



ELSEVIER

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Belief consistency and trade consistency [☆]Ehud Lehrer ^{a,b,1}, Dov Samet ^{c,*,2}^a The School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel^b INSEAD, Boulevard de Constance, 77305 Fontainebleau, France^c Faculty of Management, Tel Aviv University, Tel Aviv 69978, Israel

ARTICLE INFO

Article history:

Received 4 April 2013

Available online 27 November 2013

JEL classification:

C70

D82

Keywords:

Type spaces

Common prior

No-trade theorems

Agreeing to disagree

Belief consistency

Dutch books

ABSTRACT

Interpersonal consistency can be described in epistemic terms as a property of beliefs, or in economic terms as the impossibility of certain trades. The existence of a common prior from which all agents' beliefs are derived is of the first kind. The non-existence of an *agreeable bet*, that is, a contingent zero-sum trade which is always favorable to all agents, is of the second kind. It is well established that these two notions of consistency are equivalent for finite type spaces but not for countable ones. We present three equivalences of epistemic consistency and economic consistency conditions for countable type spaces, defining in this way three levels of consistency of type spaces: weak consistency, consistency, and strong consistency. These three levels coincide in the finite case. We fully analyze the level of consistency of type spaces based on the knowledge structure of Rubinstein's email game. The new notion of *belief consistency* introduced here helps to justify the requirement of boundedness of payoff functions in countable type spaces by showing that in a large class of spaces there exists an agreeable unbounded bet even when a common prior exists.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

1.1. Interpersonal consistency

Interpersonal consistency or inconsistency are easily identified when agents are equally informed. In such a case, if agents' beliefs are different, then, obviously, their beliefs are inconsistent. This inconsistency is expressed in purely *epistemic* terms, that is, in terms of the agents' beliefs and knowledge. But it can also be expressed in *economic* terms, i.e., in terms of behavior and choices made. We can expect agents whose beliefs are inconsistent to act and make different economic choices, and, moreover, we can even envisage certain economic interactions between them that one might call inconsistent.

Matters are more complicated when agents are differently informed. In such cases differences of beliefs may reflect the difference in information and not necessarily inconsistency of the beliefs. The question of describing consistency in epistemic and economic terms in the case of differential information is the problem we deal with in this paper.

[☆] A helpful discussion with Oyama Daisuke is acknowledged.

* Corresponding author.

E-mail addresses: lehrer@post.tau.ac.il (E. Lehrer), dovsamet@gmail.com (D. Samet).

¹ Lehrer acknowledges support of the ISF through Grant #538/11.

² Samet acknowledges support of the ISF through Grant #1517/11.

1.2. Models of knowledge and belief

In the model we use for studying interpersonal consistency, beliefs are expressed by probability distributions. Harsanyi (1967–1968) modeled environments of probabilistic beliefs of many players for studying games with incomplete information. In this model each player has *types* each of which 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 and refer to it as a type space.

1.3. The history of the concepts of interpersonal consistency

1.3.1. Common priors

The first and most straightforward definition of consistency is in epistemic terms. We say that the beliefs of agents who are differently informed at present are consistent if these beliefs are derived, by conditioning on their information, from consistent beliefs they held in a previous period, when their knowledge was still the same. Being consistent when the knowledge is the same means having the *same* beliefs. Thus, beliefs are consistent when they are derived from beliefs previously shared by the agents.

This definition of consistency is known as the common prior assumption or Harsanyi's doctrine, because he made this assumption explicitly in his model (for a discussion and justification of the common prior assumption, see Aumann, 1987 and Bonanno and Nehring, 1996, 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.

1.3.2. The agreement theorem and no-trade theorems

Aumann's (1976) agreement theorem states that 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, expressed in *epistemic* terms, for epistemic consistency. 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, 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 (1994), 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, which we call *trade consistency*, says that there is no *agreeable bet*, namely, a state dependent transfer of money between the players such that at each state 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).

1.3.4. Dutch books

The inconsistency of a *single* agent's probabilistic beliefs can be described in economic terms by Dutch books: The inconsistent agent is willing to take a bet offered to her by another agent who is guaranteed to gain at the expense of the inconsistent agent. The inconsistency of Dutch books seems to be well described by the English idiom "A fool and his money are soon parted". The role of the external agent is essential in parting the inconsistent fool from his money. The economic description of interpersonal consistency, in contrast, does not require the exploitation of the agents by an external agent. Here inconsistency seems to be related to the violation of Parmenides' *maxim ex nihilo nihil fit*: In an agreeable bet, positive expected gains for all agents are guaranteed to come from nothing. However, it is simple to see that the existence of an agreeable bet is equivalent to the existence of a Dutch book. An agreeable bet exists if and only if an external agent can obtain the consent of each agent to trade with her contingently on the state, in a way that her gains are guaranteed. Thus, interpersonal inconsistency of beliefs can be used by an external agent to exploit the agents.

1.4. Equivalence in countable spaces

We now describe the results obtained in this paper for countable state spaces. Such 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 between trade consistency and belief consistency, stated above for finite and compact type spaces, fails for countable spaces. The existence of a common prior implies trade consistency for such spaces, but the opposite entailment does not hold, as was demonstrated by Feinberg (2000). Consistency in countable state spaces proved to be a complicated phenomenon, not fully understood, and no equivalence of economic and epistemic consistency was offered for such spaces (see, for example, Feinberg, 2000; Heifetz, 2006; Hellman, 2010).

The richer structure of countable space provides variants of epistemic and economic consistency that do not exist in the finite case. We exploit these variants to define new consistency conditions and classify countable state spaces as weakly consistent, consistent, and strongly consistent. Each class is characterized in economic and epistemic terms. The three equivalence theorems we state boil down in the finite case to the unique equivalence theorem that holds in this case.

1.4.1. Weak 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.

Weak belief consistency: A type space is weakly belief consistent if there exists a common ε -prior for each ε .

