

Rationalizable Partition-Confirmed Equilibrium

Drew Fudenberg and Yuichiro Kamada*

First Version: January 29, 2011; This Version: August 11, 2011

Abstract

Rationalizable partition-confirmed equilibrium (RPCE) describes the steady state outcomes of rational learning in extensive games, when rationality is almost common knowledge and players observe a partition of the terminal nodes. RPCE allows players to make inferences about unobserved play by others; we discuss the implications of this using numerous examples. We identify the ways in which the RPCE outcomes depend on terminal node partitions, provide conditions under which they are invariant with respect terminal node partitions, and discuss the relationship of RPCE to other solution concepts in the literature.

*Fudenberg: Department of Economics, Harvard University, Cambridge, MA 02138, e-mail: dfudenberg@harvard.edu; Kamada: Department of Economics, Harvard University, Cambridge, MA 02138, e-mail: ykamada@fas.harvard.edu; We are grateful to Pierpaolo Battigalli, Eddie Dekel, Ignacio Esponda, and David K. Levine for helpful comments. NSF grants SES-0646816 and SES-0951462 provided financial assistance.

1 Introduction

Observing actual play in a game need not reveal the actions that would be taken at information sets that have never been reached. Consequently, players with incorrect beliefs about off-path play need not learn that their beliefs are in error, and so the process may not lead to a Nash equilibrium. The simplest version of this idea is self-confirming equilibrium (SCE), which corresponds to cases where players observe the terminal node of the game each time it is played, and the only restrictions placed on the players' beliefs is that they be consistent with the objective distribution on terminal nodes.¹ However, there are many applications of interest in which players do not observe the exact terminal node that is reached. For example, consider a sealed-bid uniform-price k -unit auction for a good of known value. Here the terminal node is the entire vector of submitted bids, but agents might only observe the winning price and the identity of the winning bidders. Alternatively, this information might only be made available to those who submitted nonzero bids, with the others only told that their bid was not high enough. Moreover, the players might have private information about their values for the good; since the terminal node encodes Nature's move as well as the moves of the players, the assumption that the terminal node is observed at the end of the game implies that players observe one another's types.

This leads us to replace the assumption that terminal nodes are observed with the assumption that each player has a partition over the terminal nodes, and that at the end of each play of the game each player observes the corresponding element from their terminal node partition, resulting in a generalization of self-confirming equilibrium we call "partition-confirmed equilibrium."² However, our primary interest is in cases where the rationality of the players and the observation structure are both (almost) common knowledge, so that even when there is not a commonly known outcome, players can and do reason about what their opponents might be observing. For example, we will require that player 1's conjecture about how player 2 thinks player 3 is playing must be consistent

¹We do not explicitly study dynamics here, but the solution concepts we propose are motivated by the idea that agents play the game repeatedly, with no strategic links between repetitions, and that their play is "asymptotically empirical" and "asymptotically myopic" in the sense of Fudenberg and Kreps (1995). See that paper and Fudenberg and Levine (1993b) for non-equilibrium learning models with these properties, and Fudenberg and Levine (2009) for a survey of related work.

²Battigalli (1987) and Battigalli and Guatoli (1997) introduced the concept of terminal node partitions, and used them to define conjectural equilibrium, which is closely related to SCE and to the partition-confirmed equilibrium defined in this paper. Rubinstein and Wolinsky (1994) use terminal node partitions in their definition of rationalizable conjectural equilibrium (RCE). Dekel, Fudenberg, and Levine (2004) study the implications of various terminal node partitions for SCE in static Bayesian games. Esponda (2011) extends RCE to games with moves by Nature in an epistemic framework, and also defines weaker versions that relax common knowledge of rationality and consistency to k -th level belief.

with player 1’s information about what player 2 observes.

This leads us to develop the concept of “rationalizable partition-confirmed equilibrium” (RPCE) which generalizes the “rationalizable self-confirming equilibrium” (RSCE) of Dekel, Fudenberg, and Levine (1999, DFL). A key feature of both solution concepts is that a player may be able to use observations of play at some information sets to make inferences about play at information sets the player does not observe. As with RSCE, one motivation for RPCE is that it refines the predictions of both rationalizability (Bernheim (1984) and Pearce (1984)) and of SCE: rationalizability by requiring that players’ beliefs are consistent with their observations and their knowledge of the other players’ observation structures, and SCE by requiring players use knowledge about opponents’ payoff functions to refine their beliefs.

