

Learning under Diverse World Views: Model-Based Inference

George J. Mailath* Larry Samuelson†

December 12, 2019

Online Appendices

A What’s Wrong with Different Priors?

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.¹ In contrast, model-based reasoning ensures that agents’ beliefs about the events they deem relevant are 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 (under natural conditions) 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 group of agents with different priors as it does on a group of model-based reasoners. To investigate this, we examine a sequence in which agents 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

*Department of Economics, University of Pennsylvania, and Research School of Economics, Australian National University; gmailto@econ.upenn.edu

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

¹Di Tillio et al. (2019) come to a similar conclusion.

consider a sequence in which every agent's information set I_n^i in the n^{th} term is 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 (i.e., is an oracle), 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, let $(\beta_\infty^{i,n})_{i=1,n=0}^{K,\infty}$ be the sequence of induced limiting beliefs, for each agent, 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 $(\beta_\infty^{i,n})_{i=1,n=1}^{K,\infty}$, even though the agents are oracles.

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 $(\beta_\infty^{i,n})_{i=1,n=0}^{K,\infty}$ 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

$$\beta_\infty^{i,n+1}(\omega_{I_{n+1}}) < \beta_\infty^{i,n}(\omega_{I_n}), \quad (\text{A.1})$$

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

$$\beta_\infty^{i,n+1}(\omega'_{I_{n+1}}) > \beta_\infty^{i,n}(\omega_{I_n}). \quad (\text{A.2})$$

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

$$\lim_n \beta_\infty^{i,n}(\omega_{I_n}) > 0 \implies \exists \omega \in [\omega_{\cup_n I_n}] \text{ s.t. } f(\omega) = 1 \quad (\text{A.3})$$

$$\text{and } \lim_n \beta_\infty^{i,n}(\omega_{I_n}) < 1 \implies \exists \omega \in [\omega_{\cup_n I_n}] \text{ s.t. } f(\omega) = 0. \quad (\text{A.4})$$

Requirements (A.3) and (A.4) 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 $\beta_\infty^{i,n}(\omega_{I_n}) = \mathbb{E}_{\rho^i}[f \mid \omega_{I_n}]$.

Proposition A.1 *Consider a sequence of groups of agent oracles indexed by $n = 0, \dots$, with each agent's information set in group n given by I_n ,*

where the sequence $(I_n)_{n=0}^\infty$ is increasing. Suppose the sequences $(\beta_\infty^{i,n})_{i=1,n=0}^{K,\infty}$ satisfy the martingale property and (A.3) and (A.4). Then there exists a vector of prior beliefs (ρ^1, \dots, ρ^K) generating the limiting posterior beliefs $(\beta_\infty^{i,n})_{i=1,n=0}^{K,\infty}$, i.e., $\beta_\infty^{i,n}(\omega_{I_n}) = \mathbb{E}_{\rho^i}[f \mid \omega_{I_n}]$.

Before proving this result, we make three observations. First, if $\cup_{n=0}^\infty I_n = \Omega$, then since beliefs are a martingale, $\beta_\infty^{i,n} \rightarrow f$ ρ^i -almost surely. For states with 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 $(\beta_\infty^{i,n})$ satisfying the martingale property given in (A.1)–(A.2), there is a prior rationalizing $(\beta_\infty^{i,n})$. 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 $\cup_{n=0}^\infty I_n \subsetneq \Omega$, then 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 $\beta_\infty^{i,n}$. 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 $\beta_\infty^{i,0}$ 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 $\beta_\infty^{i,0} \in (0, 1)$. By assumption, the measure $\beta_\infty^{i,1}$ attaches conditional probabilities to a collection of cylinder sets of the form $[\omega_{I_1}]$, with some of these values larger than $\beta_\infty^{i,0}$ and some smaller. Assign probabilities $\rho^i([\omega_{I_1}])$ to these sets so that the average of the conditional probabilities is $\beta_\infty^{i,0}$. 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, $(\beta_\infty^{i,n})$ is a martingale with respect to ψ^i , and so converges ψ^i -almost surely to some $\beta_\infty^{i,\infty}$ (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 (A.3) and (A.4) imply that $\beta_\infty^{i,\infty} = f$ almost surely: If $\cup_{n=0}^\infty I_n = N$, set $\rho^i = \psi^i$ and we have

$\beta_\infty^{i,n}(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$. 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 $\beta_\infty^{i,n}(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$.

