

Agreeing to Disagree and Dilation

Jiji Zhang

Lingnan University (Hong Kong)

JIJIZHANG@LN.EDU.HK

Hailin Liu

Sun Yat-sen University, Guangzhou (China)

LIUHLIN3@MAIL.SYSU.EDU.CN

Teddy Seidenfeld

Carnegie Mellon University, Pittsburgh (USA)

TEDDY@STAT.CMU.EDU

Abstract

We consider Geanakoplos and Polemarchakis's generalization of Aumann's famous result on "agreeing to disagree", in the context of imprecise probability. The main purpose is to reveal a connection between the possibility of agreeing to disagree and the interesting and anomalous phenomenon known as dilation. We show that for two agents who share the same set of priors and update by conditioning on every prior, it is impossible to agree to disagree on the lower or upper probability of a hypothesis unless a certain dilation occurs. With some common topological assumptions, the result entails that it is impossible to agree not to have the same set of posterior probabilities unless dilation is present. This result may be used to generate sufficient conditions for guaranteed full agreement in the generalized Aumann-setting for some important models of imprecise priors, and we illustrate the potential with an agreement result involving the density ratio classes. We also provide a formulation of our results in terms of "dilation-averse" agents who ignore information about the value of a dilating partition but otherwise update by full Bayesian conditioning.

Keywords: agreeing to disagree; common knowledge; dilation; imprecise probability.

1. Introduction

In a simple but insightful paper, Aumann (1976) famously showed that two (Bayesian) agents who start with the same (precise) prior cannot agree to disagree on their posteriors of a hypothesis, in the sense that if the posteriors of the hypothesis (as well as the structures of their respective information partitions) are common knowledge, then the posteriors must be equal. This result has been generalized in at least two ways. First, Aumann's result applies only to those events whose posteriors happen to be common knowledge. Geanakoplos and Polemarchakis (1982) generalized the framework to a communication setting where the agents are invited to repeatedly make their credences public via announcements and update by conditioning on the announced credences, until no new information is conveyed. They showed that for any hypothesis/event, this communication procedure is guaranteed to lead to an agreement on the probability of the hypothesis, if the agents start with the same (precise) prior (and each agent's information partition is finite).

Second, Kajii and Ui (2005, 2009) and Carvajal and Correia-da-Silva (2010) generalized Aumann's result in the setting of multiple priors. In this line of work, "agreement" is taken to mean "partial agreement", in the sense that two sets of probabilities agree if they have a non-empty intersection. These authors established several sufficient conditions under which two agents who (partially) agree on their priors are guaranteed to (partially) agree on their posteriors of a hypothesis if these posteriors are common knowledge.

In this paper, we combine the two more general settings and establish a connection between the possibility of agreeing to disagree and the interesting and anomalous phenomenon known as dilation (Good, 1974; Seidenfeld, 1981; Walley, 1991; Seidenfeld and Wasserman, 1993; Herron et al., 1997). Dilation occurs when conditioning on each element of a partition, the lower and upper probabilities of a hypothesis become more divergent than the unconditional ones. In such a case, for agents who use full Bayesian conditioning as the updating rule, their credences on a hypothesis become less precise or determinate after learning the value of the dilating partition, no matter which value they learn! This counterintuitive phenomenon is often interpreted as a distinctive challenge to the orthodox Bayesian doctrine on the value of information and to the Bayesian merging of opinions, but as far as we know, it has never been discussed in connection to Aumann’s result. We shall show that it is the key obstacle to reaching agreements via communicating posteriors by Bayesian agents with imprecise priors.

We will establish the following. After introducing the setting and reviewing the special case of precise probability in Section 2, we show in Section 3 that dilation is the *only* obstacle for agents with the same (imprecise) prior to reaching agreements on lower and upper probabilities of a hypothesis by communicating their posteriors on the hypothesis. Without dilation, the two agents in our setting are guaranteed to end up agreeing on lower and upper probabilities of the hypothesis of interest. An immediate consequence of this result, as we note in Section 4, is that under common topological assumptions, dilation is the *only* obstacle to reaching a full agreement, *full* in the sense that the sets of probability values representing the agents’ credences on the hypothesis of interest are identical. This result opens the door to generating sufficient conditions for reaching full consensus in the generalized Aumann-setting by plugging in sufficient conditions for the absence of dilation in common and important models of imprecise probabilities. As an example, we include a corollary about density ratio classes, which are shown to be dilation-immune by Seidenfeld and Wasserman (1993). In Section 5, we provide another perspective on our results and reformulate the theorems in terms of “dilation-averse” agents, who update by full Bayesian conditioning unless the information is about the value of a dilating partition (in which case they ignore the information). For such agents, they are guaranteed to end up agreeing on lower and upper probabilities, and, under some common assumptions, end up fully agreeing.

2. A Procedure of Communicating Posteriors

In Geanakoplos and Polemarchakis (1982)’s setup, two agents share a common measurable space (Ω, \mathcal{A}) and have possibly different information partitions of Ω , \mathcal{P}^1 and \mathcal{P}^2 , which are assumed to be finite. Henceforth we use $i \in \{1, 2\}$ to index the two agents, and when i is used in a statement we always intend that the statement is true for both $i = 1$ and $i = 2$. For any $w \in \Omega$, let $\mathcal{P}^i(w)$ denote the member of \mathcal{P}^i that contains w ; intuitively, $\mathcal{P}^i(w)$ represents agent i ’s initial information at state w . Both the space and the partitions are assumed to be common knowledge, in the standard sense of the term used in game theory: some proposition is common knowledge just in case agent i knows it, agent j (where $j = 3 - i$) knows that agent i knows it, agent i knows that agent j knows that agent i knows it, ... and so on. Let $\mathcal{P} = \mathcal{P}^1 \wedge \mathcal{P}^2$ be the meet of the two partitions (i.e., the finest common coarsening of \mathcal{P}^1 and \mathcal{P}^2). As Aumann (1976, p. 1237) explained, at state w , $\mathcal{P}(w)$ — the member of \mathcal{P} that contains w — is the finest event in \mathcal{A} that is common knowledge: any event that is common knowledge is a superset of $\mathcal{P}(w)$. In Geanakoplos and Polemarchakis’s setting, common

knowledge may grow as the agents communicate their posteriors of a hypothesis. So we call $\mathcal{P}(w)$ the initial common knowledge and denote it by \mathcal{C}_0 .

