Agreeing to disagree with conditional probability systems*

Elias Tsakas†

Maastricht University

December 13, 2016

Abstract

In this note, we extend Aumann's agreement theorem to a framework where beliefs are modelled by conditional probability systems á la Battigalli and Siniscalchi (1999). We prove two independent generalizations of the agreement theorem, one where the agent's share some common conditioning event, and one where they may not.

1. Introduction

According to the famous agreement theorem, if two agents have a common prior, and their posteriors for an event are commonly believed at some state that receives positive probability by the prior, then they necessarily agree on the same posterior beliefs (Aumann, 1976).[1] The main contribution of this result is twofold: on the one hand, it has been been crucial for epistemically characterizing solution concepts such as for instance Nash equilibrium (Aumann and Brandenburger, 1995), while at the same time, it has helped us to understand the role of asymmetric information on betting, trading and speculation (Milgrom and Stokey, 1982; Sebenius and Geanakoplos, 1983).[2]

Recent advancements in the theory of games with incomplete information, as well as in the epistemic approach to game theory, have recognized that it is often important to take into account the players' beliefs conditional on "unlikely" – i.e., null – events. For instance, the standard characterizations

*This note supersedes two previous papers by the same author, titled "Strong belief and agreeing to disagree" and "Hierarchies of conditional beliefs derived from commonly known priors" respectively.

†Department of Economics (AE1), Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands; Homepage: http://www.elias-tsakas.com/home.html; E-mail: e.tsakas@maastrichtuniversity.nl

[1] Notice that Aumann's original result is slightly weaker, as it is in fact stated in terms of knowledge which is stronger than belief.

[2] For a survey on the agreement theorem, see Bonanno and Nehring (1997).

of iterated admissibility in normal form games (Brandenburger et al., 2008), or rationalizability in extensive form games (Battigalli and Siniscalchi, 2002) rely on this type of beliefs. In fact, there are two ways of modeling beliefs given null events in game theory – lexicographic probability systems and conditional probability systems – both of them generalizing the standard probabilistic beliefs. In a recent paper, Bach and Perea (2013) prove Aumann's agreement theorem with lexicographic beliefs. In this paper, we complete the analysis by extending the agreement theorem to the framework of conditional probability systems.

We prove two results, both of them assuming a common prior. First, we focus on a single conditioning event which is shared by the two agents, and we show that if the posterior belief *given this conditioning event* are commonly believed, they necessarily coincide, as long as this conditioning event receives positive probability by the common prior. Then, we switch attention to cases where the two agents have different collections of conditioning events, and their posterior belief *given each of these conditioning events* are commonly believed. In this case we show that the two agents will agree on their posteriors if their collections of conditioning events are balanced (Geanakoplos, 1989) and cover the same states of nature.

2. Preliminaries

2.1. Conditional probability systems

Fix an underlying measurable space of uncertainty $(X, \mathcal{A})$ and a collection of nonempty conditioning events $\mathcal{B} \subseteq \mathcal{A}$ (not necessarily an algebra). A mapping $\mu : \mathcal{A} \times \mathcal{B} \to [0, 1]$ is called a ***conditional probability system (CPS)*** on $(X, \mathcal{A}, \mathcal{B})$ whenever it satisfies the following conditions:

(C_1) $\mu(B|B) = 1$, for all $B \in \mathcal{B}$,

(C_2) $\mu(\cdot|B) \in \Delta(X, \mathcal{A})$, for all $B \in \mathcal{B}$,

(C_3) $\mu(A|C) = \mu(A|B) \cdot \mu(B|C)$, for all $A \in \mathcal{A}$ and $B, C \in \mathcal{B}$ with $A \subseteq B \subseteq C$.