As for economic consistency, *weak trade consistency* holds when there is no agreeable bet which is bounded away from zero uniformly over states.

Weak trade consistency: A type space is weakly trade consistent if there is no $\delta > 0$ such that each agent's expected payoff at each state exceeds δ .

Our first result is:

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

1.4.2. Consistency

We call the condition of no agreeable bet *trade consistency*. As the existence of a common prior implies trade consistency and the reverse implication does not hold, we refer to the existence of a common prior as *strong belief consistency*. Finding a condition of belief consistency which is equivalent to trade consistency is a long standing open problem. In light of the weak consistency equivalence, the sought after condition of belief consistency must be stronger than weak belief consistency and strictly weaker than strong belief consistency. We next describe this condition.

A weakly consistent type space has a common ε -prior for every ε . However, if there is no common prior, the probabilities assigned by these common ε -priors to any given element of the partitions vanish when ε vanishes. The speed of this vanishing 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 vanishing ε 's that assign probabilities to the partition element which vanish infinitely more slowly than ε .

With this notion of consistency we can state our second major result:

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

1.4.3. Strong consistency

Strong trade consistency is the economic consistency condition which is equivalent to strong belief consistency, namely the existence of a common prior. It is defined in terms of ε -bets which are bets in which each player's loss in each state is smaller than ε .

Strong trade consistency: A type space is strongly trade consistent if the possible gains of a certain agent in a certain state in ε -bets vanish with ε .

Our third equivalence theorem states:

Strong consistency equivalence: A type space is strongly belief consistent if and only if it is strongly trade consistent.

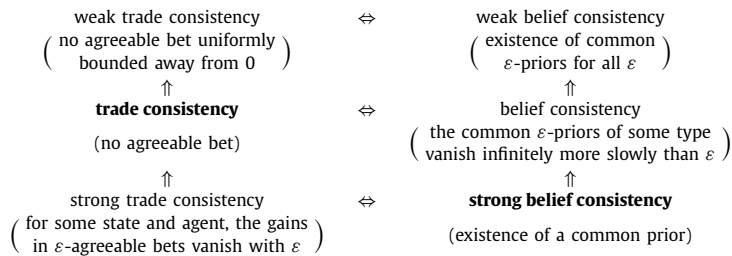


Diagram 1. Three equivalence theorems.

1.4.4. Summary

Diagram 1 summarizes succinctly the three equivalence theorems. The two conditions in boldface are the ones shown in the literature to be equivalent in finite spaces (see Theorem 1). The other four conditions are introduced here. In finite type spaces, the three economic conditions in the left column are equivalent and the three epistemic conditions in the right column are equivalent. Thus, the three equivalences across the columns coincide in the finite case.

1.5. The duality of beliefs and payoffs

Two classes of entities are studied here: beliefs, expressed as probability distributions over the states of the world, and contingent assets, expressed as state dependent payoff functions. Beliefs and contingent assets join together in generating expectation. Thus, we can view a probabilistic belief as a functional operating on state dependent payoff functions, forming their expectations and conversely, view a state dependent payoff function as a functional on belief distributions, forming the various expectations of the payoff function.

Formally, we consider probabilistic beliefs as elements of the normed space $l_1(\Omega)$ of summable functions on the state space Ω , and payoff functions as elements of $l_\infty(\Omega)$, the normed space of bounded functions on Ω . The space $l_\infty(\Omega)$ is the dual of $l_1(\Omega)$ and hence a probability distribution is a continuous functional on payoff functions. Conversely, $l_1(\Omega)$ is a subset of the dual of $l_\infty(\Omega)$, and thus a payoff function is a continuous functional on probability distributions.

The mutual duality of beliefs and payoffs is used here extensively in separating convex sets. Convexity permeates the model of beliefs as the prior beliefs of an agent form a convex set. Trades, which are payoff functions of the agents, emerge in the equivalence theorems of consistency and of weak consistency, as continuous functionals that separate convex sets defined in terms of probability distributions. The difficulty here is the condition required for separation that at least one of the convex sets has a non-empty interior. Convexity appears in the model of trades, since the set of differences between net trade and total expectation is a convex set. The existence of an agreeable bet is expressed in terms of separation of this set from the negative orthant of $l_\infty(\Omega)$. Common priors emerge here, in the equivalence theorem of strong consistency, as continuous functionals that separate these convex sets. The difficulty here is the lack of reflexivity of $l_1(\Omega)$; there are continuous functionals on $l_\infty(\Omega)$ that are finitely additive probability distributions but not sigma additive. The first are not of much use in our context, and one has to guarantee that the separating functionals are sigma additive.

1.6. Applications

We establish the level of consistency of type spaces of interest. In particular we use the new notion of belief consistency in cases where we have been unable to directly verify the existence or non-existence 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 consistency 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 strongly consistent. Thus, the existence of unbounded agreeable bets has no bearing on belief consistency.³

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

³ Bhattacharyya and Lipman (1995) look 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.

1.7. The paper's plan

Section 2 presents the model. Section 3 introduces the various kinds of belief consistency and their trade consistency 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 did not appear in the previous sections.

2. Preliminaries

Partition spaces. A *partition space* for a set $N = \{1, \dots, n\}$ of 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*.

Type spaces. A type space is a partition space where agents' probabilistic beliefs about the states vary with the state. Formally, a *type space* is a tuple $(\Omega, (\Pi_i)_{i \in N}, (t_i)_{i \in N})$, where $(\Omega, (\Pi_i)_{i \in N})$ is a partition space. For each i , $t_i : \Omega \rightarrow \Delta(\Omega)$ is i 's *type function*, where $\Delta(\Omega)$ is the set of all probability distributions on Ω . We write t_i^ω instead of $t_i(\omega)$, and call it the *type of i at ω* . 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 $t_i^\omega(\pi) = 1$. 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 . As π is the event that i 's type is t_i^π , we refer to π as a *type of i* . We consider $\Delta(\Omega)$ as a subset of $l_1(\Omega)$, in which $\Delta(\Omega)$ is closed.

