mathcal{G}^t is the σ -algebra generated by the t coordinate of $\Omega = \{0, 1\}^N$, and set $\mathcal{F}^t := \sigma(\mathcal{G}^1, \dots, \mathcal{G}^t)$. Since f is measurable with respect to $\sigma(\mathcal{G}^1, \mathcal{G}^2, \dots)$, we have

$$\mathbb{E}[f \mid \mathcal{F}^t] \rightarrow f \quad \text{a.s.}[\rho].$$

Egorov's theorem implies that for all $\varepsilon > 0$, there exists T_ε such that on an event Ω_ε , with $\rho(\Omega_\varepsilon) \geq 1 - \varepsilon/4$,

$$|\mathbb{E}[f \mid \mathcal{F}^t](\omega) - f(\omega)| < \varepsilon^2/4 \quad \forall t \geq T_\varepsilon^*, \forall \omega \in \Omega_\varepsilon. \quad (12)$$

Set $K_\varepsilon := \{1, \dots, T_\varepsilon\}$, so that $\mathcal{G}^{K_\varepsilon} = \mathcal{F}^{T_\varepsilon}$.

Claim 1 *On a full probability subset of $\Omega_\varepsilon \cap F$,*

$$\Pr\{\mathbb{E}[f \mid \mathcal{G}^{K_\varepsilon}, \mathcal{H}] \leq 1 - \varepsilon \mid \mathcal{G}^{K_\varepsilon}\} < \varepsilon/4 \quad (13)$$

and on a full probability subset of $\Omega_\varepsilon \setminus F$,

$$\Pr\{\mathbb{E}[f \mid \mathcal{G}^{K_\varepsilon}, \mathcal{H}] \geq \varepsilon \mid \mathcal{G}^{K_\varepsilon}\} < \varepsilon/4. \quad (14)$$

Proof. We prove (13); the proof of (14) follows similar lines. Define $g^\dagger(\omega) := \mathbb{E}[f \mid \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega)$, and $g(\omega) := \Pr\{g^\dagger \leq 1 - \varepsilon \mid \mathcal{G}^{K_\varepsilon}\}(\omega)$. Note that g is only measurable with respect to $\mathcal{G}^{K_\varepsilon}$ (so in particular, the inequality in (13) is measurable with respect to $\mathcal{G}^{K_\varepsilon}$), while g^\dagger is measurable with respect to the finer $\sigma(\mathcal{G}^{K_\varepsilon}, \mathcal{H})$.

Recalling that f is the indicator function of the event F , for $\omega \in \Omega^* \cap F$, (12) is

$$1 - \varepsilon^2/4 < \mathbb{E}[f \mid \mathcal{G}^{K_\varepsilon}](\omega),$$

and so (13) is implied by for ρ -almost all $\omega \in \Omega^* \cap F$,

$$\mathbb{E}[f \mid \mathcal{G}^{K_\varepsilon}](\omega) \leq 1 - \varepsilon g(\omega).$$

Since the left and right sides of the above inequality are measurable with respect to $\mathcal{G}^{K_\varepsilon}$, if the inequality does not hold, there is a positive ρ -probability event $B \in \mathcal{G}^{K_\varepsilon}$ such that,

$$\mathbb{E}[f \mid \mathcal{G}^{K_\varepsilon}](\omega) > 1 - \varepsilon g(\omega) \quad \forall \omega \in B. \quad (15)$$

Since $B \in \mathcal{G}^{K_\varepsilon}$, where χ_A is the indicator function of the event A , and the first and last (respectively, third) equalities hold because the integrating events are measurable with respect to $\mathcal{G}^{K_\varepsilon}$ (resp., $\sigma(\mathcal{G}^{K_\varepsilon}, \mathcal{H})$),