Conditional probability systems were originally introduced by Rênyi (1955). Let $\mathcal{C}(X, \mathcal{A}, \mathcal{B})$ denote the space of all CPS's on $(X, \mathcal{A}, \mathcal{B})$. Whenever it is obvious which σ-algebra we use, we omit $\mathcal{A}$ and we simply write $\mathcal{C}(X, \mathcal{B})$. For an arbitrary measurable space Y, let $\mathcal{C}(X \times Y, \mathcal{B}) := \mathcal{C}(X \times Y, \mathcal{B} \times Y)$, where $\mathcal{B} \times Y := \{B \times Y | B \in \mathcal{B}\}$ contains the cylinders in $X \times Y$ generated by $\mathcal{B}$.

We say that a CPS $\mu \in \mathcal{C}(X, \mathcal{A}, \mathcal{B})$ is ***derived from the prior*** $p \in \Delta(X, \mathcal{A})$ whenever,

$$\mu(A|B) = \frac{p(A \cap B)}{p(B)} \tag{1}$$

for all $A \in \mathcal{A}$, and for all $B \in \mathcal{B}$ with $p(B) > 0$. Notice that not every CPS can be derived from a prior, as illustrated by the following example, which is similar to the ones in Halpern (2002, Ex. 2.2) and Halpern (2010, Ex. 3.8).

Example 1. Take $X = \{x_1, x_2, x_3\}$, $\mathcal{A} = 2^X$ and $\mathcal{B} = \{\{x_1, x_2\}, \{x_2, x_3\}, \{x_3, x_1\}\}$. Define the CPS, $\mu \in \mathcal{C}(X, \mathcal{B})$ by $\mu(\{x_1\}|\{x_1, x_2\}) = 1$, $\mu(\{x_2\}|\{x_2, x_3\}) = 1$ and $\mu(\{x_3\}|\{x_3, x_1\}) = 1$. Suppose – contrary to what we want show – that μ is derived from some $p \in \Delta(X)$. Observe that $p(\{x_k\}) > 0$ for at least one $k \in \{1, 2, 3\}$, and without loss of generality let $p(\{x_1\}) > 0$. The latter implies $p(\{x_1, x_2\}) > 0$, and since μ is derived from p, it follows that $p(\{x_2\}) = \mu(\{x_2\}|\{x_1, x_2\}) \cdot p(\{x_1, x_2\}) = 0$. Now, suppose that $p(\{x_3\}) > 0$. Then, using the same argument as above, we conclude that $p(\{x_1\}) = 0$, which contradicts our initial assumption. Hence, it is necessarily the case that $p(\{x_3\}) = 0$, implying $\mu(\{x_3\}|\{x_3, x_1\}) = p(\{x_3\})/p(\{x_3, x_1\}) = 0$, which contradicts $\mu(\{x_3\}|\{x_3, x_1\}) = 1$. Therefore, there is no prior from which this CPS can be derived. $\triangleleft$

We say that $\mathcal{B}$ is ***treelike*** whenever the following conditions hold (Halpern, 2010):

(T_1) For every $B \in \mathcal{B}$ there exists a partition $\{B_1, \dots B_n\} \subseteq \mathcal{B}$ of B such that for any $C \in \mathcal{B}$ with $B \cap C \neq \emptyset$, either $B \subseteq C$ or $C \subseteq B_k$ for some $k \in \{1, \dots, n\}$.

(T_2) There is a partition $\mathcal{P} \subseteq \mathcal{B}$ such that for all $B \in \mathcal{B}$ there is some $P \in \mathcal{P}$ with $B \subseteq P$.

If $\mathcal{B}$ is treelike, every $\mu \in \mathcal{C}(X, \mathcal{A}, \mathcal{B})$ is also called treelike. Observe that in the previous example, $\mathcal{B}$ is not treelike, as there are conditioning events that are neither disjoint nor subset of one another.