The main difference is that in RSCE all players observe the distribution over terminal nodes, while RPCE allows each player to have a different partition of the terminal nodes. In this case there is no longer a publicly observed outcome path, so the implications of common knowledge of the observation structure are less immediate.³ Roughly speaking, RPCE describes situations where players know that the outcome of play has converged, even when they do not observe all aspects of this outcome themselves.

The RPCE concept is of interest in its own right; it also serves to provide additional support for the use of Nash and subgame perfect equilibrium in games where it coincides with one or the other. In particular, we will see that players can do a fair bit of reasoning about play they do not observe, even when we do not assume that players know one another’s strategies.

Many of our points can be made using “participation games,” which are one-shot simultaneous move games in which some or all of the players have the option of an action called “Out”: If a player plays Out, his payoff is 0 regardless of the play of the others, and he observes only his own action and payoff. Roughly speaking, the idea of RPCE is that if player 1 (say) always plays Out, but knows that players 2 and 3 play every period and observe the terminal node at the end of each round, and player 1 believes that play has converged, then she can use her knowledge of the payoff functions and observation structure to place restrictions on the (unobserved) play of her opponents. In particular when players 2 and 3 observe the terminal node each period, 1’s belief about their play must be concentrated on the set of Nash equilibria of the “subgame” between them. In

³A similar complication arises in “heterogeneous” equilibria, in which different agents in the same player role take different actions and thus have different observations, which is why DFL restricted attention to the “unitary” case where all agents in a given player role use the same strategy and have the same beliefs. This paper restricts attention to unitary beliefs; we plan to extend RPCE to the heterogeneous case in future research.

contrast, if player 1's choice of "Out" ends the game and prevents players 2 and 3 from acting, then when player 1 always plays Out players 2 and 3 do not have the chance to learn; here the only restriction on 1's belief when he plays Out is that the play of 2 and 3 is rationalizable.

In addition to the partition over terminal nodes, this paper differs from DFL by allowing players to have correlated beliefs about their opponents' play. As we argue by example, restricting to independent beliefs is less natural here, for when a player knows that her opponents have repeatedly played a coordination game, but has not seen their actions, it seems odd to require that the player's beliefs about the opponents correspond to a product distribution. Put differently, with partitions on terminal nodes, play of the game on its own may provide some of the players access to a common signal that is not observed by others.

In general the set of RPCE outcomes depends on the terminal node partitions, because coarser partitions provide less information and so generate fewer restrictions on the allowed beliefs. We identify four different reasons that the dependence can arise, and then give conditions under which RPCE is invariant to changes of partitions. Roughly speaking, the set of player i 's RPCE strategy is invariant to changes in her terminal node partition in simultaneous-move games when the actions of each opponent has no effect on what that opponent observes.

The paper is organized as follows. In Section 2, we provide an example to illuminate our point. In the example, we discuss two related extensive-form games, and explain why we might expect them to have different long-run outcomes under rational learning. Section 3 defines a model of extensive-form games with terminal node partitions. Section 4 revisits the example in Section 2, and analyzes other examples to show the implications of RPCE. Section 5 motivates the RPCE definition by exploring the consequences of alternative specifications. In Section 6 we analyze how the changes in the terminal node partition affect the outcomes of games with terminal node partitions. Section 7 explains the connection PCE and RPCE with other concepts from the literature, notably the rationalizable conjectural equilibrium (RCE) of Rubinstein and Wolinsky (1994). Unless stated otherwise, proofs can be found in Appendix B.

2 An Illustrative Example (Example 1)

This section provides an informal illustration of the effect of the players' knowledge about the observation structure. To do this, we compare three extensive-form games that are similar to each other, and identify an outcome that seems a plausible long-run consequence

of rational learning in one of the games but not the others. We formalize the intuition after we provide a formal model; here we give only an informal argument.