$$\begin{aligned}
\int_B \mathbb{E}[f | \mathcal{G}^{K_\varepsilon}] d\rho &= \int_B f d\rho \\
&= \int_{B \cap \{g^\dagger \leq 1-\varepsilon\}} f d\rho + \int_{B \cap \{g^\dagger > 1-\varepsilon\}} f d\rho \\
&= \int_{B \cap \{g^\dagger \leq 1-\varepsilon\}} g^\dagger d\rho + \int_{B \cap \{g^\dagger > 1-\varepsilon\}} g^\dagger d\rho \\
&\leq (1-\varepsilon) \int_B \chi_{\{g^\dagger \leq 1-\varepsilon\}} d\rho + \int_B 1 - \chi_{\{g^\dagger \leq 1-\varepsilon\}} d\rho \\
&= \int_B 1 - \varepsilon \chi_{\{g^\dagger \leq 1-\varepsilon\}} d\rho \\
&= \int_B 1 - \varepsilon g d\rho,
\end{aligned}$$

contradicting (15). □

Defining

$$B' := \{\omega : |\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) - f(\omega)| \geq \varepsilon\}$$

and

$$F' := \{\omega : \Pr\{\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}] \leq 1 - \varepsilon | \mathcal{G}^{K_\varepsilon}\}(\omega) < \varepsilon/4\}$$

we have (since, up to a zero probability event, $\Omega_\varepsilon \cap F \subseteq F'$ and $F' \in \mathcal{G}^{K_\varepsilon}$)

$$\begin{aligned}
\Pr(B' \cap F) &= \Pr\{\{\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) \leq 1 - \varepsilon\} \cap F\} \\
&\leq \Pr\{\{\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) \leq 1 - \varepsilon\} \cap F'\} \\
&= \mathbb{E}[\mathbb{E}[\chi_{\{\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) \leq 1-\varepsilon\} \cap F'} | \mathcal{G}^{K_\varepsilon}]] \\
&= \mathbb{E}[\mathbb{E}[\chi_{\{\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) \leq 1-\varepsilon\}} | \mathcal{G}^{K_\varepsilon}] \chi_{F'}] \\
&\leq \int_{\Omega^*} \Pr\{\mathbb{E}[f | \mathcal{G}^{K_\varepsilon}, \mathcal{H}](\omega) \leq 1 - \varepsilon\} | \mathcal{G}^{K_\varepsilon} \chi_{F'} d\rho + \rho(\Omega \setminus \Omega_\varepsilon) \\
&\leq \varepsilon/4 + \varepsilon/4.
\end{aligned}$$

Applying a similar argument to $B' \setminus F$, we obtain

$$\Pr(B' \setminus F) \leq \varepsilon/2,$$

so that $\rho(B') \leq \varepsilon$.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, AND M. YILDIZ (2016): “Fragility of Asymptotic Agreement under Bayesian Learning,” *Theoretical Economics*, 11. [10](#)
- AL-NAJJAR, N. I. (2009): “Decision Makers as Statisticians: Diversity, Ambiguity, and Learning,” *Econometrica*, 77(5), 1371–1401. [9](#)
- ARIELI, I., Y. BABICHENKO, AND R. SMORODINSKY (2018): “When is the Crowd Wise?” unpublished, Faculty of Industrial Engineering and Management, Technion. [37](#)
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *Annals of Statistics*, 4, 1236–1239. [24](#)
- BERG, J., R. FORSYTHE, F. NELSON, AND T. RIETZ (2008): “Results from a Dozen Years of Election Futures Markets Research,” in *Handbook of Experimental Economics Results*, ed. by C. R. Plott, and V. L. Smith, vol. 1, pp. 742–751. Elsevier. [1](#)
- BILLINGSLEY, P. (2012): *Probability and Measure*. John Wiley and Sons, New York, anniversary edn. [42](#)
- BOGACHEV, V. I. (2007): *Measure Theory*, vol. 2. Springer-Verlag. [47](#)
- BOHREN, J. A. (2016): “Informational Herding with Model Misspecification,” *Journal of Economic Theory*, 163, 222–247. [38](#)
- BOHREN, J. A., AND D. N. HAUSER (2018): “Social Learning with Model Misspecification: A Framework and a Robustness Result,” PIER Working Paper 18-017, University of Pennsylvania. [38](#)
- BRANDENBURGER, A., AND E. DEKEL (1987): “Common Knowledge with Probability 1,” *Journal of Mathematical Economics*, 16(3), 237–245. [24](#), [47](#)
- BRANDENBURGER, A., E. DEKEL, AND J. GEANAKOPLOS (1992): “Correlated Equilibrium with General Information Structures,” *Games and Economic Behavior*, 4, 182–201. [2](#)
- EYSTER, E., AND M. PICCIONE (2013): “An Approach to Asset Pricing Under Incomplete and Diverse Perceptions,” *Econometrica*, 81(4), 1483–1506. [10](#)