Instead of a common precise prior, we assume that the two agents have a common, (possibly) imprecise prior, i.e., a common, non-empty set of priors, denoted by \mathbf{Q} . Let $\mathcal{P}^1 \vee \mathcal{P}^2$ denote the join (i.e., the coarsest common refinement) of \mathcal{P}^1 and \mathcal{P}^2 . We assume that every member of $\mathcal{P}^1 \vee \mathcal{P}^2$ receives a positive probability under every measure in \mathbf{Q} , so that all the relevant conditional probabilities are well defined as ratios of unconditional probabilities. Let $H \in \mathcal{A}$ be a hypothesis of interest. Henceforth by credences or posteriors we mean the agents' credences or posteriors of H . Let $\mathbf{Q}(H)$ denote the set of prior probabilities of H : $\mathbf{Q}(H) = \{p(H) \mid p \in \mathbf{Q}\}$. For any $E \in \mathcal{A}$ such that $p(E) > 0$ for every $p \in \mathbf{Q}$, let $\mathbf{Q}(H|E) = \{p(H|E) = p(H \cap E)/p(E) \mid p \in \mathbf{Q}\}$. Unless otherwise noted (in Section 5), we assume that the agents update their credences by full Bayesian conditioning, where each and every prior in \mathbf{Q} is updated by conditioning.

Suppose the true state is w . At step 0, agent i 's information is $\mathcal{P}^i(w) \cap \mathcal{C}_0 = \mathcal{P}^i(w)$. Thus agent i updates her credence of H to $\mathbf{Q}_0^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w))$. Let $\mathcal{P}_0^i = \{E \in \mathcal{P}^i \mid E \cap \mathcal{C}_0 \neq \emptyset\}$, which is the set of those members of \mathcal{P}^i that are not ruled out by the initial common knowledge.

At step 1, the agents announce $\mathbf{Q}_0^1(H)$ and $\mathbf{Q}_0^2(H)$, respectively.¹ Consider $\mathcal{N}_1^i = \{E \in \mathcal{P}_0^i \mid \mathbf{Q}(H|E) = \mathbf{Q}_0^i(H)\}$. Intuitively, \mathcal{N}_1^i is the set of those members of \mathcal{P}_0^i that are compatible with $\mathbf{Q}_0^i(H)$, and the effect of agent i 's announcement of $\mathbf{Q}_0^i(H)$ is that it becomes common knowledge that $\mathcal{P}^i(w) \in \mathcal{N}_1^i$, or that $w \in \bigcup \mathcal{N}_1^i$ (where $\bigcup \mathcal{N}_1^i$ denotes the union of all the sets in \mathcal{N}_1^i). Therefore, after the announcements at this step, $\mathcal{C}_1 = \bigcup \mathcal{N}_1^1 \cap \bigcup \mathcal{N}_1^2$ becomes common knowledge. Let $\mathcal{P}_1^i = \{E \in \mathcal{N}_1^i \mid E \cap \mathcal{C}_1 \neq \emptyset\}$, which is the set of those members of \mathcal{N}_1^i that are not ruled out by the common knowledge at this step. Clearly $\mathcal{P}_1^i \subseteq \mathcal{N}_1^i \subseteq \mathcal{P}_0^i$ and $\mathcal{C}_1 = \bigcup \mathcal{P}_1^1 \cap \bigcup \mathcal{P}_1^2$. Now, if $\mathcal{P}_1^i = \mathcal{P}_0^i$, or equivalently, if $\mathcal{C}_1 = \mathcal{C}_0$, neither agent learns new information and their credences will stay the same no matter how many more exchanges take place; so the procedure stops. Otherwise, agent i updates credence of H to $\mathbf{Q}_1^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_1)$, and enters the next step.

In general, at step $n + 1$, the agents announce $\mathbf{Q}_n^1(H)$ and $\mathbf{Q}_n^2(H)$, respectively. Let

$$\begin{aligned} \mathcal{N}_{n+1}^i &= \{E \in \mathcal{P}_n^i \mid \mathbf{Q}(H|E \cap \mathcal{C}_n) = \mathbf{Q}_n^i(H)\} \\ \mathcal{C}_{n+1} &= \bigcup \mathcal{N}_{n+1}^1 \cap \bigcup \mathcal{N}_{n+1}^2 \\ \mathcal{P}_{n+1}^i &= \{E \in \mathcal{N}_{n+1}^i \mid E \cap \mathcal{C}_{n+1} \neq \emptyset\}. \end{aligned}$$

Again, \mathcal{N}_{n+1}^i is the set of those members of \mathcal{P}_n^i that are compatible with $\mathbf{Q}_n^i(H)$.² Hence, after the announcements at this step, \mathcal{C}_{n+1} becomes common knowledge, and \mathcal{P}_{n+1}^i is the set of those members of \mathcal{N}_{n+1}^i that are not ruled out by \mathcal{C}_{n+1} . Clearly, $\mathcal{P}_{n+1}^i \subseteq \mathcal{N}_{n+1}^i \subseteq \mathcal{P}_n^i$ and $\mathcal{C}_{n+1} = \bigcup \mathcal{P}_{n+1}^1 \cap \bigcup \mathcal{P}_{n+1}^2$. If $\mathcal{P}_{n+1}^i = \mathcal{P}_n^i$, or equivalently, if $\mathcal{C}_{n+1} = \mathcal{C}_n$, neither agent learns new information and the procedure stops; otherwise, agent i updates credence of H to $\mathbf{Q}_{n+1}^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_{n+1})$, and enters the next step.

We will refer to this procedure as the (*Bayesian*) *procedure of communicating posteriors (of H)*. Obviously, since \mathcal{P}^1 and \mathcal{P}^2 are assumed to be finite, the procedure is guaranteed to stop at step $m + 1$ for some $m \geq 0$. Aumann's original setting — where $\mathbf{Q}_0^1(H)$ and $\mathbf{Q}_0^2(H)$ are assumed to be common knowledge at step 0 (i.e., it is assumed that $\mathcal{N}_1^i = \mathcal{P}_0^i$) — is a special case in which the

1. In Geanakoplos and Polemarchakis's design, at each step, agent 2 announces her prior *after* agent 1's announcement, already taking into account whatever information is conveyed in agent 1's announcement. This feature is immaterial, at least for the purpose of this paper.

2. Note that the definition of \mathcal{N}_{n+1}^i also applies to $n = 0$, as for every $E \in \mathcal{P}_0^i$, $E \cap \mathcal{C}_0 = E$.

procedure stops at step 1. In general, the procedure stops at step $m + 1$ if and only if both $\mathbf{Q}_m^1(H)$ and $\mathbf{Q}_m^2(H)$ are already common knowledge at step m (i.e., before they are announced).