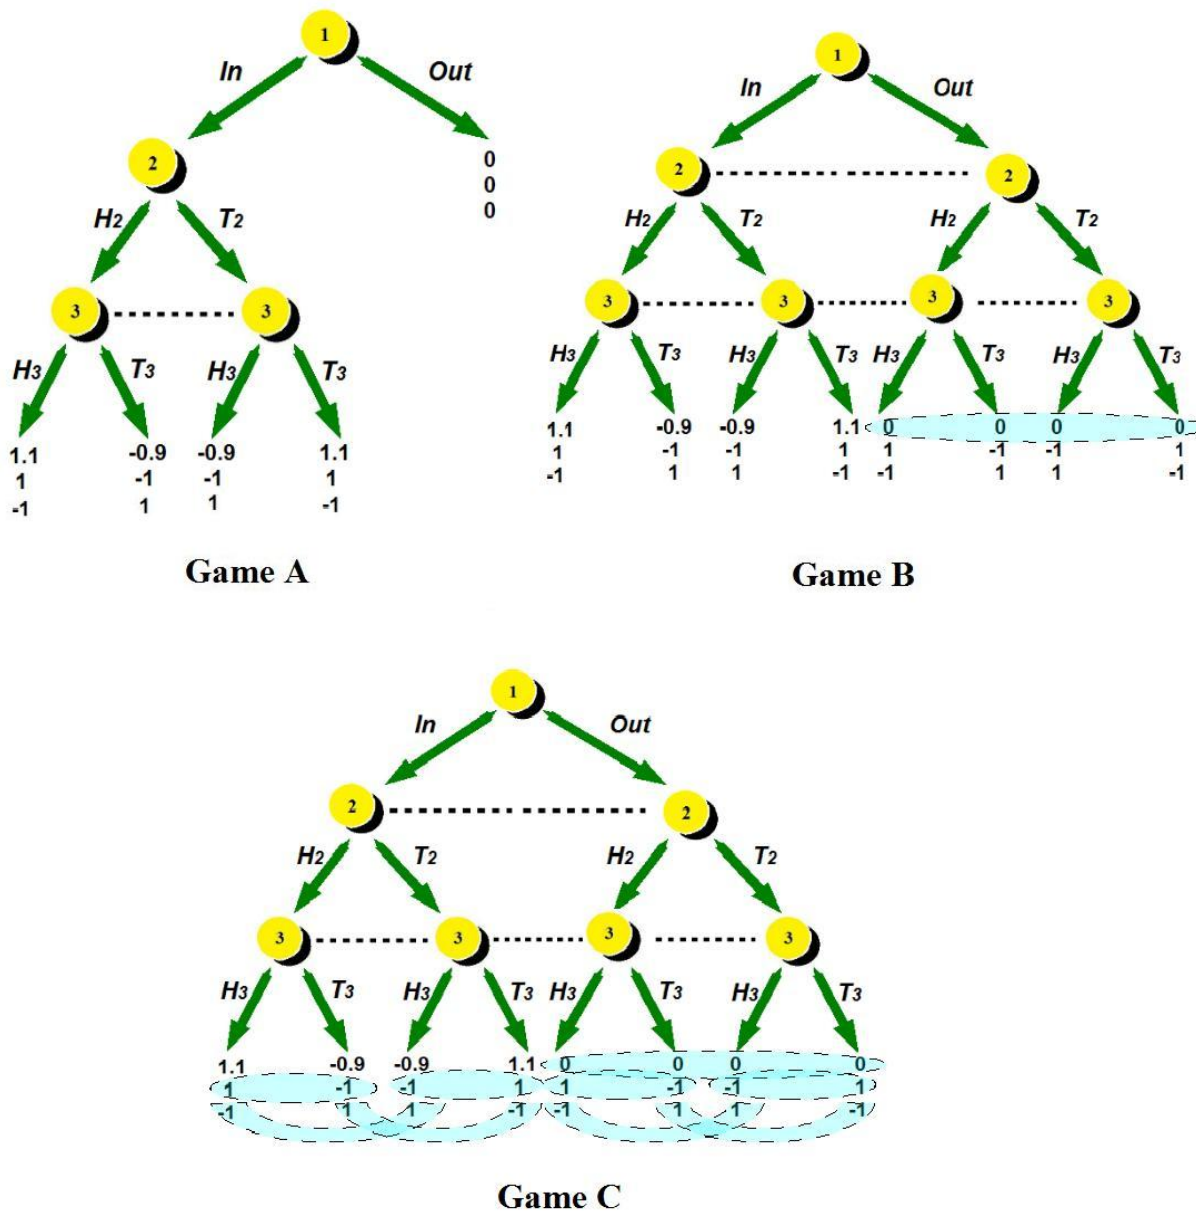


Figure 1

The three games, games A, B, and C, are depicted in Figure 0. In game A, player 1 moves first, choosing between *In* and *Out*. If he chooses *In*, players 2 and 3 play matching pennies with player *i* choosing between *H_i* and *T_i*, if players 2 and 3 match 2 gets 1 and 3 gets -1, if they do not match then player 3 gets 1 and player 2 gets -1. Player 1's

payoffs are the amount that player 2 gets plus an “extra” of 0.1, if player 1 plays *In*. Thus for example if player 2 and 3’s actions match then player 1 gets $1 + 0.1 = 1.1$. When player 1 plays *Out*, all players obtain the payoff of 0. At the end of each play of the game, players observe the exact terminal node that is reached, as in self-confirming equilibrium.

In game B, player 1 moves first, again choosing between *In* and *Out*. Instead of 2 and 3 only acting when 1 plays *In*, now they play the matching pennies game regardless of 1’s action. The map from action profiles to payoffs is exactly the same as in game A. In particular, if 1 plays *Out* she gets 0 regardless of the actions of players 2 and 3. The important assumption is that if 1 plays *Out* she observes only her own action and payoff but not the action of the other player: the corresponding cell of her terminal node partition contains four elements corresponding to the four possible choices of players 2 and 3. In Figure 1, this is denoted by connecting player 1’s payoffs at corresponding terminal nodes by dots. Players 2 and 3 observe the exact terminal nodes. Note that the observation structures for player 1 are the same in games A and B.

Finally, game C has the same extensive form and payoffs as in game B, while the terminal node partitions for players 2 and 3 are different: here we assume that they do not observe each other’s action.⁴ This unobservability is represented by connecting each player’s corresponding payoffs by dots.