Priors. Fix an agent i . The following three conditions on a probability distribution $p \in \Delta(\Omega)$ are equivalent:

1. For each $\pi \in \Pi_i$, if $p(\pi) > 0$, then $t_i^\pi(\cdot) = p(\cdot | \pi)$;
2. for each $\pi \in \Pi_i$, the equality $p = p(\pi)t_i^\pi$ holds on π ;
3. p is in the closed convex hull of i 's types $\{t_i^\omega | \omega \in \Omega\}$, or equivalently $\{t_i^\pi | \pi \in \Pi_i\}$.

A probability distribution that satisfies these conditions is called a *prior* for i . The set of all i 's priors, which by the third condition is the closed convex hull of i 's types, is denoted by P_i . A probability distribution in the set $P = \bigcap_i P_i$ is a *common prior*.

Expectations. A *payoff function* is a real valued function f on Ω . For a payoff function f and $p \in \Delta(\Omega)$, we denote by $Ex^p(f)$ the expectation of f with respect to p , when it exists. If $Ex^{t_i^\omega}(f)$ exists for each ω , we define the function $E_i(f)$ on Ω by $E_i(f)(\omega) = Ex^{t_i^\omega}(f)$. As $E_i(f)$ is constant on each $\pi \in \Pi_i$, we can write for $\pi = \Pi_i(\omega)$, $E_i(f)(\pi)$ instead of $E_i(f)(\omega)$. For $E_i(f)$ to exist it is enough that f is bounded on each $\pi \in \Pi$. This holds, of course, when f is bounded, i.e., $f \in l_\infty(\Omega)$.

3. Consistency of type spaces

Economic consistency of type spaces is expressed in terms of bets, which are state contingent zero-sum transfers among the agents. Formally,

Definition 1. A *bet* on a type space is an n -tuple of bounded payoff functions (f_1, \dots, f_n) on Ω , such that $\sum_i f_i = 0$.⁴

Definition 2. A bet (f_1, \dots, f_n) is *agreeable* if for each i , $E_i(f_i) > 0$.

In a bet (f_1, \dots, f_n) , we think of $f_i(\omega)$ as being agent's i payoff at state ω . The bet is agreeable if each player's expected payoff at each state is positive.

Definition 3. A type space is *trade consistent* if it has no agreeable bet.

It is possible to describe trade consistency in terms of Dutch books rather than agreeable bets. A *Dutch book* against the agents is an n -tuple of bounded functions (g_1, \dots, g_n) , such that $\sum_i g_i < 0$ and for each i , $E_i(g_i) > 0$. We think of g_i as a gamble offered by a bookie to agent i , in which the bookie pays $g_i(\omega)$ to i , at state ω . The first condition means that the bookie is guaranteed to gain, the second, that the agents are willing to take the gamble.

Claim 1. A type space has an agreeable bet if and only if it has a Dutch book.

⁴ The boundedness requirement on bets, which necessarily holds for finite spaces, is essential for infinite spaces. We show in Section 4.2 that unbounded agreeable bets exist independently of the consistency of the type space.

Indeed, suppose that (f_1, \dots, f_n) is an agreeable bet and define $g_i = f_i - E_i(f_i)/2$. Then, $\sum_i g_i = \sum_i -E_i(f_i)/2 < 0$, and $E_i(g_i) = E_i(f_i) - E_i(f_i)/2 > 0$. Thus, (g_1, \dots, g_n) is a Dutch book. For a Dutch book (g_1, \dots, g_n) , let $f_i = g_i - (\sum_j g_j)/n$. Then, $\sum_i f_i = 0$, and as $f_i > g_i$, $E_i(f_i) > E_i(g_i) > 0$, and hence (f_1, \dots, f_n) is an agreeable bet.

Our starting point for studying consistency of type spaces is the well-known equivalence theorem for finite type spaces (Morris, 1994; Samet, 1998b; Feinberg, 2000).⁵

Theorem 1. *A finite type space is trade consistent if and only if it has a common prior.*

For countable type spaces this equivalence does not hold. The existence of a common prior implies trade consistency, but as was shown by Feinberg (2000), there are type spaces which are trade consistent but do not have a common prior. As the condition of having a common prior is stronger than trade consistency, we adopt the following terminology:

Definition 4. *A type space is strongly belief consistent if it has a common prior.*

No reformulation of Theorem 1 was suggested in the literature that also extends to countable spaces. Here, we show that countable spaces have three levels of consistency: *weak consistency*, *consistency*, and *strong consistency*. We show how each level of consistency can be equivalently described in economic and epistemic terms. First, we define two new notions, *weak belief consistency* and *weak trade consistency*, and state their equivalence. Next we define *belief consistency* and show that it is equivalent to trade consistency in Definition 3. Finally, we define *strong trade consistency* and show its equivalence to strong belief consistency in Definition 4. In the finite case, the three conditions of economic consistency are equivalent and the three conditions of epistemic consistency are equivalent. Hence, the restriction of the three new equivalence theorems to finite spaces is Theorem 1.

3.1. Weak consistency

Weak trade consistency requires the non-existence of agreeable bets in which the expectation of each agent is not only positive, but also bounded away from 0 uniformly in all states. Formally,

Definition 5. *A type space is weakly trade consistent if there is no agreeable bet (f_1, \dots, f_n) on it for which $\inf E_i(f_i) > 0$ for each i .*

Weak belief consistency weakens the requirement of a common prior by requiring the existence of probabilities which are arbitrarily close to the sets of priors of all agents. For $\varepsilon \geq 0$, a probability distribution p is an ε -prior for i if for some $p_i \in P_i$, $\|p - p_i\| \leq \varepsilon$. The set of all ε -priors of i is denoted by P_i^ε . An element of $P^\varepsilon = \bigcap_{i \in N} P_i^\varepsilon$ 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 6. *The type space is weakly belief consistent if it has a common ε -prior for each $\varepsilon > 0$.*

To justify the term weak in this definition, note that a common prior is in particular a common ε -prior for each $\varepsilon > 0$, and therefore strong belief consistency implies weak belief consistency.

