Common Priors under Endogenous Uncertainty

Pierfrancesco Guarino* & Elias Tsakas†

23.VII.2020

Abstract

For a fixed game and a type structure that admits a common prior, Action Independence states that the conditional beliefs induced by the common prior do not depend on the player's own strategy. It has been conjectured that Action Independence can be behaviorally characterized by means of a suitable no-betting condition (Dekel & Siniscalchi, 2015), but whether this is indeed the case remains an open problem. In this paper, we prove this conjecture true by focusing on strategy-invariant bets, which are bets that cannot be manipulated by the players. In particular, first we show that at least one of the common priors satisfies Action Independence if and only if there exists no mutually acceptable strategy-invariant bet among the players. Second we show that, all common priors satisfy Action Independence if and only if there exists no mutually acceptable strategy-invariant bet among the players and an outside observer. These results allow us a deeper understanding of existing foundations of solution concepts using only epistemic conditions that are expressed in terms of type structures and are therefore elicitable.

Keywords: Common Prior, No-Betting Condition, Endogenous Uncertainty, Action Independence, Strategy-Invariant Bets.
JEL Classification Number: C70, D82.

1. Introduction

The epistemic characterizations of various well-celebrated equilibrium concepts, such as Nash equilibrium (Aumann & Brandenburger, 1995) or objective correlated equilibrium (Aumann, 1987), rely on the common prior assumption. However, as it has been recently pointed out, common priors may not always be innocent, especially in presence of endogenous uncertainty (Dekel & Siniscalchi, 2015). In particular, common priors often implicitly introduce additional beliefs beyond the ones that the analyst can elicit (viz., beyond the beliefs that are represented in a standard type structure), which may often have awkward implications. For instance, a common prior may postulate that a player's beliefs depend on this same player's own strategy, which is clearly at odds with the predominant Bayesian view that beliefs are updated only when new information arrives.

Along these lines, Dekel & Siniscalchi (2015) proposed a condition that the common prior must satisfy in order to rule out such phenomena, called *Action Independence* or *Aumann Independence* or simply *AI condition*.[1] Accordingly, the beliefs that the player inherits from the type structure are the same as the conditional beliefs induced by the common prior, irrespective of the strategy chosen by the player. Then they go on to emphasize the importance of providing behavioral foundations for the AI condition and conjecture that this should be possible by means of a suitable no-betting condition. Since then, this has remained an open problem.

In this paper, we address this question and provide an affirmative answer. In particular, we classify each type structure that admits a (not necessarily unique) common prior into one of three categories: either all common priors satisfy the AI condition, or some common priors satisfy the AI condition and some do not, or none of the common priors satisfies the AI condition. We do so by means of our two main results. First, we prove that there exists some common prior satisfying the AI condition if and only if there exists no mutually acceptable bet among the players (Theorem 1). Second, assuming that there exists at least one common prior satisfying the AI condition, we prove

*Alpen-Adria-Universität Klagenfurt (Institut für Volkswirtschaftslehre); *E-mail:* pf.guarino@hotmail.com.

†Department of Microeconomics and Public Economics (MPE), Maastricht University; *Homepage:* www.elias-tsakas.com; *E-mail:* e.tsakas@maastrichtuniversity.nl.

[1] "Action Independence" is the way in which this condition is called in Battigalli et al. (Work in Progress), while "Aumann Independence" is the name used in Dekel & Siniscalchi (2015).

that all common priors satisfy the AI condition if and only if there exists no mutually acceptable bet among the players and an outside observer (Theorem 2). Combined, the two results allow us to pin down each type structure into one of the three categories mentioned above. As a side remark, both our results restrict attention to strategy-invariant bets, i.e., bets that pay for each type the same expected payoff irrespectively of the strategy chosen by this type. In this way we avoid providing incentives that may affect the behavior of the players during the upcoming game and—as a consequence—the bets we employ do not indirectly affect the underlying type structures that aim to characterize.

Besides providing foundations for the AI condition per se, our results provide deeper insights into existing epistemic foundations of solution concepts using standard type structures.[2] Focusing on objective correlated equilibrium, a recent result shows that for every type structure admitting a common prior that satisfies AI, rationality and common belief in rationality imply a correlated equilibrium distribution (Dekel & Siniscalchi, 2015, Theorem 12.7). However, this result relies on the specific common prior employed and—in this sense—said epistemic conditions are not stated only in terms of the underlying type structure. Indeed, as we show, if a type structure belongs to our second category (i.e., it admits some common priors that satisfy the AI condition and some that do not), then it can be the case that some common prior induces a correlated equilibrium distribution, while some other common prior does not. This is exactly where our results come in handy: using Theorem 2, we can conclude that, if all common priors that a type structure admits satisfy the AI condition, then *a fortiori* all of them induce a correlated equilibrium distribution. Hence, the conditions of Dekel & Siniscalchi (2015, Theorem 12.7) can now be expressed entirely in the language of type structures without any reference to a specific common prior.

This paper belongs to a rather rich literature on the behavioral foundations of epistemic assumptions. Of course, so far, this literature has mostly focused on behavioral characterizations of the common prior assumption (Aumann, 1976; Sebenius & Geanakoplos, 1983; Milgrom & Stokey, 1982; Samet, 1998; Feinberg, 2000).[3] Notably, these earlier contributions dealt with exogenous uncertainty, i.e., the epistemic models they employed represented belief hierarchies about exogenous parameters, as opposed to our case where belief hierarchies are defined on the players' strategies. Differences between the two frameworks have been highlighted by Dekel & Siniscalchi (2015) and Battigalli et al. (Work in Progress).

The paper is structured as follows. Section 2 introduces the building blocks of our analysis, i.e., type structures, common priors, the AI condition, and bets on state spaces. In Section 3, we obtain our characterization results. In Section 4, we analyze the implications of our results in the study of objective correlated equilibrium. Finally, in Section 5, we address various issues related with our analysis. Proofs are relegated to Appendix A.

2. Theoretical Framework

2.1 Epistemic Type Structures

We begin by fixing a game in its strategic form $\Gamma := \langle I, (S_i, u_i)_{i \in I} \rangle$ (henceforth, the *game*). As usual, I is the set of players, which without loss of generality is assumed to contain only two players, viz., $I := \{\text{Ann}(a), \text{Bob}(b)\}$. For every $i \in I$, we let S_i be player i's (finite) set of strategies and $u_i : S_a \times S_b \to \Re$ the payoff function.

Given a game, a standard *epistemic type structure* (henceforth, type structure) attached on it is a tuple

$$\mathscr{T} := \langle I, (T_i, \beta_i)_{i \in I} \rangle$$

where T_i is player i's finite set of *types* and $\beta_i : T_i \to \Delta(S_j \times T_j)$ is her *belief function*. As it is well-known, every type $t_i \in T_i$ of every player $i \in I$ (inductively) encodes an infinite hierarchy of beliefs (Brandenburger & Dekel, 1993, Section 2). We say that the type structure is *non-redundant*

[2] Using standard type structures allows us to express epistemic conditions using properties that can be elicited. This is in contrast to alternative epistemic models (e.g., Aumann's partitional model) where types are attached to strategies and are—therefore—unelicitable. For a more detailed discussion of this issue, see Dekel & Siniscalchi (2015).

