Agreeing to Disagree and Lexicographic Probability Systems*

Christian W. Bach† and Jérémie Cabessa‡

EPICENTER Working Paper No. 19 (2019)

Abstract. In this note we explore agreeing to disagree with lexicographic probability systems. By means of a counterexample, it is shown that agents can agree to lexicographically disagree on their posteriors. Based on this observation, we propose the same excluding condition which essentially states that agents synchronically either neglect or consider their private information. A lexicographic agreement theorem ensues with equal posteriors at every level.

Keywords: agreeing to disagree; agreement theorem; Aumann structure; common prior assumption; interactive epistemology; lexicographic beliefs; lexicographic probability systems; locally meet constant condition; mutual absolute continuity; same excluding condition; strong agreement; universal agreement; weak agreement.

1 Introduction

The impossibility for two agents to agree to disagree is established by Aumann (1976)'s seminal agreement theorem. More precisely, if two Bayesian agents equipped with a

*A preliminary version of this work was presented at the Thirteenth Conference on Logic and the Foundations of Game and Decision Theory (LOFT13), Milan, July 2018. We are grateful to Amanda Friedenberg, Burkhard Schipper, Andrés Perea, and three anonymous referees at LOFT13 for useful and constructive comments.

†Department for Economics, Finance and Accounting, University of Liverpool Management School, Chatham Street, Liverpool, L69 7ZH, UNITED KINGDOM; EPICENTER, School of Business and Economics, Maastricht University, 6200 MD Maastricht, THE NETHERLANDS. Email: c.w.bach@liverpool.ac.uk

‡Laboratory of Mathematical Economics and Applied Microeconomics (LEMMA), University of Paris 2 – Panthéon-Assas, 4 Rue Blaise Desgoffe, 75006 Paris, FRANCE. Email: jeremie.cabessa@u-paris2.fr

common prior belief receive private information and have common knowledge of their posterior beliefs, then these posteriors must be equal. In other words, distinct posteriors cannot be common knowledge among Bayesian agents with a common prior. In this sense, agents cannot agree to disagree. A rather extensive literature on agreeing to disagree has emerged. Most contributions reconsider Aumann's impossibility theorem in more general frameworks. Notably, Bonanno and Nehring (1997) as well as Ménager (2012) provide comprehensive surveys on this literature. More recent contributions to the agreeing to disagree literature include Dégrement and Roy (2012), Hellman and Samet (2012), Bach and Perea (2013), Heifetz et al. (2013), Hellman (2013), Demey (2014), Lehrer and Samet (2014), Chen et al. (2015), Bach and Cabessa (2017), Gizatulina and Hellman (2018), Pacuit (2018), Tarbush (2018), as well as Tsakas (2018).

In game theory lexicographic beliefs play a prominent role. Intuitively, they permit to model cautious reasoning of players that do not exclude any choice from consideration yet deem some choices much more – indeed infinitely more – likely than others. This kind of inclusion-exclusion challenge is identified by Samuelson (1992), when showing that the solution concept of iterated weak dominance can be inconsistent with common knowledge. Notably lexicographic beliefs are used by Brandenburger et al. (2008) to epistemically characterize iterated weak dominance, thereby blocking inclusion-exclusion type problems. Further applications of lexicographic beliefs to games comprise Blume et al. (1991b), Brandenburger (1992), Börgers (1994), Schumacher (1999), Asheim (2002), as well as Asheim and Perea (2005). Formally, lexicographic beliefs are modelled in their most general form by lexicographic probability systems due to Blume et al. (1991a).

The natural question then emerges how the agreement theorem is affected if standard probabilities are replaced by lexicographic probability systems. Intuitively, the issue becomes whether it is impossible for cautious agents to agree to disagree or not. In fact, agreeing to disagree with a notion of lexicographic beliefs – simpler than lexicographic probability systems – is considered by Bach and Perea (2013). They employ a particular framework of cautious reasoning that admits lexicographic priors yet only unique posteriors. Consequently, agreeing to disagree with both lexicographic priors *and* lexicographic posteriors still remains to be investigated.

In this note we explore agreeing to disagree in generalized Aumann structures with lexicographic probability systems, where *all* beliefs are expressed lexicographically. More precisely, we formalize agents with a sequence of priors on the basis of which they compute a sequence of posteriors in the style of Blume et al. (1991a). A weak agreement theorem in the sense of merely identical first level posteriors obtains.

However, we show by means of a counterexample that agents can agree to disagree on their posteriors beyond the first lexicographic level. Thus, the impossibility results – neither of Aumann (1976) in the standard framework nor of Bach and Perea (2013) with lexigraphic priors – do not directly generalize to full-fledged cautious reasoning. Based on this observation, we introduce the same excluding condition which essentially states that agents synchronically either neglect or consider their private information. In fact, this condition can be seen as a variant of mutual absolute continuity from probability theory. With the same excluding condition we provide a strong agreement theorem which establishes the impossibility of agreeing to lexicographically disagree.

Besides, agreement not only at but across all lexicographic posterior levels is considered. The locally meet constant condition, which basically ensures that the agent's full lexicographic reasoning is rigid in terms of the common prior and which is stronger than a localized same excluding condition, gives rise to a universal agreement theorem. Accordingly, agents cannot only agree to lexicographically disagree but they also entertain identical posteriors across all levels.

2 Preliminaries

In set-based interactive epistemology knowledge and beliefs are modelled within the framework of Aumann structures. Formally, an Aumann structure

$$\mathcal{A} = \big(\Omega, (\mathcal{I}_i)_{i \in I}, p\big)$$

consists of a finite set Ω of possible worlds, a finite set I of agents, a possibility partition $\mathcal{I}_i$ of Ω for every agent $i \in I$, and a common prior $p : \Omega \to [0,1]$ such that $\sum_{\omega \in \Omega} p(\omega) = 1$. The cell of $\mathcal{I}_i$ containing the world ω is denoted by $\mathcal{I}_i(\omega)$ and contains those worlds deemed possible by agent i at world ω. It is also assumed that no information is excluded a priori, i.e. $p\big(\mathcal{I}_i(\omega)\big) > 0$ for all $i \in I$ and for all $\omega \in \Omega$.