Theorem 2. *A type space is weakly trade consistent if and only if it is weakly belief consistent.*

3.2. Consistency

We define belief consistency which implies weak belief consistency and is implied by strong belief consistency. To motivate this definition we first provide a necessary and sufficient condition for the existence of a common prior, in terms of common ε -priors. Obviously, weak belief consistency is a necessary condition, because a common prior is in particular a common ε -prior for each ε . The existence of a common prior is equivalent to further requiring that $p^\varepsilon(\pi)$, for common ε -priors p^ε , do not vanish for all π when ε vanishes.

In what follows we adopt the convention that the supremum over an empty set is $-\infty$. Note also that as $P^{\varepsilon'} \subseteq P^\varepsilon$ for $\varepsilon' < \varepsilon$, $\sup_{p^\varepsilon \in P^\varepsilon} p^\varepsilon(\pi)$ is an increasing function of ε , and therefore it has a limit in the extended real line when ε approaches 0.

⁵ This theorem also holds for infinite type spaces when the type space is Hausdorff compact and bets are continuous functions (Feinberg, 2000; Heifetz, 2006).

Proposition 1. A type space is strongly belief consistent, that is, it has a common prior, if and only if for some $\pi \in \bigcup_i \Pi_i$,

$$\lim_{\varepsilon \rightarrow 0^+} \sup_{p^\varepsilon \in P^\varepsilon} p^\varepsilon(\pi) > 0. \tag{1}$$

Since belief consistency should be stronger than weak belief consistency it should imply the existence of a common ε -prior for each $\varepsilon > 0$. Since it should fall short of requiring the existence of a common prior, it must allow for the vanishing of $p^\varepsilon(\pi)$ when ε vanishes, for each π . That is, it should allow for the limit in (1) to be 0 for each π . The way belief consistency strengthens weak belief consistency is by restricting the rate of the vanishing $p^\varepsilon(\pi)$'s. This rate is measured by the limit properties of $p^\varepsilon(\pi)/\varepsilon$. We require that for some π and a vanishing sequence of ε 's, $p^\varepsilon(\pi)$ vanishes infinitely more slowly than ε . Formally,

Definition 7. The type space is belief consistent if for some $\pi \in \bigcup_i \Pi_i$,

$$\sup_{\varepsilon > 0} \sup_{p^\varepsilon \in P^\varepsilon} p^\varepsilon(\pi)/\varepsilon = \infty. \tag{2}$$

Claim 2. Belief consistency implies weak belief consistency and is implied by strong belief consistency.

Indeed, if the type space is not weakly belief consistent, then for small enough positive ε , $P^\varepsilon = \emptyset$, which means that the limit in (2) is $-\infty$. 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 (2) holds.

Theorem 3. A type space is trade consistent if and only if it is belief consistent.

3.3. Strong consistency

Strong belief consistency, namely the existence of a common prior, is stronger than trade consistency. We now present an economic consistency condition which is stronger than trade consistency and is equivalent to strong belief consistency. Thus, this condition must prohibit certain bets even when they are not agreeable.

A bet (f_1, \dots, f_n) is ε -agreeable if for each i , $E_i(f_i) > -\varepsilon$. The set of all ε -agreeable bets is denoted by \mathcal{F}_ε . Obviously, these sets are homogeneous in ε , that is for each ε , $\mathcal{F}_\varepsilon = \varepsilon \mathcal{F}_1$. To measure agent i 's gains at state ω when expectation losses of ε are allowed, we consider $d_i^\omega(\varepsilon) = \sup_{(f_1, \dots, f_n) \in \mathcal{F}_\varepsilon} E_i(f_i)(\omega)$. By the homogeneity of the sets \mathcal{F}_ε it follows that $d_i^\omega(\varepsilon) = \varepsilon d_i^\omega(1)$. Therefore, there are two possibilities: either $d_i^\omega(1) = \infty$, in which case $d_i^\omega(\varepsilon) = \infty$ for each $\varepsilon > 0$, or $d_i^\omega(1) < \infty$, in which case $\inf_{\varepsilon > 0} d_i^\omega(\varepsilon) = 0$. Strong trade consistency requires that for at least one agent and one state, the second possibility holds. We define formally,

Definition 8. A type space is strongly trade consistent if for some i and ω ,

$$\inf_{\varepsilon > 0} \sup_{(f_1, \dots, f_n) \in \mathcal{F}_\varepsilon} E_i(f_i)(\omega) = 0,$$

or alternatively, $\sup_{(f_1, \dots, f_n) \in \mathcal{F}_1} E_i(f_i)(\omega) < \infty$.

Strong trade consistency implies trade consistency. To see this, suppose that trade consistency does not hold. Then there exists an agreeable bet (f_1, \dots, f_n) . As $(f_1, \dots, f_n) \in \mathcal{F}_\varepsilon$ for all $\varepsilon > 0$, it follows that for each i and ω , $0 < E_i(f_i)(\omega) \leq \inf_{\varepsilon > 0} \sup_{(f_1, \dots, f_n) \in \mathcal{F}_\varepsilon} E_i(f_i)(\omega)$.

Theorem 4. A type space is strongly belief consistent if and only if it is strongly trade consistent.

In light of Theorems 2, 3, and 4 we refer henceforth to weak consistency, consistency, and strong consistency of type spaces without the qualification of trade or belief.