[3] A notable exception is the work of Barelli (2009), where there is a behavioral characterization of a weakening of the common prior assumption called action consistency.

if every type encodes a different infinite hierarchy of beliefs. Throughout the paper, without loss of generality, we focus on non-redundant type structure (see Section 5.2).

Let $\Omega := S_a \times S_b \times T_a \times T_b$ denote the state space induced by the type structure. For each $s_i \in S_i$ and $t_i \in T_i$, we define the events $[\![s_i]\!] := \{s_i\} \times T_i \times S_j \times T_j$ and $[\![t_i]\!] := S_i \times \{t_i\} \times S_j \times T_j$ respectively, with $[\![s_i, t_i]\!] := [\![s_i]\!] \cap [\![t_i]\!]$.

Definition 1 (Common Prior). *A type structure $\mathscr{T}$* admits a common prior *if there exists a $\pi \in \Delta(\Omega)$ such that, for every $i \in I$ and every $t_i \in T_i$,*

i) $\pi([\![t_i]\!]) > 0$,

ii) $\beta_i(t_i)(s_j, t_j) = \pi([\![s_j, t_j]\!] | [\![t_i]\!])$, *for all* $(s_j, t_j) \in S_j \times T_j$.

If there exists such π, then π is deemed a common prior *of $\mathscr{T}$, with $\mathscr{T}$ admitting it. The set of all common priors of $\mathscr{T}$ is denoted by $\Pi_{\mathscr{T}} \subseteq \Delta(\Omega)$.*

Obviously, not all type structures admit a common prior. Nevertheless, throughout this paper we are only interested in type structures that admit at least one common prior. It is not difficult to see that there may exist multiple common priors.

Example 1. Let $S_a = \{U, D\}$ and $S_b = \{L, R\}$, and consider the type structure with type spaces $T_a = \{t_a\}$ and $T_b = \{t_b\}$, and belief mappings illustrated below.

	(L, t_b)	(R, t_b)
$\beta_a(t_a)$	$1/2$	$1/2$

	(U, t_a)	(D, t_a)
$\beta_b(t_b)$	$1/2$	$1/2$

Note that this type structure admits multiple common priors. Indeed, verify that the probability measure in the following table constitutes a common prior for every $\varepsilon \in [0, \frac{1}{2}]$.

	(L, t_b)	(R, t_b)
(U, t_a)	$1/2 - \varepsilon$	ε
(D, t_a)	ε	$1/2 - \varepsilon$

Notably, even this very simple type structure has uncountably many common priors. $\diamond$

It has to be emphasized that we do not take π as a primitive of our epistemic model, as we do not consider a (fictitious) ex-ante stage from which the players' beliefs are updated conditioning on the realization of the types.[4] Instead, the only primitives of our model are the belief hierarchies that are encoded in the type structure. In this sense, the common prior enriches our model by inserting additional information to our model in terms of beliefs that are not described within the type structure. Take for instance the previous example with $\varepsilon = 0$ and observe that the common prior essentially introduces new beliefs for Ann, viz., according to this common prior, once Ann has chosen U, she believes that Bob will choose L with probability 1. Clearly, we have to be careful in how we interpret and use these additional beliefs, and—ideally—we want them to be inconsequential for our game-theoretic analysis.

2.2 The AI Condition

As we have discussed in the previous section, a common prior often contains additional information—beyond what is described by the type structure—due to the fact that players form conditional beliefs given each of their own strategies. As a result, we have to be aware of the fact that, upon having introduced a common prior, inconsistencies may arise. We illustrate this point in the following example.

[4]For a discussion of this issue, see Aumann (1976, 1987, 1998), Gul (1998), Bonanno & Nehring (1999).

Example 1 (cont). First, recall that Ann's only type puts equal probability to (L, t_b) and (R, t_b) according to her belief function. Take the common prior induced by setting $\varepsilon := 0$. Now, observe that, if Ann chooses to play U, her beliefs induced by the common prior attach probability 1 to L. However, if she chooses to play D, her beliefs induced by the common prior attach probability 1 to R. $\diamond$

The inconsistency in the previous example is due to a conceptually awkward property of the common prior. Namely, according to the chosen common prior, Ann updates her beliefs without having received any new information. She does so after having conditioned with respect to an endogenous variable, viz., her own planned strategy. This is clearly at odds with the standard Bayesian view according to which updating is the result of taking into account new information. This discrepancy was first identified by Dekel & Siniscalchi (2015, Example 12.4, pp.642–643), who went on to propose the following property that a common prior must satisfy in order to avoid such inconsistencies.

Definition 2 (AI Condition). *The common prior $\pi \in \Pi_{\mathscr{T}}$ satisfies the* AI condition *if, for all $i \in I$ and all $(s_i, t_i), (s'_i, t_i) \in S_i \times T_i$ with $\pi([\![s_i, t_i]\!]) > 0$ and $\pi([\![s'_i, t_i]\!]) > 0$,*

$$\pi([\![s_j, t_j]\!] | [\![s_i, t_i]\!]) = \pi([\![s_j, t_j]\!] | [\![s'_i, t_i]\!]), \tag{2.1}$$

for every $(s_j, t_j) \in S_j \times T_j$, with $\Pi^{AI}_{\mathscr{T}} \subseteq \Pi_{\mathscr{T}}$ denoting the set of common priors that satisfy the AI condition.

The acronym AI can stand for "Action Independence" or "Aumann Independence". It is not difficult to verify that the only common prior in Example 1 that satisfies the AI condition is the one obtained by setting $\varepsilon = \frac{1}{4}$. Indeed, for any other ε, it is the case that $\pi([\![L, t_b]\!] | [\![U, t_a]\!]) = \frac{1}{2} - \varepsilon \neq \varepsilon = \pi([\![L, t_b]\!] | [\![D, t_a]\!])$, implying that Ann's beliefs depend on her own strategy.

An important remark is warranted here. Namely, the AI condition has a bite when the type structure represents belief hierarchies about endogenous variables, viz., the players' own strategies. Whenever the type structure represents beliefs about exogenous variables (e.g., about the preferences of the players), it is clearly the case that the choice of a player's own strategy does not affect her own beliefs about said exogenous variables.

Remark 1. *A common prior $\pi \in \Pi_{\mathscr{T}}$ satisfies the AI condition if and only if, for all $(s_j, t_j) \in S_j \times T_j$,*

$$\pi([\![s_j, t_j]\!] \mid [\![s_i, t_i]\!]) = \beta_i(t_i)(s_j, t_j)$$

for every $i \in I$ and $(s_i, t_i) \in S_i \times T_i$ with $\pi([\![s_i, t_i]\!]) > 0$, i.e., whenever conditional beliefs are derived from a common prior that satisfies the AI condition, said conditional beliefs coincide with those inherited from the type structure.