We adapt an example from Geanakoplos and Polemarchakis (1982) to illustrate this procedure.

Example 1 Let $\Omega = \{w_1, w_2, w_3, w_4, w_5, w_6, w_7, w_8, w_9\}$ and \mathcal{A} be the power set of Ω . Let $\mathcal{P}^1 = \{\{w_1, w_2, w_3\}, \{w_4, w_5, w_6\}, \{w_7, w_8, w_9\}\}$ and $\mathcal{P}^2 = \{\{w_1, w_2, w_3, w_4\}, \{w_5, w_6, w_7, w_8\}, \{w_9\}\}$. Let $H = \{w_3, w_4\}$, and suppose the true state of the world is w_1 . For the common set of priors, suppose \mathbf{Q} is a density ratio class (Seidenfeld and Wasserman, 1993; see also Section 4):

$$\mathbf{Q} = \{(q_1, q_2, q_3, q_4, q_5, q_6, q_7, q_8, q_9) \mid \sum_{1 \leq j \leq 9} q_j = 1, \text{ and } \frac{1}{2} \leq \frac{q_k}{q_l} \leq 2, 1 \leq k, l \leq 9\}.$$

It is easy to calculate that the lower probability of H is: $\underline{\mathbf{Q}}(H) = \inf_{p \in \mathbf{Q}} p(H) = 1/8$ (obtained at $(1/8, 1/8, 1/16, 1/16, 1/8, 1/8, 1/8, 1/8, 1/8)$), and the upper probability of H is: $\overline{\mathbf{Q}}(H) = \sup_{p \in \mathbf{Q}} p(H) = 4/11$ (obtained at $(1/11, 1/11, 2/11, 2/11, 1/11, 1/11, 1/11, 1/11, 1/11)$). Since \mathbf{Q} is closed and connected, $\mathbf{Q}(H) = [1/8, 4/11]$.

Suppose the two agents in this example carry out the procedure of communicating posteriors. Here is a summary of the execution:

Step 0 $\mathcal{C}_0 = (\mathcal{P}^1 \wedge \mathcal{P}^2)(w_1) = \Omega$. $\mathcal{P}^1(w_1) = \{w_1, w_2, w_3\}$ and $\mathbf{Q}_0^1(H) = \mathbf{Q}(H|\mathcal{P}^1(w_1)) = [1/5, 1/2]$; $\mathcal{P}^2(w_1) = \{w_1, w_2, w_3, w_4\}$ and $\mathbf{Q}_0^2(H) = \mathbf{Q}(H|\mathcal{P}^2(w_1)) = [1/3, 2/3]$.

Step 1 Agent i announces $\mathbf{Q}_0^i(H)$. $\mathcal{N}_1^1 = \{\{w_1, w_2, w_3\}, \{w_4, w_5, w_6\}\}$ (for $\mathbf{Q}(H|\{w_7, w_8, w_9\}) = \{0\} \neq \mathbf{Q}_0^1(H)$.) $\mathcal{N}_1^2 = \{\{w_1, w_2, w_3, w_4\}\}$ (for $\mathbf{Q}(H|\{w_5, w_6, w_7, w_8\}) = \mathbf{Q}(H|\{w_9\}) = \{0\} \neq \mathbf{Q}_0^2(H)$.) Thus $\mathcal{C}_1 = \bigcup \mathcal{N}_1^1 \cap \bigcup \mathcal{N}_1^2 = \{w_1, w_2, w_3, w_4\}$, and $\mathcal{P}_1^i = \mathcal{N}_1^i$. $\mathbf{Q}_1^1(H) = \mathbf{Q}(H|\mathcal{P}^1(w_1) \cap \mathcal{C}_1) = [1/5, 1/2]$; $\mathbf{Q}_1^2(H) = \mathbf{Q}(H|\mathcal{P}^2(w_1) \cap \mathcal{C}_1) = [1/3, 2/3]$.³

Step 2 Agent i announces $\mathbf{Q}_1^i(H)$. $\mathcal{N}_2^1 = \{\{w_1, w_2, w_3\}\}$ (for $\mathbf{Q}(H|\{w_4, w_5, w_6\} \cap \mathcal{C}_1) = \{1\} \neq \mathbf{Q}_1^1(H)$.) $\mathcal{N}_2^2 = \mathcal{P}_1^2$. Thus $\mathcal{C}_2 = \{w_1, w_2, w_3\}$, and $\mathcal{P}_2^i = \mathcal{N}_2^i$. $\mathbf{Q}_2^1(H) = \mathbf{Q}(H|\mathcal{P}^1(w_1) \cap \mathcal{C}_2) = [1/5, 1/2]$; $\mathbf{Q}_2^2(H) = \mathbf{Q}(H|\mathcal{P}^2(w_1) \cap \mathcal{C}_2) = [1/5, 1/2]$.

Step 3 Agent i announces $\mathbf{Q}_2^i(H)$. $\mathcal{N}_3^i = \mathcal{P}_2^i$, and so $\mathcal{C}_3 = \mathcal{C}_2$. The procedure stops.

In this example, the communication ends up making each agent's private information public. This is not always the case, as later examples will illustrate. When (at least one agent's) private information remains private, it is in general possible to agree to disagree. However, in the case of a precise prior, that is, if $\mathbf{Q} = \{\tilde{p}\}$ is a singleton, Geanakoplos and Polemarchakis (1982, Proposition 1) showed that when the procedure stops at step $m + 1$, it is necessarily the case that $\mathbf{Q}_m^1(H) = \mathbf{Q}_m^2(H)$. We present a version of the argument here that will facilitate our subsequent discussion. Suppose the procedure stops at step $m + 1$. It means that $\mathcal{P}_{m+1}^i = \mathcal{P}_m^i$ (for both $i = 1, 2$, as we always intend). This entails, by the definition of \mathcal{P}_{m+1}^i , that

$$\forall E \in \mathcal{P}_m^i, \mathbf{Q}(H|E \cap \mathcal{C}_m) = \mathbf{Q}_m^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m). \quad (1)$$

Since $\mathbf{Q} = \{\tilde{p}\}$, $\mathbf{Q}(H|E \cap \mathcal{C}_m) = \{\tilde{p}(H|E \cap \mathcal{C}_m)\}$ and $\mathbf{Q}_m^i(H) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m) = \{\tilde{p}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)\}$. It follows that

$$\forall E \in \mathcal{P}_m^i, \tilde{p}(H|E \cap \mathcal{C}_m) = \tilde{p}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m). \quad (2)$$

3. Although $\mathbf{Q}_1^i(H) = \mathbf{Q}_0^i(H)$, the procedure goes on, because some agent still acquires new information in this step.