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 (A.3) and (A.4) imply that we can choose $\rho^i \in \Delta(\Omega)$ so that its marginal on $X^{\cup I_n}$ agrees with ψ^i and $\beta_\infty^{i,\infty}(\omega_{\cup I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{\cup I_n})}[f(\omega)]$. This then implies $\beta_\infty^{i,n}(\omega_{I_n}) = \mathbb{E}_{\rho^i(\cdot|\omega_{I_n})}[f(\omega)]$. ■

We now contrast this result with a group of model-based reasoners. We again consider a sequence that receives 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.

Corollary A.1 *Consider a sequence $n = 1, \dots$, of groups of model-based reasoners, with agent i 's model given by M^i , and each agent i 's information set I_n^i in group 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 priors, the only sources of disagreement among agents with different models arise out of the different ways agents interpret information they think irrelevant.

B Subsequent Updating in Example 2

We complete the discussion of updating in Example 2. The second round 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.) 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.

C 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 variable 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:

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

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 1 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 $\{3, 5, 7, 9, \dots\}$. ◆

D 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.

We first discuss the case where 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 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 D.1 *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 1.4, 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. ■

The common knowledge of limit beliefs implies an agreement theorem.

Proposition D.2 *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). ■

When the models are infinite, as in Example C, the belief revision process may continue without end. At no stage during the belief-revision process in Appendix C are the beliefs common knowledge. Despite this difficulty,

there is an appropriate notion of common knowledge when the 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.² 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 $\bigwedge \mathcal{G}^i$ such that $\omega \in G'$ and $\rho^i(\{\omega' \in G' : \omega' \notin G\} | \omega'') = 0$ for all $\omega'' \in \Omega$.³ The last requirement is simply that G' is a subset of G , up to a zero measure set, under each agent's 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 D.3 *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 $\beta_{n+1}^i = \mathbb{E}[f | \mathcal{I}^i, \mathcal{B}_n]$ are a martingale and converge almost surely to $\mathbb{E}[f | \mathcal{I}^i, \mathcal{B}_\infty] =: \beta_\infty^i$. 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

²Bogachev (2007, Corollary 10.4.10) ensures the existence of such conditional probabilities.

³This 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.

there exists $\tilde{\omega} \in \cap_n A_n \setminus A$. But then $\beta_n^i(\omega) = \beta_n^i(\tilde{\omega})$ for all n , and since the beliefs converge,⁴ $\beta_\infty^i(\tilde{\omega}) = b^i$, a contradiction. ■

Green (2012) presents an agreeing-to-disagree result for infinite models that would allow us to extend Proposition D.2 to this case.

E Implications of Necessary Agreement

Suppose X and N are finite, $M^2 \subsetneq M^1$, $I^1 = M^1$, and $I^2 = \emptyset$, that is, player 1 has full information, and player 2 observes nothing, but thinks only a subset of the variables in 1's model are relevant. We can, without loss of generality, write $X^{M^1} = X_1 \times X_2$, $X^{M^2} = X_2$, and suppose all states $(\omega_1, \omega_2) \in X_1 \times X_2$ have positive probability. Since player 2 has no information, there is no updating after 2 has updated from the announcement of 1's full information beliefs. Necessary agreement in this context means that for all $(\omega_1, \omega_2) \in X_1 \times X_2$, if $f^1(\omega_1, \omega_2) = \alpha$, then

$$\begin{aligned} \alpha &= \mathbb{E}[f^2(\omega_2) \mid b^1 = \alpha] \\ &= \mathbb{E}[\mathbb{E}[f^1(\omega_1, \omega_2) \mid \omega_2] \mid b^1 = \alpha]. \end{aligned}$$

We now argue that necessary agreement implies that the first coordinate of 1's model is redundant for 1, that is, f^1 is independent of ω_1 .