Following the previous remark—while assuming that the type structure $\mathscr{T}$ admits some (not necessarily unique) common prior—we can classify $\Pi_{\mathscr{T}}$ into one of the following categories:

(Π_1) all common priors satisfy the AI condition, i.e., $\Pi_{\mathscr{T}} = \Pi^{AI}_{\mathscr{T}} \neq \emptyset$;

(Π_2) there are multiple common priors, some satisfying the AI condition and some not, i.e., $\Pi_{\mathscr{T}} \supsetneq \Pi^{AI}_{\mathscr{T}} \neq \emptyset$;

(Π_3) there exists no common prior satisfying the AI condition, i.e., $\Pi^{AI}_{\mathscr{T}} = \emptyset$.

Since the AI condition allows us to epistemically characterize equilibrium concepts, such as objective correlated equilibrium (Dekel & Siniscalchi, 2015, Theorem 12.4) or Nash equilibrium (Dekel & Siniscalchi, 2015, Theorem 12.7), we would like to know in which of the aforementioned categories each type structure is classified. This is actually the main research question of this paper, which we address in the following sections.

2.3 Bets on the State Space

It is common practice to characterize (epistemic) properties in terms of appropriate no-betting conditions in light of the fact that a player's willingness to accept a bet is observable. Therefore, these characterizations allow us to build testable hypotheses about the property of interest. This approach dates back to the early contributions of Milgrom & Stokey (1982) and Sebenius & Geanakoplos (1983), and the subsequent work of Samet (1998) and Feinberg (2000) among others, which have led to a full characterization of common priors in models with exogenous uncertainty. In their recent review article, Dekel & Siniscalchi (2015) conjectured that a similar characterization of the AI condition in terms of a suitable no-betting condition should be possible in type structures that represent beliefs about endogenous variables.

A *bet* on the state space Ω is a profile of random variables $g := (g_a, g_b)$, with $g_i \in \Re^\Omega$ for each $i \in I$, such that $g_a + g_b = 0$, i.e., simply put, a bet is a zero-sum contingent claim.

Definition 3 (Willingness to bet). *Player $i \in I$* accepts *g at some state in $[\![s_i, t_i]\!]$ if*

$$\mathbb{E}[g_i|s_i, t_i] := \sum_{s_j \in S_j} \sum_{t_j \in T_j} g_i(s_i, t_i, s_j, t_j) \cdot \beta_i(t_i)(s_j, t_j) \geq 0.$$

Player i strictly accepts *the bet at some state in $[\![s_i, t_i]\!]$ if the previous inequality is strict.*

Crucially, note that a player's willingness to accept (resp., strictly accept) a bet at some state depends on the beliefs she inherits from the type structure. That is, even if we fix a common prior $\pi \in \Pi_{\mathscr{T}}$, the player does not evaluate the bet using the conditional beliefs $\pi(\cdot|[\![s_i, t_i]\!])$, but rather using the conditional beliefs $\pi(\cdot|[\![t_i]\!])$ which always coincide with $\beta_i(t_i)$. This assumption follows naturally from the fact that our primitive concept is a type structure and not a common prior that is admitted by it. We illustrate this point in the context of Example 1.

Example 1 (cont). Take the bet that pays Ann the following amounts at each state.

	(L, t_b)	(R, t_b)
(U, t_a)	20	-10
(D, t_a)	-10	10

It is clear that according to the beliefs that Ann inherits from the type structure, she (weakly) accepts the bet at every state in $[\![D, t_a]\!]$, as she attaches probability $\frac{1}{2}$ to each of Bob's strategy-type pair thus yielding zero expected payoff at both states in $[\![D, t_a]\!]$. If we instead used the common prior to evaluate the bet, Ann's willingness to accept it would depend on which—among the multiple common priors—we would fix, viz., for $\varepsilon = \frac{1}{4}$ her willingness to bet would be the same as above, for $\varepsilon < \frac{1}{4}$ she would reject the bet in $[\![D, t_a]\!]$, whereas for $\varepsilon > \frac{1}{4}$ she would strictly accept it in $[\![D, t_a]\!]$. But, again, which common prior is fixed from the set $\Pi_{\mathscr{T}}$ is an arbitrary modelling choice that the analyst makes and it is not based on parameters one can elicit. ◇

Going a step further, even if we use the beliefs that come from the type structure to evaluate the bet at each state, players can influence the outcome of the bet by suitably choosing their own strategy. For instance, in the previous example, it is clear that Ann would prefer to choose U to get access to the good payoff of 20. However, this would have implications for her behavior in the underlying game Γ and—*a fortiori*—on the beliefs in our type structure. This would be clearly undesirable, in the sense that the bet that we would be using to test a property of the type structure (viz., the AI condition) would affect the type structure itself, i.e., we would be having a Heisenberg type of effect. To avoid such a phenomenon, we restrict attention to bets that cannot be manipulated by the players.

Definition 4 (Strategy-Invariant Bet). *A bet g is called* strategy-invariant *(henceforth, SI)* for type *$t_i \in T_i$ of player $i \in I$ if*

$$\mathbb{E}[g_i|s_i, t_i] = \mathbb{E}[g_i|s'_i, t_i]$$

for every $s_i, s_i' \in S_i$. A bet is SI if it is SI with respect to every $t_i \in T_i$ and every $i \in I$.

The underlying idea is that, conditional on a type t_i, player i receives the same expected payoff irrespective of the strategy choosn. Hence, the bet does not interfere with the incentives that the player faces in the underlying game and—given that this is the case for every type of every player—the introduction of the bet does not affect the underlying type structure. For instance, in the previous example, to obtain an SI bet we can replace Ann's payoff of 20 with a payoff of 10. Indeed, notice that in such a case, Ann would receive 0 in expectation irrespective of whether she chooses U or D and the same is true for Bob irrespective of whether he chooses L or R. Throughout the rest of the paper, we focus exclusively on SI bets.

Definition 5 (No-Betting Condition). *An SI bet g is* mutually acceptable *if, for every $i \in I$ and every $(s_i, t_i) \in S_i \times T_i$,*

$$\mathbb{E}[g_i | s_i, t_i] \geq 0,$$

with at least one inequality being strict.

That is, we call an SI bet *mutually acceptable* if every player accepts it at every state and there exists at least one player who strictly accepts it at some state. Our main research question then becomes whether we can classify type structures that admit a common prior into one of the categories $(\Pi_1), (\Pi_2), (\Pi_3)$ based on the existence of mutually acceptable bets. The remaining of the paper addresses this question.

3. Characterization Results

3.1 Is there a Common Prior satisfying the AI Condition?

We begin by identifying the type structures $\mathscr{T}$ which admit a common prior that satisfies the AI condition. Formally, what we do first is to provide necessary and sufficient conditions in terms of the existence of mutually acceptable SI bets so that $\Pi^{AI}_{\mathscr{T}} \neq \emptyset$. In other words, our first result indicates if the type structure belongs into one of the first two categories (viz., (Π_1) or (Π_2)) or whether it belongs to the third category (viz., (Π_3)). This result already answers affirmatively the conjecture of Dekel & Siniscalchi (2015).