Since all members of \mathcal{P}_m^i are mutually disjoint, (2) entails that

$$\tilde{p}(H|\bigcup \mathcal{P}_m^i \cap \mathcal{C}_m) = \tilde{p}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m). \quad (3)$$

Recall that $\mathcal{C}_m = \bigcup \mathcal{P}_m^1 \cap \bigcup \mathcal{P}_m^2$. Hence $\bigcup \mathcal{P}_m^i \cap \mathcal{C}_m = \mathcal{C}_m$. It then follows from (3) that

$$\tilde{p}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m) = \tilde{p}(H|\mathcal{C}_m). \quad (4)$$

Therefore, $\tilde{p}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \tilde{p}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$; that is, the two agents end up agreeing.

Two comments are in order. First, equation (4) shows that the two agents are driven to the same posterior because when the communication stops, the resulting common knowledge (\mathcal{C}_m) renders each agent's private information ($\mathcal{P}^i(w)$) irrelevant to H (even if $\mathcal{P}^i(w)$ remains private). However, it does not follow that $\mathcal{P}^1(w)$ and $\mathcal{P}^2(w)$ are jointly irrelevant to H given \mathcal{C}_m . As Geanakoplos and Polemarchakis (1982, Proposition 3) observed, the consensus reached via the procedure of communicating posteriors can be different from the consensus that would result from directly exchanging private information. Clearly, they are different if and only if $\mathcal{P}^1(w)$ and $\mathcal{P}^2(w)$ are jointly relevant to H given \mathcal{C}_m , even though each is marginally irrelevant given \mathcal{C}_m (see Example 3 in Section 6).

Second, and more importantly for the purpose of this paper, a crucial step in the above argument is the move from (2) to (3), where what is needed is the following fact: if all members of a (finite) set of events \mathcal{E} are mutually disjoint, and for every $E \in \mathcal{E}$, $\tilde{p}(H|E) = q$, then $\tilde{p}(H|\bigcup \mathcal{E}) = q$. An analogous condition for imprecise probabilities would be the following: if all members of a (finite) set of events \mathcal{E} are mutually disjoint, and for every $E \in \mathcal{E}$, $\mathbf{Q}(H|E) = Q$ (where Q is a set of real numbers), then $\mathbf{Q}(H|\bigcup \mathcal{E}) = Q$. This condition does not hold in general for sets of probabilities.⁴

3. Dilation and Agreeing to Disagree on Lower and Upper Probabilities

We borrow a simple example from Carvajal and Correia-da-Silva (2010) to illustrate the failure of the said condition for sets of probabilities.

Example 2 Let $\Omega = \{w_1, w_2, w_3, w_4\}$ and \mathcal{A} the power set of Ω . Let $\mathcal{P}^1 = \{\{w_1, w_2\}, \{w_3, w_4\}\}$ and $\mathcal{P}^2 = \{\Omega\}$. Suppose $\mathbf{Q} = \{(1/2, 0, 1/2, 0), (0, 1/2, 0, 1/2)\}$; that is, the common set of priors consists of just two probability measures, represented by the two probability vectors.⁵ Let $H = \{w_2, w_3\}$, and suppose the true state of the world is w_1 .

This is an Aumann case in that the agents' posteriors on H are already common knowledge at the beginning; the procedure of communicating posteriors stops at step 1, for $\mathcal{C}_1 = \mathcal{C}_0 = \Omega$. However, $\mathbf{Q}_0^1(H) = \{0, 1\}$ and $\mathbf{Q}_0^2(H) = \{1/2\}$. Not only are the sets non-identical, they are in full

4. Even in the case of precise probability, it is well known that this condition, as a special case of conglomerability, can fail for finitely but not countably additive probability measures (de Finetti, 1972; Schervish et al., 1984; Hill and Lane, 1985). This does not matter in the setup we are considering, for the partitions are assumed to be finite. However, the original setup in Aumann (1976) seems to allow denumerable infinite partitions, in which case Aumann's result does not necessarily hold for merely finitely additive probabilities. More generally, Schervish et al. (2016) showed that conglomerability can fail in a partition of cardinality κ for a probability measure that is not κ -additive. Thus, if uncountable partitions are allowed, Aumann's result may fail even for countably additive measures.

5. In case readers are concerned that the two probabilities are not positive and are mutually singular, these special features are not essential. We can also use $\mathbf{Q} = \{(1/2 - \epsilon, \epsilon, 1/2 - \epsilon, \epsilon), (\epsilon, 1/2 - \epsilon, \epsilon, 1/2 - \epsilon)\}$, $0 < \epsilon < 1/4$, to make the same point.

disagreement in the sense that they do not even intersect and have different lower and upper probabilities. The agents agree to fully disagree. The condition we highlighted at the end of Section 2 fails dramatically in this case for agent 1's partition \mathcal{P}_1^1 (which is identical to \mathcal{P}^1 in this case): $\mathbf{Q}(H|\{w_1, w_2\}) = \mathbf{Q}(H|\{w_3, w_4\}) = \{0, 1\}$, while $\mathbf{Q}(H|\{w_1, w_2, w_3, w_4\}) = \{1/2\}$.

This dramatic failure of the condition is known as dilation (Good, 1974; Seidenfeld, 1981; Walley, 1991; Seidenfeld and Wasserman, 1993; Herron et al., 1997). No matter which member of \mathcal{P}_1^1 is the case, the resulting conditional probability is less precise than the probability conditional on $\bigcup \mathcal{P}_1^1$. Given a non-empty set of probabilities \mathbf{R} , let $\underline{\mathbf{R}}(A|E) = \inf_{p \in \mathbf{R}} p(A|E)$ denote the lower probability of A conditional on E , and $\overline{\mathbf{R}}(A|E) = \sup_{p \in \mathbf{R}} p(A|E)$ denote the upper probability of A conditional on E . Here is a definition of dilation that suits the present purpose.

