Belief consistency and trade consistency*

Ehud Lehrer[†] and Dov Samet[‡]

November 20, 2011

Abstract

Interpersonal consistency can be defined in epistemic terms as consistency of beliefs, and in economic terms as the impossibility of certain trades. More specifically, beliefs are consistent when they have a common prior, that is, when they are derived by Bayesian updating from a prior belief common to all agents. Trade consistency requires that there is no contingent zero-sum trade which is always commonly known to be favorable to all agents. It is well established that in finite, and more generally in compact type spaces, the two notions of consistency are equivalent. However, for countable type spaces trade consistency may hold even when there is no common prior. We map the relations between various notions of belief and economic consistency in countable type spaces. Our main result is an equivalence theorem for finite and infinite countable type spaces between trade consistency and a new notion of belief consistency. This equivalence is a powerful tool that enables us to fully analyze the consistency of type spaces based on the knowledge structure of Rubinstein's email game. It also helps to justify the requirement of boundedness of trade in countable type spaces by showing that in a large class of spaces there exists an agreeable bet which is possibly unbounded even when a common prior exists.

Keywords: Type spaces, Common prior, No-trade theorems, Agreeing to disagree, Belief consistency.

Journal of Economic Literature classification numbers: C70, D82

---

*A helpful discussion with Oyama Daisuke is acknowledged.

†The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel and INSEAD, Boulevard de Constance, 77305 Fontainebleau, France. e-mail: `lehrer@post.tau.ac.il`

‡Faculty of Management, Tel Aviv University, Tel Aviv 69978. e-mail: `dovsamet@gmail.com`

Lehrer acknowledges support of the ISF through Grant #538/11.
Samet acknowledges support of the ISF through Grant #1517/11.

1 Introduction

1.1 Interpersonal consistency

Interpersonal consistency or inconsistency are sometimes easily identified. For example, suppose that Alice and Bob are exposed to exactly the same economic data. If despite of having the same information, Alice believes that interest rates will rise in the near future and Bob believes that they will fall, then, obviously, their beliefs are inconsistent. This inconsistency is expressed in purely *epistemic* terms, that is, in terms of the players' beliefs and knowledge . But it can be also expressed in *economic* terms, i.e., in terms of behavior, preferences, and choices. We can expect Alice and Bob to act and make their economic choices differently, and, moreover, we can even envisage certain economic interactions between them that one might dub inconsistent.

Matters are more complicated when agents differ in their knowledge. In such cases differences of beliefs do not necessarily imply inconsistency. As in the simple case of having the same information in the case of differential knowledge we also expect a description of consistency in epistemic terms and economic terms.

1.2 Models of knowledge and belief

Here we study interpersonal consistency in situations of differential knowledge where beliefs are expressed by probability distributions. Environments of probabilistic beliefs of many players were modeled by Harsanyi (1967-68) for studying games with incomplete information. In this model each player has *types* where each of a player's type is characterized by the player's beliefs about the types of the other players.

In a more general model of knowledge and belief, introduced in Aumann (1976) in the framework of a general state space, the knowledge of each player is defined by a partition of the state space, and her beliefs in each state are given by a probability distribution over the state space. We adopt this model here, refer to it as a type space and interchangeably use the words belief, type and posterior belief to describe the beliefs of a player at a state.

1.3 The history of the concept of consistency

1.3.1 Common priors

The first and most straightforward definition of consistency is in epistemic terms. We go back to a previous period, before the players acquired the differential information described by the partitions; a period when their knowledge was the same. We assume that their beliefs at the present period were obtained by conditioning their prior beliefs on their acquired knowledge.[1] We further assume that their beliefs in the prior period *were consistent*. As their knowledge in this prior period was the same, being consistent means having the *same* beliefs. Thus, consistency at the present time means consistency in a previous period which, in turn, means having the same beliefs.

This definition of consistency is known as the common prior assumption or Harsanyi's doctrine, because he made this assumption explicitly in his model (see Aumann (1987) and Bonanno and Nehring (1996) for a discussion and justification of the common prior assumption, and the debate over the validity of this assumption between Gul (1998) and Aumann (1998)). The common prior is a probability distribution over the state space from which all types, or posterior beliefs, are derived by conditioning on acquired knowledge.

Unless a player's beliefs are the same in all states, there are many prior probability distributions from which her posterior probabilities can be derived. The set of all such priors is the closed convex hull of the types of the player (see Samet (1998b)). The question of consistency in terms of a common prior boils down to the question of whether the sets of the players' priors have a point in common.

1.3.2 The agreement theorem and no-trade theorems

Aumann's (1976) agreement theorem states that when beliefs are consistent, that is, when there exists a common prior, agents cannot agree to disagree. That is, they cannot have common knowledge of a disagreement on the posterior probability of a given event. Thus, the impossibility of agreeing to disagree is a necessary condition for epistemic consistency, expressed in epistemic terms. Feinberg (2000), Halpern (2002), and Heifetz (2006) provided a necessary and *sufficient* condition in epistemic terms for the existence of a common prior. Samet (1998a) formulated a necessary and sufficient condition, for finite spaces, in terms of iterated expectations, and Hellman (2011) extended this result for topological compact spaces.

Other researchers extended Aumann's agreement theorem to yield further necessary conditions for the common prior assumption in economic rather than epistemic terms. They showed that a necessary condition for the existence of a common prior is that certain trades cannot be commonly known to be beneficial to all traders (see, for example,

[1] The assumption of Bayesian updating also reflects a notion of consistency of one's beliefs.

Milgrom and Stokey (1982), Sebenius and Geanakoplos (1983), Rubinstein and Wolinsky (1990)).

1.3.3 Equivalence for finite and compact spaces

Following these works, Morris (1995), Bonanno and Nehring (1996), Samet (1998b), Feinberg (2000), and Halpern (2002) studied a simple no-trade condition which is not only a necessary condition for the existence of a common prior but is *equivalent* to it. Thus, an equivalence was established between an epistemic condition of consistency and an economic one. The economic condition of consistency says that there is no *agreeable bet*, namely, a state dependent transfer of money between the players such that at each state it is common knowledge among them that each player's expected gains are positive. This equivalence was established first for finite state spaces, and then extended by Feinberg (2000) to infinite compact spaces, in which case the sets of priors are also compact. Heifetz (2006) provided a short proof for the compact case along the lines of the proof in Samet (1998b).

1.4 Equivalence in countable spaces

Countable spaces play an important role in game theory and economic theory. It suffices to mention Rubinstein's (1989) email game with its countable state space, which opened the door to a rich literature on the nature of common knowledge and its role in game theory. However, the equivalence stated above for finite and compact type spaces fails for countable spaces, and the consistency in such spaces proved to be a complicated phenomenon, not fully understood (see, for example, Feinberg (2000), Heifetz (2006), and Hellman (2010)).

In the countable case both epistemic consistency and economic consistency have variants that do not exist in the finite case. Consider, first, epistemic consistency. A type space may fail to have a common prior, yet the set of agents' priors may be of distance zero from each other. We say in this case that the type space is *weakly belief consistent*. Formally, we say that a probability distribution is an *ε-prior* for an agent if its distance from the agent's set of priors does not exceed ε. A *common ε-prior* is a probability distribution which is an ε-prior for each agent. The type space is weakly belief consistent if there exists an ε-prior for each ε. Weak belief consistency is equivalent to the existence of a common prior in the case of finite and the compact cases, but not so in the case of countable spaces.

