Dynamics for Probabilistic Common Belief

Eric Pacuit*

December 23, 2015

Abstract

Robert Aumann's agreeing to disagree theorem shows that if two agents have the same prior probability and update their probability of an event E with private information by conditioning, then if the posterior probabilities of E are common knowledge, then the posteriors must be the same. Dov Monderer and Dov Samet prove a generalization of Aumann's result involving a probabilistic variant of common knowledge. In this paper, I use various methods from probabilistic and dynamic-epistemic logics to explore a dynamic characterization of the Monderer and Samet result. The main goal is to develop a model that describes the evolution of the agents' knowledge and (probabilistic) beliefs as they interact with each other and the environment. I will show how the logical frameworks are rich and flexible tools that can be used to study many dynamic processes of interactive social belief change.

*Department of Philosophy, University of Maryland, College Park, *Email*: epacuit@umd.edu, *Web*: `http://pacuit.org`. I would like to thank Alexandru Baltag, Jan van Eijk and, especially, Johan van Benthem for comments on an earlier draft of this paper that greatly improved both the content and readability of this paper.

1 Introduction

In 1976, Robert Aumann proved a fascinating result [3]. Suppose that two agents have the same prior probability and update their probability of an event E with private information by conditioning. Aumann showed that if the posterior probabilities of E are common knowledge, then the agents must assign the *same* posterior to E. This is true even if the agents receive different information. In other words, if agents have the same prior probability and update by conditioning, then the agents cannot "agree to disagree" about their posterior probabilities. This seminal result has been generalized in many ways [4, 8, 39, 40] and is still the subject of much discussion in Economics [31, 29, 7], Logic [10, 11] and, to a lesser extent, Philosophy [28, Section 3].

In this paper, I discuss a generalization of Aumann's theorem proved by Dov Monderer and Dov Samet [32]. Monderer and Samet define a probabilistic variant of common knowledge, which can serve as a natural "approximation" to common knowledge. If there is not common knowledge about the posteriors, then Aumann's theorem does not say anything about how the posteriors may be related. Monderer and Samet prove a generalization of Aumann's result: If there is common p-belief of the posteriors of an event for values of p close to 1, then the posteriors must all be very close together (a more formal statement is given in Section 2.2). This is one way in which common p-belief is a natural approximation of common knowledge. In this paper, I explore a dynamics characterization of the Monderer and Samet result, involving update actions, along the lines of Geanakoplos and Polemarchakis's well-known dynamics characterization of Aumann's Theorem [19]. Toward this end, various methods will be discussed from probabilistic and dynamic-epistemic logics.

2 Knowledge, Belief and Probability

2.1 Knowledge and Belief in Probabilistic Models

Suppose that $\mathcal{A} = \{1, \ldots, n\}$ is a non-empty set of agents. Each agent $i \in \mathcal{A}$ is endowed with a probability measure μ_i on a common probability space $(W, \mathcal{F})$, where W is a non-empty set and $\mathcal{F}$ is a σ-algebra. [1] To simplify the discussion, assume that W is finite (in this case, we can let $\mathcal{F} = \wp(W)$). [2] We say that there is a **common prior** among the agents $\mathcal{A}$ provided that for all $i, j \in \mathcal{A}$, $\mu_i = \mu_j$.

The agents' private information is represented by a partition on W. For each $i \in \mathcal{A}$, let Π_i be i's partition on W. For each $w \in W$, let $\Pi_i(w)$ be the element of Π_i containing w. So, if w is the "actual" state, then $\Pi_i(w)$ is the information that i receives at w (i.e., it is i's *information cell* at w). It is assumed that for each $w \in W$, $\Pi_i(w)$ is assigned a positive probability. For each $i \in \mathcal{A}$, i's **posterior probability at state** w, $\mu_{i,w} : \mathcal{F} \to [0, 1]$, is defined by conditioning i's prior probability on i's information cell at state w. Thus, for each set $E \in \mathcal{F}$, we have that

$$\mu_{i,w}(E) = \mu_i(E \mid \Pi_i(w)) = \frac{\mu(E \cap \Pi_i(w))}{\mu_i(\Pi_i(w))}.$$

These models have been used to characterize solution concepts in games [1, 21, 2] and as a semantics for a modal language with epistemic and probability operators [24, 20, 45] (this theme will be discussed in more detail in Section 4). The probability measure $\mu_{i,w}$ is agent i's posterior probability given i's current information (at state w). We are interested in the agents' posterior probabilities for some fixed event E. It is convenient to think of the agents'

[1] That is, $\mathcal{F} \subseteq \wp(W)$ is a collections of subsets of W such that 1) $W \in \mathcal{F}$; 2) $\mathcal{F}$ is closed under finite intersections; and 3) $\mathcal{F}$ is closed under countable unions.

[2] Much of what follows can be generalized to infinite state spaces. See [39] for a discussion of Aumann's theorem on infinite state spaces.

posteriors of a fixed event as a random variable.[3] This motivates the following notation: For an event $E \subseteq W$ and $r \in [0,1]$, let $[E_i = r] = \{w \mid \mu_{i,w}(E) = r\}$.

The agents' "knowledge" and "beliefs" are represented by operators on $\wp(W)$ (the set of subsets of $\wp(W)$). (Here we largely follow the notational conventions of epistemic game theory.) For each $i \in \mathcal{A}$, let $K_i : \wp(W) \rightarrow \wp(W)$ be defined as follows: For all $E \subseteq W$,

$$K_i(E) = \{w \mid \Pi_i(w) \subseteq E\}.$$

Then, $K_i(E)$ is an event consisting of all states w in which i's information at w implies E (the standard convention in the literature is to refer to the event $K_i(E)$ as "agent i knows that E"). It is well known that each K_i satisfies the so-called **S5** axioms: For all events E and F, $K_i(E \cap F) = K_i(E) \cap K_i(F)$, $K_i(E) \subseteq E$, $K_i(E) \subseteq K_i(K_i(E))$, and $\overline{K_i(E)} \subseteq K_i(\overline{K_i(E)})$ ($\overline{X}$ denotes the complement of X in W). The agents' beliefs are defined in an analogous way. For each $i \in \mathcal{A}$ and $p \in [0,1]$, let $B_i^p : \wp(W) \rightarrow \wp(W)$ be defined as follows: For $E \subseteq W$,

$$B_i^p(E) = \{w \mid \mu_{i,w}(E) = \mu_i(E \mid \Pi_i(w)) \geq p\},$$

where μ_i is i's prior probability on W. Thus, i *believes that* E (to degree at least p) provided that i's posterior of E is above the given threshold.

The following Lemma gathers together a number of important facts about the B_i^p operators. The proofs will not be repeated here, since they can be found in [32].

Lemma 2.1 *Suppose that* $(W, \mathcal{F})$ *is a probability space;* μ_i *is a probability measure on* $(W, \mathcal{F})$*;* Π_i *is a partition on* W*; and* $B_i^p : \wp(W) \rightarrow \wp(W)$ *is the belief operator defined above, where* $p \in [0,1]$*. Then,*

[3] Recall that, in probability theory, a **random variable** is a measurable function from the probability space $(W, \mathcal{F})$ to some other measurable space (typically the real numbers). The idea is that, for each event E and agent $i \in \mathcal{A}$, there is a function $f_{i,E} : W \rightarrow \mathbb{R}$, where $w \mapsto \mu_{i,w}(E)$.