Theorem 1. *Given a type structure $\mathscr{T}$, the following hold.*

i) If there exists some common prior that satisfies the AI condition (i.e., if $\Pi^{AI}_{\mathscr{T}} \neq \emptyset$), then there exists no mutually acceptable SI bet.

ii) If there exists no common prior satisfying the AI condition (i.e., if $\Pi^{AI}_{\mathscr{T}} = \emptyset$), then there exists a mutually acceptable SI bet.

The full proof of the result is relegated to Appendix A. For the time being, we provide some intuition together with an illustration by means of examples.

The underlying idea behind the proof of part (i) is similar to the one in Sebenius & Geanakoplos (1983). In particular, we begin with the observation (first made in Remark 1) that we can arbitrarily replace the beliefs $\beta_i(t_i)$ inherited from the type structure with the conditional beliefs $\pi(\cdot | [\![s_i, t_i]\!])$ given by a common prior that satisfies the AI condition. Hence, we can construct an auxiliary Aumann model over the same state space Ω with the common prior π such that, for every SI bet, the willingness to accept the bet using the beliefs $\beta_i(t_i)$ is the same as the willingness to accept it using the beliefs $\pi(\cdot | [\![s_i, t_i]\!])$ at every state and for every player. But then, by Sebenius & Geanakoplos (1983, Proposition 2), no mutually acceptable bet exists in the auxiliary Aumann model and—*a fortiori*—not in our original type structure either. Let us provide an illustration.

Example 1 (cont). Consider the unique common prior that satisfies the AI condition, viz., let $\varepsilon = \frac{1}{4}$, implying that $\pi \in \Pi^{AI}_{\mathscr{T}}$ is uniformly distributed in Ω. Take an SI bet that pays to Ann the following amounts at each state.

	(L, t_b)	(R, t_b)
(U, t_a)	v_1	v_2
(D, t_a)	v_3	v_4

Notice that in order for this bet to mutually acceptable, it must be the case that the following two inequalities hold for Ann:

$$\begin{aligned} v_1 + v_2 &\geq 0, \\ v_3 + v_4 &\geq 0. \end{aligned}$$

Likewise, it must be the case that the following two inequalities hold for Bob:

$$\begin{aligned} -v_1 - v_3 &\geq 0, \\ -v_2 - v_4 &\geq 0. \end{aligned}$$

At the same time, one of the previous four inequalities must be strict. Clearly, this cannot occur. If this was the case, by adding the respective sides of the four inequalities, we would obtain $0 > 0$, which is an obvious contradiction. ◇

Now, let us switch to part (ii) of the theorem. Here, the argument is slightly more involved. In the first step we make use of earlier results by Samet (1998) and Feinberg (2000). In particular, once again we construct an auxiliary Aumann model over Ω with the conditional beliefs given each information set $[\![s_i, t_i]\!]$ being set the same as the ones given by $\beta_i(t_i)$. The key observation is that, since $\Pi^{AI}_{\mathscr{T}} = \emptyset$, there exists no common prior in our auxiliary Aumann model generating these conditional beliefs. Hence, by Samet (1998, Claim, p.173) and Feinberg (2000, Theorem 2, p.146), there exists a mutually acceptable bet in the Aumann model. However, this bet is not necessarily SI. Thus, with a sequence of transformations, we obtain a new SI bet which keeps the willingness of every player to bet at every state unchanged. As a result, since the original bet was mutually acceptable, so is this new SI bet, thus completing the proof. We illustrate the key ideas with an example.

Example 2. Once again consider the strategy sets $S_a = \{U, D\}$ and $S_b = \{L, R\}$, and take the type structure with type spaces $T_a = \{t_a\}$ and $T_b = \{t_b, t'_b\}$ and belief mappings illustrated below.

	(L, t_b)	(R, t'_b)
$\beta_a(t_a)$	1/2	1/2

	(U, t_a)	(D, t_a)
$\beta_b(t_b)$	1	0
$\beta_b(t'_b)$	0	1

Observe that the only common prior that this type structure admits is the one below.

	(L, t_b)	(R, t'_b)
(U, t_a)	1/2	0
(D, t_a)	0	1/2

Clearly, this common prior does not satisfy the AI condition. Thus, let us begin by taking the following auxiliary Aumann model over Ω with the following information partitions (with Ann's partition on the left and Bob's on the right) and the corresponding conditional beliefs given each cell of the partition.

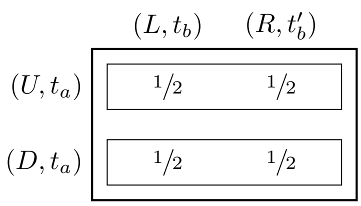

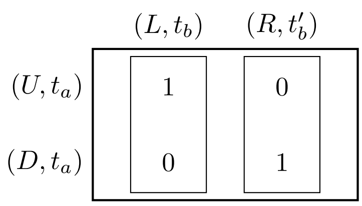

Notice that the conditional beliefs given each $[\![s_i, t_i]\!]$ coincide with those given by the type structure. However, these conditional beliefs cannot have been derived from a common prior, because, had such a prior existed, it would have belonged to $\Pi^{AI}_{\mathscr{T}}$, which we know is empty. Therefore, we can find a mutually acceptable bet in this Aumann space. Say this bet is the one that pays Ann the amounts shown below.

	(L, t_b)	(R, t'_b)
(U, t_a)	v_1	v_2
(D, t_a)	v_3	v_4

Of course, if this bet is SI, then we are done. Thus, suppose it is not, implying—without loss of generality—that

$$v_1 + v_2 > v_3 + v_4 \geq 0.$$

Then take the constant

$$c := \frac{1}{2}(v_1 + v_2) - \frac{1}{2}(v_3 + v_4) > 0$$

and subtract it from every payment in the upper information set of Ann. This will yield a new transformed bet that pays Ann the amounts shown below.

	(L, t_b)	(R, t'_b)
(U, t_a)	$v_1 - c$	$v_2 - c$
(D, t_a)	v_3	v_4

For starters, observe that the new bet is SI, as the expected payoff of Ann is the same at all states within $[\![t_a]\!]$. At the same time, Bob has become better off, as our transformation only subtracted payoffs from Ann and—*a fortiori*—added payoffs to Bob. Finally, observe that this transformation does not affect whether the bet is SI for Bob's types, i.e., the transformed bet is SI for any given type of Bob if and only if the original one was SI (which—incidentally—it trivially was). This is because we add the same constant payment to all states that correspond to any given strategy-type pair of Ann. As a result, the new transformed bet that we obtain is both mutually acceptable and SI. For the sake of illustration, such a bet could be one that pays $g_a(U, t_a, L, t_b) = g_a(D, t_a, R, t'_b) = -1$ and $g_a(D, t_a, R, t_b) = g_a(U, t_a, L, t'_b) = 1$. $\diamond$

3.2 Do all Common Priors satisfy the AI Condition?

In the previous section we provided necessary and sufficient conditions that identify whether a type structure admits a common prior that satisfies the AI condition. In particular, if the answer is positive, we can conclude that the type structure is classified in category (Π_1) or (Π_2), i.e., either all common priors satisfy the AI condition, or there are additional common priors that do not satisfy the AI condition. In this section, we provide a second result that identifies—again in terms of a no-betting condition—which of the two it is the case.