As for economic consistency, *weak trade consistency* holds when there is no *strong agreeable bet*, which is an agreeable bet, such that for some $\delta > 0$, it is common knowledged in each state that each agent's expected payoff exceeds δ. Again, in contrast to the countable case, in the finite and the compact cases, weak trade consistency and trade consistency are the same.

Our first result is:

> **Weak consistency equivalence**: *A type space is weakly belief consistent if and only if it is weakly trade consistent.*

This equivalence generalizes the equivalence theorems of the finite and compact cases, as in these cases the weak notions of consistency coincide with the unqualified notions of consistency. However, in the countable case, trade consistency is not equivalent to the existence of a common prior. Finding a notion of belief consistency which is equivalent to trade consistency is a long standing open problem. In light of the weak consistency equivalence, the sought for notion of belief consistency must be stronger than weak belief consistency. It must also be implied by the existence of a common prior, as the latter implies trade consistency also in countable spaces. However, it should be weaker than the existence of a common prior, as follows from an example of Feinberg (2000) in which a type space does not have a common prior and yet is trade consistent, as it does not admit an agreeable bet. We next describe this condition.

1.5 Belief consistency

A weakly belief consistent type space has a common ε-prior for every ε. The probabilities assigned by these priors to a given element of the partitions may change with ε. The speed of this change relative to ε, determines belief consistency. When there is a common prior, for instance, then it is an ε-prior for each ε. There must be a partition element to which the common prior assigns positive probability. The ratio of this fixed probability to ε converges to infinity as ε converges to zero. It turns out that this property of the common prior is the one that guarantees belief consistency.

> **Belief consistency**: *A type space is belief consistent if there exists a partition element and a sequence of ε-priors with ε converging to zero, such that the ratio between the probability assigned to the partition element by an ε-prior in the sequence and ε converges to infinity.*

With this notion of consistency we can state the main result of this paper:

> **Consistency equivalence**: *A type space is belief consistent if and only if it is trade consistent.*

1.6 Applications

The condition of belief consistency is a powerful tool. We use it here to establish the consistency or inconsistency of some of the type spaces we study, where we have been unable to directly verify the existence or nonexistence of an agreeable bet.

We first study the question of boundedness of bets. In finite type spaces bets are bounded by definition, and in compact spaces continuous bets are bounded by virtue of the compactness. In countable type spaces, in contrast, the boundedness of bets has to be assumed. We show that this assumption is not just a mathematical convenience. Using the equivalence theorem, we prove that a large family of type spaces admit *unbounded* agreeable bets even when a common prior exists, that is, even when players' beliefs are consistent in the strongest possible sense. Thus, the existence of unbounded agreeable bet has no bearing on belief consistency.[2]

Next, in order to illustrate the usefulness of our results, we apply the equivalence theorem to study type spaces with a knowledge structure similar to that of Rubinstein's (1989) email game, which are the simplest non-trivial infinite countable type spaces. We fully characterize the consistency of such type spaces, in terms of the improper priors they necessarily have, when types are positive.

1.7 The paper's plan

Section 2 presents the model. Section 3 introduces the various kinds of belief consistency and their trade consistency notions counterparts. Section 4 applies the results to fully characterize consistency in type spaces based on the knowledge structure of Rubinstein's email game, and show the existence of unbounded agreeable bets on a large class of type spaces. In Section 5 we discuss our results and present some open questions. Section 6 provides the proofs that were omitted in the previous sections.

[2] Bhattacharyya and Lipman (1995) looked at the question of unbounded agreeable bets the other way around. They show that for any unbounded bet on an infinite countable space there is a type structure for which this bet is agreeable.

2 Preliminaries

Partition spaces. A *partition space* for a set N of n agents, is a tuple $(\Omega, (\Pi_i)_{i\in N})$, where Ω is a countable state space, and for each i, Π_i is a partition of Ω representing i's knowledge. For each $\omega \in \Omega$ we denote by $\Pi_i(\omega)$ the element of Π_i that contains ω. Subsets of Ω are called *events*. The *meet* is the finest common coarsening of the partitions $(\Pi_i)_{i\in N}$. An event E is *common knowledge* in a state ω if the element of the meet that contains ω is a subset of E.

Type spaces. The set of all probability distributions on Ω is denoted by $\Delta(\Omega)$. We consider $\Delta(\Omega)$ as a subset of $l_1(\Omega)$, in which $\Delta(\Omega)$ is closed. A type space for the partition space $(\Omega, (\Pi_i)_{i\in N})$ is a tuple $(\Omega, (\Pi_i)_{i\in N}, (t_i)_{i\in N})$, where for each i, $t_i \colon \Omega \to \Delta(\Omega)$ is *i's type function*. We write t_i^ω instead of $t_i(\omega)$. We require that the function t_i satisfies for each i, $\pi \in \Pi_i$ and states ω and ω' in π, $t_i^{\omega'} = t_i^\omega$ and $t_i^\omega(\pi) = 1$. The *type of i at* ω is t_i^ω. Since i's type is the same in all the states in $\pi = \Pi_i(\omega)$, we also write t_i^π for t_i^ω.

Priors. The set of i's *priors*, denoted by P_i, is the closed convex hull of $\{t_i^\omega \mid \omega \in \Omega\}$, the set of all of i's types. Equivalently, a prior for i is a probability distribution $p \in \Delta(\Omega)$ such that for each i, $\pi \in \Pi_i$, and $\omega \in \pi$, $t_i^\omega(\omega)p(\pi) = p(\omega)$. A probability distribution in the set $P = \cap_i P_i$ is a *common prior*. When there exists a common prior we say that the type space is *belief consistent*.

Expectations. For a real valued function f on Ω and $p \in \Delta(\Omega)$, we denote by $\mathrm{Ex}^p(f)$ the expectation of f with respect to p, when it exists. If $\mathrm{Ex}^{t_i^\omega}(f)$ exists for each ω, we define the function $E_i(f)$ on Ω by $E_i(f)(\omega) = \mathrm{Ex}^{t_i^\omega}(f)$. As $E_i(f)$ is measurable with respect to Π_i we can write for $\pi = \Pi_i(\omega)$, $E_i(f)(\pi)$ instead of $E_i(f)(\omega)$.

Bets. A *bet* on the type space is an N-tuple of real valued functions $(f_i)_{i\in N}$ on Ω, such that $\sum_i f_i = 0$, and $E_i(f)$ is defined for each i. The bet is *agreeable* if $E_i(f_i)(\omega) > 0$ for each ω, or equivalently, if the event $\{\omega \mid E_i(f_i)(\omega) > 0, \text{ for each } i\}$ is common knowledge in each state. A type space is *trade consistent* when there is no agreeable bet on it.

For finite spaces, the two kinds of consistency are equivalent (Morris (1995), Samet (1998b), and Feinberg (2000)).

Theorem. *A finite type space is trade consistent if and only if it is belief consistent.*

This theorem also holds for infinite type spaces when the type space is Hausdorff compact and bets are continuous functions. However, its extension to countable infinite type spaces remained open (see, Feinberg (2000), Heifetz (2006)).

3 Consistent type spaces

3.1 Trade consistency

Trade consistency in terms of bets for finite spaces is extended to countable type spaces by adding a boundedness caveat.

Definition 1. *A type space is* trade consistent *if there is no agreeable bet $(f_i)_{i\in N}$ such that for each i, f_i is uniformly bounded (i.e., $f_i \in l_\infty(\Omega)$).*

The boundedness requirement, which obviously holds for finite spaces, is essential for infinite spaces. We show in Subsection 4.2 that without this requirement agreeable bets exist on a large family of type spaces, including ones that have a common prior.