1. *For all events E, $B_i^p(E)$ is a union of elements from Π_i.*

2. *For all events E, if E is a union of elements from Π_i, then $B_i^p(E) = E$.*

3. *For all events E, $B_i^p(B_i^p(E)) = B_i^p(E)$.*

4. *For all events E and F, if $E \subseteq F$, then $B_i^p(E) \subseteq B_i^p(F)$.*

5. *If (E_n) is a decreasing sequence of events then*

$$B_i^p\left(\bigcap_n E_n\right) = \bigcap_n B_i^p(E_n).$$

This Lemma demonstrates that the belief operators B_i^p have much in common with the knowledge operators K_i. Note that each K_i operator also satisfies items 1-5.

There are two crucial differences between knowledge and belief in this framework. The first follows from the laws of probability. Since the probability of the intersection of two events may be smaller than the probability of either event, it is easy to find events E and F such that $B_i^p(E) \cap B_i^p(F) \not\subseteq B_i^p(E \cap F)$. The second difference is that the agents' beliefs do not satisfy the *truth axiom* (while for the knowledge operator, we have that, for all events E, $K_i(E) \subseteq E$). However, the B_i^p operators satisfy a natural generalization of the truth axiom. While $B_i^p(E)$ does not necessarily imply that E is true, it does ensure that E has a probability of at least p. I will provide the proof of this fact here since it helps to illustrate the above definitions of the basic epistemic attitudes.

Lemma 2.2 *Suppose that $\langle W, \mathcal{F} \rangle$ is a probability space with probability measure μ_i, Π_i is a partition on W, and $B_i^p : \wp(W) \to \wp(W)$ is a p-belief operator. For all $E \subseteq W$ and $p \in [0,1]$, we have $\mu_i(E \mid B_i^p(E)) \geq p$.*

Proof. Suppose that E is an event and $p \in [0,1]$. Let $\mathcal{P}_i = \{P \mid P \in \Pi_i \text{ and } \mu_i(E \cap P) \geq p\mu_i(P)\}$. Then, the $\mathcal{P}_i$ contains all the partition cells from Π_i such that i's posterior for E is at least p. Note that $B_i^p(E) = \bigcup \mathcal{P}_i$ (cf. Clause 1 in Lemma 2.1). Since the elements of $\mathcal{P}_i$ are disjoint, we have $\mu_i(B_i^p(E)) = \sum_{P\in\mathcal{P}_i} \mu_i(P)$. Furthermore, since $E\cap B_i^p(E) = \bigcup_{P\in\mathcal{P}_i}(E\cap P)$, and each of the sets in $\mathcal{P}_i$ are disjoint, $\mu_i(E\cap B_i^p(E)) = \sum_{P\in\mathcal{P}_i} \mu_i(E\cap P)$. Finally, since, for each $P \in \mathcal{P}_i$, $\mu_i(E \mid P) \geq p$ (i.e., $\mu_i(E\cap P) \geq p\mu_i(P)$), we have

$$\mu_i(E\cap B_i^p(E)) = \sum_{P\in\mathcal{P}_i} \mu_i(E\cap P) \geq \sum_{P\in\mathcal{P}_i} p\mu_i(P) = p\sum_{P\in\mathcal{P}_i} \mu_i(P) = p\mu_i(B_i^p(E))$$

Thus, $\mu_i(E \mid B_i^p(E)) \geq p$, as desired. QED

2.2 Common Knowledge and Common p-Belief

The game theory and epistemic logic literature contains many notions of *group* knowledge. It is beyond the scope of this article to discuss all these concepts (see, for instance, [44, 15]). Instead, I will define common knowledge and the probabilistic variant of common knowledge proposed by Monderer and Samet [32].

Given knowledge operators K_i for each agent $i \in \mathcal{A}$, let $K : \wp(W) \to \wp(W)$ be $K(E) = \bigcap_{i\in\mathcal{A}} K_i(E)$. [4] Thus, $K(E)$ is the event that "*everyone knows that* E". An event is **common knowledge** if everyone knows that E and this fact is completely transparent to all the agents. This can be made precise as follows. An event E is said to be **self-evident for** i provided that $E \subseteq K_i(E)$. An event is self-evident if it is self-evident for i for each $i \in \mathcal{A}$. Thus, E is self-evident iff E is closed with respect to the agents' information partition: for all $i \in \mathcal{A}$, for all $w, v \in W$, if $w \in E$ and $v \in \Pi_i(w)$, then $v \in E$. A **common knowledge operator**, $C : \wp(W) \to \wp(W)$, is then defined as follows: For each $E \subseteq W$,

$$C(E) = \{w \mid \text{there is a self-evident set } F \text{ such that } w \in F \text{ and } F \subseteq E\}.$$

[4] This definition can be relativized to any subset of agents. I.e., for $G \subseteq \mathcal{A}$, let $K_G(E) = \bigcap_{i\in G} K_i(E)$. In order to keep notation to a minimum, I will not pursue this more general approach in this paper.

That is, E is commonly known provided that there is a true self-evident event that implies E. It is well known that, in many settings, this definition is equivalent to the more standard "infinite iterative" definition of common knowledge: [5] [6] For all events E,

$$C(E) = E \cap K(E) \cap K(K(E)) \cap K(K(K(E))) \cap \cdots$$

The definition of common p-belief is similar. Let $B^p : \wp(W) \to \wp(W)$ be the operator associating with every event E, the event $\bigcap_{i \in \mathcal{A}} B_i^p(E)$. Thus, $B^p(E)$ is the event in which every agent believes E to degree at least p. An event $E \subseteq W$ is an **evident p-belief for i** provided that $E \subseteq B_i^p(E)$ and is an **evident p-belief** if is an evident p-belief for all agents $i \in \mathcal{A}$. The **common p-belief operator**, $C^p : \wp(W) \to \wp(W)$, is defined as follows:

For all sets $E \subseteq W$,

$$\begin{aligned} C^p(E) \quad = \quad & \{w \mid \text{there is an event } F \text{ such that (i) } w \in F\text{; (ii) } F \text{ is an evident } p\text{-belief} \\ & \text{for every } i \in \mathcal{A}\text{; and (iii) } F \subseteq B_i^p(E) \text{ for all } i \in \mathcal{A}\} \ . \end{aligned}$$

Thus, an event E is commonly p-believed if there is a true self-evident p-belief that implies that everyone p-believes E. There is an equivalent definition in terms of iterations of the p-belief operators, though some care is needed since the p-belief operators are not closed under intersections. One can consult [32, 27, 34, 35, 20] for a discussion of the preceding definition of common p-belief and its relations with alternative definitions. The following example illustrates how the definition given here works.

[5] In a well-known paper, Jon Barwise [6] discusses three main different approaches to defining common knowledge: (i) the iterated view, (ii) the fixed-point view, and (iii) the shared situation view.

[6] It is worth pointing out that David Lewis had a somewhat different approach to common knowledge [30]. For Lewis, the infinite conjunction is a necessary but not a sufficient condition for common knowledge. See [9] for an illuminating discussion and a reconstruction of Lewis' notion of common knowledge. Nonetheless, following [3], the definition given in this section has become standard in game theory and epistemic logic.

Example 2.3 Suppose that there are two coins, each sitting in different drawers, and two agents, Ann (a) and Bob (b). For $i = 1, 2$, let H_i denote the event "the coin in drawer i is facing heads up" and let T_i denote the event "the coin in drawer i is facing tails up". Ann looks at the coin in drawer 1 and Bob looks at the coin in drawer 2. Suppose that after observing their respective coins, there is an announcement over a loudspeaker that both coins are facing heads up ($H_1 \cap H_2$). Assume that Ann was listening closely, but the announcement was not perfectly clear, and so the probability that she heard correctly is 0.9. Bob was not paying as close attention, and so the probability that he heard correctly is 0.8. Given that these probabilities are commonly known by Ann and Bob, the initial probability of state w_1 (where H_1 and H_2 are both true) is $0.9 \times 0.8 = 0.72$. Similar calculations for the remaining states give us the following diagram (I draw an i-labeled line between two states when the two states are in the same information cell for agent i). Furthermore, note that both agents have the same prior probability, so only one number is assigned to each state.

