International Journal of Approximate Reasoning 101 (2018) 150–162

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Agreeing to disagree and dilation ☆

Jiji Zhang [a], Hailin Liu [b,*], Teddy Seidenfeld [c]

[a] *Department of Philosophy, Lingnan University, Hong Kong, China*
[b] *Institute of Logic and Cognition, Department of Philosophy, Sun Yat-sen University, Guangzhou, China*
[c] *Department of Philosophy, Carnegie Mellon University, Pittsburgh, USA*

ARTICLE INFO

Article history:
Received 5 November 2017
Received in revised form 11 July 2018
Accepted 16 July 2018
Available online 18 July 2018

Keywords:
Agreeing to disagree
Common knowledge
Dilation
Imprecise probability

ABSTRACT

We consider Aumann's famous result on "agreeing to disagree" in the context of imprecise probabilities. Our primary aim is to reveal a connection between the possibility of agreeing to disagree and the interesting and anomalous phenomenon known as dilation. For such a purpose it is convenient to use Geanakoplos and Polemarchakis' communication setting, where agents repeatedly announce and update credences until no new information is conveyed by the announcements. We show that for agents who share the same set of priors and update by conditioning on every prior, once the procedure of communicating credences stops, 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 probability values unless dilation is present. This result may be used to generate sufficient conditions for guaranteed full agreement for some important models of imprecise priors, and we illustrate the potential with an agreement result involving 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.

1. Introduction

In a simple but insightful paper, Aumann [1] 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. The agreement obtains even if their respective (private) information is not shared. This influential result has triggered a sizable literature.[1] In this paper we consider generalizations in the context of imprecise probabilities. In particular, Kajii and Ui [19,20] and Carvajal and Correia-da-Silva [6] proved some agreement results in the setting of multiple priors. In their work, "agreement" is primarily taken to mean *partial agreement*: two sets of probabilities are said to agree if they have a non-empty intersection. These authors then established sufficient conditions under which two agents who partially agree

☆ This paper is part of the Virtual special issue on Tenth International Symposium on Imprecise Probability: Theories and Applications (ISIPTA '17), Edited by Alessandro Antonucci, Giorgio Corani, Inés Couso and Sébastien Destercke.

* Corresponding author.
E-mail addresses: jijizhang@ln.edu.hk (J. Zhang), liuhlin3@mail.sysu.edu.cn (H. Liu), teddy@stat.cmu.edu (T. Seidenfeld).

[1] Google Scholar estimates approximately 2,500 citations by March, 2018. Rubinstein and Wolinsky [29], Bonanno and Nehring [5], for example, contain useful surveys.

https://doi.org/10.1016/j.ijar.2018.07.006
0888-613X/

on their imprecise priors are guaranteed to partially agree on their imprecise posteriors of a hypothesis if these posteriors are common knowledge.

In this paper, we build on this line of work and establish a connection between the possibility of agreeing to disagree and the interesting and anomalous phenomenon known as dilation for sets of probabilities [12,33,39,36,15]. 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 conditioning as their updating rule, their credences on a hypothesis become less precise or more indeterminate 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 [28,13,24]. Moreover, if we understand the indeterminacy in question as disagreement about the hypothesis of interest among multiple agents, then dilation is naturally seen as an obstacle to decreasing the disagreement by learning in the short run. However, as far as we know, it has not been discussed in connection to Aumann's agreement result. We show that dilation is also the key to the possibility of agreeing to disagree in Aumann's sense by Bayesian agents with common, imprecise priors.

We adapt the dynamics of Geanakoplos and Polemarchakis' [11] communication setting to present our results. While Aumann's original result applies only to "those (rare) events whose posteriors happen to be common knowledge", Geanakoplos and Polemarchakis devised a communication protocol that generates posteriors for an arbitrary event that eventually converge. In this communication setting, the agents are invited to repeatedly announce their credences of an event and update by conditioning on these announced credences, until no relevant new information is conveyed. Assuming that each agent's information partition is finite, the communication procedure is guaranteed to halt and the resulting posteriors of the event become common knowledge, to which Aumann's argument can be applied. We use this communication setting as we find it illuminating for some of our purposes, especially when we formulate our results in terms of "dilation-averse" agents in Section 5. Another reason for our preference for the communication setting is that there are other dynamics for reaching a consensus on subjective probabilities through iteratively announcing and updating individual probabilities, such as DeGroot's [9] non-Bayesian method. It will be interesting in future work to contrast (our generalizations of) Geanakoplos and Polemarchakis' (Bayesian) dynamics and (generalizations of) alternative, non-Bayesian dynamics.

The rest of the paper is organized as follows. In Section 2, we review Geanakoplos and Polemarchakis' dynamics in the setting of imprecise probabilities. In Section 3 we illustrate the failure of Aumann's result when the agents' uncertainty is represented by sets of probabilities rather than precise probabilities. We show that dilation is necessary for the failure of Aumann's result. 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 for Bayesian agents to reaching a *full* agreement. Full, in the sense that the sets of probability values representing their credences on the hypothesis of interest are identical. This result opens the door to generating sufficient conditions for reaching full Aumann-consensus, 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 [36]. This to our knowledge is the first attempt to generalize Aumann's result to the context of imprecise probabilities that is concerned with "full agreement" (as opposed to mere partial agreement in the sense of non-empty intersection of sets of probabilities). 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 the result 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, when lower and upper probabilities are sufficient to identify the set of probabilities, end up fully agreeing.

2. A procedure of communicating posteriors

In Geanakoplos and Polemarchakis' [11] setup, two[2] agents share a common measurable space $(\Omega, \mathcal{A})$ 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 ([1], 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' 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 $\mathcal{P}^1$

[2] Nothing is special about the number "two" here. As is the case with many contributions to this literature, although we set up the scene with two agents, everything we say also applies to situations with any (finite) number of agents. We thank an anonymous referee for reminding us of this point.