Definition 1 (Dilation) *Let \mathbf{R} be a non-empty set of probability measures on (Ω, \mathcal{A}) . Let \mathcal{E} be a finite, non-empty set of mutually disjoint events. \mathcal{E} is said to dilate an event A with respect to \mathbf{R} (or $\mathbf{R}(\bullet|\bigcup \mathcal{E})$) if for every $E \in \mathcal{E}$, the interval $[\underline{\mathbf{R}}(A|E), \overline{\mathbf{R}}(A|E)]$ strictly contains the interval $[\underline{\mathbf{R}}(A|\bigcup \mathcal{E}), \overline{\mathbf{R}}(A|\bigcup \mathcal{E})]$.*

This is a slight generalization of the standard definition of dilation (Seidenfeld and Wasserman, 1993, p. 1141)⁶, for it considers dilation in a subspace $\bigcup \mathcal{E}$ (the definition reduces to the standard one when $\bigcup \mathcal{E} = \Omega$), but the idea and the anomalous feature are exactly the same. Again, in example 2, \mathcal{P}_1^1 , which happens to be the same as $\{E \cap \mathcal{C}_1 \mid E \in \mathcal{P}_1^1\}$, dilates the hypothesis of interest with respect to the given prior. This is not a coincidence, as Theorem 3 below shows. It is a simple consequence of the following lemma, which is a straightforward generalization of Lemma 1 in Carvajal and Correia-da-Silva (2010; also see Kajji and Ui, 2005, Proposition 3).

Lemma 2 *Suppose the procedure of communicating posteriors stops at step $m + 1$. Then*

$$\mathbf{Q}(H|\mathcal{C}_m) \subseteq [\underline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m), \overline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)]$$

for both $i = 1, 2$.

Proof As already mentioned, when the procedure stops at step $m + 1$, we have equation (1), namely,

$$\forall E \in \mathcal{P}_m^i, \mathbf{Q}(H|E \cap \mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m).$$

Consider $i = 1$ first. Let $\mathcal{P}_m^1 = \{E_1, \dots, E_k\}$. Notice that $\{E_1 \cap \mathcal{C}_m, \dots, E_k \cap \mathcal{C}_m\}$ forms a partition of $\bigcup \mathcal{P}_m^1 \cap \mathcal{C}_m$. Hence, for every $p \in \mathbf{Q}$, by the law of total probability

$$p(H|\mathcal{C}_m) = p(H|\bigcup \mathcal{P}_m^1 \cap \mathcal{C}_m) = \sum_{1 \leq j \leq k} p(H|E_j \cap \mathcal{C}_m)p(E_j \cap \mathcal{C}_m|\bigcup \mathcal{P}_m^1 \cap \mathcal{C}_m). \quad (5)$$

Given equation (1), we have that for every $1 \leq j \leq k$, $p(H|E_j \cap \mathcal{C}_m) \in \mathbf{Q}(H|E_j \cap \mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m)$. It follows that for every $1 \leq j \leq k$,

$$\underline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) \leq p(H|E_j \cap \mathcal{C}_m) \leq \overline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m). \quad (6)$$

6. Herron et al. (1997, p. 412) gave a weaker definition of dilation, requiring only that all conditional intervals contain and some of them *strictly* contain the unconditional interval. This definition (similarly generalized) would work equally well for our purpose. We thank an anonymous referee for this point.

Equation (5) and (6) together entail that

$$p(H|\mathcal{C}_m) \geq \underline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) \sum_{1 \leq j \leq k} p(E_j \cap \mathcal{C}_m | \bigcup \mathcal{P}_m^1 \cap \mathcal{C}_m) = \underline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m), \quad (7)$$

and

$$p(H|\mathcal{C}_m) \leq \overline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) \sum_{1 \leq j \leq k} p(E_j \cap \mathcal{C}_m | \bigcup \mathcal{P}_m^1 \cap \mathcal{C}_m) = \overline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m). \quad (8)$$

Since (7) and (8) hold for every $p \in \mathbf{Q}$, the desired conclusion is established for $i = 1$. The case of $i = 2$ is of course entirely parallel. \blacksquare

Lemma 2 shows that although equation (1) does not entail that $\mathbf{Q}(H|\mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$, it does entail that $\mathbf{Q}(H|\mathcal{C}_m)$ is bounded by the infimum and supremum of $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$. The following theorem is then immediate.

Theorem 3 *Suppose the procedure of communicating posteriors stops at step $m + 1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H , then $\underline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \underline{\mathbf{Q}}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$ and $\overline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \overline{\mathbf{Q}}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$.*

Proof Lemma 2 entails that for both $i = 1, 2$,

$$\underline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m) \leq \underline{\mathbf{Q}}(H|\mathcal{C}_m), \overline{\mathbf{Q}}(H|\mathcal{C}_m) \leq \overline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m).$$

Since $\mathcal{C}_m = \bigcup \mathcal{P}_m^i \cap \mathcal{C}_m$, if either of the inequality is strict, then $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ dilates H , because of equation (1). Therefore, if $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H , then

$$\underline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m) = \underline{\mathbf{Q}}(H|\mathcal{C}_m), \overline{\mathbf{Q}}(H|\mathcal{C}_m) = \overline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m).$$

The desired conclusion follows. \blacksquare

Thus, the two agents can agree to disagree on the lower or upper probability of a hypothesis only if a certain dilation takes place. Without dilation, the two agents are guaranteed to reach consensus on lower and upper probabilities by communicating posteriors.

It is worth noting that for the argument for Theorem 3 to go through, it is not necessary to require the agents to communicate their sets of posteriors. It is sufficient to ask them to communicate lower and upper probabilities at each step. Consider the procedure of communicating lower and upper posteriors: at step $n + 1$, agent i announces $\underline{\mathbf{Q}}_n^i(H)$ and $\overline{\mathbf{Q}}_n^i(H)$. Let

$$\begin{aligned} \mathcal{N}_{n+1}^{i\dagger} &= \{E \in \mathcal{P}_n^{i\dagger} \mid \underline{\mathbf{Q}}(H|E \cap \mathcal{C}_n) = \underline{\mathbf{Q}}_n^i(H) \text{ and } \overline{\mathbf{Q}}(H|E \cap \mathcal{C}_n) = \overline{\mathbf{Q}}_n^i(H)\} \\ \mathcal{C}_{n+1}^\dagger &= \bigcup \mathcal{N}_{n+1}^{1\dagger} \cap \bigcup \mathcal{N}_{n+1}^{2\dagger} \\ \mathcal{P}_{n+1}^{i\dagger} &= \{E \in \mathcal{N}_{n+1}^{i\dagger} \mid E \cap \mathcal{C}_{n+1}^\dagger \neq \emptyset\} \end{aligned}$$

If $\mathcal{P}_{n+1}^{i\dagger} = \mathcal{P}_n^{i\dagger}$, or equivalently, if $\mathcal{C}_{n+1}^\dagger = \mathcal{C}_n^\dagger$, the procedure stops; otherwise, agent i updates credence to $\underline{\mathbf{Q}}_{n+1}^i(H) = \underline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_{n+1}^\dagger)$, and enters the next step.

