

Conservative Belief and Rationality

Joseph Y. Halpern and Rafael Pass
Department of Computer Science
Cornell University
Ithaca, NY, 14853, U.S.A.
e-mail: halpern@cs.cornell.edu, rafael@cs.cornell.edu

October 3, 2012

Abstract

Brandenburger and Dekel have shown that common belief of rationality (CBR) characterizes rationalizable strategies, which are also characterized by a refinement of subjective correlated equilibrium called *a posteriori* equilibrium. It is possible that players' beliefs are incompatible, in the sense that player i can assign probability 1 to an event E to which player j assigns probability 0. One way to block incompatibility is to assume a common prior. We consider here a different approach: we require players' beliefs to be *conservative*, in the sense that all players must ascribe the actual world positive probability. Aumann has shown that, under the *common prior assumption* (CPA), common belief of rationality characterizes strategies in the support of an objective correlated equilibrium. Under the CPA, without loss of generality, all players' beliefs can be assumed to be conservative. We consider the consequences of common conservative belief of rationality (CCBR), without the common prior assumption. We show that CCBR characterizes strategies in the support of a subjective correlated equilibrium where all players' beliefs have common support. We also define a notion of strong rationalizability, and show that this is characterized by CCBR.

1 Introduction

In the Bayesian view of the world, each player has a subjective probability distribution (describing her beliefs) over the states of the world. In general, player’s beliefs may be incompatible, in the sense that one player might believe that an event E has probability 1, while another might ascribe E probability 0. (Technically, this says that player’s beliefs may not be *absolutely continuous* with respect to each other.) One way to assure that players’ beliefs are compatible is to make the *common prior assumption* (i.e., each player’s beliefs are generated by conditioning a common prior on the player’s information). Here we consider a different approach to ensuring that players’ beliefs are compatible: we restrict to beliefs that are *conservative*, where player i ’s beliefs are conservative if player i ascribes positive probability to the actual state.¹ If a player starts off with a subjective prior assigning non-zero probability to all states of the world, and updates by conditioning on information received, then his beliefs will always be conservative. Note that with conservative beliefs, players will never be “surprised”—they will never find themselves in a position to which they initially assigned probability 0.

Tan and Werlang [1988] and Brandenburger and Dekel [1987] show that common belief of rationality (CBR) characterizes rationalizable strategies [Bernheim 1984; Pearce 1984]. Brandenburger and Dekel [1987] further show that rationalizable strategies are exactly those in the support of an *a posteriori equilibrium*, a special case of a *subjective correlated equilibrium* (SCE) [Aumann 1974].² Aumann [1987] shows that, under the *common prior assumption* (CPA), CBR is characterized by strategies in the support of an objective correlated equilibrium (CE).³ Here we will be interested in what happens to these results if we restrict to conservative beliefs; more specifically, we will be interested in the consequences of *common conservative belief of rationality* (CCBR). Under the CPA, players’ beliefs regarding rationality can always be taken to be conservative, so under the CPA, CBR and CCBR coincide (see Appendix B for a formalization of this statement). Without the CPA, this is no longer the case: as is well known, in the normal-form version of the centipede game [Rosenthal 1982], all pure strategies are rationalizable, and can thus be played in a world where CBR holds; in contrast, CCBR implies the backward induction solution (even without a common prior).

Our main result is a characterization of CCBR without the common prior assumption. We first show that CCBR characterizes strategies that satisfy a notion of *strong rationalizability*; roughly speaking, this notion is identical to the traditional notion of rationalizability with an additional absolute continuity condition—if one player assigns positive probability to some strategy profile, then all players assign positive probability to this profile. We next establish an analogue of the Brandenburger-Dekel result, showing that strongly rationalizable strategies are exactly those in the support of a SCE where all players’ beliefs have common support; we call this a *common-support subjective correlated equilibrium* (CS-SCE). Interestingly, Aumann already alluded to such a notion in his original paper defining correlated equilibrium [Aumann 1974], but, as far as we know, it has not been investigated further. Since every CE is a CS-SCE, our characterization of CCBR in terms of CS-SCE shows that strategies

¹Actually, this is the case only if the state space is finite. In general, we require that the actual state be in the support of player i ’s beliefs.

²Roughly speaking, a subjective correlated equilibrium of a game G is an equilibrium resulting from augmenting G with a mediator who gives players “advice”, where players may have different subjective probabilities on the likelihood of various pieces of advice being given.

³That is, the players can now be viewed as playing with a mediator who gives advice, and have common beliefs about the likelihood of various pieces of advice being given.

consistent with CCBP always exists in finite games. So, to informally summarize this discussion, we have three distinct “epistemic” solution concepts—“CBP+CPA = CCBP+CPA < CCBP < CBP”—corresponding, respectively, to the equilibrium notions of CE, CS-SCE, and a posteriori equilibrium.

The work closest to ours is that of Stuart [1997], who considers the consequences of CBP under the assumption that players’ beliefs satisfy a variant of the notion of mutual absolute continuity (MAC). This condition implies that it is a common belief that players assign positive probability to the true state of the world. That is, it is common belief that players’ beliefs are conservative (although this may not be true). Stuart shows that in the finitely repeated prisoner’s dilemma, CBP+MAC suffices to get the backward induction outcome, but, in general, CBP+MAC does not imply that a Nash equilibrium outcome is played in games of perfect information. Stuart, however, does not provide a characterization of the set of strategies consistent with CBP+MAC.