For countable spaces we can distinguish between agreeable bets as defined above, and strongly agreeable bet. The bet $(f_i)_{i\in N}$ is *strongly agreeable* if for some $\delta > 0$, $E_i(f_i)(\omega) > \delta$ at each state ω. The lack of a strong agreeable bet defines a weak version of trade consistency.

Definition 2. *The type space is* weakly trade consistent *if there is no strongly agreeable bet $(f_i)_{i\in N}$ such that for each i, $f_i \in l_\infty(\Omega)$.*

For finite, and more generally compact spaces, a bet is agreeable if and only if it is strongly agreeable, and hence trade consistency and weak trade consistency coincide.

3.2 Weak belief-consistency

When the sets of priors are arbitrarily close in a finite or compact space, then there exists a common prior. However, an infinite countable type space can have arbitrarily close sets of priors and still not have a common prior. This gives rise to the next definition.

For $\varepsilon \geq 0$, a probability distribution p is an *ε-prior* for i if for some $p_i \in P_i$, $||p-p_i||_1 \leq \varepsilon$. The set of all ε-priors of i is denoted by P_i^ε. Let $P^\varepsilon = \cap_{i\in N} P_i^\varepsilon$. An element of P^ε is called a *common ε-prior*. Obviously, a 0-prior is a prior and a common 0-prior is a common prior, that is, $P_i^0 = P_i$ and $P^0 = P$.

Definition 3. *The type space is* weakly belief consistent *if it has a common ε-prior for each $\varepsilon > 0$.*

For this consistency the following equivalence holds.

Theorem 1. *A countable type space (infinite or finite) is weakly trade consistent if and only if it is weakly belief consistent.*

Since a finite type space is weakly belief consistent if and only if it has a common prior, this theorem generalizes the equivalence theorems of the finite, as stated in Section 2. However defining belief consistency for infinite type spaces, equivalent to trade consistency, is subtle. Obviously, type spaces that have a common prior should be considered belief consistent, as is it is straightforward to show that such spaces are trade consistent. In contrast, type spaces that are not weakly belief consistent should not be considered belief consistent, as by Theorem 1 they are not weakly trade consistent and a fortiori not trade consistent. The twilight zone of type spaces that are weakly belief consistent but do not have a common prior is divided. Some such spaces are trade consistent as was shown by Feinberg (2000) and some are not. Thus belief consistency should hold for all spaces that have a common prior and for some spaces that are weakly belief consistent and do not have a common prior.

3.3 Belief consistency

When the type space is weakly belief consistent, the common ε-priors may spread thin and the probability of all the types, namely the partitions' elements, may go to zero rapidly as $\varepsilon \to 0$. To have consistent beliefs, the probability of some types must remain relatively high when $\varepsilon \to 0$, which makes the proximity of the priors significant. The magnitude of importance in our case is the ratio $p(\pi)/\varepsilon$ for a common ε-prior p and an element π of one the partitions.

Proposition 1. *For each $\pi \in \cup_i \Pi_i$, $\lim_{\varepsilon \to 0^+} \left(\sup\{p(\pi)/\varepsilon \mid p \in P^\varepsilon\} \right)$ exists.*[3]

Definition 4. *The type space is* belief consistent *if for some $\pi \in \cup_i \Pi_i$,*

$$\lim_{\varepsilon \to 0^+} \left(\sup\{p(\pi)/\varepsilon \mid p \in P^\varepsilon\} \right) = \infty. \tag{1}$$

The following claims show that this definition generalizes the consistency defined for finite spaces as the existence of a common prior.

Claim 1.

(a) If the type space has a common prior then it is belief consistent.

(b) If the type space is not weakly belief consistent then it is belief inconsistent.

[3]We assume the extended real line here. The supremum of an empty set is defined as usual to be $-\infty$, and the limit in the proposition can be ∞ or $-\infty$.

Indeed, if there is a common prior p, then for each $\varepsilon > 0$, p is a common ε-prior, and therefore for π with $p(\pi) > 0$, the nominator in $p(\pi)/\varepsilon$ is fixed and therefore (1) holds.

If the type space is not weakly belief consistent, then there exists ε_0 such that for each $\varepsilon < \varepsilon_0$, $P^\varepsilon = \emptyset$, which means that the limit in (1) is $-\infty$, and therefore in this case belief consistency does not hold. As finite spaces either have a common prior or are not weakly belief consistent, it follows that:

Claim 2. *A finite type space is belief consistent if and only if it has a common prior.*

3.4 The main theorem

Theorem 2. *A countable type space (infinite or finite) is trade consistent if and only if it is belief consistent.*

In light of Theorems 1 and 2 we refer henceforth to weak consistency and consistency of type spaces without the qualification of trade and belief.

4 Applications

4.1 Common improper priors

A *common improper prior* for a type space is a measure $\mu \neq 0$ on Ω such that for each i, $\pi \in \Pi_i$, and $\omega \in \pi$, $\mu(\pi) < \infty$, and $t_i^\pi(\omega)\mu(\pi) = \mu(\omega)$. Heifetz (2006) conjectured that a type space is consistent if and only if it does not have a common improper prior. Hellman (2010) proved half of this conjecture: a consistent type space must have a common improper prior. However the converse is not true. We give here a sufficient condition for type spaces that have a common improper prior to be inconsistent. We prove the inconsistency by showing that this condition implies belief inconsistency.

Obviously, a type space with a meet that consists of finite sets has a common improper prior if and only if it has a common prior and it is therefore consistent. Thus, we restrict ourselves to infinite type spaces that have an infinite element in the meet. In order to simplify the formulation of the condition we assume that the type space is *connected*, that is, the meet is a singleton. For simplicity, we further assume positive type functions. A type function t_i is *positive* if for each $\pi \in \Pi_i$, t_i^π is positive on π, that is, for each $\omega \in \pi$, $t_i^\pi(\omega) > 0$.

Proposition 2. *If a connected infinite type space with positive type functions has a common improper prior with a positive lower bound, then it is inconsistent.*

4.2 Unbounded agreeable bets

Using the result in the previous subsection we are able to demonstrate the importance of the boundedness condition in Definition 1 of trade consistency. We show that for a large family of type spaces, which is far from being exhaustive, there is always an agreeable bet, possibly unbounded. This family also includes type spaces that have a common prior.

A type space is *locally finite* if each of the partitions' elements is finite.

Proposition 3. *On each locally finite type space with positive type functions there exists an agreeable bet.*

4.3 The basic infinite partition space

In this subsection we fully characterize, in terms of common improper priors, the various kinds of consistency of the simplest infinite type spaces, using Theorems 1 and 2.

Starting with a state ω_0 we define an increasing sequence of events, $(E_k)_{k=0}^{\infty}$, by $E_0 = \{\omega_0\}$ and for $k \geq 1$, $E_k = \cup_{\omega \in E_{k-1}} \cup_i \Pi_i(\omega)$. The union $E = \cup_{k \geq 0} E_k$ is the element of the meet that contains ω_0. If for some $k \geq 0$, $E_k = E_{k+1}$, then $E_k = E_m$ for all $m \geq k$. In this case the event E is the event that "all know that all know ... that E", where "all know" is iterated k times, and it is also the event that E is common knowledge.[4] Obviously, the interesting infinite cases are partition spaces where the sequence is strictly increasing. The simplest partition space of this kind is one with two agents where for each k, $E_{k+1} \setminus E_k$ is the smallest possible set, that is, a singleton.