Agents reason about events which are defined as sets of possible worlds. The common prior p naturally extends to a measure $p : \mathcal{P}(\Omega) \to [0,1]$ on the event space by setting $p(E) = \sum_{\omega \in E} p(\omega)$ for all $E \in \mathcal{P}(\Omega)$. Agents are Bayesians and consequently update the common prior with their private information as follows: the posterior belief of agent i in event E at world ω is given by

$$p\big(E \mid \mathcal{I}_i(\omega)\big) = \frac{p\big(E \cap \mathcal{I}_i(\omega)\big)}{p\big(\mathcal{I}_i(\omega)\big)}.$$

Knowledge is formalized in terms of events. The event of agent i knowing event E, denoted by $K_i(E)$, is defined as

$$K_i(E) := \{\omega \in \Omega : \mathcal{I}_i(\omega) \subseteq E\}.$$

If $\omega \in K_i(E)$, then i is said to know E at ω. Mutual knowledge of E is given by $K(E) := \bigcap_{i \in I} K_i(E)$. Setting $K^0(E) := E$, higher-order mutual knowledge is inductively defined by

$$K^m(E) := K\big(K^{m-1}(E)\big)$$

for all $m > 0$. Accordingly, mutual knowledge can also be denoted as 1-order mutual knowledge. The conjunction of all higher-order mutual knowledge yields common knowledge, which is formally defined as

$$CK(E) := \bigcap_{m > 0} K^m(E)$$

for all $E \in \mathcal{P}(\Omega)$. An equivalent formulation of common knowledge due to Aumann (1976) is based on the meet of the agents' possibility partitions.[1] Accordingly, common knowledge is defined as

$$CK(E) := \{\omega \in \Omega : (\bigwedge_{i \in I} \mathcal{I}_i)(\omega) \subseteq E\}$$

for all $E \in \mathcal{P}(\Omega)$.

Lexicographic beliefs are modelled in line with Blume et al. (1991a)'s notion of lexicographic probability systems. The following definition provides a direct adaptation of Blume et al. (1991a, Definition 3.1) to the interactive setting with multiple agents.

Definition 1. *Let Ω be a finite set of possible worlds and I be a finite set of agents. A* lexicographic probability system for agent $i \in I$ (i-LPS) *is a tuple $\rho_i = (p_i^1, \ldots, p_i^{M_i})$, for some $M_i \in \mathbb{N}$, where $p_i^m \in \Delta(\Omega)$ for all $m \in \{1, \ldots, M_i\}$.*

Lexicographic beliefs are thus sequences of standard beliefs. The index numbers of a lexicographic probability system are also referred to as lexicographic levels.

Injecting lexicographic probability systems into Aumann structures gives rise to the notion of lexicographic Aumann structures.

Definition 2. *A* lexicographic Aumann structure *is a tuple $\mathcal{A}_{LPS} = \big(\Omega, I, (\mathcal{I}_i)_{i \in I}, (\rho_i)_{i \in I}\big)$, where*

- *Ω is a set of possible worlds,*
- *I is a set of agents,*
- *$\mathcal{I}_i \subseteq 2^{\Omega}$ is a possibility partition of Ω for every agent $i \in I$,*
- *ρ_i is an i-LPS for evey agent $i \in I$,*
- *for every agent $i \in I$ and for every world $\omega \in \Omega$, there exists a lexicographic level $m \in \{1, \ldots, M_i\}$ such that $p_i^m\big(\mathcal{I}_i(\omega)\big) > 0$.*

The fifth item of Definition 2 ensures that no information is excluded a priori, and formally reflects the idea of caution. Actually, this condition can be seen as the lexicographic analogue to Aumann (1976)'s requirement for all information cells to be non-null events in the framework of Aumann structures. Besides, caution could also be modelled as follows: for all $i \in I$ and for all $\omega \in \Omega$ there exists $m_i > 0$ such that $p_i^{m_i}(\omega) > 0$. Since such a condition is stronger than the fifth item of Definition 2, the latter is preferable.

Agents use their information to reason lexicographically about events. Formally, Blume et al. (1991a, Definition 4.2) is adjusted to the context of lexicographic Aumann structures as follows.

[1] Given two partitions $\mathcal{P}_1$ and $\mathcal{P}_2$ of some set S, the partition $\mathcal{P}_1$ is called *finer* than the partition $\mathcal{P}_2$ (or $\mathcal{P}_2$ *coarser* than $\mathcal{P}_1$), if each cell of $\mathcal{P}_1$ is a subset of some cell of $\mathcal{P}_2$. Given n partitions $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_n$ of S, the finest partition that is coarser than $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_n$ is called the *meet* of $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_n$ and is denoted by $\bigwedge_{i=1}^n \mathcal{P}_i$. Moreover, given $x \in S$, the cell of the meet $\bigwedge_{i=1}^n \mathcal{P}_i$ containing x is denoted by $(\bigwedge_{i=1}^n \mathcal{P}_i)(x)$.