3.4. Finite type spaces

If for a finite space Ω , there exists a common ε -prior for each $\varepsilon > 0$, then by the compactness of $\Delta(\Omega)$ there is an accumulation point of these probability distributions and it is a common prior. Thus, for finite spaces weak belief consistency and strong belief consistency are equivalent. Hence, the three epistemic conditions of consistency in Theorems 2, 3, and 4, are equivalent for finite spaces. It follows by these theorems that the three economic conditions of consistency also coincide, and the three theorems state the same equivalence for finite spaces: that of Theorem 1.

4. Applications

4.1. Common improper priors

A common improper prior for a type space is a measure μ on Ω , which is not identically null, 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 has a common improper prior. Hellman (2010) proved half of this conjecture: a consistent type space must have a common improper prior. He also provided an example that shows that 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. Agreeable unbounded bets

Using the result in the previous subsection we are able to demonstrate the importance of the boundedness condition on payoff functions in Definition 1 of bets. We call an n -tuple of real valued functions (f_1, \dots, f_n) , which are not necessarily bounded, an *integrable bet* if $\sum_i f_i = 0$, and for each i , $E_i(f_i)$ is well defined. As in the case of bets, an integrable bet (f_1, \dots, f_n) is *agreeable* if for each i , $E_i(f_i) > 0$. A type space is *locally finite* if each of the partitions' elements is finite.

We show that agreeable integrable bets exist on a large family of type spaces that also include spaces that have a common prior.

Proposition 3. *On each locally finite type space with positive type functions there exists an agreeable integrable 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 for two agents, using the equivalence theorems of Section 3.

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 $k \geq 1$, $E_k = \bigcup_{\omega \in E_{k-1}} \bigcup_i \Pi_i(\omega)$. The union $E = \bigcup_{k \geq 0} E_k$ 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.⁶ 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 $\Pi_1 = \{\{1\}, \{2, 3\}, \dots, \{2k, 2k+1\}, \dots\}$, and $\Pi_2 = \{\{1, 2\}, \dots, \{2k-1, 2k\}, \dots\}$. The basic partition space can model the sequential exchange of messages between the two agents 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' 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, the type space is*

- (1) *strongly consistent if and only if $\sum_m \mu_m < \infty$;*
- (2) *consistent if and only if $\liminf_m \mu_m = 0$;*
- (3) *weakly consistent if and only if $\liminf_m \mu_m / \sum_{k=1}^m \mu_k = 0$.*

⁶ 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, common knowledge is still finitely generated in E .

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 $m \rightarrow \infty$, 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 strongly consistent. 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 not strongly consistent.

Consistency in the exponential family depends on α as follows.

- $\alpha < 1$. Strong consistency holds, since $\sum_{m=1}^\infty \alpha^m = \alpha/(1 - \alpha) < \infty$.
- $\alpha = 1$. Weak consistency holds but not consistency, since $\mu_m / \sum_{k=1}^m \mu_k = 1/m \rightarrow 0$, but $\mu_m = 1$ for each m .
- $\alpha > 1$. Even weak consistency does not hold, 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.

- $\beta < -1$. Strong consistency holds, since $\sum_{m=1}^\infty m^\beta < \infty$.
- $-1 \leq \beta < 0$. Consistency holds but not strong consistency, since $m^\beta \rightarrow 0$, but the series $\sum_{m=1}^\infty m^\beta$ majorizes the harmonic series obtained for $\beta = -1$, and therefore diverges.
- $\beta \geq 0$. Weak consistency holds but not consistency, since for $x \geq 0$, x^β is increasing, and thus $m^\beta / \sum_{k=1}^m k^\beta \leq m^\beta / \int_0^m x^\beta dx = (\beta + 1)/m \rightarrow 0$. However, for each m , $m^\beta \geq 1$.

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 part (2) of Proposition 4 but not the one in part (1).

5. Open problems

5.1. Constructing bets

In order to show that a type space is not consistent we can either show that it is not belief consistent, or, alternatively, construct an agreeable bet. Theorem 3 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 (3) of Proposition 4. However, in order to prove inconsistency for more general spaces that have a common improper prior, as in Proposition 2, 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 constructively describe agreeable bets in countable spaces in terms of a common improper prior, when such bets exist.

5.2. Improper priors and consistency

To some extent, the existence of a common improper prior signifies 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. If there exists a common prior p , then for some π , $p(\pi) > 0$, and as $p \in P^\varepsilon$ for all $\varepsilon > 0$, (1) holds. If (1) holds, then there exists a sequence of probability distributions (p^m) , and a sequence of priors of i , (p_i^m) , for each i , such that $\lim_m \|p^m - p_i^m\| = 0$, and $\lim_m p^m(\pi_0) > 0$ for some π_0 . Using a diagonal argument we can assume without loss of generality that $\lim_m p^m(\omega)$ exists for each ω and denote it by $p(\omega)$. Since for each ω , $\lim_m |p^m(\omega) - p_i^m(\omega)| = 0$, it follows that $\lim_m p_i^m(\omega) = p(\omega)$ for each ω . The prior p_i^m is a convex combination of i ’s types: $p_i^m = \sum_{\pi \in \Pi_i} c_i^m(\pi) t_i^\pi$, and by the convergence of the priors at each state, $c_i(\pi) = \lim_m c_i^m(\pi)$ exists for each $\pi \in \Pi_i$, and $p = \sum_{\pi \in \Pi_i} c_i(\pi) t_i^\pi$ on π . Hence, $p = \sum_{\pi \in \Pi_i} c_i(\pi) t_i^\pi$. For each m , $1 = \|p_i^m\| = \sum_{\pi \in \Pi_i} c_i^m(\pi)$. Therefore, $p(\Omega) = \sum_{\pi \in \Pi_i} c_i(\pi) \leq 1$. By our assumption $p(\pi_0) > 0$, thus p can be normalized to a probability distribution which is a convex combination of i ’s types for each i . □