Proposition 1. *If $\mu \in \mathcal{C}(X, \mathcal{A}, \mathcal{B})$ is treelike, then μ is derived from some prior $p \in \Delta(X, \mathcal{A})$.*

We say that $\mathcal{B}$ is ***balanced*** whenever there exists a collection $(\lambda_B)_{B \in \mathcal{B}}$ of real numbers such that

$$\sum_{B \in \mathcal{B}} \lambda_B \mathbf{1}_{[B]}(x) = 1 \tag{2}$$

for all $x \in [B^*] := \bigcup_{B \in \mathcal{B}} [B]$ (e.g., see Geanakoplos, 1989; Tsakas and Voorneveld, 2011), where $\mathbf{1}_{[B]} : X \to \{0, 1\}$ is the indicator function, i.e., $\mathbf{1}_{[B]}(x) = 1$ if $\omega \in [B]$, and $\mathbf{1}_{[B]}(x) = 0$ otherwise.

Remark 1. Obviously, whenever $\mathcal{B}$ is treelike, it is also balanced. Indeed, if we take the partition $\mathcal{P}$ from Condition (T_2) and we set $\lambda_B = 1$ for all $B \in \mathcal{P}$ and $\lambda_B = 0$ otherwise, Equation (2) will be trivially satisfied. However, the converse is not necessarily true, e.g., the collection $\mathcal{B}$ from Example 1 is balanced – by taking $\lambda_B = 1/2$ for all $B \in \mathcal{B}$ – but not treelike. The latter also implies that balancedness does not guarantee the existence of a prior. $\triangleleft$

2.2. Hierarchies of conditional beliefs

Following Battigalli and Siniscalchi (1999) we extend conditional probability systems to an interactive setting. Accordingly, let Θ be a finite underlying space of uncertainty (together with the discrete algebra) and $I = \{a, b\}$ be the set of our agents.[3] For each $i \in I$, take a collection of conditioning events $\mathcal{B}_i$. In dynamic games, Θ is often seen as the set of terminal histories and $\mathcal{B}_i$ is interpreted as the collection of i's information sets (e.g., see Alós-Ferrer and Ritzberger, 2016). Then, we consider a ***type-based (epistemic) model*** $\big(\Theta, (\mathcal{B}_i)_{i\in I}, (T_i)_{i\in I}, (\lambda_i)_{i\in I}\big)$, where T_i is a finite space of types and $\lambda_i : T_i \to \mathcal{C}(\Theta \times T_j, \mathcal{B}_i)$ a conditional belief mapping.[4] For notation simplicity, $\lambda_{t_i}(\cdot|B) := \lambda_i(t_i)(\cdot|B)$ denotes t_i's beliefs over $\Theta \times T_j$ given the conditioning event $B \times T_j$. If $\mathcal{B}_i$ is treelike for every $i \in I$, the type-based model is also called treelike. Standard belief hierarchies are encoded in a type-based model with a single conditioning event $\mathcal{B}_i = \{\Theta\}$ and $\lambda_i : T_i \to \Delta(\Theta \times T_j)$ (Harsanyi, 1967-68).

We define the product space $\Omega := \Theta \times T_a \times T_b$, henceforth called the state space. We naturally define $\theta(\omega) := \mathrm{Proj}_\Theta\, \omega$ and $t_i(\omega) := \mathrm{Proj}_{T_i}\, \omega$. Then, for an element (resp., subset) of a space in the product, we define its extension in Ω, by putting the corresponding element (resp., subset) within brackets, e.g., for each $\theta \in \Theta$ we write $[\theta] := \{\omega \in \Omega : \theta(\omega) = \theta\}$, and likewise for each $F \subseteq \Theta \times T_j$ we write $[F] := \{\omega \in \Omega : (\theta(\omega), t_j(\omega)) \in F\}$.