A strategy is consistent with CBP+MAC iff it is consistent with CCBP (see Appendix B for a formalization of this statement). Thus, in a sense, conservative beliefs and mutually absolutely continuous beliefs, while different ways of characterizing compatibility of players’ beliefs, turn out to be essentially equivalent in the context of common belief of rationality. Our results thus also provide a characterization of strategies consistent with CBP+MAC.

Ben-Porath [1997] considers conservative beliefs in the context of extensive-form games; roughly speaking, he shows (in our language) that if CCBP holds and, in addition, the support of each player’s beliefs is common belief, then the outcome of play is a Nash equilibrium outcome.

The rest of the paper is organized as follows. In the next section, we define a logical language and models for reasoning about (conservative) belief and common (conservative) belief in games. In Section 3, we define strong rationalizability, show that it is characterized by CCBP, and, finally, show that strong rationalizability characterizes common-support subjective correlated equilibrium. To simplify the exposition, we consider only finite models here. This means that we are considering only finitely many types of players.⁴

2 Kripke structures for games

We consider normal-form games with n players. Given a (normal-form) n -player game Γ , let $\Sigma_i(\Gamma)$ denote the strategies of player i in Γ . To reason about the game Γ , we consider a class of probability structures corresponding to Γ . A *finite probability structure* M appropriate for Γ is a tuple $(\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, where Ω is a finite set of states; \mathbf{s} associates with each state $\omega \in \Omega$ a pure strategy profile $\mathbf{s}(\omega)$ in the game Γ ; and, for each player i , \mathcal{PR}_i associates with each state $\omega \in \Omega$ a probability distribution $\mathcal{PR}_i(\omega)$ on Ω such that

1. $\mathcal{PR}_i(\omega)(\llbracket \mathbf{s}_i(\omega) \rrbracket_M) = 1$, where for each strategy σ_i for player i , $\llbracket \sigma_i \rrbracket_M = \{\omega : \mathbf{s}_i(\omega) = \sigma_i\}$, and $\mathbf{s}_i(\omega)$ denotes player i ’s strategy in the strategy profile $\mathbf{s}(\omega)$,
2. $\mathcal{PR}_i(\omega)(\llbracket \mathcal{PR}_i(\omega), i \rrbracket_M) = 1$, where for each probability measure π and player i , $\llbracket \pi, i \rrbracket_M = \{\omega : \mathcal{PR}_i(\omega) = \pi\}$.

⁴The restriction to finite models is essentially without loss of generality if we assume that the support of all probability measures in the model has probability 1. While this assumption does not hold in general, it is satisfied under a number of standard topological assumptions; see [Halpern and Pass 2011] for more details.

These assumptions say that player i assigns probability 1 to his actual strategy and beliefs. (We have implicitly assumed that all sets in finite structures are measurable.)

Given a probability structure $M = (\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, let

- $play_i(\sigma_i)$ (“player i plays σ_i ”) be the set of states $\omega \in \Omega$ such that $\sigma_i = \mathbf{s}_i(\omega)$;
- RAT_i (“player i is rational”) be the set of states $\omega \in \Omega$ such that $\mathbf{s}_i(\omega)$ is a best response, given player i ’s beliefs on the strategies of other players induced by $\mathcal{PR}_i(\omega)$; and let
- $\mathbf{B}_i(E)$ (“player i believes E ”) be the set of states $\omega \in \Omega$ such that $\mathcal{PR}_i(\omega)(E) = 1$.

Let $RAT = RAT_1 \cap \dots \cap RAT_n$ and $play(\vec{\sigma}) = play_1(\sigma_1) \cap \dots \cap play_n(\sigma_n)$.

We next define the *conservative belief* operator \mathbf{B}^* . We want to capture the intuition that if a player has a conservative belief, the player should not rule out anything he has not observed evidence against. This, in particular, means that the player should assign positive probability to all events that occur in the true state of the world (as he could never have observed evidence against them). One way to formalize this would be to restrict to structures where all players’ beliefs satisfy this condition (as in Aumann’s [1974] model). In order to consider statements involving both belief and conservative belief (and to be able to compare CCBR with CPA+CBR), we want to allow models where worlds can have probability 0. Roughly speaking, we say that conservative belief of event E holds in a world ω iff (a) the agent believes E at ω (i.e., the agent’s probability measure at ω assigns E probability 1) and (b) the agent assigns ω positive probability.⁵ Formally, let $Supp(\pi)$ denote the support of the probability measure π . Say that a state ω is *i -conservative* if $\omega \in Supp(\mathcal{PR}_i(\omega))$ —roughly speaking, agent i ’s probability distribution at ω assigns positive probability to ω . Say that ω is *conservative* if ω is *i -conservative* for all agents i .

- Let $\mathbf{B}_i^*(E)$ (“player i conservatively believes E ”) be the set of states $\omega \in \mathbf{B}_i(E)$ such that ω is *i -conservative*.