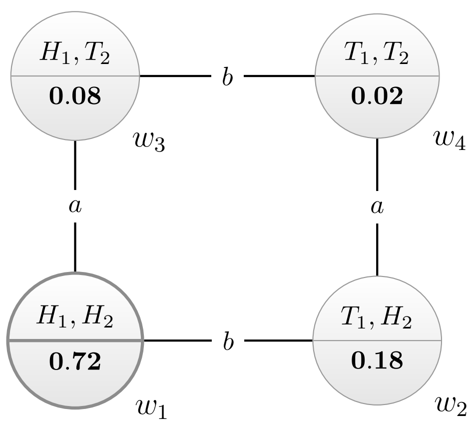

In particular, we have that $w_1 \in B_a^{0.9}(H_1 \cap H_2) \cap B_b^{0.8}(H_1 \cap H_2)$. This means that $X = \{w_1\}$ is an evident 0.8-belief for both Ann and Bob. Since $X \subseteq B_a^{0.8}(H_1 \cap H_2) = \{w_1, w_2, w_3, w_4\} = B_b^{0.8}(H_1 \cap H_2)$, we have that $w_1 \in C_{a,b}^{0.8}(H_1 \cap H_2)$. That is, while it is not common knowledge that the coins are both facing heads up, it is common 0.8-belief that they are. Finally, note that the agents' posteriors for the event $H_1 \cap H_2$ are 0.9 and 0.8.

2.3 The Agreement Theorems

Using the above definitions, I can now state Aumann's Theorem and Monderer and Samet's generalization involving common p-belief. I start with the statement of Aumann's Theorem.

Theorem 2.4 (Aumann, [3]) *Suppose that $(W, \mathcal{F})$ is a finite probability space and μ is a probability measure on $(W, \mathcal{F})$. Further, suppose that there is a common prior: for each $i \in \mathcal{A}$, $\mu_i = \mu$ is i's prior probability measure. For all real numbers $r_1, \ldots, r_n \in [0, 1]$ and events $E \in \mathcal{F}$, if $C(\bigcap_{i\in\mathcal{A}}[E_i = r_i]) \neq \emptyset$, then for all $i, j \in \mathcal{A}$, $r_i = r_j$.*

The proof of this theorem is not difficult once the definitions are in place (see [3]). Next, the Monderer and Samet result identifies the consequence of assuming that the agents' posterior probabilities of an event E are commonly p-believed (rather than commonly known).

Theorem 2.5 (Monderer and Samet, [32]) *Suppose that $(W, \mathcal{F})$ is a finite probability space and μ is a probability measure on $(W, \mathcal{F})$. Further, suppose that there is a common prior: for each $i \in \mathcal{A}$, $\mu_i = \mu$ is i's prior probability measure. For all real numbers $r_1, \ldots, r_n \in [0, 1]$ and events $E \subseteq W$, if $C^p(\bigcap_{i\in\mathcal{A}}[E_i = r_i]) \neq \emptyset$, then for all $i, j \in \mathcal{A}$, $|r_i - r_j| < 1 - p$.*

The original result in [32] had a different bound on the differences between the posteriors. The bound used in the above version was provided by Zvika Neeman [36], who also showed that this bound cannot be improved. In the special case when $p = 1$ and the posteriors are commonly p-believed, then the differences between the posteriors is $1 - p = 0$ (i.e., the posteriors must be the same). Thus, the above theorem generalizes Aumann's Theorem. [7]

[7] Note that common 1-belief is not necessarily equivalent to common knowledge. The two may come apart when there are states in an information cell that are assigned probability 0. Dealing with probability 0 events is a very interesting topic, although I do not focus on it in this paper.

Example 2.6 To further illustrate this theorem, consider the model depicted in Example 2.3, and let $E = H_1 \cap H_2$. Then, the prior probability of E is $\mu_a(E) = \mu_b(E) = 0.72$, while the posterior probability of E is $\mu_{a,w_1}(E) = 0.9$ and $\mu_{b,w_1}(E) = 0.8$. Furthermore, we can easily see that $[E_a = 0.9] = \{w_1, w_3\}$ and $[E_b = 0.8] = \{w_1, w_2\}$. In this model, $C^{0.8}_{\{a,b\}}([E_a = 0.9] \cap [E_b = 0.8]) = C^{0.8}_{\{a,b\}}(\{w_1\}) = \{w_1\}$. Indeed, we have that $|\mu_{a,w_1}(E) - \mu_{b,w_1}(E)| = 0.9 - 0.8 = 0.1 < 1 - 0.8$, as expected according to Theorem 2.5.

3 Dynamics

Aumann's theorem raises a natural question: How, exactly, do the agents come to agree on their posteriors? *If* there is common knowledge of the posteriors of an event, then these posteriors must be equal. However, in general, common knowledge will be achieved only after the agents gather and exchange information.

3.1 Processes of Inquiry and Information Exchange

There are two main stages to this dynamic process of information exchange. Starting with a common probability space $(W, \mathcal{F}, \mu)$:

1. Each agent i receives private information represented by a partition Π_i on W. The idea is that the agents each ask questions, or perform experiments, the accumulated answers (or results) of which result in a partition over the set of states. It is assumed that the agents know the question that the other agents ask (i.e., each agent i knows the partitions Π_j for each agent J), but not the answers that they receive (i.e., for each agent i and state w, i does not, in general, know the event $\Pi_j(w)$).

2. The agents exchange some information based on their current knowledge and beliefs.

> The information exchanged could be the answer to the question that they asked [43], the posteriors of some fixed event E [19, 3], whether the agents would accept a bet based on E [41], or, more generally, the output of some discrete random variable [33, 38]. Based on the information received from the other agent(s), the agents (further) *refine* their initial partition. This may result in new information that the agents can exchange.

This dynamic process of information exchange *converges* when the agents cannot learn anything new. That is, there is convergence when no further exchange of information results in any change to the agents' knowledge or beliefs.

Many interesting questions can be asked about information exchange in general, beyond the issue of convergence. For instance, what is commonly known (or commonly believed) among the agents after a process converges? Or, have the agents reached *consensus* about some previous disagreement? For instance, is there agreement about the posterior probability of some fixed event, or about whether to accept a bet based on some fixed event?

3.2 Announcing Posterior Probabilities and Model Update

Geanakoplos and Polemarchakis develop a dynamic characterization of Aumann's theorem by defining a general process of information exchange among agents with different information along the lines of the process described above [19]. [8] The authors show that when the agents exchange their current probability for some fixed event E, convergence is reached after a finite number of exchanges. Furthermore, when the process of information exchange converges, there is common knowledge of the current probabilities of E. Aumann's theorem, then, implies that these probabilities must be the same. Thus, the exchange of information results not only in common knowledge, but also in *consensus* about the probability of E.

[8]See [10, 11] for a discussion of dynamic modal logics of beliefs motivated by this work.

This result is a dynamic characterization of Aumann's agreeing to disagree theorem. To make this statement more precise, we give some formal definitions.

An **epistemic-probability model** is a tuple $\mathcal{M} = \langle W, \{\Pi_i\}_{i\in\mathcal{A}}, \mu\rangle$, where $W \neq \emptyset$ is a finite set of states, for each $i \in \mathcal{A}$, Π_i is a partition on W, and μ is a probability measure on[9] W (μ is the **common prior**). In the remainder of this paper, we will mostly be interested in pairs $(\mathcal{M}, w)$, called **pointed epistemic-probability models**, where $\mathcal{M}$ is an epistemic probability model and w is a state from $\mathcal{M}$. As the agents exchange information, this model is *transformed* into a new one – a common idea in the literature on information dynamics [5, 44, 37].