Now, let us extend t_i's conditional beliefs from $\lambda_{t_i} \in \mathcal{C}(\Theta \times T_j, \mathcal{B}_i)$ to $\pi_{t_i} \in \mathcal{C}(\Omega, \mathcal{B}_i)$, where for an arbitrary $E \subseteq \Omega$ and an arbitrary $B \in \mathcal{B}_i$ we obtain

$$\pi_{t_i}(E|B) := \lambda_{t_i}\big(\{(\theta, t_j) \in \Theta \times T_j : (\theta, t_i, t_j) \in E\} \;\big|\; B\big). \tag{3}$$

Let us now define two events of particular interest. Fix an arbitrary event $E \subseteq \Omega$ and a probability $q_i \in [0, 1]$. Then, for each $B \in \mathcal{B}_i$, we define the event

$$[q_i]^B := \{\omega \in \Omega : \pi_{t_i(\omega)}(E|B) = q_i\}. \tag{4}$$

Moreover, we define the event

$$[q_i] := \bigcap_{B \in \mathcal{B}_i} [q_i]^B. \tag{5}$$

We denote the event that i ***(fully) believes*** in E by

$$K_i(E) := \bigcap_{B \in \mathcal{B}_i} \left\{ \, \omega \in \Omega : \pi_{t_i(\omega)}(E|B) = 1 \, \right\} . \tag{6}$$

[3] Our entirely analysis can be generalized to a compact metrizable space of uncertainty and a finite set of agents.

[4] Tsakas (2014) extends Battigalli and Siniscalchi's (1999) construction to also accommodate uncertainty about the collection of conditioning events. To do so, we introduce a function that associates each type with a collection of conditioning events.

Obviously, by definition $K_i([t_i]) = [t_i]$, i.e., our agents satisfy the usual introspection axioms. Now, we say that E is mutually believed at the states in $K(E) := K_a(E) \cap K_b(E)$. Moreover, for each $m > 0$ we inductively define the event that E is m-fold mutually believed by $K^{m+1}(E) := K(K^m(E))$ with $K^1(E) := K(E)$. Finally, we define the event that E is ***commonly (fully) believed*** by

$$CK(E) := \bigcap_{k=1}^{\infty} K^m(E). \tag{7}$$

Now, generalizing the definition we introduced in the single-agent case, we say that i's (conditional) beliefs are ***derived from a prior*** $p \in \Delta(\Omega)$ if, for all $t_i \in T_i$ and for all $F \subseteq \Omega$

$$\pi_{t_i}(F|B) = \frac{p\big(F \cap [B] \cap [t_i]\big)}{p\big([B] \cap [t_i]\big)}, \tag{8}$$

for all $B \in \mathcal{B}_i$ with $p\big([B] \cap [t_i]\big) > 0$. We say that $p \in \Delta(\Omega)$ is a ***common prior*** if the conditional beliefs of every $i \in I$ are derived from p.

Remark 2. Obviously, if $T_i = \{t_i\}$ is a singleton, (8) reduces to (1), i.e., the two (seemingly) distinct definitions of a prior are consistent with each other. Furthermore, if $\mathcal{B}_i = \{\Theta\}$, our definition of a common prior reduces to the standard one, which has been studied extensively in the literature (e.g., see Aumann, 1976; Bonanno and Nehring, 1999; Feinberg, 2000; Halpern, 2002). ◁

3. Generalized agreement theorem

According to Aumann's famous agreement theorem, in the framework with standard beliefs, if the two agents share a common prior and their posteriors about some given event are commonly believed at some state that receives positive probability by the prior, then these posteriors necessarily coincide. In this section we extend Aumann's idea to a setting with conditional probability systems.

We prove two (independent) extensions of Aumann (1976). First, we show that whenever the two agents share some conditioning event and their conditional beliefs *given this event* are commonly believed, then they necessarily coincide. Notice that we do not require common belief in the posteriors given other conditional events.

Theorem 1. *Fix a type-based model* $\big(\Theta, (\mathcal{B}_i)_{i\in I}, (T_i)_{i\in I}, (\lambda_i)_{i\in I}\big)$ *and take some* $B \in \mathcal{B}_a \cap \mathcal{B}_b$*. Let* $p \in \Delta(\Omega)$ *be a common prior and take* $q_a \in [0,1]$ *and* $q_b \in [0,1]$ *such that* $CK([q_a]^B \cap [q_b]^B) \cap \mathrm{Supp}(p) \cap [B] \neq \emptyset$*. Then,* $q_a = q_b$*.*

Our assumption of the conditioning event B receiving positive probability by the common prior is a crucial one, as illustrated by the following example.