Proof of Theorem 2. Suppose first that the type space is not weakly trade consistent. That is, there exists an agreeable bet (f_1, \dots, f_n) and $\delta > 0$, such that $E_i(f_i)(\omega) > \delta$ for each i 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 $\|p^\varepsilon - p_i\| < \varepsilon$. $\text{Ex}^{p^\varepsilon}(f_i) = \text{Ex}^{p^\varepsilon}(f_i) - \text{Ex}^{p_i}(f_i) + \text{Ex}^{p_i}(f_i) \geq -\varepsilon + \text{Ex}^{p_i}(E_i(f_i)) \geq -\varepsilon + \delta$. However, $\sum f_i = 0$ implies $\sum \text{Ex}^{p^\varepsilon}(f_i) = 0$. 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, $g = (g_1, \dots, g_n)$, 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 $\hat{g} = (\hat{g}_1, \dots, \hat{g}_n)$ 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 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 $\hat{h} = (\hat{h}_1, \dots, \hat{h}_n)$ 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 $h_i = \hat{h}_i + b_i$. The functional $h = (h_1, \dots, h_n)$ 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. \square

Proof of Theorem 3. 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 $\text{Ex}^{q_i}(\|t_i - s_i\|) \leq 2\varepsilon$.

Proof. Since i is fixed we suppress the index i in the proof. By the definition of types, $\text{Ex}^q(\|t - s\|) = \sum_{\pi \in \Pi} q(\pi) \|t(\pi) - s(\pi)\| \leq \sum_{\pi \in \Pi} \|q(\pi)t(\pi) - p(\pi)t(\pi)\| + \sum_{\pi \in \Pi} \|p(\pi)t(\pi) - q(\pi)s(\pi)\|$. 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 $\sum_{\pi \in \Pi} \sum_{\omega \in \pi} |p(\omega) - q(\omega)| = \|q - p\| \leq \varepsilon$. \square

Lemma 2. If $p \in P^\varepsilon$, then for any bet (f_1, \dots, f_n) , $\text{Ex}^p(\sum_i E_i(f_i)) \leq 2n(\sum_i \|f_i\|_\infty)\varepsilon$.

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

Assume that the type space is belief consistent, and (2) holds for $\pi = \Pi_i(\omega)$. Suppose to the contrary that (f_1, \dots, f_n) is a bounded agreeable bet. Since for each j , $E_j(f_j) > 0$, it follows that for any p , $\text{Ex}^p(\sum_j E_j(f_j)) \geq \text{Ex}^p(E_i(f_i)) \geq p(\pi)E_i(f_i)(\omega)$. 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, $\text{Ex}^{p^\varepsilon}(\sum_j E_j(f_j)) > 2n\varepsilon$, 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 $(f_1^\pi, \dots, f_n^\pi)$ such that $E_i(f_i^\pi)(\pi) > 0$, and for each j , $E_j(f_j^\pi) \geq 0$. Let π_1, π_2, \dots be an enumeration of the elements of $\bigcup_i \Pi_i$. The bet (f_1, \dots, f_n) , where $f_i = \sum_{k \geq 1} 2^{-k} f_i^{\pi_k}$, is a bounded agreeable bet.

The bet $(f_1^\pi, \dots, f_n^\pi)$ 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 $t_{i_0}(\pi)(\omega) > 0$. 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 $l_1(\Omega)$, $C_i^{2,\delta}$ is the closed convex hull of $\{t_i^{\omega'} \mid t_i^{\omega'} \neq t_i^\omega\}$, and C_i^δ is the closed convex hull of $C_i^{1,\delta} \cup C_i^{2,\delta}$. Suppose that a probability distribution $p \in \Delta(\Omega)$ is in $\bigcap_i C_i^\delta$. 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, $p_j(\Pi_j(\omega)) = \alpha_j t_j^\omega(\Pi_j(\omega)) + (1 - \alpha_j)q_j(\Pi_j(\omega)) = \alpha_j$. Hence, $p(\Pi_j(\omega)) \geq \alpha_j - \alpha_j \delta$, and therefore

$$p(\Pi_j(\omega))/(\alpha_j \delta) \geq 1/\delta - 1. \tag{3}$$

Suppose that for each $\delta > 0$, $(\bigcap_i C_i^\delta) \cap \Delta(\Omega) \neq \emptyset$. Since there are finitely many agents, there is an agent j such that (3) holds for arbitrarily small δ . This implies that (2) holds, contrary to our assumption. We conclude that for some $\delta > 0$, $(\bigcap_i C_i^\delta) \cap \Delta(\Omega) = \emptyset$. 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 $g = (g_1, \dots, g_n)$ on $l_0(\Omega)^N$, where for each i , $g_i \in l_\infty(\Omega)$, 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 c be the constant function on Ω that takes the value c . Define $\hat{g} = (\hat{g}_1, \dots, \hat{g}_n)$ by $\hat{g}_i = g_i - c/n$. 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 proof of Theorem 2, implies that $\sum_i \hat{g}_i \leq 0$.

Define $\hat{h} = (\hat{h}_1, \dots, \hat{h}_n)$ 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 $h_i = \hat{h}_i + b_i$. The functional (h_1, \dots, h_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_1, \dots, h_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_1, \dots, h_n) is the required bet (f_j^π) . Otherwise, we increase i_0 's payoffs at ω at the expense of j 's, as follows. Choose a 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 $(f_1^\pi, \dots, f_n^\pi)$ coincides with (h_1, \dots, h_n) . Since $t_{i_0}(\omega)(\omega) > 0$, $E_j(f_{i_0}^\pi)(\omega) > 0$ as required. \square

Proof of Theorem 4. Suppose that there is a common prior p . Choose ω with $p(\omega) > 0$ and fix i . Assume that strong trade consistency does not hold. Then, for all $\varepsilon > 0$, $\sup_{(f_i) \in \mathcal{F}_\varepsilon} E_i(f_i)(\omega) = \infty$. Thus, for $\varepsilon = p(\omega)/(2n)$ there exists an ε -agreeable bet (f_i) such that $E_i(f_i)(\omega) > 1$. Now, $\text{Exp}^p(\sum_j f_j) = \sum_j \text{Exp}^p(f_j) = \sum_j \text{Exp}^p(E_j(f_j)) \geq -n\varepsilon + p(\omega) > 0$. But, as $\sum_j f_j = 0$, $\text{Exp}^p(\sum_j f_j) = 0$.