The type of transformation studied by Geanakoplos and Polemarchakis and subsequent researchers involves a *refinement* of the agents' partitions. A partition Π is a **refinement** of Π', provided that for all $X \in \Pi'$, there are $Y_1, \ldots, Y_k \in \Pi_i$ such that $X = Y_1 \cup \cdots \cup Y_k$. We also say that Π' is a **coarsening** of Π. Given an event $E \subseteq W$ and a partition Π_i of W for agent i, let Π_i^E be the coarsest partition that refines Π_i and for all $X \in \Pi_i^E$, either $X \subseteq E$ or $E \subseteq X$. If Π_i represents i's current information, or current question, then Π_i^E represents i's information after updating with E. Given a sequence of events $\overline{E} = (E_1, \ldots, E_n)$, let $\mathcal{M}^{\overline{E}} = \langle W, \{\Pi_i^{E_i}\}_{i\in\mathcal{A}}, \mu\rangle$. Then, $\mathcal{M}^{\overline{E}}$ describes the agents' knowledge and beliefs after every $i \in \mathcal{A}$ updates their partition with the event E_i. Using this notation, we can describe the information exchanges used in a dynamic characterization of Aumann's Theorem.

Suppose that $\mathcal{M} = \langle W, \{\Pi_i\}_{i\in\mathcal{A}}, \mu\rangle$ is an epistemic-probability model, $F \subseteq W$ is an event, and $v \in W$ is the "actual world". For each $i \in \mathcal{A}$, let $r_i = \mu_{i,v}(F)$. Then, the information exchanged at state v by all the agents is the sequence $\overline{E} = (E_1, \ldots, E_n)$, where

[9] Since the set set of states W is finite, we can assume that the σ-algebra is $\wp(W)$. Furthermore, it is often convenient, when the set of states is finite, to view the probability measure as a function $\mu : W \to [0, 1]$ with $\sum_{w\in W} \mu(w) = 1$. Then, the probability of any set $X \subseteq W$ is $\mu(X) = \sum_{w\in X} \mu(w)$.

for each $i \in \mathcal{A}$, $E_i = \bigcap_{j\in\mathcal{A}, j\neq i}\{w \mid \mu_{j,w}(F) = r_j\}$. That is, each agent i learns the other agents' current probability of the event F. The agents "indirectly" share their information by repeatedly reporting their posteriors of the event F at state v. This sequence of information exchanges is described by a sequence of model transformations:

$$(\mathcal{M}_0, v) \overset{\overline{E}_0}{\Longrightarrow} (\mathcal{M}_1, v) \overset{\overline{E}_1}{\Longrightarrow} (\mathcal{M}_2, v) \overset{\overline{E}_2}{\Longrightarrow} \cdots \overset{\overline{E}_{m-1}}{\Longrightarrow} (\mathcal{M}_m, v),$$

where for $k = 0, \ldots, m-1$, $\overline{E}_k$ is the information about the agents' posterior probabilities of F from model $\mathcal{M}_k$ at state v, and for $k = 0, \ldots, m-1$, $\mathcal{M}_{k+1} = \mathcal{M}_k^{\overline{E}_k}$.

A sequence of model transformations is said to *converge* when no further information exchanges change the model. That is, a sequence converges when there is a model $\mathcal{M}_m$ such that $\mathcal{M}_m^{\overline{E}_m} = \mathcal{M}$. Geanakoplos and Polemarchakis prove that for any (finite) epistemic probability model $\mathcal{M}$, event F and state v from $\mathcal{M}$, the sequence of model transformations generated by F and $(\mathcal{M}, v)$ is guaranteed to converge. For an epistemic-probability model $\mathcal{M}$, event F and state v, let $\mathcal{M}[F, v]$ be the epistemic-probability model that is generated after a sequence of information exchanges about the posteriors of F at v converges.

Since the posteriors of the event F in $\mathcal{M}[F, v]$ are common knowledge, Aumann's theorem then guarantees that the posteriors must be the same in $\mathcal{M}[F, v]$. [10]

The convergence of the agents' opinions about an event F depends on the assumption that there is a common prior *and* that the agents know the other agents' possible information cells. I.e., the agents know which question was asked, or experiment performed, by the other agents. If an agent i announces that her posterior for F is r, then the other agents can rule out any state in which conditioning the common prior on i's information at that state does not assign probability r to F. Thus, even if all the agents announce the same posterior for

[10] Actually, one can see that the posteriors for the event F must be equal without appealing to Aumann's Theorem. That is, it can be shown that the agents' posteriors for F must be equal in $\mathcal{M}[F, v]$, cf. [11].

an event, they may still learn something. This is illustrated by the following example.

Example 3.1 Recall the situation described in Example 2.3. Two (fair) coins are flipped and placed in two different boxes. Ann observes the coin in the first box and Bob observes the coin in the second box. Let $\mathcal{M}_0$ represent Ann and Bob's beliefs before they observe the coins. Assume that both Ann and Bob initially believe that the coins are fair. Then, $\mathcal{M}_0 = \langle W, \{\Pi_i\}_{i\in\{a,b\}}, \mu\rangle$, where $\Pi_a = \Pi_b = \{W\}$ and for $j = 1, \ldots, 4$, $\mu(w_j) = \frac{1}{4}$. We are interested in 4 events: $H_1 = \{w_1, w_3\}$ (the coin in box 1 is lying heads up); $H_2 = \{w_1, w_2\}$ (the coin in box 2 is lying heads up); $T_1 = \{w_2, w_4\}$ (the coin in box 1 is lying tails up); and $T_2 = \{w_3, w_4\}$ (the coin is box 2 is lying tails up).

The first "learning event" is that Ann observes the coin in box 1 and Bob observes the coin in box 2. The agents' observations transform the model $\mathcal{M}_0$ into $\mathcal{M}_1 = \langle W, \{\Pi_a^{H_1}, \Pi_b^{H_2}\}, \mu\rangle$, in which Ann knows whether coin 1 is heads up or tails up $(K_a(H_1) \cup K_a(T_1))$ and Bob knows whether coin 2 is heads up or tails up $(K_b(H_2) \vee K_b(T_2))$. We are assuming here that the observations are "out in the open", so that Ann knows that Bob is observing the coin in box 2 and Bob knows that Ann is observing the coin in box 1. [11] After the agents observe the coins in their respective boxes, they exchange information. Suppose that both coins are lying heads up. There are two situations to contrast.

In the first situation, the agents exchange their current probabilities of the event $E = H_1 \cap H_2$. Note that the agents' posterior probabilities are the same for E: $\mu_{a,w_1}(E) = \mu_{b,w_1}(E) = \frac{1}{2}$. However, the agents do learn something from this exchange of information. Since $\mu_{a,w_2}(E) = \mu_{a,w_4}(E) = 0$ and $\mu_{b,w_2}(E) = \mu_{b,w_4}(E) = 0$, it becomes common knowledge that both coins are lying heads up. Thus, the agents become certain of event E (i.e., both Ann and Bob assign probability 1 to E after exchanging their information). Indeed, note

[11] This assumption can be dropped using ideas from *dynamic epistemic logic*: see [44, 37] for a discussion.

that we get the same result even if the coins are biased (so that the initial probability is not uniform). In general, the agents will become certain of the event E in any model in which $\mu(E \mid \Pi_a^{H_1}(w_1)) \neq \mu(E \mid \Pi_a^{T_1}(w_1))$ or $\mu(E \mid \Pi_b^{H_2}(w_1)) \neq \mu(E \mid \Pi_b^{T_2}(w_1))$. So, for example, the agents will become certain of E after exchanging their probabilities for E in the model that was described in Example 2.3.