Let $\mathbf{EB}(E)$ (“everyone believes E ”) be defined as $\mathbf{B}_1(E) \cap \dots \cap \mathbf{B}_n(E)$; and define $\mathbf{EB}^k(E)$ for all k inductively by letting $\mathbf{EB}^1(E) = \mathbf{EB}(E)$ and $\mathbf{EB}^{k+1}(E) = \mathbf{EB}(\mathbf{EB}^k(E))$. Let $\mathbf{EB}^*(E)$ (“everyone conservatively believes E ”) be $\mathbf{B}_1^*(E) \cap \dots \cap \mathbf{B}_n^*(E)$, and define $(\mathbf{EB}^*)^k$ analogously to \mathbf{EB}^k . Let $\mathbf{CB}(E)$ (“common belief of E holds”) be $\bigcap_{k>1} \mathbf{EB}^k(E)$, and let $\mathbf{CB}^*(E)$ (“common conservative belief of E holds”) be $\bigcap_{k>1} (\mathbf{EB}^*)^k(E)$. The following observation will prove useful.

Lemma 2.1: *If Γ is a game, $M = (\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ is a finite structure that is appropriate for Γ , and $\omega \in \mathbf{CB}^*(RAT)$, then for all players j , $Supp(\mathcal{PR}_j(\omega)) \subseteq \mathbf{CB}^*(RAT)$.*

Proof: Suppose that $\omega \in \mathbf{CB}^*(RAT)$ and $\omega' \in Supp(\mathcal{PR}_j(\omega))$. Since $\omega \in (\mathbf{EB}^*)^k(RAT)$, it follows that $\omega \in B_j((\mathbf{EB}^*)^{k-1}(RAT))$, so $\omega' \in (\mathbf{EB}^*)^{k-1}(RAT)$. Thus, $\omega' \in \mathbf{CB}^*(RAT)$. Hence, $Supp(\mathcal{PR}_j(\omega)) \subseteq \mathbf{CB}^*(RAT)$. ■

⁵In an earlier version of this paper [Halpern and Pass 2011], we used a logic-based approach, and defined events E based on formulas φ in some underlying logic. We took conservative belief of φ to hold if φ was *valid*—true in all states in all structures. This meant, for example, that an agent i would conservatively believe *true* in all states, even ones that he did not assign positive probability. While this made the resulting logic more similar in spirit to standard modal logics, it was an *ad hoc* assumption and made the semantics more complicated. Since the assumption played no role in any of our results, we have not made it here.

3 Characterizing CCBR

To put our result into context, we first restate the characterizations of rationalizability given by Tan and Werlang [1988] and Brandenburger and Dekel [1987] in our language. We first recall Pearce’s [1984] definition of rationalizability.

Definition 3.1: A strategy σ_i for player i in game Γ is *rationalizable* if, for each player j , there is a set $\mathcal{Z}_j \subseteq \Sigma_j(\Gamma)$ and, for each strategy $\sigma'_j \in \mathcal{Z}_j$, a probability measure $\mu_{\sigma'_j}$ on $\Sigma_{-j}(\Gamma)$ whose support is a subset of \mathcal{Z}_{-j} such that

- $\sigma_i \in \mathcal{Z}_i$; and
- for each player j and strategy $\sigma'_j \in \mathcal{Z}_j$, strategy σ'_j is a best response to (the beliefs) $\mu_{\sigma'_j}$.

■

Theorem 3.2: [Brandenburger and Dekel 1987; Tan and Werlang 1988] σ_i is a rationalizable strategy for i in a game Γ iff there exists a finite structure M that is appropriate for Γ and a state ω such that $\omega \in \mathbf{CB}(\text{RAT}) \cap \text{play}_i(\sigma_i)$.

It is well-known [Pearce 1984] that a strategy is rationalizable iff it survives iterated deletion of strongly dominated strategies. Thus, the characterization of rationalizability in Theorem 3.2 is also a characterization of strategies that survive iterated deletion of strongly dominated strategies. As a result, common belief of rationality is always consistent—that is, every finite game has a structure where common belief of rationality holds. Furthermore, since iterated deletion of strongly dominated strategies can be done in polynomial time, there exists an efficient procedure that for every player finds a strategy that is consistent with common belief of rationality.

Clearly CCBR implies CBR, so in a state satisfying CCBR, the strategies used must be rationalizable. However, not all rationalizable strategies can be played in a state satisfying CCBR (see Appendix A for an example). We now introduce a strengthening of the notion of rationalizability called *strong rationalizability*, which characterizes strategies consistent with CCBR. Roughly speaking, this notion is defined just as the traditional notion of rationalizability, but with an additional “absolute continuity” condition.

Definition 3.3: A strategy σ_i for player i in game Γ is *strongly rationalizable* if, for each player j , there is a set $\mathcal{Z}_j \subseteq \Sigma_j(\Gamma)$ and, for each strategy $\sigma'_j \in \mathcal{Z}_j$, a probability measure $\mu_{\sigma'_j}$ on $\Sigma_{-j}(\Gamma)$ whose support is a subset of \mathcal{Z}_{-j} such that

1. $\sigma_i \in \mathcal{Z}_i$;
2. for all players j and strategies $\sigma'_j \in \mathcal{Z}_j$, σ'_j is a best response to (the beliefs) $\mu_{\sigma'_j}$; and
3. for all players j, h and all strategy profiles $\vec{\sigma}' \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$, if $\mu_{\sigma'_j}(\sigma'_{-j}) > 0$, then $\mu_{\sigma'_h}(\sigma'_{-h}) > 0$.