This results in the *basic partition space*, the state space of which is the set of natural numbers $\mathbb{N} = \{1, 2, 3, \dots\}$, and the partitions of agents 1 and 2 are $\Pi_1 = \big\{\{1\}, \{2, 3\}, \dots, \{2k, 2k+1\} \dots\big\}$, and $\Pi_2 = \big\{\{1, 2\}, \dots, \{2k-1, 2k\} \dots\big\}$. The basic partition space can model the sequential exchange of messages between the two agents, starting with agent 1. State k describes the knowledge after k messages have been sent. It was used in Rubinstein (1989) for the analysis of the electronic email game.

For simplicity we consider type spaces on the basic partition space with positive type functions. Each such type space has a common improper prior $\mu = (\mu_1, \mu_2, \dots)$. Indeed, define $\mu_1 > 0$ arbitrarily and for each $k \geq 2$ let $\mu_k = \mu_{k-1} t_i^k(k)/t_i^k(k-1)$, where $i = 1$ for odd k and $i = 2$ for even k. Then μ is a common improper prior, and a measure μ'

[4]When the type space is locally finite, then $E_k = E_{k+1}$ implies that the element of the meet that contains π is finite. When there are infinite partition elements then E can be infinite, but still, common knowledge is finitely generated in E.

is a common improper prior if and only if $\mu' = c\mu$ for some positive c.

Proposition 4. *Consider a type space on the basic partition space with positive type functions and a common improper prior μ. Then,*

(1) *the type space is weakly consistent if and only if* $\liminf_m \mu_m / \sum_{k=1}^{m} \mu_k = 0$*;*

(2) *the type space is consistent if and only if* $\liminf_m \mu_m = 0$*;*

(3) *the type space has a common prior if and only if* $\sum_m \mu_m < \infty$*.*

We illustrate Proposition 4 with two families of type spaces, defined by their common improper priors, on the basic partition space: *exponential*, where $\mu_m = \alpha^m$ for $\alpha > 0$, and *polynomial*, where $\mu_m = m^\beta$. The two families share one type space in common, when $\alpha = 1$ and $\beta = 0$. For this type space, in each partition element $\{m, m+1\}$ the type is $(1/2, 1/2)$.

In the exponential family, the ratio between the probability of two consecutive states is fixed, $(1/\alpha)$. Thus, in each partition element $\{m, m+1\}$ the type is $(1/(1+\alpha), \alpha/(1+\alpha))$. In the polynomial family, the ratio between the probability of two consecutive states converges to 1, as $m \to \infty$, and thus the types on $\{m, m+1\}$ converge to $(1/2, 1/2)$.

We show below, that in the exponential family, type spaces with large α are not even weakly consistent, and the only consistent spaces are those that have a common prior. In the polynomial family, in contrast, all type spaces, even those with large β, are weakly consistent, and there are type spaces that are consistent but do not have a common prior.

Consistency in the exponential family depends on α as follows.

- $\alpha < 1$. *There exists a common prior*, since $\sum_{m=1}^{\infty} \alpha^m = \alpha/(1-\alpha) < \infty$. .
- $\alpha = 1$. *The type space is weakly consistent but not consistent.* Indeed, $\mu_m = 1$ for each m, but $\mu_m / \sum_{k=1}^{m} \mu_k = 1/m \to 0$.
- $\alpha > 1$. *The type space is not weakly consistent*, since $\alpha^m / \sum_{k=1}^{m} \alpha^k = (\alpha^m - 1)/(\alpha - 1) \geq 1$.

Consistency in the polynomial family depends on β as follows.

- $\beta < -1$. *There exists a common prior*, since $\sum_{m=1}^{\infty} m^\beta < \infty$.

- $-1 \leq \beta < 0$. *The type space is consistent but does not have a common prior.* The series $\sum_{m=1}^{\infty} m^{\beta}$, that majorizes the harmonic series which is obtained for $\beta = -1$, diverges. However, $m^{\beta} \to 0$.

- $\beta \geq 0$. *The type space is weakly consistent but not consistent.* For each m, $m^{\beta} \geq 1$. However, since for $x \geq 0$, x^{β} is increasing, it follows that $m^{\beta} / \sum_{k=1}^{m} k^{\beta} \leq m^{\beta} / \int_{0}^{m} x^{\beta} dx = (\beta + 1)/m \to 0$.

It is easy to see that in Feinberg's (2000) example of a consistent type space that does not have a common prior, there is an improper common prior that satisfies the condition in **(2)** but not the one in **(3)**.

5 Comments and open problems

5.1 Constructing bets

In order to show that a type space has an agreeable bet we can either show that it is not belief consistent, using our main result, Theorem 2, or alternatively, construct the agreeable bet. Theorem 2 does not help us with such a construction, since the proof that an agreeable bet exists when the space is not belief consistent uses a separation theorem which by nature is not constructive.

Analyzing the simplest type spaces based on the basic partition model, we constructed an agreeable bet, using the common improper prior, to prove inconsistency in part (1) of Proposition 4. However in order to prove inconsistency for more general spaces that have a common improper prior, as in Propositions 2 and 4, we showed that the type spaces are not belief consistent. We do not know how to construct agreeable bets for the spaces in these propositions.

The question that arises is whether it is possible to describe constructively agreeable bets in countable spaces in terms of a common improper prior, when such bets exist.

5.2 Consistency in non-compact uncountable spaces

The analysis of consistency for non-compact *uncountable* spaces is a challenging problem. Our analysis for countable spaces as well as the analysis for compact spaces in Heifetz (2006) is based on the fact that the dual of the linear space of countably additive probability distributions consists of functions on the space. In the countable case the linear space of probability distributions is $l_1(\Omega)$ and its dual is the linear space of bounded functions on Ω, $l_\infty(\Omega)$. For a compact space, the linear space of countably additive probability distributions is, by the Riesz representation theorem, the space of continuous linear functionals on $C(\Omega)$—the set of continuous functions on Ω. When this space of functionals is endowed with the weak* topology, its topological dual is $C(\Omega)$. In these two cases, agreeable bets are obtained as separating functionals between sets of probability distributions.

However, when Ω is neither countable nor compact (either because it is not endowed with a topology or because it is a topological space but not compact), a continuous functional on the linear space of countably additive probability distributions is typically not a function on Ω. Suppose for example that Ω is the interval $[0,1]$. Here, countably additive probability distributions cannot be viewed as functions on Ω. There is no natural way to define the space of the all countably additive measures as a norm space and obtain agreeable bets as linear functionals over this space.

One way to bypass this difficulty is to take a dual approach and to start up with the space of bounded functions. This space also contains the agreeable bets and is relatively easy to handle. However, the dual of this space is the set of all finitely additive measures. Thus, using this approach to obtain a common prior as a separating functional would result in a finitely additive prior.

Allowing finitely additive rather than countably additive common priors is problematic. Consider, for instance, a locally finite countable space and a finitely additive measure that assigns probability 0 to every partition element and probability 1 to Ω. This measure is a common prior, because it is consistent with every player's posterior on each partition element. Thus, no matter what posteriors the players have, in a locally finite countable space there always exists a finitely additive common prior. In summary, eliminating the σ-additivity restriction over common priors renders the notion of a common prior somewhat vacuous.