We begin with the premise that $\Pi^{AI}_{\mathscr{T}} \neq \emptyset$, i.e., there exists a common prior satisfying the AI condition, implying (by Theorem 1) that there exists no mutually acceptable SI bet between Ann and Bob. Now, we extend our type structure by introducing a dummy player (henceforth, the *outside observer*). The outside observer has a single strategy that is completely inconsequential to anybody and his beliefs are given by a common prior $\pi \in \Pi_{\mathscr{T}}$.

Formally, define the extended game $\overline{\Gamma} := \langle \overline{I}, (\overline{S}_i, \overline{u}_i)_{i \in \overline{I}} \rangle$ (henceforth, the *extended game*), where $\overline{I} := I \cup \{d\}$ is the original set of players augmented by the outsider observed, $\overline{S}_i := S_i$ for each $i \in I$ and $\overline{S}_d := \{s_d\}$ for the outside observer, and finally $\overline{u}_i(s_a, s_b, s_d) := u_i(s_a, s_b)$ for each original player $i \in I$ and $\overline{u}_d(s_a, s_b, s_d) := 0$ for the outside observer, for each extended strategy profile (s_a, s_b, s_d).

Now, take a standard type structure $\mathscr{T}$ (over the original game) admitting common priors and fix an arbitrary $\pi \in \Pi_{\mathscr{T}}$. Extend the type structure to

$$\overline{\mathscr{T}}_\pi := \langle \overline{I}, (\overline{T}_i, \overline{\beta}_i)_{i \in \overline{I}} \rangle,$$

such that $\overline{T}_i := T_i$ for every original player $i \in I$ and $\overline{T}_d := \{t_d\}$ for the outside observer, and $\overline{\beta}_i(t_i)(s_j, t_j, s_d, t_d) := \beta_i(t_i)(s_j, t_j)$ for each $i \in I$ and $\overline{\beta}_d(t_d)(s_i, t_i, s_j, t_j) := \pi(s_i, t_i, s_j, t_j)$ for the outside observer, for each original state $(s_i, t_i, s_j, t_j) \in \Omega$. That is, the original players inherit the beliefs from the original standard type structure, while the outside observer adopts the beliefs that are given by the common prior.

It is straightforward to verify that the extended state space $\overline{\Omega} = \overline{S}_a \times \overline{T}_a \times \overline{S}_b \times \overline{T}_b \times \overline{S}_d \times \overline{T}_d$ is homeomorphic to the original state space Ω. Indeed, we simply need to take each $(s_a, t_a, s_b, t_b) \in \Omega$ and augment it with the outside observer's unique strategy-type pair to obtain the corresponding state $(s_a, t_a, s_b, t_b, s_d, t_d) \in \overline{\Omega}$. In this sense, with a slight abuse of notation and without loss of generality, we keep using the original state space to also identify states in the extended state space.

An extended bet on the state space Ω is a profile of random variables $\overline{g} := (\overline{g}_a, \overline{g}_b, \overline{g}_d)$, with $\overline{g}_i \in \Re^\Omega$ for each $i \in \overline{I}$, such that $\overline{g}_a + \overline{g}_b + \overline{g}_d = 0$. We keep restricting attention to SI bets and notice that every extended bet is trivially SI for the outside observer: hence, we only need to make sure that an extended bet is SI for the original players. Willingness (resp. strict willingness) to accept a bet is naturally extended to the present framework, viz., player $i \in \overline{I}$ accepts $\overline{g}$ at some state in $[\![s_i, t_i]\!]$ if $\mathbb{E}[\overline{g}_i | s_i, t_i] \geq 0$ and strictly accepts it if the inequality is strict. Then we can naturally state the no-betting condition appropriate for this framework as follows: an extended SI bet $\overline{g}$ is mutually acceptable if every player $i \in \overline{I}$ accepts it at every state and there exists at least one player who strictly accepts it at some state.

Now, let us state our second result, which allows us to identify whether the original type structure $\mathscr{T}$ is classified in the category (Π_1) or in the category (Π_2), by means of the existence of a mutually acceptable extended SI bet.

Theorem 2. *Fix a standard type structure that admits a common prior that satisfies the AI condition, i.e.,* $\Pi^{AI}_{\mathscr{T}} \neq \emptyset$. *Then the following hold.*

i) If all common priors in $\Pi_{\mathscr{T}}$ *satisfy the AI condition (i.e., if* $\Pi^{AI}_{\mathscr{T}} = \Pi_{\mathscr{T}}$*), then there exists no mutually acceptable extended SI bet.*

ii) If there exists a common prior in $\Pi_{\mathscr{T}}$ *that does not satisfy the AI condition (i.e., if* $\Pi^{AI}_{\mathscr{T}} \subsetneq \Pi_{\mathscr{T}}$*), then there exists a mutually acceptable extended SI bet.*

Once again, we relegate the formal proof to Appendix A and here we only provide some intuition.

Proving part (i) of the theorem follows directly from Theorem 1. In particular, pick an arbitrary common prior (which, by hypothesis, satisfies the AI condition) and—once again—construct an auxiliary Aumann model in which the conditional beliefs given each information set are derived from said common prior. Hence, there is no mutually acceptable extended SI bet. Finally, observe that the willingness to accept any extended SI bet in the auxiliary Aumann model is the same as the willingness to accept the same bet using the beliefs that are derived from the extended type structure, which completes our argument. We now provide an illustration of what above.

Example 3. Once again, let $S_a = \{U, D\}$ and $S_b = \{L, R\}$, and consider the type structure with type spaces $T_a = \{t_a, t'_a\}$ and $T_b = \{t_b, t'_b\}$, and belief mappings illustrated below.

	(L, t_b)	(R, t'_b)
$\beta_a(t_a)$	1	0
$\beta_a(t'_a)$	0	1

	(U, t_a)	(D, t'_a)
$\beta_b(t_b)$	1	0
$\beta_b(t'_b)$	0	1

Notice that the distribution below is a common prior admitted by the type structure for every $\delta \in (0, 1)$.

	(L, t_b)	(R, t'_b)
(U, t_a)	δ	0
(D, t'_a)	0	$1 - \delta$

In fact, these are all the common priors that our type structure admits. Now, it is clear that they all satisfy the AI condition and it is not difficult to see that there exists no extended SI bet which is mutually acceptable. Indeed, take an arbitrary bet that pays Ann and Bob at each state the following amounts.