Note that even though player 1 receives the same information in these games, the observation structures of players 2 and 3 differ. In game A, players 2 and 3 do not observe each other’s play when 1 plays *Out*, so there is no reason for player 1 to expect their play to resemble a Nash equilibrium. Consequently, an impatient player 1 might choose to play *Out*, fearing that player 2 would lose to player 3. In game B, on the other hand, players 2 and 3 observe each other’s play, whatever player 1’s action is. Thus they should be playing as in the Nash equilibrium of the matching pennies game, and 1 knows this, so she should play *In*. In game C, even though 2 and 3 play repeatedly, they do not observe each others’ play. Hence 1 can expect that 2 is consistently losing to player 3, and so she can play *Out*.

In Section 4 we formalize this intuition. Before doing so, we provide a formal model of extensive-form games with terminal node partitions in the next section.

⁴This means that their payoffs are not measurable with respect to the partitions. This specification is made only to make our point in a simple example. Example 5 illustrates much the same point with payoffs that are measurable with respect to the partitions.

3 The Model

3.1 Extensive-Form Games with Terminal Node Partitions, Strategies, and Beliefs

X is the finite set of nodes, with $Z \subseteq X$ being the set of terminal nodes. The distribution over Nature N 's moves is known to all players. The set of players is $I = \{1, \dots, n\}$. H_i is the collection of player i 's information sets. Set $H = \bigcup_{i \in I} H_i$ and $H_{-i} = H \setminus H_i$. Let $A(h)$ be the set of available actions at $h \in H$, $A_i = \bigcup_{h \in H_i} A(h)$, and $A = \times_{i \in I} A_i$ and $A_{-i} = \times_{j \neq i} A_j$. For each $z \in Z$, player i 's payoff is $u_i(z)$.