5.3 Improper priors and consistency

The existence of a common improper prior signifies, to some extent, consistency among the players' beliefs. In the case that a common improper prior exists, what might prevent the existence of a regular common prior, is not any sort of discrepancy between the players' beliefs, but rather the infinite structure of the space and the ratios, induced by these beliefs, between the probabilities of different states.

Proposition 4 fully characterized, in terms of an improper prior, if it exists, the various kinds of consistency in the case of the basic partition model. The relation in a general model between an improper prior, if it exists, and consistency of all kinds remains open.

6 Proofs

Proof of Proposition 1: Fix $\pi \in \cup_i \Pi_i$ and let $\varphi(\varepsilon) = \sup\{p(\pi) \mid p \in P^\varepsilon\}$. We show that φ is increasing and concave on an interval of 0, which in turn implies that $\lim_{\varepsilon \to 0^+} \varphi(\varepsilon)/\varepsilon$ exists. If $P^\varepsilon = \emptyset$ for some ε, then $\varphi = -\infty$ on the interval $(0, \varepsilon]$, in particular φ is increasing and concave. We may therefore assume that $\varphi > -\infty$ on an interval around 0. Fix positive numbers $\varepsilon_1 < \varepsilon_2$. Since $P^{\varepsilon_1} \subseteq P^{\varepsilon_2}$, φ is increasing. We need to show that for $\alpha \in (0,1)$, $\varphi(\alpha\varepsilon_1 + (1-\alpha)\varepsilon_2) \geq \alpha\varphi(\varepsilon_1) + (1-\alpha)\varphi(\varepsilon_2)$. We first show that $\alpha P^{\varepsilon_1} + (1-\alpha)P^{\varepsilon_2} \subseteq P^{\alpha\varepsilon_1 + (1-\alpha)\varepsilon_2}$. It is enough to show that for each i, $\alpha P_i^{\varepsilon_1} + (1-\alpha)P_i^{\varepsilon_2} \subseteq P_i^{\alpha\varepsilon_1+(1-\alpha)\varepsilon_2}$. To see this let, for $k = 1, 2$, $p^{\varepsilon_k} \in P_i^{\varepsilon_k}$, and $p_k \in P_k$, such that $||p^{\varepsilon_k} - p_k|| \leq \varepsilon_k$. Then, $p = \alpha p_1 + (1-\alpha)p_2 \in P_i$, and $||\alpha p^{\varepsilon_1} + (1-\alpha)p^{\varepsilon_2} - p||_1 \leq \alpha||p^{\varepsilon_1} - p_1||_1 + (1-\alpha)||p^{\varepsilon_2} - p_2||_1 \leq \alpha\varepsilon_1 + (1-\alpha)\varepsilon_2$. Thus, $\alpha p^{\varepsilon_1} + (1-\alpha)p^{\varepsilon_2} \in P_i^{\alpha\varepsilon_1+(1-\alpha)\varepsilon_2}$.

Now, $I_k = \{p(\pi) \mid p \in P^{\varepsilon_k}\}$, for $k = 1, 2$, and $I = \{p(\pi) \mid p \in P^{\alpha\varepsilon_1+(1-\alpha)\varepsilon_2}\}$ are intervals, and by what we have shown, $\alpha I_1 + (1-\alpha)I_2 \subseteq I$. Thus, $\alpha \sup I_1 + (1-\alpha)\sup I_2 = \sup \alpha I_1 + \sup(1-\alpha)I_2 \leq \sup I$, which completes the proof of concavity. ∎

Proof of Theorem 1: Suppose first that the type space is not weakly trade consistent. That is, there exist an agreeable bet $(f_i)_{i \in N}$ and $\delta > 0$, such that $E_i(f_i)(\omega) > \delta$ at each state ω. We may assume, without loss of generality, that $||f_i||_\infty \leq 1$ for each $i \in N$. Suppose that p^ε is a common ε-prior and $p_i \in P_i$ satisfies $\|p^\varepsilon - p_i\| < \varepsilon$. $\mathrm{Ex}^{p^\varepsilon}(f_i) = \mathrm{Ex}^{p^\varepsilon}(f_i) - \mathrm{Ex}^{p_i}(f_i) + \mathrm{Ex}^{p_i}(f_i) \geq -\varepsilon + \mathrm{Ex}^{p_i}(E_i(f_i)) \geq -\varepsilon + \delta$. However, $\sum f_i = 0$ implies $\sum \mathrm{Ex}^{p^\varepsilon}(f_i) = 0$. Thus, $0 \geq -\varepsilon + \delta$, or equivalently, $\varepsilon \geq \delta$. Thus, ε-priors exist only for ε's which are bounded away from 0, and therefore the space is not weakly belief consistent.

Assume now that the space is not weakly belief consistent. Then, there is $\varepsilon > 0$ such that no common $n\varepsilon$-prior exists. Let $Q = \times_{i \in N} P_i$, and $C = Q + \varepsilon B^N$, where B is the unit ball in $l_1(\Omega)$. Denote by D the diagonal in $\Delta(\Omega)^N$, i.e., the set of $(p_i)_{i \in N}$ such that for all i and j, $p_i = p_j$. Then, C and D are disjoint. Moreover, both sets are convex and C has a non-empty interior. Thus, the two sets can be separated by a non-trivial continuous functional (see Dunford and Schwartz (1957)). That is, there is a continuous functional $g \neq 0$ on $l_1(\Omega)^N$ and a constant c such that $gx \geq c \geq gy$ for each $x \in C$ and

$y \in D$. The dual of $l_1(\omega)^N$ is $l_\infty(\omega)^N$, and thus, $g = (g_i)_{i\in N}$, where $g_i \in l_\infty(\Omega)$, and for $x = (x_i) \in l_1(\omega)^N$, $gx = \sum_i g_i x_i$.

We may assume without loss of generality that $||g|| > 1$ and therefore there exists $z \in B^N$ such that $gz = 1$. Then, for each $x \in Q$, $g(x - \varepsilon z) = gx - \varepsilon gz \geq c$. Thus, for each $x \in Q$, $gx \geq c + \varepsilon$.

Let $\mathbf{c}$ be the constant function on Ω that takes the value c. Define $\hat{g} = (\hat{g}_i)_{i\in N}$ by $\hat{g}_i = g_i - \mathbf{c}/n$. Since a constant functional returns the constant on probability distributions, it follows that for all $x \in Q$ and $y \in D$, $\hat{g}x \geq \varepsilon > 0 \geq \hat{g}y$. The inequality $\hat{g}y \leq 0$ for each $y \in D$ means that that for each $p \in \Delta(\Omega)$, $\sum_i \hat{g}_i p \leq 0$. That is, $(\sum_i \hat{g}_i)p \leq 0$. This implies that $\sum_i \hat{g}_i \leq 0$. Define $\hat{h} = (\hat{h}_i)_{i\in N}$ by $\hat{h}_i = \hat{g}^i - (\sum_i \hat{g}_i)/n$. Then, for each i, $\hat{h}_i \geq \hat{g}_i$, and therefore $\hat{h}x \geq \varepsilon$ for each $x \in Q$. Moreover, $\sum_i \hat{h}_i = 0$.