■

We now show that a strategy is strongly rationalizable iff it is played in a world where CCBR holds.

Theorem 3.4: σ_i is a strongly rationalizable strategy for i in a game Γ iff there exists a finite structure M that is appropriate for Γ and a state ω such that $\omega \in \mathbf{CB}^*(\text{RAT}) \cap \text{play}_i(\sigma_i)$.

Proof: For the “only if” direction, suppose that σ_i is strongly rationalizable. Choose $\mathcal{Z}_j \subseteq \Sigma_j(\Gamma)$ and measures $\mu_{\sigma'_j}$ for each strategy $\sigma'_j \in \mathcal{Z}_j$ guaranteed to exist by Definition 3.3. Define an appropriate structure $M = (\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$, where

- $\Omega = \{\vec{\sigma}' \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n : \text{there exists a player } j \text{ such that } \mu_{\sigma'_j}(\sigma'_{-j}) > 0\}$;
- $\mathbf{s}(\vec{\sigma}') = \sigma'$;
- $\mathcal{PR}_j(\vec{\sigma}')(\vec{\sigma}'') = 0$ if $\sigma'_j \neq \sigma''_j$, and $\mathcal{PR}_j(\vec{\sigma}')(\vec{\sigma}'') = \mu_{\sigma'_j}(\sigma''_{-j})$ otherwise. (Note that this is a well-defined probability distribution since, if $\sigma''_{-j} \in \text{Supp}(\mu_{\sigma'_j})$, then $(\sigma'_j, \sigma''_{-j}) \in \Omega$.)

Since $\sigma_i \in \mathcal{Z}_i$, there exists some state $\omega \in \Omega$ such that $\mathbf{s}_i(\omega) = \sigma_i$. By Condition 2 in the definition of strong rationalizability, each player is best responding to his beliefs at every state, so we have that $\vec{\sigma}' \in \text{RAT}$ for all states $\vec{\sigma}'$. We next show that for every player j and every strategy profile $\vec{\sigma}' \in \Omega$, we have that $\mathcal{PR}_j(\vec{\sigma}')(\{\vec{\sigma}'\}) > 0$. This will show that every state in Ω is conservative. For suppose that $\vec{\sigma}' \in \Omega$. By definition, $\mathcal{PR}_j(\vec{\sigma}')(\{\vec{\sigma}'\}) = \mu_{\sigma'_j}(\sigma'_{-j})$. Furthermore, recall that $\vec{\sigma}' \in \Omega$ iff there exists some h such that $\mu_{\sigma'_h}(\sigma'_{-h}) > 0$. By Condition 3 in the definition of strong rationalizability, we have that if $\mu_{\sigma'_h}(\sigma'_{-h}) > 0$, then $\mu_{\sigma'_j}(\sigma'_{-j}) > 0$, and thus $\mathcal{PR}_j(\vec{\sigma}')(\{\vec{\sigma}'\}) > 0$. It now easily follows that $\vec{\sigma}' \in \mathbf{CB}^*(\text{RAT})$ for all $\vec{\sigma}' \in \Omega$. This proves the only if direction.

For the “if” direction, suppose that $M = (\Omega, \mathbf{s}, \mathcal{PR}_1, \dots, \mathcal{PR}_n)$ is appropriate for Γ , and $\omega \in \mathbf{CB}^*(\text{RAT})$. We want to show that $\sigma_i = \mathbf{s}_i(\omega)$ is strongly rationalizable. Note that by Lemma 2.1, we have that for every j , $\omega' \in \mathbf{CB}^*(\text{RAT})$, the support of $\mathcal{PR}_j(\omega')$ is a subset of $\mathbf{CB}^*(\text{RAT})$.

For each player j , let $\mathcal{Z}_j = \{\mathbf{s}_j(\omega') : \omega' \in \mathbf{CB}^*(\text{RAT})\}$. By definition, $\sigma_i \in \mathcal{Z}_i$. For each j and each $\sigma'_j \in \mathcal{Z}_j$, we now define $\mu_{\sigma'_j}$. For $\sigma'_j \in \mathcal{Z}_j$, let $\Omega_{\sigma'_j} = \{\omega' \in \mathbf{CB}^*(\text{RAT}) : \mathbf{s}_j(\omega') = \sigma'_j\}$. Let $\mu_{\sigma'_j}$ be the result of projecting $\sum_{\omega' \in \Omega_{\sigma'_j}} \mathcal{PR}_j(\omega') / |\Omega_{\sigma'_j}|$ onto $\Sigma_{-j}(\Gamma)$. Since, for every $\omega' \in \mathbf{CB}^*(\text{RAT})$, the support of $\mathcal{PR}_j(\omega')$ is a subset of $\mathbf{CB}^*(\text{RAT})$, we have that the support of $\mu_{\sigma'_j}$ is a subset of \mathcal{Z}_{-j} . Furthermore, since j is rational at each state in $\mathbf{CB}^*(\text{RAT})$, it follows that σ'_j is a best response to the projection of $\mathcal{PR}_j(\omega')$ onto $\Sigma_{-j}(\Gamma)$ for each $\omega' \in \Omega_{\sigma'_j}$; thus, σ'_j is also a best response to a convex combination of these projections.