Definition 3. *Let $\mathcal{A}_{LPS}$ be a lexicographic Aumann structure, $\omega \in \Omega$ be some world, and $i \in I$ be some agent. The* conditional lexicographic probability system of agent i given his information at world ω (ω-conditional i-LPS) *is the tuple*

$$\rho_i^{\omega} = \left(p_i^{m_{i,\omega}^1}\left(\cdot \mid \mathcal{I}_i(\omega)\right), \ldots, p_i^{m_{i,\omega}^{L_{i,\omega}}}\left(\cdot \mid \mathcal{I}_i(\omega)\right)\right)$$

where

- *the sequence $(m_{i,\omega}^l)_{l=1}^{L_{i,\omega}}$ of indices is given by $m_{i,\omega}^0 = 0$, $m_{i,\omega}^l = \min\left\{m \in \{1, \ldots, M_i\} : p_i^m\left(\mathcal{I}_i(\omega)\right) > 0 \text{ and } m > m_{i,\omega}^{l-1}\right\}$ as well as $L_{i,\omega} = length\left((m_{i,\omega}^l)\right)$,*
- *$p_i^{m_{i,\omega}^l}\left(E \mid \mathcal{I}_i(\omega)\right) = \frac{p_i^{m_{i,\omega}^l}\left(E \cap \mathcal{I}_i(\omega)\right)}{p_i^{m_{i,\omega}^l}\left(\mathcal{I}_i(\omega)\right)}$ for all $E \subseteq \Omega$ and for all $l \in \{1, \ldots, L_{i,\omega}\}$.*

An essential difference between lexicographic Aumann structures and the standard framework resides in the former equipping agents with multiple levels of – and not unique – posteriors beliefs.

The common prior assumption in Aumann structures can be directly generalized to the lexicographic setting as follows.

Definition 4. *Let $\mathcal{A}_{LPS} = \left(\Omega, I, (\mathcal{I}_i)_{i \in I}, (\rho_i)_{i \in I}\right)$ be a lexicographic Aumann structure. The lexicographic Aumann structure $\mathcal{A}_{LPS}$ satisfies the* common prior assumption (CPA)*, if there exists $M \in \mathbb{N}$ such that $M_i = M_j = M$ and $p_i^m = p_j^m = p^m$ for all $i, j \in I$ and for all $m \in \{1, \ldots, M\}$. In this case, the tuple $\rho = (p^1, \ldots, p^M)$ is called* common prior *and $\mathcal{A}_{LPS}^{CPA} = \left(\Omega, I, (\mathcal{I}_i)_{i \in I}, \rho\right)$ is called* lexicographic Aumann structure with a common prior*.*

3 Weak Agreement

Since the agents hold levels of posterior beliefs, agreement becomes a multifarious notion. In fact, it is now shown that common knowledge of lexicographic posteriors ensures the agents' first level posterior beliefs to coincide.

Theorem 1 (WAT). *Let $\mathcal{A}_{LPS}^{CPA}$ be a lexicographic Aumann structure with a common prior ρ. Let $E \subseteq \Omega$ be some event and $\omega \in \Omega$ be some world. If*

$$CK\Big(\bigcap_{i \in I} \bigcap_{l \in \{1, \ldots, L_{i,\omega}\}} \{\omega' \in \Omega : p^{k_{i,\omega'}^l}\left(E \mid \mathcal{I}_i(\omega')\right) = p^{k_{i,\omega}^l}\left(E \mid \mathcal{I}_i(\omega)\right)\}\Big) \neq \emptyset,$$

then

$$p^{k_{i,\omega}^1}\left(E \mid \mathcal{I}_i(\omega)\right) = p^{k_{j,\omega}^1}\left(E \mid \mathcal{I}_j(\omega)\right)$$

for all $i, j \in I$.

Proof. Let $i^* \in I$ be some agent, $A_{i^*} \subseteq \Omega$ be some set such that $(\bigwedge_{i \in I} \mathcal{I}_i)(\omega) = \bigcup_{\omega' \in A_{i^*}} \mathcal{I}_{i^*}(\omega')$ and $\mathcal{I}_{i^*}(\omega_1) \cap \mathcal{I}_{i^*}(\omega_2)$ for all $\omega_1, \omega_2 \in A_{i^*}$, as well as $m^* \in \{1, \ldots, M\}$ be the first lexicographic level with $p^{m^*}\big((\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big) > 0$. Consider some world $\omega^* \in A_{i^*}$. Hence, $p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) \geq 0$.

If $p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) > 0$, then $k^1_{i^*,\omega^*} = m^*$, and by Bayesian updating,

$$p^{k^1_{i^*,\omega^*}}\big(E \mid \mathcal{I}_{i^*}(\omega^*)\big)\cdot p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) = p^{m^*}\big(E\cap\mathcal{I}_{i^*}(\omega^*)\big)$$

holds. Alternatively, if $p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) = 0$, then $p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big)$, thus $p^{m^*}\big(E\cap\mathcal{I}_{i^*}(\omega^*)\big) = 0$. Since $L_{i^*,\omega^*} \geq 1$, the first lexicographic level posterior measure $p^{k^1_{i^*,\omega^*}}$ exists and $k^1_{i^*,\omega^*} > m^*$. Consequently,

$$p^{k^1_{i^*,\omega^*}}\big(E \mid \mathcal{I}_{i^*}(\omega^*)\big)\cdot p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) = p^{m^*}\big(E\cap\mathcal{I}_{i^*}(\omega^*)\big)$$

holds trivially. Therefore,

$$p^{k^1_{i^*,\omega^*}}\big(E \mid \mathcal{I}_{i^*}(\omega^*)\big)\cdot p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) = p^{m^*}\big(E\cap\mathcal{I}_{i^*}(\omega^*)\big)$$

obtains for all $\omega^* \in A_{i^*}$.