	(L, t_b)	(R, t'_b)
(U, t_a)	v^a_1, v^b_1	v^a_2, v^b_2
(D, t_a)	v^a_3, v^b_3	v^a_4, v^b_4

Notice that, in order for Ann to (weakly) accept this bet, it must be the case that $v^a_1 \geq 0$ and $v^a_4 \geq 0$ and, likewise for Bob, to (weakly) accept this bet we must have $v^b_1 \geq 0$ and $v^b_4 \geq 0$. But then it is the case that the outside observer's expected payoff is equal to

$$-\delta(v^a_1 + v^b_1) - (1-\delta)(v^a_4 + v^b_4) \leq 0,$$

with equality holding if and only if $v^a_1 = v^a_4 = v^b_1 = v^b_4 = 0$. Hence, there exists no mutually acceptable extended SI bet. ◇

Concerning part (ii) of the theorem, here our proof is constructive and proceeds as follows. For starters, since there exists a common prior that satisfies the AI condition, there exists no mutually acceptable SI bet between Ann and Bob (by Theorem 1). Nevertheless, since there also exists a common prior π that does not satisfy the AI condition, it is necessarily the case that some player (among the original players) evaluates a bet at some state using the beliefs that are inherited from the type structure, which differ from the conditional beliefs given from the common prior π. On the other hand, the outside observer does use π to evaluate a bet. This discrepancy allows us to construct a bet that pays Ann and Bob 0 in expectation at every state, and yields a strictly positive expected payoff to the outside observer. This is clearly a mutually acceptable extended SI bet. Let us provide an illustration.

Example 1 (cont). Consider a common prior that does not satisfy the AI condition, e.g., let $\varepsilon = 0$. Then consider the bet that pays Bob 0 at every state and pays Ann the amounts depicted in the following table.

	(L, t_b)	(R, t_b)
(U, t_a)	-1	1
(D, t_a)	1	-1

Obviously, Ann receives 0 in expectation at all states, as she uses the beliefs she inherits from the type structure, which distribute probability uniformly across 1 and -1. On the other hand, the outside observer uses the beliefs that come from the common prior, which put probability 1 to the two states where Ann receives -1, and—as a result—the outside observer gets 1 in expectation, i.e., he strictly accepts the bet. Therefore, this is a mutually acceptable extended SI bet. ◇

4. Objective Correlated Equilibrium & the AI Condition

Let us now focus on how our results allow us to deeper understand (the role of AI condition for) existing epistemic characterizations of objective correlated equilibrium (henceforth, *correlated equilibrium*). According to the original characterization of Aumann (1987), common knowledge of rationality together with a common prior yield a correlated equilibrium distribution. However, Aumann's model has the feature that each type (viz., each partition cell in Aumann's language) is exogenously endowed with a strategy. As a result, the AI condition is automatically satisfied (which is why we do not need to explicitly postulate it in Aumann's theorem), but—at the same time—bundling each type with a given strategy leads to an undesirable situation, viz., types in an Aumann model cannot be always elicited.

To deal with this issue, Dekel & Siniscalchi (2015) minimally modify Aumann's conditions to obtain the following result: a type structure expressing rationality and common belief in rationality together with a common prior that satisfies the AI condition lead to a correlated equilibrium distribution. Formally, take a type structure $\mathscr{T}$ that admits a common prior $\pi \in \Pi^{AI}_{\mathscr{T}}$, such that for every $i \in I$ and every $(s_i, t_i) \in S_i \times T_i$ with $\pi([\![s_i, t_i]\!]) > 0$, we have

$$\sum_{s_j \in S_j} \pi([\![s_j]\!] \,|\, [\![s_i, t_i]\!]) u_i(s'_i, s_j) \geq \sum_{s_j \in S_j} \pi([\![s_j]\!] \,|\, [\![s'_i, t_i]\!]) u_i(s'_i, s_j),$$

for every $s'_i \in S_i$. Then $\mathrm{marg}_{S_a \times S_b} \pi$ is a correlated equilibrium distribution (Dekel & Siniscalchi, 2015, Theorem 12.4).[5]

Here we should crucially note that the choice of the common prior (among the multiple common priors that exist) plays a significant role in the previous result. This is easily verified by the fact that the common prior is directly used to obtain the correlated equilibrium distribution. Thus, it might be the case that the same type structure $\mathscr{T}$ induces a correlated equilibrium for some common priors in $\Pi_{\mathscr{T}}$ and not for some others. We now illustrate that this is indeed the case.

Example 1 (cont). Consider a standard symmetric coordination game:

	L	R
U	$1, 1$	$0, 0$
D	$0, 0$	$1, 1$

Then consider the common prior induced by $\varepsilon = \frac{1}{2}$. Clearly, this common prior does not satisfy the AI condition. Moreover, it does not induce a correlated equilibrium distribution, as it leads the players to miscoordinate with probability 1. If instead we had chosen another common prior that satisfied the AI condition (i.e., if we had chosen $\varepsilon = \frac{1}{4}$), then the resulting distribution of strategy profiles would have constituted a correlated equilibrium. $\diamond$

The conclusion is that the epistemic conditions for correlated equilibrium cannot be expressed entirely within a type structure, but reference to the choice of the common prior has to be made. This is exactly where our results come in handy: following Theorem 2, if the type structure under scrutiny does not allow for side-betting with an outside observer, then all common priors satisfy the AI condition and—as a result—all common priors that the type structure $\mathscr{T}$ admits do induce a correlated equilibrium distribution. Therefore, the result of Dekel & Siniscalchi (2015) would only make reference to the type structure, implying that we obtain sufficient conditions for correlated equilibrium using only properties that can be elicited.

5. Discussion

5.1 Rationality and Common Belief in Rationality

As we have already mentioned, the reason behind our choice of focusing on SI bets is that they allow us to isolate the betting element that could potentially arise in a game from the actual play in the game. It is essentially due to this point that in this paper there is no actual reference to specific games: by focusing on SI bets, we make irrelevant the specific game upon which the betting can potentially take place, thus ending up to be in position to focus on the specific property of the type structures (and infinite hierarchies of beliefs) we are actually interested in, namely, the AI condition.

Indeed, the importance of our results lies in the fact that sidebetting on games is strictly related to the presence of endogenous uncertainty, i.e., the players' behavior has an impact on the realization of the state of nature (i.e., the outcome). Focusing on SI bets prevent us from the need to create a larger game to incorporate potential side bets in the actual game under scrutiny.

[5] To be more precise, Dekel & Siniscalchi (2015) also impose a minimality condition, which is not relevant for our discussion, and therefore it is without loss of generality to omit explicit reference to it. As a side remark, this minimality condition is equivalent to our assumption of working with type structures that are non-redundant.

Indeed, SI sidebets on a game do not alter the players' incentives to choose a strategy instead of another one in the actual play the game. That is, SI bets make players indifferent between different strategies.

However, it is possible to argue that in an actual game, not all strategies are 'equal', in the sense that a rational player should not choose some strategies. As a matter of fact, we feel this would betray the exercise, since the rationality of a player in a game should not have any bite in the analysis of her infinite hierarchies of beliefs by means of sidebets. Nevertheless, the question stands if our results are true when the focus is on strategies that are consistent with the rationality of the players and their mutual beliefs on their rationality. It is clear that all our results naturally extend to a setting where the focus is on those strategies that correspond to the strategies chosen in a type structure where the event *Rationality and Common Belief of Rationality* is nonempty.