In the second situation, the agents exchange their current probabilities of the event $F = (H_1 \cap H_2) \cup (T_1 \cap T_2)$. That is, F is the event "the coins in the boxes match". As in the first situation, both Ann and Bob assign probability 1/2 to this event: $\mu_{a,w_1}(F) = \mu_{b,w_1}(F) = \frac{1}{2}$. However, in this case, exchanging their current probabilities of F does not result in any changes to the model. This is because $\mu(F \mid H_1) = \mu(F \mid T_1) = \mu(F \mid H_2) = \mu(F \mid T_2) = \frac{1}{2}$. Note that the same phenomenon occurs if we assume that the coins are perfectly correlated (so that, in $\mathcal{M}_0$, we have $\mu(w_1) = \mu(w_4) = \frac{1}{2}$ and $\mu(w_2) = \mu(w_3) = 0$), and the agents exchange their current probabilities of E.

In both cases, the agents' initial observations result in the same probability of the event in question. In the first situation, sharing these probabilities causes the agents to become certain of the event. In the second situation, however, sharing the probabilities does not result in any changes to the agents' knowledge or beliefs. The crucial difference is the agents' higher-order information—i.e., what the agents know about the other agent's knowledge and beliefs. It is the higher-order information that drives the changes in the agents' beliefs about an event E. Suppose that agent j announces that the probability of some event E is q. Given j's current information partition and this announcement, the other agents can partition the set of states into those in which j's announcement is true (i.e., j's current probability of E is q) and those in which the announcement is false. The agents then use this new information to refine their partitions.

3.3 Information Exchanges for Common p-Belief

What type of information exchanges should be used in a dynamic characterization of Monderer and Samet's generalization of Aumann's Theorem? That is, for an event F and an epistemic-probability model $\mathcal{M}$, what dynamic process will converge on a model in which there is common p-belief of the agents' current probability of F?

I start by considering a specific example. The goal is to explain how we account for the model from Example 2.3. That is, if the above-discussed two situations from Example 3.1 represent two extreme points on a scale of different types of dynamic processes of information exchanges, then is there an intermediate type of process that results in the model from Example 2.3?

One way to answer this question is to use a "softer" type of belief change to represent the agents' initial observation. The idea is that situations such as the one described in the model from Example 2.3 arise because the agents receive *uncertain evidence* about the coins. The standard way to update beliefs given uncertain evidence is to use an approach pioneered by Richard Jeffrey [26]. Here are some relevant details for our purpose.

In the dynamic process of information exchange described in this section, initially, each agent receives some evidence. In the situations described above, each agent receives different evidence (Ann receives evidence that the coin in box 1 is heads up, while Bob receives evidence that the coin in box 2 is heads up), but both are certain that the evidence is correct. Thus, after taking this evidence into account, Ann is *certain* that the coin is lying heads up in box 1, and Bob is *certain* that the coin is lying heads up in box 2. Jeffrey is interested in situations in which the agents *shift* their probabilities in response to some evidence, though they do not necessarily become *certain* that the evidence is correct. For instance, Ann and Bob may observe the coins in the boxes under a dim light. For a set of states W, Jeffrey calls a sequence $(Q_1 : r_1, \ldots, Q_m : r_m)$ a *learning experience*, where $\{Q_i\}$ is a partition on W, for each $j = 1, \ldots, m$, $r_j \in [0, 1]$, and $\sum_j r_j = 1$. The intuition is that each r_j is the new posterior of the partition cell Q_j.

Then here is the new update rule. Given a probability μ on W, the new probability given a learning experience $(Q_1 : r_1, \ldots, Q_m : r_m)$, denoted by μ_{new}, of an event E is:

$$\mu_{new}(E) = \sum_{i=1}^{n} r_i \mu(E \mid Q_i).$$

It is not hard to see that Jeffrey updating is a generalization of conditioning: The above equation reduces to standard conditioning with a learning sequence in which exactly one element has a weight of 1, while all the others have weights of 0.

Starting with the initial model $\mathcal{M}_0$, suppose that Ann's learning experience is $(H_1 : 0.9, T_1 : 0.1)$ and Bob's is $(H_1 : 0.8, T_1 : 0.2)$. Then, the model from Example 2.3 results from Ann and Bob applying the Jeffrey update rule *and assuming that these learning experiences are common knowledge.* This type of multi-agent Jeffrey update can be represented using the probabilistic event models and update rule from [45]. [12]

However, this is not the end of the story. As noted above, if Ann and Bob share their probabilities of the event E in the model from Example 2.3, then they will become certain in the event E. This implies that the model from Example 2.3 can be maintained as long as Ann and Bob do not communicate. A complete analysis should allow exchanges of information to be intermixed with learning about an event, either by conditioning or by Jeffrey update.

We conclude with a brief discussion of some of the issues that arise when moving from

[12]The details of the probabilistic event models and update rules are beyond the scope of this article. One may consult [45] for a discussion. An interesting question for future research is to explore the types of epistemic-probability models that arise from to look at the long-run dynamics of applying of updating an epistemic-probability model with a sequence of probabilistic event models.

common knowledge to common p-belief. I will not arrive at a complete solution, but we will see what I take to be the main challenge. Again, it helps to illustrate our theme with a specific example, this time, a scenario based on the proof of Proposition 1 in [19].

Example 3.2 Suppose that $\mathcal{M}_0 = \langle W, \{\Pi_a, \Pi_b\}, \mu_0 \rangle$, where

- $W = \{w_1, w_2, w_3, w_4, w_5, w_6, w_7, w_8, w_9\}$;
- $\Pi_a = \{\{w_1, w_2, w_3\}, \{w_4, w_5, w_6\}, \{w_7, w_8, w_9\}\}$ and

  $\Pi_b = \{\{w_1, w_2, w_3, w_4\}, \{w_5, w_6, w_7, w_8\}, \{w_9\}\}$;
- for all $j = 1, \ldots, 9$, $\mu(w_j) = \frac{1}{9}$.

Moreover, suppose that the event that we are interested in is $E = \{w_1, w_5, w_9\}$. The relevant model (without the probabilities) is pictured as follows – where agent a's partition is the solid line; agent b's partition is the dashed line; and E is the gray ellipse:

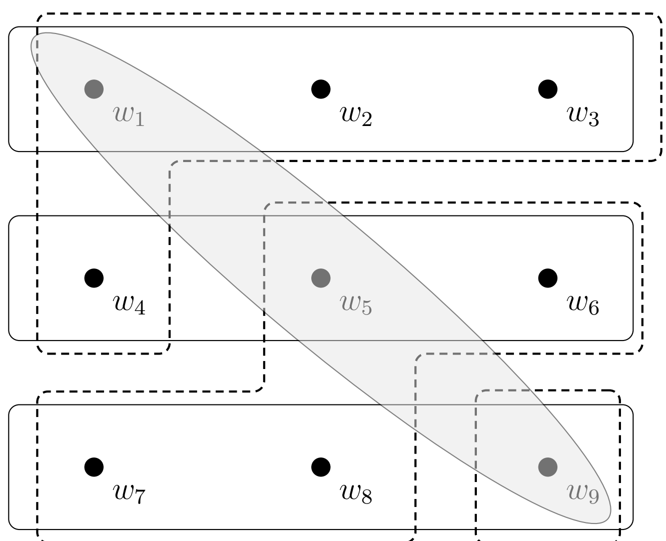

Starting with this initial model, the process of information exchanges that was described in this section proceeds as follows:

- Round 1: Ann announces that her probability of E is $\frac{1}{3}$ and Bob announces that his is $\frac{1}{4}$. After the agents share these probabilities, the probabilities in E do not change. However, Ann's partition is refined so that $\Pi_a(w_7) = \{w_7, w_8\}$ and $\Pi_a(w_9) = \{w_9\}$. Bob's partition is not changed either. Thus, this refinement does not change Ann's and Bob's probability for E.