Finally, we need to show that for every strategy profile $\vec{\sigma}' \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ and all players j and h , if $\mu_{\sigma'_j}(\sigma'_{-j}) > 0$ then $\mu_{\sigma'_h}(\sigma'_{-h}) > 0$. So suppose that $\vec{\sigma}' \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ and $\mu_{\sigma'_j}(\sigma'_{-j}) > 0$; that is, there exists some state $\omega' \in \Omega_{\sigma'_j}$ such that $\mathcal{PR}_j(\omega')([\sigma'_{-j}]_M) > 0$. Since M is a structure appropriate for Γ , we must have that $\mathcal{PR}_j(\omega')([\sigma'_j]_M) = 1$. Thus, $\mathcal{PR}_j(\omega')([\sigma']_M) > 0$, and there must be a state ω'' such that $\mathcal{PR}_j(\omega')(\omega'') > 0$ and $\mathbf{s}(\omega'') = \vec{\sigma}'$. Since the support of $\mathcal{PR}_j(\omega')$ is a subset of $\mathbf{CB}^*(\text{RAT})$, $\omega'' \in \mathbf{CB}^*(\text{RAT})$, and thus $\omega'' \in \Omega_{\sigma'_h}$. Since $\omega'' \in \mathbf{CB}^*(\text{RAT})$, we have that $\mathcal{PR}_h(\omega'')(\omega'') > 0$. It follows that $\mu_{\sigma'_h}$ assigns positive probability to σ'_{-h} . Thus, σ_i is strongly rationalizable. ■

We next show that strongly rationalizable strategies are in the support of a subjective correlated equilibrium where the players’ beliefs have common support. We first recall the definitions of objective and subjective correlated equilibrium.

Definition 3.5: An (*objective*) *correlated equilibrium (CE)* in a game Γ is a distribution π on $\Sigma_1(\Gamma) \times \dots \times \Sigma_n(\Gamma)$ such that, for all players i and strategies $\sigma_i \in \Sigma_i(\Gamma)$ such that $\pi(\sigma_i) > 0$, we have that for all $\sigma'_i \in \Sigma_i(\Gamma)$,

$$\sum_{\sigma_{-i} \in \Sigma_{-i}(\Gamma)} \pi(\sigma_{-i} | \sigma_i) u(\sigma) \geq \sum_{\sigma_{-i} \in \Sigma_{-i}(\Gamma)} \pi(\sigma_{-i} | \sigma_i) u(\sigma'_i, \sigma_{-i}).$$

■

In a *subjective correlated equilibrium*, rather than there being one common probability measure π , each player i has her own probability measure π_i on $\Sigma_1(\Gamma) \times \dots \times \Sigma_n(\Gamma)$.

Definition 3.6: A *subjective correlated equilibrium (SCE)* in a game Γ is a profile $\vec{\pi}$ of probability measures on $\Sigma_1(\Gamma) \times \dots \times \Sigma_n(\Gamma)$ such that for all players i and strategies $\sigma_i \in \Sigma_i$ such that $\pi_i(\sigma_i) > 0$, we have that for all $\sigma'_i \in \Sigma_i(\Gamma)$,

$$\sum_{\sigma_{-i} \in \Sigma_{-i}(\Gamma)} \pi_i(\sigma_{-i} | \sigma_i) u(\sigma) \geq \sum_{\sigma_{-i} \in \Sigma_{-i}(\Gamma)} \pi_i(\sigma_{-i} | \sigma_i) u(\sigma'_i, \sigma_{-i}).$$

■

As pointed out by Aumann, in an SCE each player i assigns (subjective) probability 0 to the event that she wishes to deviate after receiving its signal, but some other player j might assign positive probability to the event that i wishes to deviate. Aumann [1974] suggested two method for circumventing this problem. The first was to require that all the π_i 's have the same support; we call such an SCE a *common-support subjective correlated equilibrium*.

Definition 3.7: $\vec{\pi}$ is a *common-support subjective correlated equilibrium (CS-SCE)* if it is a SCE and, for all players i and j , the probability measures π_i and π_j have the same support. ■

Clearly, if π is an objective correlated equilibrium, then (π, \dots, π) is a CS-SCE. As the next example shows, the converse does not hold in general.

Example 3.8: Consider the simple 2-player game where $\Sigma_1 = \{L, R\}$ and $\Sigma_2 = \{\ell, r\}$. The payoffs in the game are given in the first table, π_1 is given in the second table, and π_2 is given in the third table. We leave it to the reader to check that (π_1, π_2) is a CS-SCE. Intuitively, player 1 justifies his moves by believing that player 2 is more likely to play ℓ when he plays L and that player 2 is more likely to play r when he plays R ; player 2 has symmetric beliefs. Since $\pi_1 \neq \pi_2$, (π_1, π_2) is not a CE.