5.2 Non-Redundancy and Elicitability

In Section 2.1, we explicitly assumed that we work with type structures that are non-redundant. First let us stress that our results do not formally rely on the non-redundancy assumption. So, why do we impose it? The reason is conceptual. In particular, bets are contingent claims that allocate payments to each player conditional on strategy profiles and belief hierarchies. This is a natural implicit assumption, as belief hierarchies can be elicited and strategies are directly observed. Therefore, allowing for redundancies would lead to multiple types representing the same hierarchy and—as a result—whenever we elicit said hierarchy it would be ambiguous which type should be used to determine the payments of bet. Hence, in such a case, we would need to also require bets to be measurable with respect to the belief hierarchies, i.e., in presence of redundancies two different types that induce the same belief hierarchy should induce the exact same payments. However, this would make our model unnecessarily more complex, which is why we rule out redundancies in the first place.

Appendix

A. Proofs

Proof of Theorem 1. *Part (i):* Assume that there exists a common prior $\pi \in \Pi^{AI}_{\mathscr{T}}$. By the fact that $\pi \in \Pi^{AI}_{\mathscr{T}}$, it is the case that $\beta_i(t_i)(s_j, t_j) = \pi([\![s_j, t_j]\!] | [\![s_i, t_i]\!])$ for every $(s_i, t_i) \in S_i \times T_i$ and every $(s_j, t_j) \in S_j \times T_j$ (see Remark 1). Hence, for every SI bet $\widetilde{g}$, we can rewrite i's expected payoff at each state in $[\![s_i, t_i]\!]$ as

$$\mathbb{E}[\widetilde{g}_i | s_i, t_i] = \sum_{s_j \in S_j} \sum_{t_j \in T_j} \widetilde{g}_i(s_i, t_i, s_j, t_j) \cdot \pi([\![s_j, t_j]\!] | [\![s_i, t_i]\!]).$$

Now, we proceed by contradiction and assume that there exists a mutually acceptable SI bet g, implying that $\mathbb{E}[g_i | s_i, t_i] \geq 0$ for every $(s_i, t_i) \in S_i \times T_i$ and every $i \in I$, with at least one inequality being strict. Hence, by the fact that $\{[\![s_i, t_i]\!] | (s_i, t_i) \in S_i \times T_i\}$ is a partition of Ω, we obtain

$$\begin{aligned} 0 &\leq \sum_{s_j \in S_j} \sum_{t_j \in T_j} \pi([\![s_i, t_i]\!]) \cdot \mathbb{E}[g_i | s_i, t_i] \\ &= \sum_{s_i \in S_i} \sum_{t_i \in T_i} \pi([\![s_i, t_i]\!]) \sum_{s_j \in S_j} \sum_{t_j \in T_j} g_i(s_i, t_i, s_j, t_j) \cdot \pi([\![s_j, t_j]\!] | [\![s_i, t_i]\!]) \\ &= \sum_{s_i \in S_i} \sum_{t_i \in T_i} \sum_{s_j \in S_j} \sum_{t_j \in T_j} g_i(s_i, t_i, s_j, t_j) \cdot \pi(s_i, t_i, s_j, t_j), \end{aligned}$$

with the inequality being strict for at least one of the two players. Thus, we finally add the respective sides of the inequalities for the two players to obtain

$$0 < \sum_{s_a \in S_a} \sum_{t_a \in T_a} \sum_{s_b \in S_b} \sum_{t_b \in T_b} \Big(g_a(s_a, t_a, s_b, t_b) + g_b(s_a, t_a, s_b, t_b) \Big) \cdot \pi(s_a, t_a, s_b, t_b) = 0,$$

which is the desired contradiction. Hence, there exists no mutually acceptable SI bet.

Part (ii): We begin by defining the following auxiliary Aumann model $\langle \Omega, (\mathcal{P}_i, \pi_i)_{i \in I} \rangle$, where $\mathcal{P}_i := \{[\![s_i, t_i]\!] | (s_i, t_i) \in S_i \times T_i\}$ is the information partition of $S_i \times T_i$ cylinders and, for each player $i \in I$, $\pi_i \in \Delta(\Omega)$ is a probability measure such that the conditional beliefs given each information set $[\![s_i, t_i]\!]$ agree with the beliefs obtained from the type structure, i.e., we set

$$\pi_i([\![s_j, t_j]\!] | [\![s_i, t_i]\!]) := \beta_i(t_i)(s_j, t_j),$$

for every $(s_j, t_j) \in S_j \times T_j$. Note that it is necessarily the case that $\pi_a \neq \pi_b$, otherwise there would exist a common prior that satisfies the AI condition, which cannot be (by hypothesis). Therefore, our auxiliary Aumann structure does not admit a common prior. Hence, by Samet (1998, Claim, p.173) and Feinberg (2000, Theorem 2, p.146), there exists a mutually acceptable bet $f := (f_a, f_b)$, with $f_i \in \Re^\Omega$ for every $i \in I$, such that $f_a + f_b = 0$. In particular, for every $i \in I$ and every $[\![s_i, t_i]\!]$, it is the case that

$$\sum_{s_j \in S_j} \sum_{t_j \in T_j} f_i(s_i, t_i, s_j, t_j) \cdot \pi_i([\![s_j, t_j]\!] | [\![s_i, t_i]\!]) \geq 0,$$

with at least one inequality being strict. By definition of π_i, it follows that, for every $i \in I$ and every $[\![s_i, t_i]\!]$, it is the case that

$$\mathbb{E}[f_i | s_i, t_i] \geq 0,$$

with at least one inequality being strict. Now, if f is an SI bet, we are done. Thus, we assume that it is not and we define a new bet $g := (g_a, g_b)$ as follows. First, for each $i \in I$ and each $(s_i, t_i) \in S_i \times T_i$, we let

$$c_i(s_i, t_i) := \mathbb{E}[f_i | s_i, t_i] - \min_{s_i' \in S_i} \mathbb{E}[f_i | s_i', t_i] \geq 0.$$

Then, for each state $(s_i, t_i, s_j, t_j) \in \Omega$ and each player $i \in I$, we define

$$g_i(s_i, t_i, s_j, t_j) := f_i(s_i, t_i, s_j, t_j) - c_i(s_i, t_i) + c_j(s_j, t_j).$$

First of all, it is trivially verified that g is indeed a bet. Now, notice that, for any fixed $t_i \in T_i$, we obtain

$$\mathbb{E}[g_i | s_i, t_i] = \min_{s_i' \in S_i} \mathbb{E}[f_i | s_i', t_i] + c_j(s_j, t_j),$$

for all $s_i \in S_i$. Observe that $\mathbb{E}[g_i | s_i, t_i] = \mathbb{E}[g_i | s_i', t_i]$ for all $s_i, s_i' \in S_i$, i.e., g is an SI bet. Moreover, $\mathbb{E}[g_i | s_i, t_i] \geq 0$ for all $(s_i, t_i) \in S_i \times T_i$. Finally, observe that, since f is not SI in the first place, we obtain $c_j(s_j, t_j) > 0$ for at least one player $j \in I$ and at least one pair $(s_j, t_j) \in S_j \times T_j$. Therefore, there exists at least one pair (s_i, t_i) of j's opponent such that $\mathbb{E}[g_i | s_i, t_i] > 0$. This implies that g is mutually acceptable, which completes the proof. ■

Proof of Theorem 2. *Part (i):* The proof is almost identical to the one of the first part of Theorem 1. The only caveat is that there are three players now, Ann, Bob and the outside observer. Nevertheless, the outside observer can only have one belief, viz., it is necessarily the case that $\overline{\beta}_d(t_d) = \pi$ for some $\pi \in \Pi_{\mathscr{T}}^{AI}$. Now, suppose that there exists a mutually acceptable extended SI bet. By following the exact same steps in the proof of Theorem 1, we obtain

$$0 < \sum_{s_a \in S_a} \sum_{t_a \in T_a} \sum_{s_b \in S_b} \sum_{t_b \in T_b} \Big(g_a(s_a, t_a, s_b, t_b) + g_b(s_a, t_a, s_b, t_b) + g_d(s_a, t_a, s_b, t_b) \Big) \cdot \pi(s_a, t_a, s_b, t_b) = 0,$$

which is an obvious contradiction.