- FAMA, E. (1970): “Efficient Capital Markets: A Review of Theory and Empirical Work,” *Journal of Finance*, 25(2), 383–417. 1
- (1991): “Efficient Capital Markets: II,” *Journal of Finance*, 46(5), 1575–1617. 1
- FRIEDMAN, D., AND J. RUST (1993): *The Double Auction Market: Institutions, Theories, and Evidence*. Addison Wesley, New York. 38
- GALTON, F. (1907): “Vox Populi (The Wisdom of Crowds),” *Nature*, 1949(75), 450–451. 1, 39
- GEANAKOPOLOS, J. (1989): “Game Theory Without Partitions, and Applications to Speculation and Consensus,” Cowles foundation discussion paper no. 914, Yale University. 2
- GEANAKOPOLOS, J., AND H. POLEMARCHAKIS (1982): “We Can’t Disagree Forever,” *Journal of Economic Theory*, 27, 192–200. 12, 18, 46
- GIACOMINI, R., V. SKRETA, AND J. TURÉN (2007): “Models, Inattention and Bayesian Updates,” University College London and University of Texas at Austin. 10
- GILBOA, I., AND L. SAMUELSON (2012): “Subjectivity in Inductive Inference,” *Theoretical Economics*, 7(2), 183–216. 9
- GREEN, E. J. (2012): “Events Concerning Knowledge,” Working paper, Penn State University. 25
- HONG, H., J. C. STEIN, AND J. YU (2007): “Simple Forecasts and Paradigm Shifts,” *Journal of Finance*, 52(3), 1207–1242. 10
- KOESSLER, F., C. NOUSSAIR, AND A. ZIEGELMEYER (2008): “Parimutuel Betting under Asymmetric Information,” *Journal of Mathematical Economics*, 44(7), 733–744. 1
- MILGROM, P. R. (1979): “A Convergence Theorem for Competitive Bidding with Differential Information,” *Econometrica*, 47(3), 679–688. 1
- MILGROM, P. R., AND N. STOKEY (1982): “Information, Trade, and Common Knowledge,” *Journal of Economic Theory*, 26, 17–27. 1
- MORRIS, S. (1994): “Trade with Heterogeneous Prior Beliefs and Asymmetric Information,” *Econometrica*, 62, 1327–1347. 2

- OTTAVIANI, M., AND P. N. SØRENSEN (2010): “Noise, information, and the favorite-longshot bias in parimutuel predictions,” *American Economic Journal: Microeconomics*, 2(1), 58–85. 1
- PEARL, J. (2009): *Causality*. Cambridge University Press. 6
- ROUX, N., AND J. SOBEL (2015): “Group Polarization in a Model of Information Aggregation,” *American Economic Journal: Microeconomics*, 7(4), 202–232. 28
- SAVAGE, L. J. (1972): *The Foundations of Statistics*. Dover Publications, New York, originally 1954. 3
- SETHI, R., AND M. YILDIZ (2012): “Public Disagreement,” *American Economic Journal: Microeconomics*, 4(3), 57–95. 12
- SPIEGLER, R. (2016): “Bayesian Networks and Boundedly Rational Expectations,” *Quarterly Journal of Economics*, 131, 1243–1290. 6, 10, 38
- (2018): “News and Archival Information in Games,” unpublished, Tel Aviv and UCL. 38
- STROOCK, D. W. (2011): *Probability Theory: An Analytic View*. Cambridge university press, second edn. 5
- SUROWIECKI, J. (2004): *The Wisdom of Crowds*. Anchor Books, New York. 1, 29
- WILSON, C. (1977): “A Model of Insurance Markets with Incomplete Information,” *Journal of Economic Theory*, 16, 167–207. 1
- WOLFERS, J., AND E. ZITZEWITZ (2004): “Prediction Markets,” *Journal of Economic Perspectives*, 18(2), 107–126. 1, 38