- Round 2: Ann announces that her probability of E is $\frac{1}{3}$ and Bob announces that his is $\frac{1}{4}$. As a result of this announcement, Bob refines his partition so that $\Pi_b(w_5) = \{w_5, w_6\}$ and $\Pi_b(w_7) = \{w_7, w_8\}$. Once again, this refinement does not change Ann's and Bob's probability for E.

- Round 3: Ann announces that her probability of E is $\frac{1}{3}$ and Bob announces that his is $\frac{1}{4}$. As a result of this announcement, Ann further refines her partition so that $\Pi_a(w_4) = \{w_4\}$ and $\Pi_a(w_5) = \{w_5, w_6\}$. This refinement does not change Ann's and Bob's probability for E.

- Round 4: Ann announces that her probability of E is $\frac{1}{3}$ and Bob announces that his is $\frac{1}{4}$. As a result of this announcement, Bob further refines his partition so that $\Pi_b(w_4) = \{w_4\}$ and $\Pi_b(w_1) = \{w_1, w_2, w_3\}$. Now Ann and Bob both assign probability $\frac{1}{3}$ to the event E.

After round 4, the process converges since there is no new information that can be shared. Note that, as the reader is invited to check, $w_1 \in C^{\frac{1}{4}}(E)$, and this does not change during the information exchange – until the process converges, when we have $w \in C^{\frac{1}{3}}(E)$.

So, we have identified an obvious further question at this stage. *What type of information updates will increase the p so that there is common p-belief of the event E?* If we want to increase the common p-belief of the event E, then there needs to be a more drastic change to Ann and Bob's probabilities during the process of information exchange.

One natural approach is to assume that the agents receive information that changes the relative likelihood of their information cells. For instance, suppose that in round 1, the agents learn that Ann's partition cell $\{w_1, w_2, w_3\}$ is twice as likely as her other cells. This can be represented by a Jeffrey update along the lines of the ones used in the previous section. The problem with this approach is that this type of change in the probabilities does not shift the agents' posteriors of the event E. Indeed, a defining feature of Jeffrey updating is that an agent's conditional probabilities are left unchanged [14].

So, clearly, more is to be done in analyzing the agent dynamics that we started with. In particular, one important clue from our example is that increasing the common p-belief of the event E is closely tied to learning about E. That is, the type of updates that will increase common p-belief of E involve simply raising the (prior) probability of E.

We leave our discussion of probabilistic-epistemic update dynamics at this stage, hoping to have illustrated the main issues and challenges involved in analyzing Monderer-Samet type results and related scenarios. A complete analysis of the dynamic processes that result from combining learning directly about an event while also allowing agents to share their probabilities of that event will be left for another paper.

4 Logical Considerations

One obvious approach to studying the fine-structure of the scenarios discussed in the preceding sections is the introduction of logical languages and their associated techniques for analyzing reasoning by, or about, agents.

4.1 Epistemic Probabilistic Base Logic

Many different logical frameworks have been proposed to reason about epistemic-probability models [17, 16, 25, 46, 45]. One may consult [13] for a comparative discussion of these different logical frameworks. In this section, I highlight the formal language from [16, 17], being the most natural language to formalize the agreement theorems discussed in this paper.

Suppose that At is a countable set of atomic propositions, and then, let $\mathcal{L}$ be the smallest set of formulas generated by the following grammar:

$$p \mid \neg\varphi \mid \varphi \wedge \psi \mid \mathbf{K}_i\varphi \mid \mathbf{C}\varphi \mid p_1\mathbf{P}_i(\varphi_1) + \cdots + p_k\mathbf{P}_i(\varphi_k) \geq q$$

where $i \in \mathcal{A}$, $p \in \mathsf{At}$ and $p_1, \ldots, p_k, q$ are rational numbers in $[0, 1]$. [13] The additional Boolean connectives $(\vee, \rightarrow, \leftrightarrow)$ are defined as usual. This language can be interpreted on our earlier epistemic-probability models, when we add a valuation function to interpret the atomic propositions. Suppose that $\mathcal{M} = \langle W, \{\Pi_i\}_{i\in\mathcal{A}}, \{\mu_i\}, V\rangle$ is an epistemic-probability model where $V : \mathsf{At} \rightarrow \wp(W)$ is a valuation function. Truth of formulas $\varphi \in \mathcal{L}$ at a pointed epistemic-probability model, denoted $\mathcal{M}, w \models \varphi$, is defined by recursion:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$
- $\mathcal{M}, w \models \neg\varphi$ iff $\mathcal{M}, w \not\models \varphi$
- $\mathcal{M}, w \models \varphi \wedge \psi$ iff $\mathcal{M}, w \models \varphi$ and $\mathcal{M}, w \models \psi$
- $\mathcal{M}, w \models \mathbf{K}_i\varphi$ iff $w \in K_i([\![\varphi]\!]_{\mathcal{M}})$
- $\mathcal{M}, w \models \mathbf{C}\varphi$ iff $w \in C([\![\varphi]\!]_{\mathcal{M}})$
- $\mathcal{M}, w \models p_1\mathbf{P}_i(\varphi) + \cdots + p_k\mathbf{P}_i(\varphi) \geq q$ iff $\sum_{j=1}^{k} p_j\mu_i([\![\varphi_j]\!]_{\mathcal{M}} \mid \Pi_i(w)) \geq q$.

[13] Here the restriction to rational numbers ensures that the language is countable.

It is straightforward to show that we can now also define formulas $p_1\mathbf{P}_i(\varphi) + \cdots + p_k\mathbf{P}_i(\varphi) \leq q$ and $p_1\mathbf{P}_i(\varphi)+\cdots+p_k\mathbf{P}_i(\varphi) = q$, with the obvious interpretations. Furthermore, for rational numbers $p, q \in [0,1]$, it is useful to define $p = q$ to be the formula $p\mathbf{P}_i(\top) + -q\mathbf{P}_i(\top) = 0$ (where $\top$ is the formula $p \vee \neg p$).

Excursion. In much of the literature on epistemic-probability modal logic, a more general semantics is used, with different probability measures assigned to each state. A *general epistemic-probability model* is then a structure $\langle W, \{R_i\}_{i\in\mathcal{A}}, \{\pi_i\}_{i\in\mathcal{A}}, V\rangle$, where $W \neq \emptyset$; for each $i \in \mathcal{A}$, $R_i \subseteq W \times W$ is an equivalence relation and $V : \mathsf{At} \to \wp(W)$ is a valuation function. For each $i \in \mathcal{A}$, π_i assigns a probability function to each state: For each $i \in \mathcal{A}$, $\pi_i : W \to (W \rightharpoonup [0,1])$ such that[14] for all $w \in W$,

$$\sum_{v\in dom(\pi_i(w))} \pi_i(w)(v) = 1$$

Then, the interpretation of the earlier basic probability formula is:

> $\mathcal{M}, w \models p_1\mathbf{P}_i(\varphi) + \cdots + p_k\mathbf{P}_i(\varphi) \geq q$ iff $\sum_{i=1}^{k} p_i P_i(w)([\![\varphi]\!]_{\mathcal{M}}) \geq q$,
>
> where for any set $X \subseteq W$, $P_i(w)(X) = \sum_{v\in X} P_i(w)(v)$.