Example 2. Let $\Theta = \{\theta_1, \theta_2, \theta_3\}$ be the underlying space of uncertainty, with $\mathcal{B}_i = \{\{\theta_1\}, \{\theta_2, \theta_3\}\}$ be the common collection of conditioning events of each $i \in I = \{a, b\}$. Moreover, consider the type-based epistemic model $(\Theta, \mathcal{B}_a, \mathcal{B}_b, T_a, T_b, \lambda_a, \lambda_b)$ such that $T_a = \{t_a^0\}$, $T_b = \{t_b^1, t_b^2\}$, with $\lambda_{t_a^0}(\cdot|\{\theta_1\})$ being uniformly distributed over $\{(\theta_1, t_b^1), (\theta_1, t_b^2)\}$ and $\lambda_{t_a^0}(\cdot|\{\theta_2, \theta_3\})$ being uniformly distributed over $\{(\theta_2, t_b^1), (\theta_2, t_b^2)\}$, while at the same time $\lambda_{t_b^k}(\{\theta_1, t_a^0\}|\{\theta_1\}) = 1$ and $\lambda_{t_b^k}(\{\theta_3, t_a^0\}|\{\theta_2, \theta_3\}) = 1$ for each $k \in \{1, 2\}$. Notice that the conditional beliefs are derived from the common prior $p \in \Delta(\Theta \times T_a \times T_b)$ which is uniformly distributed over $\{(\theta_1, t_a^0, t_b^1), (\theta_1, t_a^0, t_b^2)\}$. Now, consider the event $A = \{\theta_1, \theta_2\} \times T_a \times T_b$, and observe that it is commonly believed that a puts probability 1 to A both given $\{\theta_1\}$ and given $\{\theta_2, \theta_3\}$, whereas b puts probability 1 to A given $\{\theta_1\}$ and probability 0 to A given $\{\theta_2, \theta_3\}$. Hence, a and b completely disagree on their probabilistic assessment over A given the conditioning event $\{\theta_2, \theta_3\}$. That is, disagreement occurs only given the conditioning event that receives 0 probability by the common prior, whereas the two agents agree given the conditional event that receives positive probability by the common prior. $\triangleleft$

Obviously Aumann's theorem follows directly from the previous result by setting $\mathcal{B}_a = \mathcal{B}_b = \{\Theta\}$. In this case, $CK([q_a] \cap [q_b]) \cap \mathrm{Supp}(p) \neq \emptyset$ directly implies $q_a = q_b$ as in Aumann (1976).

Second, we show that even if the agents have different conditional events, as long as the two collections cover the same states of nature and the conditional beliefs *given each event* are commonly believed, then they necessarily coincide. That is, in our second result we allow for additional classes of conditioning events, but at the same time we are more restrictive in the interactive beliefs about the posteriors.

Theorem 2. *Fix a balanced type-based model* $\big(\Theta, (\mathcal{B}_i)_{i\in I}, (T_i)_{i\in I}, (\lambda_i)_{i\in I}\big)$ *with* $B_a^* = B_b^* =: B^*$. *Let* $p \in \Delta(\Omega)$ *be a common prior and take some* $q_a \in [0, 1]$ *and* $q_b \in [0, 1]$, *such that* $CK([q_a] \cap [q_b]) \cap \mathrm{Supp}(p) \cap [B^*] \neq \emptyset$. *Then,* $q_a = q_b$.

The previous result is closely related to a version of the agreement theorem for nonpartitional information structures (Geanakoplos, 1989, Thm. 6). Moreover, note that it can be directly extended to epistemic models where different $t_i \in T_i$ may have a different (balanced) $\mathcal{B}_{t_i}$, similarly to Tsakas (2014). The only requirement in this case would be that $B_{t_i}^* = B^*$ for all $t_i \in T_i$ and all $i \in I$, i.e., all collections of conditioning events cover the same natural states in Θ.