	L	R
ℓ	(1, 0)	(0, 1)
r	(0, 1)	(1, 0)

	L	R
ℓ	.45	.05
r	.05	.45

	L	R
ℓ	.05	.45
r	.45	.05

■

The second (and seemingly preferred) method of Aumann to circumvent the problem above was to require that the equilibrium strategy remains optimal even after the players receive their signal, no matter what the signal is. Aumann called this refinement an *a posteriori equilibrium*. This notion was

further formalized by Brandenburger and Dekel [1987], who also show that it characterizes the set of rationalizable strategies (and thus also characterizes strategies consistent with CBR).

As we now show, CS-CSE characterizes strongly rationalizable strategies, and thus, by Theorem 3.4, the set of strategies consistent with CCBR.

Theorem 3.9: σ_i is a strongly rationalizable strategy for i in a game Γ iff there exists a common-support subjective correlated equilibrium $\vec{\pi}$ such that σ_i is in the support of the projection of π_i onto $\Sigma_i(\Gamma)$.

Proof: For the “only if” direction, for every player j , choose $\mathcal{Z}_j \subseteq \Sigma_j(\Gamma)$ and the probability measure $\mu_{\sigma'_j}$ for each strategy $\sigma'_j \in \mathcal{Z}_j$ guaranteed to exist by Definition 3.3; define π_j as the distribution obtained by uniformly sampling $\sigma'_j \in \mathcal{Z}_j$ and then sampling σ'_{-j} according to $\mu_{\sigma'_j}$; that is, $\pi_j(\vec{\sigma}') = \mu_{\sigma'_j}(\sigma'_{-j})/|\mathcal{Z}_j|$. Clearly $\vec{\pi}$ is a subjective correlated equilibrium and σ_i is in the support of the projection of π_i onto $\Sigma_i(\Gamma)$ (since $\sigma_i \in \mathcal{Z}_i$). To show that $\vec{\pi}$ has common support, consider any profile $\vec{\sigma}'$. For all players j and h , if $\pi_j(\vec{\sigma}') > 0$, then $\sigma'_j \in \mathcal{Z}_j$, and $\mu_{\sigma'_j}(\sigma'_{-j}) > 0$ (by construction), which means that $\sigma'_h \in \mathcal{Z}_h$ and $\mu_{\sigma'_h}(\sigma'_{-h}) > 0$ (by definition of strong rationalizability); it follows that $\pi_h(\vec{\sigma}') > 0$.

For the “if” direction, choose the profile $\vec{\pi}$ of probability measures guaranteed to exist by Definition 3.7. For every player j , define \mathcal{Z}_j as the support of the projection of π_j onto $\Sigma_j(\Gamma)$, and for every strategy $\sigma'_j \in \mathcal{Z}_j$, define $\mu_{\sigma'_j}(\sigma'_{-j}) = \pi_j(\sigma'_{-j}|\sigma'_j)$; it follows that σ'_j is a best response to $\mu_{\sigma'_j}$. Furthermore, $\sigma_i \in \mathcal{Z}_i$, since it is in the support of the projection of π_i onto $\Sigma_i(\Gamma)$. Since, for all players j and h , the measures π_j and π_h have the same support, we have that for all strategy profiles $\vec{\sigma}' \in \mathcal{Z}_1 \times \dots \times \mathcal{Z}_n$ such that $\mu_{\sigma'_j}(\sigma'_{-j}) > 0$, $\mu_{\sigma'_h}(\sigma'_{-h}) > 0$. Finally, for all players j and strategies $\sigma'_j \in \mathcal{Z}_j$, to see that the support of $\mu_{\sigma'_j}$ is a subset of \mathcal{Z}_{-j} , note that the support of $\mu_{\sigma'_j}$ is a subset of the projection of the support of π_j onto $\Sigma_{-j}(\Gamma)$. The support of $\mu_{\sigma'_j}$ is thus a subset of $\prod_{h \in [-j]} \mathcal{Z}_h^j$, where for all players j' and h , $\mathcal{Z}_h^{j'}$ is the projection of the support of $\pi_{j'}$ onto $\Sigma_h(\Gamma)$. But since $\vec{\pi}$ has common support, we have that, for all players j' and h , $\mathcal{Z}_h^{j'} = \mathcal{Z}_h^h = \mathcal{Z}_h$, and thus the support of $\mu_{\sigma'_j}$ is a subset of \mathcal{Z}_{-j} . ■

Acknowledgments: We thank Bob Aumann for asking the questions that inspired us to think about strong rationalizability (although they were asked in a completely different context!) and Adam Brandenburger for pointing out Stuart’s work [Stuart 1997]. Halpern is supported in part by NSF grants ITR-0325453, IIS-0534064, IIS-0812045, and IIS-0911036, AFOSR grants FA9550-08-1-0438 and FA9550-09-1-0266, and ARO grant W911NF-09-1-0281. Pass is supported in part by a Alfred P. Sloan Fellowship, a Microsoft New Faculty Fellowship, NSF CAREER Award CCF-0746990, AFOSR YIP Award FA9550-10-1-0093, and DARPA and AFRL under contract FA8750-11-2-0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US government.

A An example separating CBR and CCBR

As the following example shows, CBR does not imply CCBR in general.