There are two important differences between this model and the epistemic-probability models used in this paper. The first is that since agents' probability functions may vary from state to state, the agents do not necessarily know their probability function. The second difference is that the probability functions used to interpret the probability formulas are not conditioned on the agents' current information. The approach taken in this section ensures that for all formulas φ, rational numbers $p \in [0,1]$ and epistemic-probability models $\mathcal{M}$,

[14]The notation $\pi : W \rightharpoonup [0,1]$ means that π is partial function. Let $dom(\pi)$ be the domain of π. This assumption does not play an important role in this paper, though it is is needed in order to deal with states that should be ignored when assigning probability to an event. For instance, after φ is publicly announced, the states that do not satisfy φ should not figure into any probabilistic calculations.

$B_i^p([\![\varphi]\!]_{\mathcal{M}}) = [\![\mathbf{P}_i(\varphi) \geq p]\!]_{\mathcal{M}}$. In both frameworks, conditional probability is definable: For formulas φ and ψ, let $\mathbf{P}_i(\varphi \mid \psi) \geq q$ be defined as $\mathbf{P}_i(\varphi \wedge \psi) + -q\mathbf{P}_i(\psi) = 0$. [15]

In the above language $\mathcal{L}$, the agreeing to disagree theorem is expressible as follows:

$$\mathbf{C}(\mathbf{P}_a(\varphi) = p \wedge \mathbf{P}_b(\varphi) = q) \rightarrow p = q$$

Aumann's result says that the above formula is valid (true at all states) on any model in which there is a common prior. Completely formalizing this theorem would require a formula that expresses the fact that there is a common prior. One may consult [22, 23, 18] for approaches to this challenging problem, and [11] for a complete logical analysis of this approach to formalizing Aumann's Theorem.

4.2 Extension to Common p-Belief

In order to extend the above logical analysis to Monderer and Samet's generalization of Aumann's theorem, we must go beyond the above formalism, and add a common p-belief operator to the language. To do this we add, for each rational number $p \in [0, 1]$, a modal operator $\mathbf{C}^p$ to the language $\mathcal{L}$. The interpretation of these formulas is straightforward:

$$\mathcal{M}, w \models \mathbf{C}^p\varphi \text{ iff } w \in C^p([\![\varphi]\!]_{\mathcal{M}})$$

In a language with common p-belief operators defined in this way, the Monderer and Samet agreement theorem can be expressed – where the same comment as above applies about the further task of formalizing the common prior:

$$\mathbf{C}^p(\mathbf{P}_a(\varphi) = p \wedge \mathbf{P}_b(\varphi) = q) \rightarrow ((a - b < 1 - p) \wedge (b - a < 1 - p))$$

[15] One further difficulty with expressing $B_i^p([\![\varphi]\!]_{\mathcal{M}})$ precisely is in finding formulas that express agent i's information at each state. See [11] for an elegant solution to this problem when formalizing Aumann's theorem in this more general setting, as well as other relevant results: see below.

To the best of my knowledge, finding a complete logic for the language $\mathcal{L}$ with common p-belief operators is still an open problem (cf. [16] for a brief discussion).

4.3 Dynamic Logics of Probabilistic Epistemic Update

In this richer setting, the processes of information exchange discussed in Section 3 can also be subjected to a logical analysis using existing dynamic epistemic-probability logics [45, 12].

In particular, announcements of the agents' current posteriors can be described in the language using a *public announcement operator*. For each $\varphi \in \mathcal{L}$ (with common p-belief operators), let $[\varphi]\psi$ be a formula with the intended interpretation is that "after φ is publicly announced, ψ is true". Suppose that $\mathcal{M} = \langle W, \{\Pi_i\}_{i\in\mathcal{A}}, \{\mu_i\}_{i\in\mathcal{A}}, V\rangle$ is an epistemic-probability model. The model $\mathcal{M}^\varphi$ is the structure $\langle W^\varphi, \{\Pi_i^\varphi\}_{i\in\mathcal{A}}, \{\mu_i^\varphi\}_{i\in\mathcal{A}}, V^\varphi\rangle$, where $W^\varphi = W$, for all $i \in \mathcal{A}$, $\Pi^\varphi = \Pi_i^{[\![\varphi]\!]_\mathcal{M}}, V^\varphi = V$, and for all $\mu_i^\varphi(\cdot) = \mu_i(\cdot \mid [\![\varphi]\!]_\mathcal{M})$. This assumes that only formulas that are assigned non-zero probability are announced (cf. [12] for an illuminating discussion of this). Truth of public announcement formulas is defined as follows:

$$\mathcal{M}, w \models [\varphi]\psi \text{ iff, if } \mathcal{M}, w \models \varphi, \text{ then } \mathcal{M}^\varphi, w \models \psi.$$

The above formula assumes that only true statements can be announced. This assumption can be dropped (see [12] for a discussion), although this extra generality is not needed for the purposes of this section. As a basic application, the key single step in the process of information exchanges discussed in Section 3 is represented by the dynamic formula $[\mathbf{P}_a(\varphi) = p \wedge \mathbf{P}_a(\varphi) = q]\psi$. A complete logic for the dynamic processes of information exchange that underlie Aumann's Theorem can be found in [11].

4.4 Toward a Dynamic Logic for Common p-Belief

I conclude this section with a brief discussion of a key logical issue for the Monderer-Samet setting. *What are the recursion axioms for the common p-belief operators?*

A fundamental insight in the dynamic epistemic logic literature [44] is that there is a strong logical relationship between what is true before and after an announcement in the form of so-called *recursion axioms*. The recursion axioms describe the effect of an announcement (or indeed, any relevant sort of information-producing event) in terms of what is true before the announcement. For instance, the reader can easily check that the following formula is valid (i.e., true at all states) in any epistemic-probability model:

$$[\psi](\mathbf{P}_i(\varphi) \geq q) \leftrightarrow (\varphi \rightarrow \mathbf{P}_i([\psi]\varphi \mid \psi) \geq q)$$

When a logical language becomes strictly more expressive by adding public announcement operators, recursion axioms are not always available. Adding public announcement operators to epistemic logic (even without probability operators) with common knowledge is such a case. It was shown by [5] that the language of epistemic logic with common knowledge and public announcements is more expressive than epistemic logic with common knowledge. Therefore, a recursion axiom for formulas of the form $[\varphi]\mathbf{C}\psi$ does not exist.

Nonetheless, a recursion axiom-style analysis is still possible [42]. The key idea is to introduce a *conditional common knowledge* operator $\mathbf{C}(\varphi, \psi)$ saying that φ is true at all states that are reachable by finite paths of the agents' accessible relations *going through states satisfying* φ. There are recursion axioms in this more expressive language:

$$[!\psi]\mathbf{C}\varphi \leftrightarrow (\psi \rightarrow \mathbf{C}([!\psi]\varphi, \psi).$$

A natural question is whether this approach can be extended to common p-belief operators. The answer is that we can, but the definition of conditional common p-belief operators requires some care. We first define a conditional p-belief operator: $B_i^p(E,F) = \{w \mid \mu_i(E \mid \Pi_i(w) \cap F) \geq p\}$. Then, conditional common p-belief is defined as follows:

$$C^p(E,F) = \{w \mid \text{there is a } U \subseteq W \text{ such that } w \in U \cap F,\ U \subseteq B^p(U,F), \text{ and } U \subseteq B^p(E,F)\}$$

Once this definition is in place, it is not hard to see that the following formula is valid:

$$[!\psi]\mathbf{C}^p\varphi \leftrightarrow (\psi \to \mathbf{C}^p([!\psi]\varphi, \psi).$$

This section has merely given a taste of the insights that can be gained by using existing dynamic epistemic probability logics to reason about agreement theorems and the dynamic processes of information exchange that leads to agreement. A lot more lies ahead of us. [16]

5 Conclusion

This article has raised more questions than it has answered, and its main contribution is conceptual rather than technical. It has argued for a more fine-grained analysis of the dynamic process of information exchange that underlies agreeing to disagree type results. Our main goal was to develop a model that describes the evolution of the agents' knowledge and beliefs as they interact with each other and the environment. At each moment, the agents receive an input. This input might be an announcement from one or more of the agents in the group (about some property of the agent's current beliefs, such as the probability of some fixed event), or, more generally, some signal revealing the value of some unknown random variable. Based on this input, the agents update their beliefs.