Let $\bar{\alpha} := \max f^1(\omega_1, \omega_2)$, and let $(\bar{\omega}_1, \bar{\omega}_2)$ be values that achieve $\bar{\alpha}$, i.e., $f^1(\bar{\omega}_1, \bar{\omega}_2) = \bar{\alpha}$. Conditional on the announcement $\bar{\alpha}$, necessary agreement implies

$$\bar{\alpha} = \mathbb{E}[f^2(\omega_2) \mid b^1 = \bar{\alpha}],$$

which implies $\mathbb{E}[f^1(\omega_1, \omega_2) \mid \omega_2] = \bar{\alpha}$ for all ω_2 in the support of the conditional beliefs $\rho_2(\cdot \mid b^1 = \bar{\alpha}) \in \Delta(X_2)$. But this implies that for all $\omega_1 \in X_1$ and for all ω_2 in the support of the conditional beliefs $\rho_2(\cdot \mid b^1 = \bar{\alpha}) \in \Delta(X_2)$, $f^1(\omega_1, \omega_2) = \bar{\alpha}$; in particular, for such ω_2 , f^1 is independent of ω_1 .

Since $X_1 \times X_2$ is finite, we can now argue inductively. Suppose $\alpha' := \max\{f^1(\omega_1, \omega_2) < \bar{\alpha}\}$ and let (ω'_1, ω'_2) be values that achieve α' , i.e., $f^1(\omega'_1, \omega'_2) = \alpha'$. After agent 1's announcement of α' , agent 2 assigns positive probability to ω'_2 . Moreover, from the previous paragraph, for all ω_1 , $f^1(\omega_1, \omega'_2) \leq \alpha'$ (if $f^1(\omega_1, \omega'_2) = \bar{\alpha}$ for some ω_1 , then ω'_2 is in the support of $\rho_2(\cdot \mid b^1 = \bar{\alpha})$ and so $f^1(\omega'_1, \omega'_2) = \bar{\alpha} \neq \alpha'$, a contradiction). But then, for all $\omega_1 \in X_1$ and for all ω_2 in the support of the conditional beliefs $\rho_2(\cdot \mid b^1 = \alpha')$, $f^1(\omega_1, \omega_2) = \alpha'$; in particular, for such ω_2 , f^1 is independent of ω_1 . Proceeding in this way

⁴The sentence previously footnoted implies we can assume beliefs converge on $\cap_n A_n$.

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 F.1: Agreement need not imply redundancy in the presence of correlation.

for progressively lower values of beliefs of agent 1, we conclude that f^1 is independent of ω_1 for all ω_2 .

F An Example Illustrating Redundancy and Correlation

We start with the general specification given in Figure F.1. Agent 1 observes every variable in 1's model, and so never does any updating past the 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 2 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.,

$$\frac{ax+by}{a+b} \neq \frac{cz+dw}{c+d} \tag{F.1}$$

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

$$(a+c)by = ac(z-x) + b \frac{a+c}{b+d} (by+dw) \tag{F.2}$$

and

$$(b+d)cz = bd(y-w) + c \frac{b+d}{a+c} (ax+cz). \tag{F.3}$$

Setting $b = c = 0$ gives the case where the two variables 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 (F.1). 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 exist particular values of a, b, c, d satisfying the equations (F.2)–(F.3) for necessary agreement.

References

- Aumann, Robert J.** 1976. "Agreeing to Disagree." *Annals of Statistics*, 4(6): 1236–1239.
- Billingsley, Patrick.** 2012. *Probability and Measure*. . Anniversary ed., New York:John Wiley and Sons.
- Bogachev, Vladimir I.** 2007. *Measure Theory, Volume 2*. Berlin:Springer.
- Brandenburger, Adam, and Eddie Dekel.** 1987. "Common Knowledge with Probability 1." *Journal of Mathematical Economics*, 16(3): 237–245.
- Di Tillio, Alfredo, Ehud Lehrer, Herakles Polemarchakis, and Dov Samet.** 2019. "Bayesian Monologues and Dialogues." Bocconi University, University of Tel Aviv, and University of Warwick.
- Geanakoplos, John, and Heraklis Polemarchakis.** 1982. "We Can't Disagree Forever." *Journal of Economic Theory*, 27(2): 192–200.
- Green, Edward J.** 2012. "Events Concerning Knowledge." Penn State University Working Paper.