Example A.1: Consider an N -move centipede game [Rosenthal 1982] viewed as a normal-form game, where $N > 1$ is even. More precisely, the action set Σ_1 for player 1 is $\{1, 3, \dots, N-1\}$, and the action set Σ_2 for player 2 is $\{2, 4, \dots, N\}$. (Intuitively, a move of k means that the player first stops at node k .) We take $u_1(k, m)$ to be $2^k + 1$ if $k < m$ and 2^{m-1} otherwise; similarly, we take $u_2(k, m)$ to be $2^m + 1$ if $m < k$ and 2^{k-1} otherwise. Thus, for example, if player 1 stops at round 1, then player 1's utility is 3 and player 2's utility is 1; if player 1 stops first at round 3 and player 2 stops at round 2, then player 1's utility is 2 and player 2's utility is 5. It is easy to see that all strategies are rationalizable (that is, we can take $\mathcal{Z}_1 = \Sigma_1$ and $\mathcal{Z}_2 = \Sigma_2$ in Definition 3.1.) For all strategies $j \in \Sigma_2$, we can take μ_j to put probability 1 on player 1 playing 1 (stopping right away). This belief clearly justifies player 2 playing j (i.e., stopping at j), for all j . For strategy $k \in \Sigma_1$, let μ_k put probability 1 on player 2 playing $k+1$. Clearly this belief makes playing k rational. Thus, all strategies for both players are rationalizable, and therefore compatible with CBR. By way of contrast, as we now show, the only pure strategy for player 1 that is compatible with CCBR is playing 1.

For definiteness, we do this for a 20-move centipede. Define a (k, k') -state to be a state where player 1 plays k and player 2 plays k' . We first claim that there is no $(k, 20)$ -state with $k > 1$ where CCBR holds. For if ω is such a $(k, 20)$ -state, then player 2 must ascribe ω positive probability at ω . Consider the largest k' such that player 2 ascribes positive probability to a $(k', 20)$ -state at ω . By assumption, $k' > 1$. Then, given his beliefs, player 2 is strictly better off playing $k' - 1$ than playing 20. This is a contradiction. Thus, there is no $(k, 20)$ -state compatible with CCBR and $k > 1$.

It now easily follows by Lemma 2.1 that there is no $(19, k)$ -state compatible with CCBR (because playing 19 can be justified for player 1 only if he ascribes positive probability to $(19, 20)$, but, as we have just argued, there is no such state compatible with CCBR). It follows that there is no $(k, 18)$ -state compatible with CCBR with $k > 1$ (since a $(k, 18)$ state can be justified only by a $(19, k)$ state). Continuing in this way, an easy induction shows that the only state that can be compatible with CCBR have the form $(1, k)$. It is also easy to see that the strategies played in these states are in fact compatible with CCBR. ■

This example shows that CCBR implies that the backward induction outcome is played in the centipede game. It follows from Corollary B.6 below and results of Stuart [1997] that the same is true in finitely repeated prisoner's dilemma. We might hope that the same is true in all games of perfect information, but Stuart [1997, Figure3a] gives a simple example showing that this is not the case in general; indeed, in general, with CCBR, not even the Nash equilibrium outcome is played. This observation follows as an easy corollary of our characterization of CCBR in terms of correlated equilibria: all (objective) correlated equilibrium strategies satisfy CCBR, and it is already well known that correlated equilibrium outcomes are not necessarily Nash equilibrium outcomes.

B Relating CCBR to CPA and MAC

In this appendix, we compare CCBR with (a) CBR in the presence of CPA and (b) Stuart's [1997] requirement of mutual absolute continuity.

We first make precise what we mean by CPA.

Definition B.1: A structure $M = (\Omega, \mathfrak{s}, \mathcal{P}\mathcal{R}_1, \dots, \mathcal{P}\mathcal{R}_n)$ appropriate for Γ satisfies the CPA if there exists a probability π on Ω and partitions Π_1, \dots, Π_n of Ω such that, for all $\omega \in \Omega$ and players i , the

support of $\mathcal{P}\mathcal{R}_i(\omega)$ is contained in $\Pi_i(\omega)$ (where $\Pi_i(\omega)$ is the cell of the partition Π_i containing ω), $\pi(\Pi_i(\omega)) > 0$, and for all $\omega' \in \Pi_i(\omega)$, $\mathcal{P}\mathcal{R}_i(\omega)(\omega') = \pi(\omega' | \Pi_i(\omega))$. That is, $\mathcal{P}\mathcal{R}_i(\omega)$ is obtained from π by conditioning on $\Pi_i(\omega)$.⁶ ■

Proposition B.2: *Suppose that $M = (\Omega, \mathbf{s}, \mathcal{P}\mathcal{R}_1, \dots, \mathcal{P}\mathcal{R}_n)$ is a finite structure appropriate for Γ satisfying the CPA. Then for every state $\omega \in \Omega$, if ω is i -conservative for some player i , then ω is conservative. Furthermore, for every conservative state ω , we have*

- (a) $\omega \in \mathbf{B}_i(E)$ iff $\omega \in \mathbf{B}_i^*(E)$, and
- (b) $\omega \in \mathbf{CB}(E)$ iff $\omega \in \mathbf{CB}^*(E)$.