Corollary 1. *Fix a treelike type-based model* $\big(\Theta, (\mathcal{B}_i)_{i\in I}, (T_i)_{i\in I}, (\lambda_i)_{i\in I}\big)$. *Let* $p \in \Delta(\Omega)$ *be a common prior and take some* $q_a \in [0, 1]$ *and* $q_b \in [0, 1]$, *such that* $CK([q_a] \cap [q_b]) \cap \mathrm{Supp}(p) \neq \emptyset$. *Then,* $q_a = q_b$.

A. Proofs

Proof of Proposition 1. First of all observe that since $\mathcal{B}$ covers X, since $\mathcal{P} \subseteq \mathcal{B}$ is a partition of X. Then, consider an arbitrary collection $(\alpha_P)_{P\in\mathcal{P}}$ or positive reals such that $\sum_{P\in\mathcal{P}} \alpha_P = 1$. For each $A \in \mathcal{A}$ define

$$p(A) := \sum_{P\in\mathcal{P}} \alpha_P \cdot \mu(A|P). \tag{A.1}$$

Verifying that p is a probability measure in $\Delta(X, \mathcal{A})$ is trivial. Moreover, notice that for every $A \in \mathcal{A}$ and every $B \in \mathcal{B}$ with $p(B) > 0$, it is the case that

$$\begin{aligned}
\frac{p(A\cap B)}{p(B)} &= \frac{\sum_{P\in\mathcal{P}} \alpha_P \cdot \mu(A\cap B|P)}{\sum_{P'\in\mathcal{P}} \alpha_{P'} \cdot \mu(B|P')} \\
&= \mu(A|B) \cdot \frac{\sum_{P\in\mathcal{P}} \alpha_P \cdot \mu(B|P)}{\sum_{P'\in\mathcal{P}} \alpha_{P'} \cdot \mu(B|P')} \qquad (A.2)\\
&= \mu(A|B),
\end{aligned}$$

with (A.2) following from (C_3) since $A \cap B \subseteq B \subseteq P$, thus completing the proof. □

Proof of Theorem 1. Define $T_i^B := \{t_i \in T_i : t_i(\omega) = t_i$ for some $\omega \in CK([q_a]^B \cap [q_b]^B) \cap \mathrm{Supp}(p) \cap [B]\}$, which is by hypothesis nonempty. Since the event $[q_i]^B$ is measurable with respect to the partition $\{[t_i]|t_i \in T_i\}$, either $[t_i] \subseteq [q_i]^B$ or $[t_i] \cap [q_i]^B = \emptyset$. Now, for an arbitrary $t_i \in T_i^B$, it is necessarily the case that $\emptyset \neq [t_i] \cap CK([q_i]^B) \subseteq [t_i] \cap K_i([q_i]^B) \subseteq [t_i] \cap [q_i]^B$. Hence, for each $t_i \in T_i^B$,

$$\begin{aligned}
q_i &= \pi_{t_i}(E|B) \\
&= p\big(E|[B] \cap [t_i]\big).
\end{aligned}$$

The second equality follows from $p([B] \cap [t_i]) > 0$, which is by definition true for all $t_i \in T_i^B$. Now, define the set $M := B \times T_a^B \times T_b^B$. Hence,

$$\begin{aligned}
q_i &= \sum_{t_i\in T_i^B} q_i \cdot p([B] \cap [t_i]|M) \\
&= \sum_{t_i\in T_i^B} p\big(E|[B] \cap [t_i]\big) \cdot p([B] \cap [t_i]|M) \\
&= p(E|M).
\end{aligned}$$

Finally, since $p(E|M)$ does not depend on $i \in I$, it is the case that $q_a = q_b$. □