Let $a_i = \inf_{p_i \in P_i} \hat{h}_i p_i$. Then, $\sum_i a_i \geq \varepsilon$. Let $b_i = (\sum_i a_i)/n - a_i$. Then, $\sum_i b_i = 0$ and for each i, $a_i + b_i \geq \varepsilon/n$. Let $h_i = \hat{h}_i + \mathbf{b}_i$. The functional $(h_i)_{i\in N}$ satisfies $\sum_i h_i = 0$ and for each i and $p_i \in P_i$, $h_i p_i \geq a_i + b_i \geq \varepsilon/n$. In particular, for each ω, $p_i t_i^\omega \geq \varepsilon/n$. Thus, h is a strongly agreeable bet. ■

Proof of Theorem 2: We first prove two lemmas.

Lemma 1. *If $||q_i - p_i|| \leq \varepsilon$, for $p_i \in P_i$, and s_i is a type function for i for which q_i is a prior, then* $\mathrm{Ex}^{q_i}(||t_i - s_i||) \leq 2\varepsilon$.

Proof. Since i is fixed we suppress the index i in the proof. By the definition of types, $\mathrm{Ex}^q(||t-s||) = \sum_{\pi\in\Pi} q(\pi)||t(\pi) - s(\pi)|| \leq \sum_{\pi\in\Pi} ||q(\pi)t(\pi) - p(\pi)t(\pi)|| + \sum_{\pi\in\Pi} ||p(\pi)t(\pi) - q(\pi)s(\pi)||$. The first sum equals $\sum_{\pi\in\Pi} |q(\pi) - p(\pi)|\,||t(\pi)||$ and as $||t(\pi)|| = 1$, it is $\sum_{\pi\in\Pi} |q(\pi) - p(\pi)| \leq ||q - p|| \leq \varepsilon$. The second term equals $\sum_{\pi\in\Pi}\sum_{\omega\in\pi} |p(\omega)) - q(\omega)| = ||q - p|| \leq \varepsilon$. ■

Lemma 2. *If $p \in P^\varepsilon$, then for any bet $(f_i)_{i\in N}$,* $\mathrm{Ex}^p(\sum_i E_i(f_i)) \leq 2n(\sum_i ||f_i||_\infty)\varepsilon$.

Proof. Let $\hat{t}_i$ be a type function of i for which p is a prior, and let $\hat{E}_i(f_i)(\omega) = \mathrm{Ex}^{\hat{t}_i^\omega}(f_i)$. Then, as p is a common prior for $(\hat{t}_i)_{i\in N}$, $\mathrm{Ex}^p(\sum_i \hat{E}_i(f_i)) = \mathrm{Ex}^p(\sum_i f_i) = 0$. Thus, $|\mathrm{Ex}^p(\sum_i E_i(f_i))| = |\mathrm{Ex}^p(\sum_i E_i(f_i)) - \mathrm{Ex}^p(\sum_i \hat{E}_i(f_i))| \leq \sum_i \mathrm{Ex}^p(|E_i(f_i) - \hat{E}_i(f_i)|)$. Now, $|E_i(f_i)(\omega) - \hat{E}_i(f_i)(\omega)| \leq ||t_i^\omega - \hat{t}_i^\omega||_1 ||f_i||_\infty$. Thus, $\mathrm{Ex}^p(|E_i(f_i) - \hat{E}_i(f_i)|) \leq ||f_i||_\infty \mathrm{Ex}^p(||t_i^\omega - \hat{t}_i^\omega||$. The required inequality follows by Lemma 1. ■

Assume that the type space is belief consistent, and (1) holds for $\pi = \Pi_i(\omega)$. Suppose to the contrary that (f_i) is a bounded agreeable bet. Since for each j, $E_j(f_j) > 0$, it follows that for any p, $\mathrm{Ex}^p(\sum_j E_j(f_j)) \geq \mathrm{Ex}^p(E_i(f_i)) \geq p(\pi)E_i(f_i)(\omega)$. Since $E_i(f_i)(\omega) > 0$, it follows by our assumption that for some ε and $p^\varepsilon \in P^\varepsilon$, $p^\varepsilon(\pi)E_i(f_i)(\omega)/\varepsilon > 2n$. Thus, $\mathrm{Ex}^{p^\varepsilon}(\sum_j E_j(f_j)) > 2n\varepsilon$, which contradicts Lemma 2.

Conversely, suppose that the space is not belief consistent. We construct for each i and $\pi \in \Pi_i$ a bounded bet $(f_i^\pi)_{i\in N}$ such that $E_i(f_i^\pi)(\pi) > 0$, and for each j, $E_j(f_j^\pi) \geq 0$. Let $\pi_1, \pi_2, \ldots$ be an enumeration of the elements of $\cup_i \Pi_i$. The bet $(f_i)_{i\in N}$, where $f_i = \sum_{k\geq 1} 2^{-k} f_i^{\pi_k}$, is a bounded agreeable bet.

The bet (f_i^π) is constructed as a continuous functional that separates convex sets of probability distributions. However, sets of probability distributions have empty interior in $l_1(\Omega)$, and therefore separation of such sets by continuous functionals is not guaranteed. In the first stage of the proof we enlarge the sets of priors to sets with a non-empty interior. Of course, these sets should be enlarged cautiously such that they can still be separated. The ability to do so is guaranteed by the lack of belief consistency.

Fix i_0 and $\pi \in \Pi_{i_0}$. Choose $\omega \in \pi$ such that $t_{i_o}(\pi)(\omega) > 0$. Consider for each i and $\delta > 0$ the following sets in $\Delta(\Omega)$: $C_i^{1,\delta} = t_i^\omega + \delta B$, where B is the unit ball in $l_1(\Omega)$, $C_i^{2,\delta} = \mathrm{Cov}\{t_i^{\omega'} \mid t_i^{\omega'} \neq t_i^\omega\}$, and $C_i^\delta = \mathrm{Cov}(C_i^{1,\delta}, C_i^{2,\delta})$. Suppose that a probability distribution $p \in \Delta(\Omega)$ is in $\cap_i C_i^\delta$. Then, for each i there are α_i in $[0,1]$, $x_i \in B$, and $q_i \in C_i^{2,\delta}$ such that $p = \alpha_i t_i^\omega + \alpha_i \delta x_i + (1-\alpha_i)q_i$. Clearly $p_i = \alpha_i t_i^\omega + (1-\alpha_i)q_i$ is in P_i, and $||p - p_i|| = ||\alpha_i \delta x_i|| \leq \alpha_i \delta$. Let $\alpha_j = \max_i \alpha_i$. Then, $p \in P^{\alpha_j \delta}$. If $\alpha_j = 0$, then p is a common prior, which is impossible by our assumption of a lack of belief consistency. Thus $\alpha_j > 0$. Now, $p_j(\Pi_j(\omega)) = \alpha_j t_j^\omega(\Pi_j(\omega)) + (1-\alpha_i)q_j(\Pi_j(\omega)) = \alpha_j$. Hence, $p(\Pi_j(\omega)) \geq \alpha_j - \alpha_j\delta$, and therefore

$$p(\Pi_j(\omega))/(\alpha_j\delta) > 1/\delta - 1. \tag{2}$$

Suppose that for each $\delta > 0$, $(\cap_i C_i^\delta) \cap \Delta(\Omega) \neq \emptyset$. Since there are finitely many agents, there is an agent j such that (2) holds for arbitrarily small δ. This implies that (1) holds, contrary to our assumption. We conclude that for some $\delta > 0$, $(\cap_i C_i^\delta) \cap \Delta(\Omega) = \emptyset$. The last equality implies that the set $C^\delta = \times_i C_i^\delta$ and the diagonal D in $\Delta(\Omega)^N$ are disjoint. Since C^δ has a non-empty interior in $l_1(\Omega)^N$, there is a non-zero continuous functional $g = (g_i)_{i\in N}$ on $l_1(\Omega)^N$, where for each i, $g_i \in l_1(\Omega)$, and a constant c such that $gx \geq c \geq gy$ for each $x \in C^\delta$ and $y \in D$. For any $(x_i)_{i\in N} \in B^N$, $(t_i^\omega + \delta x_i)_{i\in N} \in C^\delta$. Thus, $\sum_i (g_i t_i^\omega + \delta g_i x_i) \geq c$. Suppose that $\sum_i g_i t_i^\omega = c$; then $\sum_i g_i x_i = 0$. But as $g \neq 0$ it cannot vanish on B^N. Thus, $\sum_i g_i t_i^\omega > c$. We transform the functional g into the desired bet in steps similar to the ones in the previous proof. The first two transformations of g

into $\hat{g}$ and then into $\hat{h}$ result in a bounded bet, $\hat{h}$.