Proof: Since M satisfies the CPA, there is some probability π on Ω and partitions Π_1, \dots, Π_n satisfying the conditions of Definition B.1. By definition of $\mathcal{P}\mathcal{R}_1, \dots, \mathcal{P}\mathcal{R}_n$, we have that if ω is i -conservative for some i , then it is conservative. Part (a) follows from the definition of \mathbf{B}^* , while part (b) follows by an easy induction, relying on the fact that all i -conservative states are conservative. ■

Corollary B.3: *If M is a structure that satisfies CPA and $\omega \in \mathbf{CB}(RAT) \cap \text{play}_i(\sigma_i)$ for some state ω , then σ_i is strongly rationalizable.*

Proof: Let $M = (\Omega, \mathbf{s}, \mathcal{P}\mathcal{R}_1, \dots, \mathcal{P}\mathcal{R}_n), \omega, \sigma_i$ be as in the statement of Corollary B.3. Let ω' be some state in the support of $\mathcal{P}\mathcal{R}_i(\omega)$; by definition ω' is i -conservative, and thus, by Proposition B.2, conservative. Furthermore, $\omega \in \mathbf{CB}(RAT) \cap \text{play}_i(\sigma_i)$ which by Proposition B.2 means that $\omega \in \mathbf{CB}^*(RAT) \cap \text{play}_i(\sigma_i)$. Thus, by Theorem 3.4, σ_i is strongly rationalizable. ■

Turning to mutual absolute continuity, we first give Stuart's definition in our setting.⁷

Definition B.4: A structure $M = (\Omega, \mathbf{s}, \mathcal{P}\mathcal{R}_1, \dots, \mathcal{P}\mathcal{R}_n)$ appropriate for Γ satisfies mutual absolute continuity if, for all states ω , players i and j , we have $\mathcal{P}\mathcal{R}_i(\omega)(\omega) > 0$ iff $\mathcal{P}\mathcal{R}_j(\omega)(\omega) > 0$. ■

We can now prove a result similar in spirit to Proposition B.2.

Proposition B.5: *Suppose that $M = (\Omega, \mathbf{s}, \mathcal{P}\mathcal{R}_1, \dots, \mathcal{P}\mathcal{R}_n)$ is a finite structure appropriate for Γ satisfying mutual absolute continuity. Then for every state $\omega \in \Omega$, if ω is i -conservative for some $i \in [n]$, then ω is conservative. Furthermore, for every conservative state ω , we have*

- (a) $\omega \in \mathbf{B}_i(E)$ iff $\omega \in \mathbf{B}_i^*(E)$, and
- (b) $\omega \in \mathbf{CB}(E)$ iff $\omega \in \mathbf{CB}^*(E)$.

Proof: Mutual absolute continuity is equivalent to the statement that all i -conservative states are conservative. The rest of the proof is identical to the proof of Proposition B.2. ■

⁶The requirement that $\pi(\Pi_i(\omega)) > 0$ is implicitly made by Aumann, but not by all authors that consider the CPA (c.f. [Halpern 2002]). Authors who allow $\pi(\Pi_i(\omega))$ to be 0 place no constraints on what $\mathcal{P}\mathcal{R}_i(\omega')$ should be for $\omega' \in \Pi_i(\omega)$.

⁷As Stuart [1997] points out, this definition of mutual absolute continuity is similar to, but not the same as, the standard definition of the notion.

Corollary B.6: *There exists a finite structure M satisfying mutual absolute continuity and a state ω such that $\omega \in \mathbf{CB}(RAT) \cap \text{play}_i(\sigma_i)$ iff σ_i is strongly rationalizable.*

Proof: The “only if” direction follows just as in Corollary B.3. For the “if” direction, note that the structure in the proof of Theorem 3.4 satisfies mutual absolute continuity (since all states are conservative).

■

Friedenberg and Meier [2010, Lemma 8.2] show that a structure that satisfies the CPA also satisfies mutual absolute continuity. Combining the Friedenberg-Meier result with Proposition B.5 gives Proposition B.2.

References

- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96.
- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55, 1–18.
- Ben-Porath, E. (1997). Rationality, Nash equilibrium and backward induction in perfect information games. *Review of Economic Studies* 64, 23–46.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica* 52(4), 1007–1028.
- Brandenburger, A. and E. Dekel (1987). Rationalizability and correlated equilibria. *Econometrica* 55, 1391–1402.
- Friedenberg, A. and M. Meier (2010). The context of the game. Unpublished manuscript. A preliminary version of the paper appears in *Theoretical Aspects of Rationality and Knowledge: Proc. Twelfth Conference (TARK 2009)*, 2009, pp. 134–135.
- Halpern, J. Y. (2002). Characterizing the common prior assumption. *Journal of Economic Theory* 106(2), 316–355.
- Halpern, J. Y. and R. Pass (2011). Justified belief and rationality. Unpublished manuscript; available at www.cs.cornell.edu/home/halpern/papers/cbr.pdf.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4), 1029–1050.
- Rosenthal, R. W. (1982). Games of perfect information, predatory pricing, and the chain store paradox. *Journal of Economic Theory* 25, 92–100.
- Stuart, H. W. (1997). Common belief of rationality in the finitely repeated prisoners’ dilemma. *Games and Economic Behavior* 19(1), 133–143.
- Tan, T. and S. Werlang (1988). The Bayesian foundation of solution concepts of games. *Journal of Economic Theory* 45(45), 370–391.