Proof of Theorem 2. *Step 1.* Define $T_i^* := \{t_i \in T_i : t_i(\omega) = t_i$ for some $\omega \in CK([q_a] \cap [q_b]) \cap \mathrm{Supp}(p) \cap [B^*]\}$, which is by hypothesis nonempty. Then, define the set $M^* := B^* \times T_a^* \times T_b^*$. Then, following Tsakas and Voorneveld (2011, Proof of Lemma), we define $\mathcal{Q}_i$ as the coarsest partition of B^* such that generates $\mathcal{B}_i$, i.e., formally for every $\theta \in B^*$, the element of $\mathcal{Q}_i$ that contains θ is defined by $\mathcal{Q}_i(\theta) := \bigcap_{B\in\mathcal{B}_i:\theta\in B} B$. Thus, we obtain

$$
\begin{aligned}
p(E|M^*) &= \sum_{Q\in\mathcal{Q}_i} p(E|Q\times T_a^*\times T_b^*)\cdot p(Q\times T_a^*\times T_b^*|M^*) \\
&= \sum_{Q\in\mathcal{Q}_i}\Bigg(\sum_{B\in\mathcal{B}_i:B\supseteq Q}\lambda_B\Bigg)\cdot p(E|Q\times T_a^*\times T_b^*)\cdot p(Q\times T_a^*\times T_b^*|M^*) && \text{(A.3)}\\
&= \sum_{B\in\mathcal{B}_i}\lambda_B\sum_{Q\in\mathcal{Q}_i:Q\subseteq B} p(E|Q\times T_a^*\times T_b^*)\cdot p(Q\times T_a^*\times T_b^*|M^*)\\
&= \sum_{B\in\mathcal{B}_i}\lambda_B\sum_{Q\in\mathcal{Q}_i:Q\subseteq B} p(E|Q\times T_a^*\times T_b^*)\cdot p(Q\times T_a^*\times T_b^*|B\times T_a^*\times T_b^*)\cdot p(B\times T_a^*\times T_b^*|M^*)\\
&= \sum_{B\in\mathcal{B}_i}\lambda_B\cdot p(B\times T_a^*\times T_b^*|M^*)\cdot p(E|B\times T_a^*\times T_b^*), && \text{(A.4)}
\end{aligned}
$$

with (A.3) following from the fact that $\mathcal{B}_i$ is balanced. Before moving forward, let us point out that $p(Q\times T_a^*\times T_b^*)$ or $p(B\times T_a^*\times T_b^*)$ might in principle be equal to 0 for some $Q\in\mathcal{Q}_i$ or $B\in\mathcal{B}_i$ respectively. In such cases, the conditional probabilities given these null events take arbitrary values.

Step 2. Now, similarly to the proof of Theorem 1, it is the case that $[t_i]\subseteq[q_i]$ for all $t_i\in T_i^*$, thus implying $\pi_{t_i}(E|B)=q_i$ for all $B\in\mathcal{B}_i$ and all $t_i\in T_i^*$.

Step 3. Now take an arbitrary $B\in\mathcal{B}_i$ such that $p(E|B\times T_a^*\times T_b^*)>0$. Then, there exists some $t_i\in T_i^*$ such that $p([B]\cap[t_i])>0$, and therefore since p is a common prior, we obtain $p(E|B\times T_a^*\times T_b^*)=q_i$ for all $B\in\mathcal{B}_i$ with $p(E|B\times T_a^*\times T_b^*)>0$. Hence, by (A.4) it follows that

$$
\begin{aligned}
p(E|M^*) &= q_i\sum_{B\in\mathcal{B}_i}\lambda_B\cdot p(B\times T_a^*\times T_b^*|M^*)\\
&= q_i\sum_{B\in\mathcal{B}_i}\lambda_B\sum_{Q\in\mathcal{Q}_i} p(B\times T_a^*\times T_b^*|Q\times T_a^*\times T_b^*)\cdot p(Q\times T_a^*\times T_b^*|M^*)\\
&= q_i\sum_{Q\in\mathcal{Q}_i} p(Q\times T_a^*\times T_b^*|M^*)\sum_{B\in\mathcal{B}_i}\lambda_B\cdot p(B\times T_a^*\times T_b^*|Q\times T_a^*\times T_b^*)\\
&= q_i\sum_{Q\in\mathcal{Q}_i} p(Q\times T_a^*\times T_b^*|M^*) && \text{(A.5)}\\
&= q_i,
\end{aligned}
$$