Besides, as

$$A_{i^*} \subseteq CK\Big(\bigcap_{i\in I}\bigcap_{l\in\{1,\dots,L_{i,\omega}\}}\{\omega' \in \Omega : p^{k^l_{i,\omega'}}\big(E \mid \mathcal{I}_i(\omega')\big) = p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big)\}\Big),$$

it is the case that $p^{k^l_{i,\omega^*}}(E \mid \mathcal{I}_i(\omega^*)) = p^{k^l_{i,\omega}}(E \mid \mathcal{I}_i(\omega))$ for all $\omega^* \in A_{i^*}$, for all $l \in \{1,\dots,L_{i,\omega}\}$, and for all $i \in I$. Thus,

$$p^{k^1_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big)\cdot p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) = p^{m^*}\big(E\cap\mathcal{I}_{i^*}(\omega^*)\big)$$

holds for all $\omega^* \in A_{i^*}$. It follows that

$$\sum_{\omega^*\in A_{i^*}} p^{k^1_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big)\cdot p^{m^*}\big(\mathcal{I}_{i^*}(\omega^*)\big) = \sum_{\omega^*\in A_{i^*}} p^{m^*}\big(E\cap\mathcal{I}_{i^*}(\omega^*)\big),$$

and by countable additivity,

$$p^{k^1_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big)\cdot p^{m^*}\big((\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big) = p^{m^*}\big(E\cap(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big).$$

Since i^* has been chosen arbitrarily,

$$p^{k^1_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big) = \frac{p^{m^*}\big(E\cap(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big)}{p^{m^*}\big((\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big)}$$

holds for all $i^* \in I$. Therefore, $p^{k^1_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big) = p^{k^1_{j,\omega}}\big(E \mid \mathcal{I}_j(\omega)\big)$ obtains for all $i, j \in I$. ■

Agents can thus not agree to disagree on their first lexicographic level posterior beliefs. Yet, the preceding theorem does not furnish any insight on lexicographic level posteriors deeper than level one. Accordingly, **WAT** establishes a form of weak agreement for the lexicographic framework.

For the special case of exclusively admitting the first level posteriors – formally, only considering $p_i^{m^1_{i,\omega}}\big(\cdot \mid \mathcal{I}_i(\omega)\big)$ for all $\omega \in \Omega$ and for all $i \in I$ – our framework of lexicographic Aumann structures becomes essentially equivalent to Bach and Perea (2013)'s model, which only employs a lexicographic common prior but unique posteriors. In particular, their non-overlapping support condition across lexicographic prior levels is not assumed in our framework. Actually, **WAT** can be seen as a generalization of Bach and Perea (2013, Theorem 1). Moreover, if not only the posteriors but also the common prior is restricted to a single probability measure, i.e. formally setting $M = 1$, then Aumann (1976)'s framework is recovered in effect, and **WAT** becomes the original agreement theorem.

4 Disagreement

With lexicographic probability systems weak agreement in the sense of equal first level posteriors obtains. Attention is now shifted towards the deeper lexicographic levels. In fact, it is possible for agents to agree to disagree on posteriors beyond the first lexicographic level, as the following example illustrates.

Example 1. Let $\mathcal{A}^{CPA}_{LPS} = \big(\Omega, I, (\mathcal{I}_i)_{i\in I}, \rho\big)$ be a lexicographic Aumann structure with a common prior, where

- $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$,
- $I = \{Alice, Bob\}$,
- $\mathcal{I}_{Alice} = \big\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\big\}$,
- $\mathcal{I}_{Bob} = \{\Omega\}$,
- and $\rho = (p^1, p^2, p^3)$ with $p^1(\omega_1) = 1$, $p^2(\omega_2) = \frac{1}{3}$, $p^2(\omega_3) = \frac{2}{3}$, $p^3(\omega_4) = 1$.

Consider the event $E = \{\omega_1, \omega_3\}$. Observe that

$$p^{k^1_{Alice,\omega}}\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = p^1\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = 1$$

for all $\omega \in \{\omega_1, \omega_2\}$, and

$$p^{k^1_{Alice,\omega}}\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = p^2\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = 1$$

for all $\omega \in \{\omega_3, \omega_4\}$. Consequently, $p^{k^1_{Alice,\omega}}\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = 1$ obtains at every world $\omega \in \Omega$. Also, observe that

$$p^{k^1_{Bob,\omega}}\big(E \mid \mathcal{I}_{Bob}(\omega)\big) = p^1\big(E \mid \mathcal{I}_{Bob}(\omega)\big) = 1$$

for all $\omega \in \Omega$. Therefore, *Alice*'s and *Bob*'s first level posteriors coincide.

Moreover, it is the case that

$$p^{k^2_{Alice,\omega}}\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = p^2\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = 0$$

for all $\omega \in \{\omega_1, \omega_2\}$, and

$$p^{k^2_{Alice,\omega}}\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = p^3\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = 0$$

for all $\omega \in \{\omega_3, \omega_4\}$. Hence, $p^{k^2_{Alice,\omega}}\big(E \mid \mathcal{I}_{Alice}(\omega)\big) = 0$ obtains at every world $\omega \in \Omega$. Also,

$$p^{k^2_{Bob,\omega}}\big(E \mid \mathcal{I}_{Bob}(\omega)\big) = p^2\big(E \mid \mathcal{I}_{Bob}(\omega)\big) = \frac{2}{3}$$

holds at every world $\omega \in \Omega$. Therefore, *Alice*'s and *Bob*'s second level posteriors do not coincide.

Taking $\omega = \omega_1$ guarantees that

$$CK\Big(\bigcap_{i \in I} \bigcap_{l \in \{1,\dots,L_{i,\omega}\}} \big\{\omega' \in \Omega : p^{k^l_{i,\omega'}}(E \mid \mathcal{I}_i(\omega')) = p^{k^l_{i,\omega}}(E \mid \mathcal{I}_i(\omega))\big\}\Big) = CK(\Omega) = \Omega \neq \emptyset,$$

while

$$p^{k^2_{Alice,\omega}}(E \mid \mathcal{I}_{Alice}(\omega)) = 0 \neq \frac{2}{3} = p^{k^2_{Bob,\omega}}(E \mid \mathcal{I}_{Bob}(\omega))$$

obtains. ♣

A possibility result on agreeing to lexicographically disagree thus emerges.