Suppose, conversely, that strong trade consistency holds. Then, for some i_0 and ω_0 ,

$$\inf_{\varepsilon > 0} \sup_{(f_1, \dots, f_n) \in \mathcal{F}_\varepsilon} E_{i_0}(f_{i_0})(\omega_0) = 0.$$

Since for all $\omega \in \Pi_i(\omega_0)$, $E_{i_0}(f_{i_0})(\omega) = E_{i_0}(f_{i_0})(\omega_0)$, we can assume without loss of generality that $t_{i_0}^{\omega_0}(\omega_0) > 0$.

Let $C = \{\sum_i (f_i - E_i(f_i)) \mid f_i \in L_\infty(\Omega)\}$, L_- be the negative orthant of $L_\infty(\Omega)$, and B the unit ball in $l_\infty(\Omega)$. We show that strong trade consistency implies that for some $\varepsilon > 0$, $\text{conv}(C, 1_{\omega_0} + \varepsilon B) \cap L_- = \emptyset$. That is, the convex hull of C and $1_{\omega_0} + \varepsilon B$ is disjoint of L_- . Consider $\varepsilon > 0$ for which the intersection is not empty. Then for some $\alpha \in [0, 1]$, $h \in B$, and $(f_1, \dots, f_n) : \alpha \sum_i (f_i - E_i(f_i)) + (1 - \alpha)(1_{\omega_0} + \varepsilon h) < 0$. If $\alpha = 1$, then $\sum_i (f_i - E_i(f_i)) < 0$. Let $g = -\sum_i (f_i - E_i(f_i))$ and for each i , $g_i = f_i - E_i(f_i) + g/n$. Then $\sum_i g_i = 0$ and for each i , $E_i(g_i) = E_i(g/n)$ and as $g > 0$, $E_i(f_i) > 0$. Thus, trade consistency does not hold, and a fortiori strong trade consistency does not hold, contrary to our assumption. Thus, $\alpha < 1$, and we can define a function g by $g = \alpha(1 - \alpha)^{-1} \sum_i (f_i - E_i(f_i)) + 1_{\omega_0} - \varepsilon h < 0$.

Define for each $i \neq i_0$, $g_i = \alpha(1 - \alpha)^{-1} (f_i - E_i(f_i)) + (\varepsilon h - g)/(n - 1)$, and $g_{i_0} = \alpha(1 - \alpha)^{-1} (f_{i_0} - E_{i_0}(f_{i_0})) + 1_{\omega_0}$. Then, $\sum_i g_i = 0$. Also, for all $i \neq i_0$, $E_i(g_i) \geq E_i(\varepsilon h) + E_i(-g) > E_i(\varepsilon h) \geq -\varepsilon$, and $E_{i_0}(g_{i_0}) \geq 0$. Therefore, (g_1, \dots, g_n) is an ε -agreeable bet. Note, however, that $E_{i_0}(g_{i_0})(\omega_0) = E_i(1_{\omega_0})(\omega_0) = t_{i_0}^{\omega_0}(\omega_0)$. Therefore, if $\text{conv}(C, 1_{\omega_0} + \varepsilon B) \cap L_- \neq \emptyset$ for all $\varepsilon > 0$, then $\inf_{\varepsilon > 0} \sup_{(f_1, \dots, f_n) \in \mathcal{F}_\varepsilon} E_{i_0}(f_{i_0})(\omega_0) \geq t_{i_0}^{\omega_0}(\omega_0) > 0$, contrary to our assumption.

Let $\varepsilon > 0$ be such that $\text{conv}(C, 1_{\omega_0} + \varepsilon B) \cap L_- = \emptyset$, as both sets are convex and have non-empty interior there is a non-zero continuous functional φ on $l_\infty(\Omega)$ (namely, a finitely additive measure on $l_\infty(\Omega)$) such that $\varphi(g) \geq 0 \geq \varphi(h)$ for all $g \in \text{conv}(C, 1_{\omega_0} + \varepsilon B)$ and $h \in L_-$. We use φ to construct a common prior.

We show first that for each ω , $\varphi(1_\omega) \geq 0$ and $0 < \sum_{\omega'} \varphi(1_{\omega'}) < \infty$. Indeed, for $\varepsilon > 0$, $-1_\omega - \varepsilon \in L_-$. Thus $-\varphi(1_\omega) - \varepsilon \varphi(1_\Omega) \leq 0$. Since this is true for all positive ε , $\varphi(1_\omega) \geq 0$. As $\varphi \neq 0$, there exists $h \in B$ such that $\varphi(h) > 0$. Since $1_{\omega_0} - \varepsilon h \in 1_{\omega_0} + \varepsilon B$, it follows that $\varphi(1_{\omega_0} - \varepsilon h) \geq 0$, that is, $\varphi(1_{\omega_0}) \geq \varepsilon \varphi(h) > 0$. Thus, $0 < \sum_{\omega'} \varphi(1_{\omega'})$. Finally, for any finite subset Ω' of Ω , $\sum_{\omega \in \Omega'} \varphi(1_\omega) = \varphi(1_{\Omega'}) \leq \|\varphi\| \|1_{\Omega'}\| = \|\varphi\|$. Thus $\sum_{\omega'} \varphi(1_{\omega'}) \leq \|\varphi\| < \infty$. Define a probability function p on Ω by $p(\omega) = \varphi(1_\omega) / \sum_{\omega'} \varphi(1_{\omega'})$.