Since for each i, $P_i \subset C_i^{\delta}$, g separates $Q = \times_i P_i$ from D. Let $\mathbf{c}$ be the constant function on Ω that takes the value c. Define $\hat{g} = (\hat{g}_i)_{i\in N}$ by $\hat{g}_i = g_i - \mathbf{c}/n$. Then, for all $x \in Q$ and $y \in D$, $\hat{g}x \geq 0 \geq \hat{g}y$, and $\sum_i \hat{g}_i t_i^{\omega} > 0$. The inequality $\hat{g}y \leq 0$, as in the previous proof, that $\sum_i \hat{g}_i \leq 0$.

Define $\hat{h} = (\hat{h}_i)_{i\in N}$ by $\hat{h}_i = \hat{g}^i - (\sum_i \hat{g}_i)/n$. Then, $\hat{h}_i \geq \hat{g}_i$ for each i, and therefore $\hat{h}x \geq 0$ for each $x \in Q$, and $\sum_i \hat{h}_i t_i^{\omega} \geq \sum_i \hat{g}_i t_i^{\omega} > 0$. Moreover $\sum_i \hat{h}_i = 0$.

We transform $\hat{h}$ into h to have a bet which is not rejected at any state. Let $a_i = \inf_{p_i \in P_i} \hat{h}_i p_i$. Then, $\sum_i a_i \geq 0$. Let $b_i = (\sum_i a_i)/n - a_i$. Then $\sum_i b_i = 0$ and for each i, $a_i + b_i \geq 0$. Define $h_i = \hat{h}_i + \mathbf{b}_i$. The functional $(h_i)_{i\in N}$ satisfies: $\sum_i h_i = 0$; for each i and $p_i \in P_i$, $h_i p_i \geq 0$; and $\sum_i h_i t_i^{\omega} > 0$.

Thus, $(h_i)_{i\in N}$ is a bounded bet, and for each i and ω', $E_i(h_i)(\omega') = h_i t_i^{\omega'} \geq 0$ and $\sum_i E_i(h_i)(\omega) = \sum_i h_i t_i^{\omega} > 0$. From these two inequalities it follows that for some j, $E_j(h_j)(\omega) > 0$. If $j = i_0$ then (h_i) is the required bet (f_i^{π}). Otherwise, we increase i_0's payoffs at ω on the expense of j's, as follows. Choose small enough $\varepsilon > 0$ such that $E_j(h_j)(\omega) > \varepsilon$, and let $f_j^{\pi}(\omega) = h_j(\omega) - \varepsilon$ and $f_{i_0}^{\pi}(\omega) = h_{i_0}(\omega) + \varepsilon$. For all other states and agents (f_i^{π}) coincides with (h_i). Since $t_{i_0}(\omega)(\omega) > 0$, $E_j(f_{i_0}^{\pi})(\omega) > 0$ as required. ∎

Proof of Proposition 2: Let $(\Omega, (\Pi_i)_{i\in N}, (t_i)_{i\in N})$ be a connected infinite type space. Suppose that μ is a common improper prior on this space that has a positive lower bound. It follows that every π in $\cup_i \Pi_i$ is finite, because for an infinite π, the requirement that $\mu(\pi) < \infty$ contradicts the existence of such a bound.

Consider a graph the nodes of which are the states and the edges are pairs (ω, ω') such that $\{\omega, \omega'\} \subseteq \pi$ for some $\pi \in \cup_i \Pi_i$. The graph is infinite, and since all the elements of $\cup_i \Pi_i$ are finite, each node has finitely many neighbors. Moreover, the connectedness of the type space means that the graph is connected. By König's lemma, König (1936), for each state ω_0 there exists a simple infinite path $\omega_0, \omega_1, \ldots$, where simplicity means that no state is repeated.

We first prove the proposition in the case that μ is constant. Fix ω_0, and let $\omega_0, \omega_1, \ldots$ be a simple infinite path and $i_1, i_2, \ldots$ an infinite sequence of elements in N such that $\omega_{k+1} \in \pi_{i_k}(\omega_k)$ for all $k \geq 0$. By Claim 1, it is enough to show that (1) does not hold when the space is weakly belief consistent. Let p be a common ε-prior. Then, for each i there is a prior p^i such that $||p - p^i|| \leq \varepsilon$. Denote $p_k = p(\omega_k)$ and $p_k^i = p^i(\omega_k)$. Then, $\sum_{k\geq 0} |p_{k+1} - p_k| = \sum_i \sum_{\{k|i_k=i\}} |p_{k+1} - p_k|$. Since μ is constant it follows that for $i_k = i$, $p_k^i = p_{k+1}^i$, and therefore $|p_{k+1} - p_k| \leq |p_{k+1} - p_{k+1}^i| + |p_k^i - p_k|$.

Thus, $\sum_{\{k|i_k=i\}} |p_{k+1} - p_k| \leq 2||p - p^i|| \leq 2\varepsilon$. Hence, $\sum_{k\geq 0} |p_{k+1} - p_k| \leq 2n\varepsilon$. Now, $|p_0 - p_m| \leq \sum_{k\geq 0}^{m-1} |p_{k+1} - p_k| \leq 2n\varepsilon$. As the elements of the path are distinct, $p_m \to 0$. Thus, $p_o \leq 2n\varepsilon$. Therefore for each ε, $p \in P^\varepsilon$, and ω, $p(\omega) \leq 2n\varepsilon$, and for each i, $p(\pi_i(\omega)) \leq 2n|\pi_i(\omega)|\varepsilon$, which shows that (1) does not hold.