Part (ii): The proof is constructive. Fix a common prior $\pi \in \Pi_{\mathscr{T}} \setminus \Pi_{\mathscr{T}}^{AI}$, and take the extended type structure such that the outside observer's beliefs agree with this common prior, i.e., $\overline{\beta}_d(t_d) = \pi$. Since π is a common prior that does not satisfy the AI condition, there exists a $E_j \subseteq S_j \times T_j$ and two strategy-type pairs $(s_i, t_i), (s_i', t_i) \in S_i \times T_i$ such that

$$\pi(E_j | [\![s_i, t_i]\!]) > \beta_i(t_i)(E_j) > \pi(E_j | [\![s_i', t_i]\!]).$$

Obviously, this directly implies that

$$\pi(E_j^c | [\![s_i, t_i]\!]) < \beta_i(t_i)(E_j^c) < \pi(E_j^c | [\![s_i', t_i]\!]),$$

where $E_j^c = (S_j \times T_j) \setminus E_j$. Now, we define the extended bet $\overline{g}$ as follows.

- At all states $(s_i'', t_i'', s_j'', t_j'') \notin [\![s_i, t_i]\!] \cup [\![s_i', t_i]\!]$, set $\overline{g}_k(s_i'', t_i'', s_j'', t_j'') := 0$ for every player $k \in \overline{I}$.
- At all states $(s_i'', t_i'', s_j'', t_j'') \in [\![s_i, t_i]\!] \cup [\![s_i', t_i]\!]$, set $\overline{g}_j(s_i'', t_i'', s_j'', t_j'') := 0$. Hence, for player i and the outside observer we have $\overline{g}_i(s_i'', t_i'', s_j'', t_j'') = -\overline{g}_d(s_i'', t_i'', s_j'', t_j'')$, where i's payments are shown in the table below.

	E_j	E_j^c
(s_i, t_i)	v_1	v_2
(s_i', t_i)	v_3	v_4

That is, payments are measurable with the respect to the events in $\{\{(s_i, t_i)\}, \{(s_i', t_i)\}\} \times \{E_j, E_j^c\}$. Then we set i's payments to be such that

$$\beta_i(t_i)(E_j)v_1 + \beta_i(t_i)(E_j^c)v_2 = \beta_i(t_i)(E_j)v_3 + \beta_i(t_i)(E_j^c)v_4 = 0,$$

with $v_1 < 0$ and $v_4 < 0$, and *a fortiori* $v_2 > 0$ and $v_3 > 0$.

By construction, both i's and j's expected payoffs from the bet are equal to 0 at all states, i.e.,

$$\mathbb{E}[\overline{g}_i | s_i'', t_i''] = \mathbb{E}[\overline{g}_j | s_j'', t_j''] = 0,$$

at all $(s_i'', t_i'', s_j'', t_j'') \in \Omega$. Hence, the bet $\overline{g}$ is SI for players i and j and moreover it is (weakly) acceptable for both of them. Now, observe that

$$\begin{aligned}
\mathbb{E}[\overline{g}_d | s_d, t_d] &= -\pi([\![s_i, t_i]\!])\Big(\pi(E_j | [\![s_i, t_i]\!])v_1 + \pi(E_j^c | [\![s_i, t_i]\!])v_2\Big) \\
&\quad -\pi([\![s_i', t_i]\!])\Big(\pi(E_j | [\![s_i', t_i]\!])v_3 + \pi(E_j^c | [\![s_i', t_i]\!])v_4\Big) \\
&> -\pi([\![s_i, t_i]\!])\Big(\beta_i(t_i)(E_j)v_1 + \beta_i(t_i)(E_j^c)v_2\Big) \\
&\quad -\pi([\![s_i', t_i]\!])\Big(\beta_i(t_i)(E_j)v_3 + \beta_i(t_i)(E_j^c)v_4\Big) \\
&= 0,
\end{aligned}$$

which implies that the outside observer is strictly willing to accept the extended bet $\overline{g}$. Thus, overall $\overline{g}$ is a mutually acceptable extended SI bet, which completes the proof. ■

References

Aumann, R.J. (1976). Agreeing to Disagree. *The Annals of Statistics*, **4**, 1236–1239. [2, 3]

——— (1987). Correlated Equilibrium as an Expression of Bayesian Rationality. *Econometrica*, **55**, 1–18. [1, 3, 10]

——— (1998). Common Priors: A Reply to Gul. *Econometrica*, **66**, 929–938. [3]

Aumann, R.J., Brandenburger, A. (1995). Epistemic Conditions for Nash Equilibrium. *Econometrica*, **63**, 1161–1180. [1]

Barelli, P. (2009). Consistency of Beliefs and Epistemic Conditions for Nash and Correlated Equilibria. *Games and Economic Behavior*, **67**, 363–375. [2]

Battigalli, P., Friedenberg, A., Siniscalchi, M. (Work in Progress). Epistemic Game Theory: Reasoning about Strategic Uncertainty. [1, 2]

Bonanno, G., Nehring, K. (1999). How to Make Sense of The Common Prior Assumption under Incomplete Information. *International Journal of Game Theory*, **28**, 409–434. [3]

Brandenburger, A., Dekel, E. (1993). Hierarchies of Beliefs and Common Knowledge. *Journal of Economic Theory*, **59**, 189–198. [2]

Dekel, E., Siniscalchi, M. (2015). Epistemic Game Theory. In *Handbook of Game Theory* (H.P. Young, S. Zamir, editors), volume IV. North-Holland, Amsterdam. [1, 2, 4, 5, 6, 11]

Feinberg, Y. (2000). Characterizing Common Priors in the Form of Posteriors. *Journal of Economic Theory*, **91**, 127–179. [2, 5, 7, 13]

Gul, F. (1998). A Comment on Aumann's Bayesian View. *Econometrica*, **66**, 923–927. [3]

Milgrom, P., Stokey, N. (1982). Information, Trade and Common Knowledge. *Journal of Economic Theory*, **26**, 17–27. [2, 5]

Samet, D. (1998). Common Priors and Separation of Convex Sets. *Games and Economic Behavior*, **24**, 172–174. [2, 5, 7, 13]

Sebenius, J.K., Geanakoplos, J. (1983). Don't Bet on It: Contingent Agreements with Asymmetric Information. *Journal of the American Statistical Association*, **78**, 424–426. [2, 5, 6]