To see that p is a common prior, set, for $\omega \in \pi \in \Pi_i$, $f_i = 1_\omega$ and $f_j = 0$ for $j \neq i$. Then, $1_\omega - E_i(1_\omega) \in C$. Similarly $-1_\omega - E_i(-1_\omega) \in C$. Thus, $\varphi(1_\omega - E_i(1_\omega)) = 0$. Since $E_i(1_\omega) = t_i^\pi(\omega) 1_\pi$, we have $\varphi(1_\omega) = t_i^\pi(\omega) \varphi(1_\pi)$. Summing over all ω' in π we conclude that $\sum_{\omega' \in \pi} \varphi(1_{\omega'}) = \varphi(1_\pi)$. Hence $\varphi(1_\omega) = t_i^\pi(\omega) \sum_{\omega' \in \pi} \varphi(1_{\omega'})$, which implies $p(\omega) = t_i^\pi(\omega) p(\pi)$. \square

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 $\bigcup_i \Pi_i$ 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 $\pi \in \bigcup_i \Pi_i$. The graph is infinite, and since all the elements of $\bigcup_i \Pi_i$ 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 (see Diestel, 2000), for each state ω_0 there exists a simple infinite path $\omega_0, \omega_1, \dots$, 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 $\omega_0, \omega_1, \dots$ be a simple infinite path and i_1, i_2, \dots 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 2, it is enough to show that (2) 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 $|p_{k+1} - p_k| \leq |p_{k+1} - p_{k+1}^i| + |p_{k+1}^i - p_k|$. 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, $|p_0 - p_m| \leq \sum_{k \geq 0} |p_{k+1} - p_k| \leq 2n\varepsilon$. As the elements of the path are distinct, $p_m \rightarrow 0$. Thus, $p_0 \leq 2n\varepsilon$. 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 (2) 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 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 3, there exists a bounded agreeable bet $(\hat{f}_1, \dots, \hat{f}_n)$ 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_1, \dots, f_n) by $f_i(\omega) = \hat{f}_i(\omega)/\mu(\omega)$. Since μ is bounded away from 0, this is a bounded bet. For $\pi \in \Pi_i$, $E_i(\pi)(f_i) = \sum_{\omega \in \pi} \hat{t}_i^\pi(\omega) f_i(\omega) = (1/\mu(\pi)) \sum_{\omega \in \pi} \hat{f}_i(\omega) > 0$. Thus, (f_1, \dots, f_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. \square

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 n -tuple of functions (f_1, \dots, f_n) . Since μ can approach 0, these functions are not necessarily bounded, but by local finiteness, $E_i(f_i)$ is defined for each i . Thus, (f_1, \dots, f_n) is an agreeable integrable bet. \square

Proof of Proposition 4. Part (1) is obvious. A common prior is in particular a common improper prior that sums to 1. If there is a summable common improper prior, then by normalizing it we get a common prior.

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 (2) 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.

To prove (3), 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 (3) holds then the type space is weakly belief consistent.

Conversely, suppose that the condition on μ in (3) 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 2 the type space is not weakly trade consistent. \square

References

Aumann, R., 1976. Agreeing to disagree. *Ann. Stat.* 4, 1236–1239.
 Aumann, R., 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
 Aumann, R., 1998. Common priors: A reply to Gul. *Econometrica* 66, 929–938.
 Bhattacharyya, S., Lipman, B.L., 1995. Ex ante versus interim rationality and the existence of bubbles. *Econ. Theory* 6, 469–494.
 Bonanno, G., Nehring, K., 1996. Fundamental agreement: A new foundation for the Harsanyi doctrine. Working Paper No. 96-02. University of California, Davis.
 Diestel, R., 2000. *Graph Theory*. Springer-Verlag, New York.

- Dunford, N., Schwartz, J.T., 1957. *Linear Operators*, part I. Interscience, New York.
- Feinberg, Y., 2000. Characterizing common priors in the form of posteriors. *J. Econ. Theory* 91, 127–179.
- Gul, F., 1998. A comment on Aumann's Bayesian view. *Econometrica* 66, 923–927.
- Halpern, J., 2002. Characterizing the common prior assumption. *J. Econ. Theory* 106, 316–355.
- Harsanyi, J.C., 1967–1968. Games with incomplete information played by Bayesian players, parts I, II, & III. *Manage. Sci.* 14, 159–182. 320–334, 486–502.
- Heifetz, A., 2006. The positive foundation of the common prior assumption. *Games Econ. Behav.* 56, 105–120.
- Hellman, Z., 2010. Countable spaces and common priors. DP 604. Center for the Study of Rationality, The Hebrew University, Jerusalem.
- Hellman, Z., 2011. Iterated expectations, compact spaces, and common priors. *Games Econ. Behav.* 72, 163–171.
- Milgrom, P., Stokey, N.L., 1982. Information, trade and common knowledge. *J. Econ. Theory* 26 (1), 17–27.
- Morris, S., 1994. Trade with heterogenous prior beliefs and asymmetric information. *Econometrica* 62, 1327–1347.
- Rubinstein, A., 1989. The electronic mail game: Strategic behavior under almost common knowledge. *Amer. Econ. Rev.* 79, 385–391.
- Rubinstein, A., Wolinsky, A., 1990. On the logic of “agreeing to disagree” type of results. *J. Econ. Theory* 51, 184–193.
- Samet, D., 1998a. Common priors and iterated expectations. *Games Econ. Behav.* 24 (1), 131–141.
- Samet, D., 1998b. Common priors and the separation of convex sets. *Games Econ. Behav.* 24 (1), 172–174.
- Sebenius, J., Geanakoplos, J., 1983. Don't bet on it: Contingent agreements with asymmetric information. *J. Amer. Statistical Assoc.* 78, 424–426.