As before, this modified procedure is guaranteed to stop at step $m + 1$ for some $m \geq 0$, because \mathcal{P}^1 and \mathcal{P}^2 are assumed to be finite. The version of Lemma 2 on this procedure remains valid, for equation (1) is not necessary for the argument. All that is needed is the weaker condition that

$$\forall E \in \mathcal{P}_m^i, \underline{\mathbf{Q}}(H|E \cap \mathcal{C}_m) = \underline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m) \text{ and } \overline{\mathbf{Q}}(H|E \cap \mathcal{C}_m) = \overline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$$

This weaker condition obviously remains true when \mathcal{P}_m^i is replaced by $\mathcal{P}_m^{i\dagger}$ and \mathcal{C}_m by \mathcal{C}_m^\dagger . Hence we also have the following variant of Theorem 3.

Theorem 4 *Suppose the procedure of communicating lower and upper posteriors stops at step $m + 1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m^\dagger \mid E \in \mathcal{P}_m^{i\dagger}\}$ does not dilate H , then $\underline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m^\dagger) = \underline{\mathbf{Q}}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m^\dagger)$ and $\overline{\mathbf{Q}}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m^\dagger) = \overline{\mathbf{Q}}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m^\dagger)$.*

Proof Extremely similar to that of Theorem 3. ■

Although both procedures result in consensus on lower and upper probabilities in the absence of dilation, in general the agreements they lead to may well be different, for in general communicating lower and upper posteriors conveys less information than communicating the full sets of posteriors.

4. More Agreement Results

Under some common assumptions, however, lower and upper probabilities are sufficient to identify the full set, in which case the two procedures are equivalent and, more importantly, the consensus reached in the absence of dilation will be full consensus. For example, if we follow Carvajal and Correia-da-Silva (2010) to assume that the set of priors is closed and connected (or follow Kajii and Ui (2005) to assume that the set of posteriors is a closed interval), we obtain the following result.

Theorem 5 *Suppose \mathbf{Q} is closed and connected (with respect to the total variation topology), and suppose the procedure of communicating posteriors stops at step $m + 1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H , then $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$.*

Proof Given the assumption that all the relevant conditional probabilities are well defined as ratios of unconditional probabilities, the mapping from \mathbf{Q} to $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ is continuous.⁷ Hence, since \mathbf{Q} is assumed to be closed and connected, $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ is a closed interval. Thus $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ is identified by $\overline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ and $\underline{\mathbf{Q}}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$. Then Theorem 3 entails that $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) = \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m)$. ■

To our knowledge, Theorem 5 is the first attempt to formulate a generalization of Aumann's agreement theorem in the context of imprecise probability that takes agreement to mean full agreement (identical set of posteriors). In addition to revealing a connection to the important phenomenon of dilation, Theorem 5 may also be used to generate sufficient conditions for guaranteed full agreement via communicating posteriors, for important models of imprecise probability, if sufficient conditions for the absence of dilation in those models are known. As a simple example, consider the density ratio classes for finite spaces (Wasserman, 1992; Seidenfeld and Wasserman, 1993).

Definition 6 (Density Ratio Prior) *Let $\Omega = \{w_1, \dots, w_n\}$ and \mathcal{A} the power set of Ω . A density ratio prior is defined by*

$$\mathbf{D}_{p,k} = \{(q_1, \dots, q_n) \mid \sum_{1 \leq j \leq n} q_j = 1 \text{ and } \frac{q_h}{q_j} \leq k \frac{p_h}{p_j}, \forall 1 \leq h, j \leq n\}$$

where $k \geq 1$ and (p_1, \dots, p_n) is a probability vector such that $p_j > 0$ for all $1 \leq j \leq n$.

7. The mapping is given by: $p \mapsto p(H \cap \mathcal{P}^i(w) \cap \mathcal{C}_m) / p(\mathcal{P}^i(w) \cap \mathcal{C}_m)$, which is obviously continuous, as long as the ratio is always defined.

For instance, Example 1 in Section 2 employs a density ratio prior, where p is the uniform distribution over the 9-atom algebra and $k = 2$.

Corollary 7 *If two agents start with a common density ratio prior and carry out the procedure of communicating posteriors, they are guaranteed to reach the same set of posteriors.*

Proof Seidenfeld and Wasserman (1993, Theorem 4.1) showed that the density ratio priors are dilation-immune in the sense that no finite partition of the sample space dilates any event. Note also that if \mathbf{D} is a density ratio prior on (Ω, \mathcal{A}) , then for every $E \in \mathcal{A}$, $\mathbf{D}(\bullet|E)$ remains a density ratio prior on the space restricted to E , which follows easily from Definition 6. Moreover, a density ratio prior is obviously closed and connected. Then Theorem 5 entails the desired conclusion. ■

Finally, if we consider just *partial* agreement, in the sense of a non-empty intersection of sets of posteriors, we can drop the assumption of connectedness in Theorem 5.

Theorem 8 *Suppose \mathbf{Q} is closed, and suppose the procedure of communicating posteriors stops at step $m + 1$. If for both $i = 1, 2$, $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H , then $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) \cap \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m) \neq \emptyset$.*

Proof Since \mathbf{Q} is closed, $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ is also closed, for the mapping from \mathbf{Q} to $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ is continuous. Thus, $\mathbf{Q}(H|\mathcal{P}^i(w) \cap \mathcal{C}_m)$ contains its infimum and supremum. It then follows from Theorem 3 that $\mathbf{Q}(H|\mathcal{P}^1(w) \cap \mathcal{C}_m) \cap \mathbf{Q}(H|\mathcal{P}^2(w) \cap \mathcal{C}_m) \neq \emptyset$. ■

5. Dilation-Averse Updating

The presence of dilation may alarm some agents, who may be inclined to think that they are permitted or even rationally required to ignore information about the value of a partition that dilates the hypothesis of interest (Grünwald and Halpern, 2004). Whether dilating information should be ignored is a matter of debate. For example, Kyburg's (1974) theory of "epistemological", interval-valued probability precludes altogether the possibility of dilation. However, the rule in his theory that is responsible for the impossibility of dilation was forcefully criticized by Levi (1977) on the grounds that it runs afoul of some basic Bayesian tenets even when the theory delivers precise probability values. We do not pretend to resolve the debate here, but we would like to reformulate the main ideas of this paper in terms of dilation-averse agents, which we believe provides a useful perspective to think about our results. An agent participating in the communication procedure is said to be *dilation-averse* if she does not condition on her information about the value of a partition that dilates the hypothesis of interest, but is otherwise happy to update by Bayesian conditioning. Suppose it is common knowledge that the two agents are dilation-averse. To model this situation, the procedure of communicating posteriors should be modified as follows.