We have shown how concrete results from game theory sharpen our intuitions and raise concrete issues in thinking about the preceding phenomena through the lense of probabilistic

[16] For further illustrations, the reader may consult [10] in a framework without probabilities.

update steps. We also hope to have shown, or at least made it plausible to the reader, how the logical framework discussed in Section 4 is a rich and flexible tool that can be used to study many dynamic processes of interactive social belief change.

References

[1] Robert Aumann. Correlated equilibrium as an expression of bayesian rationality. *Econometrica*, 55:1 – 18, 1987.

[2] Robert Aumann and Adam Brandenburger. Epistemic conditions for the Nash equilibrium. *Econometrica*, 63(5):1161 – 1180, 1995.

[3] Rorbert Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236 – 1239, 1976.

[4] Michael Bacharach. Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, 37:167 – 190, 1985.

[5] Alexandru Baltag, Larry Moss, and Slawomir Solecki. The logic of common knowledge, public announcements and private suspicions. In I. Gilboa, editor, *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK 98)*, pages 43 – 56, 1998.

[6] Jon Barwise. Three views of common knowledge. In *Proceedings of Theoretical Aspects of Rationality and Knowledge (TARK)*, 1987.

[7] Giacomo Bonanno and Klaus Nehring. Agreeing to disagree: A survey. Available at `http://www.econ.ucdavis.edu/faculty/bonanno/wpapers.htm`, 1997.

[8] Johanthan A.K. Cave. Learning to agree. *Economics Letters*, 12:147 – 152, 1983.

[9] Robin Cubitt and Robert Sugden. Common knowledge, salience and convention: A reconstruction. *Economics and Philosophy*, 19(2):175 – 210, 2003.

[10] Cédric Dégremont and Olivier Roy. Agreement theorems in dynamic-epistemic logic. *Journal of Philosophical Logic*, 41(4):735 – 764, 2012.

[11] Lorenz Demey. Agreeing to disagree in probabilistic dynamic epistemic logic. *Synthese*, 191(3):409 – 438, 2014.

[12] Lorenz Demey and Barteld Kooi. Logic and probabilistic update. In Alexandru Baltag and Sonja Smets, editors, *Johan van Benthem on Logic and Information Dynamics*, volume 5 of *Outstanding Contributions to Logic*, pages 381–404. Springer International Publishing, 2014.

[13] Lorenz Demey, Barteld Kooi, and Joshua Sack. Logic and probability. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Fall 2014 edition, 2014.

[14] Persi Diaconis and Sandy Zabell. Updating subjective probability. *Journal of the American Statistical Association*, 77:822 – 830, 1982.

[15] Ronald Fagin, Joseph Halpern, Yoram Moses, and Moshe Vardi. *Reasoning about Knowledge*. The MIT Press, 1995.

[16] Ronald Fagin and Joseph Halpern. Reasoning about knowledge and probability. *Journal of the ACM*, 41(2):340 – 367, 1994.

[17] Ronald Fagin, Joseph Halpern, and Nimrod Megiddo. A logic for reasoning about probabilities. *Information and Computation*, 87(1):78 – 128, 1990.

[18] Yossi Feinberg. Characterizing common priors in the form of posteriors. *Journal of Economic Theory*, 91(2):127 – 179, 2000.

[19] John Geanakoplos and Herakles Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192 – 200, 1982.

[20] Joseph Halpern and Ronald Fagin. Reasoning about knowledge and probability. *Journal of the ACM*, 41(2):340 – 367, 1994.

[21] Joseph Halpern and Yoram Moses. Characterizing solution concepts in terms of common knowledge of rationality. Manuscript, `https://www.cs.cornell.edu/home/halpern/papers/kbeq2.pdf`, 2014.

[22] Joseph Halpern. Characterizing the common prior assumption. *Journal of Economic Theory*, 106(2):316 – 355, 2002.

[23] Aviad Heifetz. The positive foundation of the common prior assumption. *Games and Economic Behavior*, 56(1):105 – 120, 2006.

[24] Aviad Heifetz and Philipe Mongin. Probability logic for type spaces. *Games and Economic Behavior*, 35(1):31 – 53, 2001.

[25] Aviad Heifetz and Philippe Mongin. Probability logic for type spaces. *Games and Economic Behavior*, 35:31 – 53, 2001.

[26] Richard Jeffrey. *The Logic of Decision*. McGraw-Hill, 1983.

[27] Atsushi Kajii and Stephen Morris. Common p-belief: The general case. *Games and Economic Behavior*, 18:73 – 82, 1997.

[28] Thomas Kelly. The epistemic significance of disagreement. *Oxford Studies of Epistemology*, 105(167 - 196), 2005.

[29] Ehud Lehrer and Dov Samet. Agreeing to agree. *Theoretical Economics*, 6:269 – 287, 2011.

[30] David Lewis. *Convention*. Cambridge University Press, 1969.

[31] Paul Milgrom and Nancy Stokey. Information, trade and common knowledge. *Journal of Economic Theory*, 26:17 – 27, 1982.

[32] Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170 – 190, 1989.

[33] Dov Monderer and Dov Samet. Stochastic common learning. *Games and Economic Behavior*, 9(2):161 –171, 1995.

[34] Philipe Mongin and Luc Lismont. Strong completeness theorems for weak logics of common belief. *Journal of Philosophical Logic*, 32:115 – 137, 2003.

[35] Stephen Morris. Approximate common knowledge revisited. *International Journal of Game Theory*, 28:385 – 408, 1999.

[36] Zvika Neeman. Approximating agreeing to disagree results with common p-beliefs. *Games and Economic Behavior*, 12:162 – 164, 1996.

[37] Eric Pacuit. Dynamic epistemic logic II: Logics of information change. *Philosophy Compass*, 8(9):815 – 833, 2013.

[38] Rohit Parikh and Paul Krasucki. Communication, consensus, and knowledge. *Journal of Economic Theory*, 52(1):178 – 189, 1990.

[39] Dov Samet. Agreeing to disagree in infinite information structures. *Games and Economic Behavior*, 21(3):213 – 218, 1992.

[40] Dov Samet. Agreeing to disagree: The non-probabilistic case. *Games and Economic Behavior*, 69(1):169 – 174, 2010.

[41] James Sebenius and John Geanakoplos. Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association*, 78(382):424 – 426, 1983.

[42] Johan van Benthem, Jan van Eijck, and Barteld Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620 – 1662, 2006.

[43] Johan van Benthem. One is a lonely number: On the logic of communication. *Logic Colloquium*, 27(96 - 129), 2002.

[44] Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2014.

[45] Johan van Benthem, Jelle Gerbrandy, and Barteld Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67 – 96, 2009.

[46] Chunlai Zhou. A complete deductive system for probability logic. *Journal of Logic and Computation*, 19(6):1427 – 1454, 2009.