Remark 1. There exist a lexicographic Aumann structure $\mathcal{A}^{CPA}_{LPS} = \big(\Omega, I, (\mathcal{I}_i)_{i \in I}, \rho\big)$ with a common prior, some event $E \subseteq \Omega$, and some world $\omega \in \Omega$, such that

$$CK\Big(\bigcap_{i \in I} \bigcap_{l \in \{1,\dots,L_{i,\omega}\}} \big\{\omega' \in \Omega : p^{k^l_{i,\omega'}}(E \mid \mathcal{I}_i(\omega')) = p^{k^l_{i,\omega}}(E \mid \mathcal{I}_i(\omega))\big\}\Big) \neq \emptyset$$

and

$$p^{k^l_{i,\omega}}(E \mid \mathcal{I}_i(\omega)) \neq p^{k^l_{j,\omega}}(E \mid \mathcal{I}_j(\omega))$$

for some $i, j \in I$ and for some $l \in \{1, \dots, \min_{i,j}\{L_{i,\omega}, L_{j,\omega}\}\}$.

Accordingly, common knowledge of the agents' lexicographic posteriors does not suffice to establish agreement at all lexicographic levels. The agents can entertain distinct posteriors at lexicographic levels beyond one, and at the same time acknowledge this divergence.

5 Strong Agreement

The impossibility result of **WAT** is weak, since it only affects the first lexicographic posterior level. A condition is now introduced that will then be used to derive a strong impossibility result in the sense of equal posteriors at every lexicographic level.

Definition 5. *Let $\mathcal{A}^{CPA}_{LPS}$ be a lexicographic Aumann structure with a common prior ρ and $\omega \in \Omega$ be some world. The common prior ρ is* same excluding, *if it is the case that*

$$p^m\big(\mathcal{I}_i(\omega)\big) = 0, \text{ if and only if, } p^m\big(\mathcal{I}_j(\omega)\big) = 0$$

for all $\omega \in \Omega$ for all $i, j \in I$, and for all $m \in \{1, \dots, M\}$.

Intuitively, the same excluding condition ensures that agents always either all consider or all neglect their respective information at any lexicographic prior level. Since the common prior assumption can be interpreted as some form of like-mindedness of the agents, the same excluding condition reflects an intensified like-mindedness.

Formally, the same excluding condition is closely related to the notion of mutual absolute continuity in probability theory. Let μ and ν be measures on some set Ω, and define $\mu << \nu$, if $\nu(F) = 0$ implies $\mu(F) = 0$ for all $F \in \mathcal{P}(\Omega)$. The two measures μ and ν are called mutually absolutely continuous, whenever $\mu << \nu$ and $\nu << \mu$. Observe that the common prior ρ induces for every level $m \in \{1, \dots, M\}$ and for every player $i \in I$ a measure $\mu_i^m : \mathcal{P}(\Omega) \to [0, 1]$ given by

$$\mu_i^m(F) := \begin{cases} 0 & \text{if } F = \emptyset \\ \sum_{\omega \in F} \frac{p^m\big(\mathcal{I}_i(\omega)\big)}{|\mathcal{I}_i(\omega)|} & \text{otherwise,} \end{cases}$$

for all $F \in \mathcal{P}(\Omega)$. Now, if $\mu_i^m(F) > 0$ for some $F \in \mathcal{P}(\Omega)$, i.e. $\sum_{\omega \in F} \frac{p^m\big(\mathcal{I}_i(\omega)\big)}{|\mathcal{I}_i(\omega)|} > 0$, then there exists $\omega^* \in F$ such that $p^m\big(\mathcal{I}_i(\omega^*)\big) > 0$. By the same excluding condition, $p^m\big(\mathcal{I}_j(\omega^*)\big) > 0$ thus holds too, and consequently $\mu_j^m(F) = \sum_{\omega \in F} \frac{p^m\big(\mathcal{I}_j(\omega)\big)}{|\mathcal{I}_j(\omega)|} > 0$. Conversely, if $p^m\big(\mathcal{I}_i(\omega)\big) > 0$ for some $\omega \in \Omega$, then $\mu_i^m(\{\omega\}) > 0$. By mutual absolute continuity, $\mu_j^m(\{\omega\}) > 0$ hence also obtains, and consequently $p^m\big(\mathcal{I}_j(\omega)\big) > 0$. Therefore, the following characterization of the same excluding condition in terms of mutual absolute continuity from probability theory ensues.

Remark 2. Let $\mathcal{A}_{LPS}^{CPA}$ be a lexicographic Aumann structure with a common prior ρ. The common prior ρ is same excluding, if and only if, μ_i^m and μ_j^m are mutually absolutely continuous for all $i, j \in I$ and for all $m \in \{1, \dots, M\}$.

Hence, the same excluding condition can be seen as a variant of mutual absolute continuity.

Moreover, the same excluding condition is conceptually similar to Stuart (1997)'s version of mutual absolute continuity.[2] Accordingly, if some agent's belief assigns a positive probability to a state (which essentially corresponds to a possible world in our framework), then so do all the other agents. Even though Stuart's model in its general form does not impose any prior, an agent's belief can be interpreted as posterior. While the underlying idea of Stuart's mutual absolute continuity and our same excluding condition is tantamount – some form of synchronicity in both consideration and omission – mutual absolute continuity concerns posterior beliefs as well as possible worlds yet same excludingness refers to prior beliefs as well as information.

It turns out that the same excluding condition together with the common prior assumption and common knowledge of posteriors implies that the agents' posterior beliefs coincide at all lexicographic levels.

[2]In fact, mutual absolute continuity plays an important role in establishing all period defection in the normal-form model of the finitely repeated prisoners' dilemma (Stuart, 1997, Proposition).

Theorem 2 (SAT). *Let $\mathcal{A}_{LPS}^{CPA}$ be a lexicographic Aumann structure with a common prior ρ that is same excluding. Let $E \subseteq \Omega$ be some event and $\omega \in \Omega$ be some world. If*

$$CK\Big(\bigcap_{i\in I}\bigcap_{l\in\{1,\dots,L_{i,\omega}\}}\{\omega'\in\Omega : p^{k^l_{i,\omega'}}\big(E \mid \mathcal{I}_i(\omega')\big) = p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big)\}\Big) \neq \emptyset,$$

then

$$p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big) = p^{k^l_{j,\omega}}\big(E \mid \mathcal{I}_j(\omega)\big)$$

for all $i, j \in I$ and for all $l \in \big\{1,\dots,\min_{i\in I}\{L_{i,\omega}\}\big\}$.