We now drop the assumption that μ is constant. If for some ω, $\mu(\omega) = 0$, then by the connectedness of space and the positivity of the the type function it follows that μ is constantly 0 which is excluded by definition. Thus, $\mu(\omega) > 0$ for each ω. Consider the type space $(\Omega, (\Pi_i)_{i\in N}, (\hat{t}_i)_{i\in N})$, where for each i, $\pi \in \Pi_i$, and $\omega \in \pi$, $\hat{t}_i^\pi(\omega) = 1/|\pi|$. The constant function on Ω is a common improper prior for this space and thus, as we have shown, it is inconsistent. By Theorem 2, there exists a bounded agreeable bet $(\hat{f}_i)_{i\in N}$ on this space. Thus, for each i and $\pi \in \Pi_i$, $\sum_{\omega\in\pi} \hat{f}_i(\omega) > 0$. Define a bet $(f_i)_{i\in N}$ by $f_i(\omega) = \hat{f}_i(\omega)/\mu(\omega)$. Since μ is bounded away from 0, this is a bounded bet. For $\pi \in \Pi_i$, $E_i(\pi)(f_i) = \sum_{\omega\in\pi} t_i^\pi(\omega) f_i(\omega) = (1/\mu(\pi)) \sum_{\omega\in\pi} \hat{f}_i(\omega) > 0$. Thus, $(f_i)_{i\in N}$ is a bounded agreeable bet on the space $(\Omega, (\Pi_i)_{i\in N}, (t_i)_{i\in N})$. Hence the space is not trade consistent. ■

Proof of Proposition 3: Let $(\Omega, (\Pi_i)_{i\in N}, (t_i)_{i\in N})$ be a locally finite type space $(\Omega, (\Pi_i)_{i\in N})$ with positive type functions. If there is no common improper prior on the space, then by Hellman (2010) there exists a bounded agreeable bet on the space. If there is a common improper prior μ on the space, then as argued in the proof of Proposition 2, it follows from the positivity of the type functions that $\mu(\omega) > 0$ for each ω. This fact and the local finiteness assumption make it possible to construct a type space $(\Omega, (\Pi_i)_{i\in N}, (\hat{t}_i)_{i\in N})$ as in the proof of 2. Following this proof we also construct the agreeable bet $(f_i)_{i\in N}$. Since μ can approach 0, the bet is not necessarily bounded. ■

Proof of Proposition 4: To prove **(1)**, define for each $m \geq 1$ a probability distribution p^m on the interval of integers $[1, m]$ by $p^m(k) = \mu_k/S_m$, where $S_m = \sum_{k=1}^m \mu_k$. For even m, p^m is a prior for 1, and for odd m, p^m is a prior for 2. For each $m \geq 1$, $||p^m - p^{m+1}|| = \sum_{k=1}^m (\mu_k/S_m - \mu_k/S_{m+1}) + \mu_{m+1}/S_{m+1} = (1/S_m - 1/S_{m+1}) \sum_{k=1}^m \mu_k + \mu_{m+1}/S_{m+1} = [(S_{m+1} - S_m)/(S_{m+1}S_m)]S_m + \mu_{m+1}/S_{m+1} = 2\mu_{m+1}/S_{m+1}$. Thus, for each m, p^m is a $2\mu_{m+1}/S_{m+1}$-prior. This shows that if the condition on μ in **(1)** holds then the type space is weakly belief consistent.

Conversely, suppose that the condition on μ in **(1)** does not hold. Then, for some $\varepsilon > 0$, $\mu_m/S_m > \varepsilon$ for all m. As for each m, $\mu_m/S_m \leq \mu_m/S_1$ it follows that for all m, $\mu_m > \varepsilon S_1 > 0$. Let g be a function on the state space defined by $g(m) = S_m/\mu_m$ for $m \geq 1$. Define a bet (f_1, f_2) by $f_1(m) = g(m)$ for all odd m, $f_1(m) = -g(m)$ for all even

m, and $f_2(m) = -f_1(m)$ for all m. The bet is bounded as $|f_i(m)| = g(m) < 1/\varepsilon$. The bet satisfies for $\pi = \{m, m+1\} \in \Pi_i$, $E_i(f_i)(m) = -\mu_m g(m) + \mu_{m+1} g(m+1) = \mu_{m+1} > \varepsilon S_1 > 0$. Thus, by Theorem 1 the type space is not weakly trade consistent.

To prove **(2)**, consider the probability distribution p^m defined above. Denote $\varepsilon_m = 2\mu_{m+1}/S_{m+1}$. Then, as was shown above, $p^m \in P_1^{\varepsilon_m}$ for all m. Now, $p^m(\{1\})/\varepsilon_m = \mu_1/(S_m \varepsilon_m) = [\mu_1 S_{m+1}]/[2\mu_{m+1} S_m] = \mu_1/[2\mu_{m+1}] + \mu_1/[2S_m]$. If $\liminf_m \mu_m = 0$ then condition 1 is satisfied and the type space is belief consistent.

Conversely, if $\liminf_m \mu_m \neq 0$ then μ is bounded away from 0 and the type space is not consistent by Proposition 2.

Part **(3)** is obvious. A common prior is in particular a common improper prior the sum of which is 1. If there is a summable common improper prior, then by normalizing it we get a common prior. ■

References

Aumann, R. (1976), Agreeing to Disagree, *The Annals of Statistics*, 4, pp. 1236-1239.

Aumann, R. (1987), Correlated Equilibrium as an Expression of Bayesian Rationality, *Econometrica*, 55, pp. 1-18.

Aumann, R. (1998), Common Priors: A Reply to Gul, *Econometrica*, 66, pp. 929-938.

Bhattacharyya, S. and B. L. Lipman (1995), Ex Ante Versus Interim Rationality and the Existence of Bubbles, *Economic Theory*, 6, pp. 469-494.

Bonanno, G. and K. Nehring (1996), Fundamental Agreement: A New Foundation for the Harsanyi Doctrine, Working Paper No. 96-02, University of California, Davis.

Dunford N. and J.T. Schwartz (1957), *Linear Operators, Part I*, Interscience, New York.

Feinberg, Y. (2000), Characterizing Common Priors in the Form of Posteriors, *Journal of Economic Theory*, 91, pp. 127-179.

Geanakoplos, J. and H. Polemarchakis (1982), We Can't Disagree Forever, *Journal of Economic Theory*, 28, pp. 192–200.

Gul, F. (1998), A Comment on Aumann's Bayesian View, *Econometrica* 66, pp. 923–927.

Halpern, J. (2002), Characterizing the Common Prior Assumption, *Journal of Economic Theory*, 106, pp. 316-355.

Harsanyi, J. C. (1967-68), Games with Incomplete Information Played by Bayesian Players, parts I,II, & III, *Management Science*, 14, pp. 159-182, 320-334, 486-502.

Heifetz, A. (2006), The Positive Foundation of the Common Prior Assumption, *Games and Economic Behavior*, 56, pp. 105-120.

Hellman, Z. (2011), Iterated Expectations, Compact Spaces, and Common Priors, *Games and Economic Behavior*, 72, pp. 163-171.

Hellman, Z. (2010), Countable Spaces and Common Priors, *manuscript.*

König, D. (1936), *Theorie der Endlichen und Unendlichen Graphen: Kombinatorische Topologie der Streckenkomplexe*, Akad. Verlag, Leipzig

Milgrom, P. and N. L. Stokey (1982), Information, Trade and Common Knowledge, *Journal of Economic Theory*, 26, No. 1, pp. 17-27.

Morris, S. (1994), Trade with Heterogenous Prior Beliefs and Asymmetric Information, *Econometrica*, 62, pp. 1327-1347.

Rubinstein, A. (1989), The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge,, *American Economic Review*, 79, pp. 385–391.

Rubinstein, A. and A. Wolinsky (1990), On the logic of "Agreeing to Disagree" Type of Results, *Journal of Economic Theory*, 51, pp. 184–193.

Samet, D. (1998a), Common Priors and Iterated Expectations, *Games and Economic Behavior*, 24, No. 1, pp. 131–141.

Samet, D. (1998b), Common Priors and the Separation of Convex Sets, *Games and Economic Behavior*, 24, No. 1, pp. 172–174.

Sebenius, J., and J. Geanakoplos (1983), Don't Bet On It: Contingent Agreements with Asymmetric Information, *Journal of the American Statistical Association*, 78, pp. 424–426.