with (A.5) following again from the fact that $\mathcal{B}_i$ is balanced. Indeed, the conditional probability can be seen as an indicator function, with $p(Q\times T_a^*\times T_b^*|M^*)=1$ if $Q\subseteq B$ and $p(Q\times T_a^*\times T_b^*|M^*)=0$ otherwise. Finally, notice that $p(E|M^*)$ does not depend on $i\in I$, thus completing the proof. $\square$

Proof of Corollary 1. Recall from Remark 1 that every treelike $\mathcal{B}_i$ is also balanced. Moreover, by (T_2) it follows that $B_a^*=B_b^*=\Theta$ for all $i\in I$. Hence, $CK([q_a]\cap[q_b])\cap\mathrm{Supp}(p)\neq\emptyset$ directly implies $CK([q_a]\cap[q_b])\cap\mathrm{Supp}(p)\cap[B^*]\neq\emptyset$. Thus, all the conditions of Theorem 2 are satisfied, and therefore $q_a=q_b$. $\square$

References

Alós-Ferrer, C. & Ritzberger, K. (2016). *The theory of extensive form games.* Springer.

Aumann, R.J. (1976). Agreeing to disagree. *Annals of Statistics* 4, 1236-1239.

Aumann, R.J. & Brandenburger, A. (1995). Epistemic conditions for Nash equilibrium. *Econometrica* 63, 1161–1180.

Bach, C.W. & Perea, A. (2013). Agreeing to disagree with lexicographic prior beliefs. *Mathematical Social Sciences* 66, 129–133.

Bonanno, G. & Nehring, K. (1997). Agreeing to disagree: a survey. *Working Paper, UC Davis.*

——— (1999). How to make sense of the common prior assumption under incomplete information. *International Journal of Game Theory* 28, 409–434.

Battigalli, P. & Siniscalchi, M. (1999). Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *Journal of Economic Theory* 88, 188–230.

——— (2002). Strong belief and forward induction reasoning. *Journal of Economic Theory* 106, 356–391.

Brandenburger, A., Friedenberg, A. & Keissler, H.J. (2008). Admissibility in games. *Econometrica* 76, 307–352.

Feinberg, Y. (2000). Characterizing common priors in the form of posteriors. *Journal of Economic Theory* 91, 127–179.

Geanakoplos, J. (1989). Game theory without partitions, and applications to speculation and consensus. *Cowles Foundation Discussion Paper* 914.

Halpern, J.Y. (2002). Characterizing the common prior assumption. *Journal of Economic Theory* 106, 316–355.

——— (2010). Lexicographic probability, conditional probability, and nonstandard probability. *Games and Economic Behavior* 68, 155–179.

Harsanyi, J. (1967-68). Games with incomplete information played by Bayesian players, I-III. *Management Science* 14, 159–182, 320–334, 486–502.

Milgrom, P. & Stokey, N. (1982). Information, trade and common knowledge. *Journal of Economic Theory* 26, 17–27.

Rênyi, A. (1955). On a new axiomatic theory of probability. *Acta Mathetica Academiae Scientiarum Hungaricae* 6, 285–335.

Sebenius, J. & Geanakoplos, J. (1983). Don't bet on it: Contingent agreements with asymmetric information. *Journal of the American Statistical Association* 78, 424–426.

Tsakas, E. (2014). Epistemic equivalence of extended belief hierarchies. *Games and Economic Behavior* 86, 126–144.

Tsakas, E. & Voorneveld, M. (2011). On consensus through communication without a commonly known protocol. *Journal of Mathematical Economics* 47, 733–739.