Proof. We first show that if ρ is same excluding, then $(k^l_{i,\omega})_{l=1}^{L_{i,\omega}} = (k^l_{j,\omega'})_{l=1}^{L_{i,\omega'}}$ for all $\omega' \in (\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$ and for all $i, j \in I$. Let $i^* \in I$ and $\omega' \in (\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$. Since $\omega' \in (\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$, the world ω' is reachable from ω, i.e., there exists a sequence $(P^k)_{k=1}^N$ of information cells such that $\omega \in P^1$, $\omega' \in P^N$, and $P^{k-1} \cap P^k \neq \emptyset$ for all $1 \leq k < N$. Since ρ is same excluding, it is the case that, $p^m(P^{k-1}) = 0$ if and only if $p^m(P^k) = 0$ for all $m \in \{1,\dots,M\}$ and for all $1 \leq k < N$. Thus, $p^m(P^1) = 0$ if and only if $p^m(P^N) = 0$ for all $m \in \{1,\dots,M\}$. Since $\omega \in \mathcal{I}_{i^*}(\omega)\cap P^1$, $\omega' \in \mathcal{I}_{i^*}(\omega')\cap P^N$ and ρ is same excluding, it follows that $p^m\big(\mathcal{I}_{i^*}(\omega)\big) = 0$ if and only if $p^m\big(\mathcal{I}_{i^*}(\omega')\big) = 0$, for all m $\in \{1,\dots,M\}$. Consequently, $(k^l_{i^*,\omega})_{l=1}^{L_{i^*,\omega}} = (k^l_{i^*,\omega'})_{l=1}^{L_{i^*,\omega'}}$. Towards a contradiction, suppose that there exist $j^* \in I$ and $l \in \big\{1,\dots,\min\{L_{i^*,\omega'}, L_{j^*,\omega'}\}\big\}$ such that $k^l_{i^*,\omega'} \neq k^l_{j^*,\omega'}$. Without loss of generality suppose that l is the least such index. Then, either $k^l_{i^*,\omega'} < k^l_{j^*,\omega'}$, in which case, $p^{k^l_{i^*,\omega'}}\big(\mathcal{I}_{i^*}(\omega')\big) > 0$ and $p^{k^l_{i^*,\omega'}}\big(\mathcal{I}_{j^*}(\omega')\big) = 0$, or $k^l_{i^*,\omega'} > k^l_{j^*,\omega'}$, in which case, $p^{k^l_{j^*,\omega'}}\big(\mathcal{I}_{i^*}(\omega')\big) = 0$ and $p^{k^l_{j^*,\omega'}}\big(\mathcal{I}_{j^*}(\omega')\big) > 0$. In both cases, this contradicts the same excludingness of ρ. Therefore, $(k^l_{i,\omega'})_{l=1}^{L_{i,\omega'}} = (k^l_{j,\omega'})_{l=1}^{L_{j,\omega'}}$ for all $i, j \in I$ and for all $\omega' \in (\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$.

Let $i^* \in I$ and $l^* \in \big\{1,\dots,\min_{i\in I}\{L_{i,\omega}\}\big\}$. Then, there exists $A_{i^*} \subseteq \Omega$ such that $(\bigwedge_{i\in I}\mathcal{I}_i)(\omega) = \bigcup_{\omega'\in A_{i^*}}\mathcal{I}_{i^*}(\omega')$ and $\mathcal{I}_{i^*}(\omega_1)\cap\mathcal{I}_{i^*}(\omega_2)$ for all $\omega_1, \omega_2 \in A_{i^*}$. Note that $A_{i^*} \subseteq (\bigwedge_{i\in I}\mathcal{I}_i)(\omega) \subseteq CK(E') \subseteq E'$, where

$$E' = \bigcap_{i\in I}\bigcap_{l\in\{1,\dots,L_{i,\omega}\}}\big\{\omega'\in\Omega : p^{k^l_{i,\omega'}}\big(E \mid \mathcal{I}_i(\omega')\big) = p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big)\big\}.$$

Since $(k^l_{j,\omega})_{l=1}^{L_{j,\omega}} = (k^l_{j',\omega'})_{l=1}^{L_{j',\omega'}}$ holds for all $\omega' \in (\bigwedge_{i\in I}\mathcal{I}_i)(\omega)$ and for all $j, j' \in I$, it follows that

$$p^{k^{l^*}_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big) = p^{k^{l^*}_{i^*,\omega'}}\big(E \mid \mathcal{I}_{i^*}(\omega')\big) = \frac{p^{k^{l^*}_{i^*,\omega'}}\big(E\cap\mathcal{I}_{i^*}(\omega')\big)}{p^{k^{l^*}_{i^*,\omega'}}\big(\mathcal{I}_{i^*}(\omega')\big)} = \frac{p^{k^{l^*}_{i,\omega}}\big(E\cap\mathcal{I}_{i^*}(\omega')\big)}{p^{k^{l^*}_{i,\omega}}\big(\mathcal{I}_{i^*}(\omega')\big)}$$

for all $\omega' \in A_{i^*}$ and for all $i \in I$. Consequently,

$$p^{k^{l^*}_{i^*,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big)\cdot p^{k^{l^*}_{i,\omega}}\big(\mathcal{I}_{i^*}(\omega')\big) = p^{k^{l^*}_{i,\omega}}\big(E\cap\mathcal{I}_{i^*}(\omega')\big),$$

for all $\omega' \in A_{i^*}$ and for all $i \in I$. Summing over all $\omega' \in A_{i^*}$, countable additivity yields

$$p^{k^{l^*}_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big) = \frac{p^{k^{l^*}_{i,\omega}}\big(E\cap(\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big)}{p^{k^{l^*}_{i,\omega}}\big((\bigwedge_{i\in I}\mathcal{I}_i)(\omega)\big)}$$

for all $i \in I$. Therefore, $p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big) = p^{k^l_{j,\omega}}\big(E \mid \mathcal{I}_j(\omega)\big)$ for all $i, j \in I$ and for all $l \in \big\{1, \ldots, \min_{i \in I}\{L_{i,\omega}\}\big\}$. ■

Consequently, it is impossible for agents to agree to lexicographically disagree whenever the same excluding condition is satisfied. In contrast to **WAT** that only ensures a weak form of agreement at the first posterior level, **SAT** establishes strong agreement in the sense of identity of posteriors at all lexicographic levels. The appropriate impossibility result for cautious reasoning is thus given by **SAT**, because the full thinking range is concerned.