At step 0, for each agent i , if $\mathcal{P}_0^i = \{E \in \mathcal{P}^i \mid E \cap \mathcal{C}_0 \neq \emptyset\}$ dilates H , she updates her credence by conditioning on the common knowledge \mathcal{C}_0 : $\mathbf{Q}_0^1(H) = \mathbf{Q}(H|\mathcal{C}_0)$; otherwise, she updates her credence in the standard way: $\mathbf{Q}_0^1(H) = \mathbf{Q}(H|\mathcal{P}^1(w))$.

At step $n + 1$, the agents announce $\mathbf{Q}_n^1(H)$ and $\mathbf{Q}_n^2(H)$, respectively. Consider the set $\tilde{\mathcal{N}}_{n+1}^i = \{E \in \mathcal{P}_n^i \mid \mathbf{Q}(H|E \cap \mathcal{C}_n) = \mathbf{Q}_n^i(H)\}$. It is easy to see that $\tilde{\mathcal{N}}_{n+1}^i = \emptyset$ if and only if there was dilation at step n . Let

$$\mathcal{N}_{n+1}^i = \begin{cases} \mathcal{P}_n^i & \text{if } \tilde{\mathcal{N}}_{n+1}^i = \emptyset, \\ \tilde{\mathcal{N}}_{n+1}^i & \text{otherwise.} \end{cases}$$

That is, when $\tilde{\mathcal{N}}_{n+1}^i = \emptyset$, no new information is conveyed by the announcement of $\mathbf{Q}_n^i(H)$. As before, let $\mathcal{C}_{n+1} = \bigcup \mathcal{N}_{n+1}^1 \cap \bigcup \mathcal{N}_{n+1}^2$, and $\mathcal{P}_{n+1}^i = \{E \in \mathcal{N}_{n+1}^i \mid E \cap \mathcal{C}_{n+1} \neq \emptyset\}$. Clearly, it remains true that $\mathcal{P}_{n+1}^i \subseteq \mathcal{N}_{n+1}^i \subseteq \mathcal{P}_n^i$ and $\mathcal{C}_{n+1} = \bigcup \mathcal{P}_{n+1}^1 \cap \bigcup \mathcal{P}_{n+1}^2$. If $\mathcal{P}_{n+1}^i = \mathcal{P}_n^i$, or equivalently, if $\mathcal{C}_{n+1} = \mathcal{C}_n$, the procedure stops; otherwise, agent i updates credence of H according to whether $\{E \cap \mathcal{C}_{n+1} \mid E \in \mathcal{P}_{n+1}^i\}$ dilates H . If it does not dilate H , the credence is updated to $\mathbf{Q}_{n+1}^i(H) = \mathbf{Q}(H \mid \mathcal{P}^i(w) \cap \mathcal{C}_{n+1})$; otherwise, the credence is updated to $\mathbf{Q}_{n+1}^i(H) = \mathbf{Q}(H \mid \mathcal{C}_{n+1})$.

For instance, if the agents in Example 2 are commonly known to be dilation-averse and follow the above procedure, then at step 0, seeing that her partition $\{\{w_1, w_2\}, \{w_3, w_4\}\}$ dilates H , agent 1 will ignore her private information (i.e., $\{w_1, w_2\}$) and go with $\mathbf{Q}_0^1(H) = \mathbf{Q}(H \mid \mathcal{C}_0) = \{1/2\}$. Then at step 1, $\mathcal{P}_1^i = \mathcal{N}_1^i = \mathcal{P}_0^i$, and the procedure stops (with a consensus).

As the original, Bayesian procedure of communicating posteriors, this dilation-averse procedure will surely stop at step $m + 1$ for some $m \geq 0$. It is then very easy to adapt the arguments for Theorems 3, 5, and 8 to show the following:

Theorem 9 *Suppose that the dilation-averse procedure of communicating posteriors stops at step $m + 1$. Then*

- 1) $\underline{\mathbf{Q}}_m^1(H) = \underline{\mathbf{Q}}_m^2(H)$ and $\overline{\mathbf{Q}}_m^1(H) = \overline{\mathbf{Q}}_m^2(H)$;
- 2) If \mathbf{Q} is closed, then $\mathbf{Q}_m^1(H) \cap \mathbf{Q}_m^2(H) \neq \emptyset$; and
- 3) If \mathbf{Q} is closed and connected, then $\mathbf{Q}_m^1(H) = \mathbf{Q}_m^2(H)$.

Proof For each i , either $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ dilates H , in which case $\mathbf{Q}_m^i(H) = \mathbf{Q}(H \mid \mathcal{C}_m)$ by the design of the procedure, or $\{E \cap \mathcal{C}_m \mid E \in \mathcal{P}_m^i\}$ does not dilate H , in which case the argument for Theorem 3 is applicable to derive that $\mathbf{Q}_m^i(H) = \mathbf{Q}(H \mid \mathcal{C}_m)$ and $\overline{\mathbf{Q}}_m^i(H) = \overline{\mathbf{Q}}(H \mid \mathcal{C}_m)$. Either way we have 1). The derivations of 2) and 3) from 1) are the same as those of Theorems 8 and 5. ■

Therefore, two agents who are commonly known to be dilation-averse cannot agree to disagree on lower or upper probabilities, and, under common assumptions, cannot agree not to fully agree.

6. Concluding Remarks

Like Aumann's original result, the results in this paper are mathematically simple once the framework is set up, but they highlight an interesting connection between the possibility of agreeing to disagree and the phenomenon of dilation. We offered two perspectives to view this connection. For Bayesian agents with a common set of priors, agreeing to disagree on lower or upper posteriors entails the presence of dilation for at least one of them. For dilation-averse (but otherwise Bayesian) agents with a common set of priors, it is impossible to agree to disagree on lower or upper posteriors.

Although the absence of dilation is sufficient for Bayesian agents to reach agreements by communicating posteriors, it is not necessary. Here is a simple example to show this.

Example 3 *Let $\Omega = \{w_1, w_2, w_3, w_4\}$ and \mathcal{A} be its power set. Suppose $\mathcal{P}^1 = \{\{w_1, w_2\}, \{w_3, w_4\}\}$ and $\mathcal{P}^2 = \{\{w_1, w_3\}, \{w_2, w_4\}\}$. Let $H = \{w_1, w_4\}$, and suppose the true state of the world is w_1 . Let \tilde{p} be the uniform distribution over the 4-atom algebra, and Λ be the set of all distributions over the 4-atom algebra. Define $\mathbf{Q} = \{(0.8\tilde{p} + 0.2q) \mid q \in \Lambda\}$.⁸*

8. This ϵ -contamination model ($\epsilon = 0.2$) can be equivalently specified as the largest set of probability measures on the 4-atom algebra satisfying the constraint that every atom receives a lower probability of 0.2.