Besides, observe that Example 1 actually suggests that **SAT** qualifies as tight with respect to the same excluding condition.[3] Indeed, the other two key assumptions, i.e. common prior as well common knowledge of posteriors, but not the same excluding condition hold, while the consequent, i.e. lexicographically identical posterior beliefs, fails.

Moreover, continuity in agreeing to lexicographically disagree follows from **SAT** in the sense that equal prior beliefs up to some lexicographic level imply equal posterior beliefs up to a corresponding lexicographic level. Suppose that the common prior assumption is weakened such that the agents' priors coincide up to some level $M^* < M$, and modify the initial lexicographic Aumann structure by truncating the agent's lexicographic priors at M^*, which is equivalent to imposing a common prior $\rho = (p^1, \ldots, p^{M^*})$. By **SAT** it follows that common knowledge of lexicographic posteriors at some world $\omega \in \Omega$ implies equal posterior measures for every level $l \in \min_{i \in I}\{L_{i,\omega}\}$ in the truncated structure, and hence also up to level $\min_{i \in I}\{L_{i,\omega}\}$ in the initial lexicographic Aumann structure. In this sense, the lexicographic impossibility result of **SAT** is continuous.

6 Universal Agreement

An utmost zealous form of agreement does require identical posteriors not only at – but also across – all lexicographic levels. In fact, the conditions of **SAT** are not sufficient for such an universal form of agreement. Consider the following slight variation of Example 1, where only ρ is replaced by $\varrho = (q^1, q^2)$ where $q^1(\omega_1) = q^1(\omega_3) = \frac{1}{2}$ and $q^2(\omega_2) = q^2(\omega_3) = \frac{1}{2}$. Note that common knowledge of posteriors holds at all worlds and the same excluding condition is satisfied, yet universal agreement does not obtain: the level 1 and level 2 posteriors are distinct. In order to exclude agreeing to disagree with lexicographic probability systems in the universal sense the following condition is needed.

Definition 6. *Let $\mathcal{A}^{CPA}_{LPS}$ be a lexicographic Aumann structure with a common prior ρ and $\omega \in \Omega$ be some world. The common prior ρ is* locally meet constant *at ω, if* $p^m(\omega') = p^m(\omega)$, for all $\omega' \in (\bigwedge_{i \in I} \mathcal{I}_i)(\omega)$ and for all $m \in \{1, \ldots, M\}$.

Intuitively, the locally meet constant condition ensures that the probabilistic reasoning of the agent about all reachable worlds remains rigid throughout the lexicographic

[3] Tightness is interpreted in the style of Aumann and Brandenburger (1995), i.e. whether dropping only one assumption of a result were to already break its conclusion.

levels. Note that if a common prior is locally meet constant at some world, then it is also (locally) same excluding at that world. In this sense, local meet constancy is stronger than same excludingness.

As a matter of fact, an impossibility result on agreeing to disagree across lexicographic levels unfolds with the locally meet constant condition.

Theorem 3 (UAT). *Let $\mathcal{A}_{LPS}^{CPA}$ be a lexicographic Aumann structure with a common prior ρ. Let $E \subseteq \Omega$ be some event and $\omega \in \Omega$ be some world. If ρ is locally meet constant at ω and*

$$CK\Big(\bigcap_{i \in I} \bigcap_{l \in \{1,\dots,L_{i,\omega}\}} \big\{\omega' \in \Omega : p^{k^l_{i,\omega'}}(E \mid \mathcal{I}_i(\omega')) = p^{k^l_{i,\omega}}(E \mid \mathcal{I}_i(\omega))\big\}\Big) \neq \emptyset,$$

then

$$p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big) = p^{k^{l'}_{j,\omega}}\big(E \mid \mathcal{I}_j(\omega)\big)$$

for all $i, j \in I$ and for all $l, l' \in \big\{1, \dots, \min_{i \in I}\{L_{i,\omega}\}\big\}$.

Proof. Consider some agents $i^*, j^* \in I$ as well as some lexicographic prior level $m \in \{1, \dots, M\}$. Observe that for all $\omega' \in (\bigwedge_{i \in I} \mathcal{I}_i)(\omega)$, by local meet constancy, it is the case that $p^m\big(\mathcal{I}_{i^*}(\omega')\big) > 0$, if and only if, $p^m\big(\mathcal{I}_{j^*}(\omega')\big) > 0$. An analogous argument as in the proof of **SAT** ensures that $(k^l_{i^*,\omega})_{l=1}^{L_{i^*,\omega}} = (k^l_{j^*,\omega'})_{l=1}^{L_{j^*,\omega'}}$ for all $\omega' \in (\bigwedge_{i \in I} \mathcal{I}_i)(\omega)$.

Consider some lexicographic posterior level $l^* \in \big\{1, \dots, \min_{i \in I}\{L_i, \omega\}\big\}$. By a similar argument as in the proof of **SAT**,

$$p^{k^{l^*}_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big) = \frac{p^{k^{l^*}_{j^*,\omega}}\big(E \cap (\bigwedge_{i' \in I} \mathcal{I}_{i'})(\omega)\big)}{p^{k^{l^*}_{j^*,\omega}}\big((\bigwedge_{i' \in I} \mathcal{I}_{i'})(\omega)\big)}$$

holds. With local meet constancy, it follows that

$$p^{k^{l^*}_{i^*,\omega}}\big(E \mid \mathcal{I}_{i^*}(\omega)\big) = \frac{\mid E \cap (\bigwedge_{i' \in I} \mathcal{I}_{i'})(\omega) \mid}{\mid (\bigwedge_{i' \in I} \mathcal{I}_{i'})(\omega) \mid}.$$

Therefore, $p^{k^l_{i,\omega}}\big(E \mid \mathcal{I}_i(\omega)\big) = p^{k^{l'}_{j,\omega}}\big(E \mid \mathcal{I}_j(\omega)\big)$ for all $i, j \in I$ and for all $l, l' \in \big\{1, \dots, \min_{i \in I}\{L_{i,\omega}\}\big\}$. ■

Accordingly, agents cannot agree to lexicographically disagree across any level whenever local meet constancy holds. Thus, **UAT** elevates the strong agreement of **SAT** to universal agreement. Furthermore, the local meet constant condition in **UAT** is tight, as in the slight variation of Example 1 all assumptions of **UAT** but local meet constancy are satisfied and the consequent of **UAT** fails.

7 Conclusion

Within the set-based epistemic framework enriched by lexicographic probability systems three types of agreement theorems emerge. If the agents' posteriors are common knowledge, weak agreement directly obtains. However, lexicographic disagreement cannot be excluded without further assumptions. Imposing the same excluding condition in addition to common knowledge of posteriors, a result of strong agreement, yielding same posteriors at all lexicographic levels, ensues. Finally, the locally meet constant condition enables a universal agreement theorem, where the posteriors coincide throughout the full reasoning range.

In terms of applications, it might be intriguing to use our strong agreement theorem for epistemic foundations of solution concepts involving cautious reasoning. For instance, Myerson (1978)'s proper equilibrium could be analyzed with analogous conditions to Aumann and Brandenburger (1995) adapted to the lexicographic framework used here, where the Strong Agreement Theorem forces equal conjectures. We leave such issues for future research.

References

Asheim, G. B. (2002): On the Epistemic Foundation for Backward Induction. *Mathematical Social Sciences* 44, 121–144.

Asheim, G. B. and Perea, A. (2005): Sequential and Quasi-Perfect Rationalizability in Extensive Games. *Games and Economic Behavior* 53, 15–42.

Aumann, R. J. (1976): Agreeing to Disagree. *Annals of Statistics* 4, 1236–1239.

Aumann, R. J. and Brandenburger, A. (1995): Epistemic Conditions for Nash Equilibrium. *Econometrica* 63, 1161–1180.

Bach, C. W. and Cabessa, J. (2017): Limit-Agreeing to Disagree. *Journal of Logic and Computation* 27, 1169–1187.

Bach, C. W. and Perea, A. (2013): Agreeing to Disagree with Lexicographic Prior Beliefs. *Mathematical Social Sciences* 66, 129–133.

Blume, L.; Brandenburger, A. and Dekel, E. (1991a): Lexicographic Probabilities and Choice under Uncertainty. *Econometrica* 59, 61–79.

Blume, L.; Brandenburger, A. and Dekel, E. (1991b): Lexicographic Probabilities and Equilibrium Refinements. *Econometrica* 59, 81–98.

Bonanno, G. and Nehring, K. (1997): Agreeing to Disagree: A Survey. Mimeo.

Börgers, T. (1994): Weak Dominance and Approximate Common Knowledge. *Journal of Economic Theory* 64, 265–276.

Brandenburger, A. (1992): Lexicographic Probabilities and Iterated Admissibility. In *Economic Analysis of Markets and Games*, ed. P. Dasgupta et al., MIT Press, 282–290.

Brandenburger, A.; Friedenberg, A. and Keisler, J. (2008): Admissibility in Games. *Econometrica* 76, 307–352.

Chen, Y.-C.; Lehrer, E.; Li, J.; Samet, D. and Shmaya, E. (2015): Agreeing to Disagree and Dutch Books. *Games and Economic Behavior* 93, 108–116.

Dégremont, C. and Roy, O. (2012): Agreement Theorems in Dynamic-Epistemic Logic. *Journal of Philosophical Logic*, 41, 735–764.

Demey, L. (2014): Agreeing to Disagree in Probabilistic Dynamic Epistemic Logic. *Synthese*, 191, 409–438.

Gizatulina, A. and Hellman, Z. (2018): Hegerogeneous Priors and Trade. Mimeo.

Heifetz, A.; Meier, M. and Schipper, B. C. (2013): Unawareness, Beliefs and Speculative Trade. *Games and Economic Behavior* 77, 100–121.

Hellman, Z. (2013): Almost Common Priors. *International Journal of Game Theory* 42, 399–410.

Hellman, Z. and Samet, D. (2012): How Common are Common Priors? *Games and Economic Behavior* 74, 517–525.

Lehrer, E. and Samet, D. (2014): Belief Consistency and Trade Consistency. *Games and Economic Behavior* 83, 165–177.

Ménager, L. (2012): Agreeing to Disagree: A Review. Mimeo.

Myerson, R. B. (1978): Refinements of the Nash Equilibrium Concept. *International Journal of Game Theory* 7, 73–80.

Pacuit, E. (2018): Agreement Theorems with Qualitative Conditional Probability. Mimeo.

Samuelson, L. (1992): Dominated Strategies and Common Knowledge. *Games and Economic Behavior* 4, 284–313.

Schuhmacher, F. (1999): Proper Rationalizability and Backward Induction. *International Journal of Game Theory* 28, 599–615.

Stuart, H. W. (1997): Common Belief of Rationality in the Finitely Repeated Prisoners' Dilemma. *Games and Economic Behavior* 19, 133–143.

Tarbush, B. (2018): Counterfactuals in "Agreeing to Disagre" Type Results. Mimeo.

Tsakas, E. (2018): Agreeing to Disagree with Conditional Probability Systems. forthcoming in *B.E. Journal of Theoretical Economics*.