Like Example 2, this is an Aumann case, where the posteriors of H are common knowledge without announcements, because $\mathbf{Q}(H|\{w_1, w_2\}) = \mathbf{Q}(H|\{w_3, w_4\}) = [1/3, 2/3]$, and $\mathbf{Q}(H|\{w_1, w_3\}) = \mathbf{Q}(H|\{w_2, w_4\}) = [1/3, 2/3]$. So the procedure of communicating posteriors stops at step 1, and $\mathcal{C}_1 = \mathcal{C}_0 = \Omega$. Dilation does occur, for both agents, because $\mathbf{Q}(H|\mathcal{C}_0) = \mathbf{Q}(H) = [0.4, 0.6]$, which is strictly contained in $[1/3, 2/3]$. Despite the presence of dilations, the two agents will still reach an agreement even if they are not dilation-averse, though the agreement is different from the one dilation-averse agents would reach.

It is also worth noting that this example is a generalization of an example from Geanakoplos and Polemarchakis (1982), which was used to illustrate the fact we mentioned in Section 2, that the consensus resulting from communicating posteriors can be different from the consensus resulting from directly exchanging private information. If both pieces of private information in the example become public, the two agents will converge on a precise, extreme probability.

We close by mentioning two ways our results may be expanded. First, when “agreement” is interpreted as partial agreement, the common prior assumption may also be relaxed to the assumption that priors (partially) agree, i.e., that the two sets of priors have a non-empty intersection. This is, for example, what Carvajal and Correia-da-Silva (2010) assume in their results. Their main agreement result about Bayesian agents (Proposition 1) is that if two Bayesian agents have closed and connected sets of priors that have a non-empty intersection, and both sets of posteriors on a hypothesis are common knowledge, then the sets of posteriors also have a non-empty intersection. This result, just like Aumann’s original result, is straightforwardly generalizable to the setting of communicating posteriors. The more interesting question, in light of our results here, is what purchase the condition of no dilation has in the context of priors that do not fully agree, or to put it differently, whether stronger agreement results are available in this context for dilation-averse agents.

Second, we have only considered the full Bayesian updating rule (and the dilation-averse variant). Other updating rules may be examined in our setting, especially the Dempster-Shafer or maximum likelihood updating considered by Kajii and Ui (2005) and Carvajal and Correia-da-Silva (2010). For Dempster-Shafer updating, Carvajal and Correia-da-Silva’s main agreement result requires each agent’s set of likelihood maximizers as well as their sets of posteriors to be common knowledge, which suggests that in general communication of posteriors alone is not enough to guarantee agreement. One natural idea is to allow also the communication of likelihood maximizers. On the other hand, Seidenfeld (1997) showed that for ϵ -contamination models (Huber, 1973; Berger, 1984) Dempster-Shafer updating is equivalent to Bayesian updating. Therefore, if we can derive a corollary about ϵ -contamination models (in the spirit of Corollary 7) from Theorem 5 and results on dilation in ϵ -contamination models, that will also apply to Dempster-Shafer updating.

Acknowledgements

This research was supported by the Research Grants Council of Hong Kong under the General Research Fund LU13600715, and by a Faculty Research Grant from Lingnan University.

References

- Aumann, R. (1976). Agreeing to disagree. *Ann. Stat.*, 4: 1236-1239.
- Berger, J. (1984). The robust Bayesian viewpoint (with discussion). In *Robustness in Bayesian Statistics* (J. Kadane, ed.), pp. 63-124. North-Holland, Amsterdam.

- Carvajal, A. and Correia-da-Silva, J. (2010). *Agreeing to disagree with multiple priors* (No. 368). Universidade do Porto, Faculdade de Economia do Porto.
- De Finetti, B. (1972). *Probability, Induction, and Statistics*. John Wiley, New York.
- Geanakoplos, J. and Polemarchakis, M. (1982). We can't disagree forever. *J. Econ. Theory*, 26: 363-390.
- Good, I. J. (1974). A little learning can be dangerous. *Br. J. Philos. Sci.*, 25: 340-342.
- Grünwald, P. D. and Halpern, J. Y. (2004). When ignorance is bliss. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, pp. 226-234. AUAI Press.
- Herron, T., Seidenfeld, T. and Wasserman, L. (1997). Divisive conditioning: Further results on dilation. *Philosophy of Science*, 64: 411-444.
- Hill, B. M. and Lane, D. (1985). Conglomerability and countable additivity. *Sankhyā: The Indian Journal of Statistics, Series A*: 366-379.
- Huber, P. J. (1973). The use of Choquet capacities in statistics. *Bull. Internat. Statist. Inst.*, 45: 181-191.
- Levi, I. (1977). Direct inference. *The Journal of Philosophy*, 74(1): 5-29.
- Kajii, A. and Ui, T. (2005). Incomplete information games with multiple priors. *JPN Econ. Rev.*, 56: 332-351.
- Kajii, A., and Ui, T. (2009). Interim efficient allocations under uncertainty. *J. Econ. Theory*, 144: 337-353.
- Kyburg H. E. (1974). *The Logical Foundations of Statistical Inference*. Reidel, Dordrecht.
- Schervish, M. J., Seidenfeld, T. and Kadane, J. B. (1984). The extent of non-conglomerability of finitely additive probabilities. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 66(2): 205-226.
- Schervish, M. J., Seidenfeld, T. and Kadane, J. B. (2016). Non-conglomerability for countably additive measures that are not κ -additive. *The Review of Symbolic Logic*, forthcoming.
- Seidenfeld, T. (1981). Levi on the dogma of randomization. In *Henry E. Kyburg and Isaac Levi* (R. Bogdan, ed.), pp. 263-291. Reidel, Dordrecht.
- Seidenfeld, T. (1997). Some static and dynamic aspects of robust Bayesian theory. In *Random Sets*, pp. 385-406. Springer, New York.
- Seidenfeld, T. and Wasserman, L. (1993). Dilation for sets of probabilities. *Ann. Stat.*, 21: 1139-1154.
- Wasserman, L. (1992). Invariance properties of density ratio priors. *Ann. Stat.*, 20: 2177-2182.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, New York.