To model what players observe at the end of each round of play, let $\mathbf{P}_i = (P_i^1, \dots, P_i^{L_i})$ be a partition over Z . Except where otherwise noted, we will require that $u_i(z) = u_i(z')$ if terminal nodes z and z' are in the same partition cell, so that payoffs are measurable with respect to terminal node partitions.⁵

Player i 's behavioral strategy π_i is a map from H_i to probability distributions over actions, satisfying $\pi_i(h) \in \Delta(A(h))$ for each $h \in H_i$. The set of all behavioral strategies for i is Π_i , and the set of behavioral strategy profiles is $\Pi = \times_{i \in I} \Pi_i$. Let $\Pi_{-i} = \times_{j \neq i} \Pi_j$ and $\Pi_{-i,k} = \times_{j \neq i,k} \Pi_j$, with typical elements π_{-i} and $\pi_{-i,j}$, respectively. For $\pi \in \Pi$ and $\pi_i \in \Pi_i$, let $H(\pi)$ and $H(\pi_i)$ be the collection of information sets reached with positive probability given π and (π_i, π'_{-i}) , respectively, where π'_{-i} is any completely mixed behavioral strategy.

A strategy profile π completely determines a probability distribution over terminal nodes; let $d(\pi)(z)$ be the probability of reaching $z \in Z$ given π . We let $D_i(\pi)(P_i^l) = \sum_{z \in P_i^l} d(\pi)(z)$ for each cell P_i^l of player i 's partition; We assume that the extensive form has perfect recall in the usual sense, and extend perfect recall to terminal node partitions by requiring that two terminal nodes must be in different cells of \mathbf{P}_i if they correspond to different actions by player i . If every terminal node is in a different cell of \mathbf{P}_i , the partition \mathbf{P}_i is said to be discrete. If the cell i observes depends only on i 's actions, the partition is called trivial.⁶

As in DFL, we will require optimality at some off-path information sets, so we will need to specify assessments at off-path information sets: Player i 's assessment a_i is a map that assigns to each of i 's information sets a probability distribution over nodes in that information set. Let \mathcal{A}_i be the set of assessments for player i ; i 's assessment a_i and his opponents' behavioral strategies π_{-i} completely determine i 's expected payoff for playing any strategy π_i , conditional on any $h \in H_i$. Denote by $\mu^i \in \Delta(\mathcal{A}_i \times \Pi_{-i})$ the **belief** held

⁵In particular, Theorem 5 will assume this condition.

⁶Even if the terminal node partition is trivial, i may be able to distinguish among some terminal nodes, because we require that i 's terminal node partition distinguishes between i 's own actions.

by player i . That is, a belief is a probability distribution over pairs of an assessment and a profile of the opponents' strategies. Notice that we allow the possibility that player 1's belief about players 2 and 3's plays be (subjectively) correlated. For example, it is possible that 1 is sure that 2 and 3 play the same action, while she is not sure whether it is L or R . This type of belief is consistent with the assumption that 1 thinks 2 and 3 play independently. We say that version v_i^k has an **independent belief** if the projection of $\mu^{i,k}$ on Π_{-i} corresponds to a product measure. In this case we associate $(a_i^k, \mu^{i,k})$ with $(a_i^k, \pi_{-i}^{i,k}) \in \mathcal{A}_i \times \Pi_{-i}$.⁷ Note that the space that beliefs lie in is not $\mathcal{A}_i \times \Delta(\Pi_{-i})$, but $\Delta(\mathcal{A}_i \times \Pi_{-i})$. This is because an assessment may be correlated with the expectation of opponents' strategies, so it may not be sufficient to postulate only a single assessment. Example 9 shows why this extra generality is desirable.

We say that $\pi_i \in \Pi_i$ is a **best response to** $\mu^i \in \Delta(\mathcal{A}_i \times \Pi_{-i})$ at $\bar{H}_i \subseteq H_i$ if the restriction of π_i to the subtree starting at $h \in \bar{H}_i$ is optimal against μ^i in that subtree. The formal definition of optimality here is made complicated by the fact that we allow player i 's assessment over off-path nodes to be correlated with his belief about the opponents' strategies. The expected payoff conditional on reaching a given information set depends on the relative weights over different assessments, and we need to specify how these weights are determined. Appendix A gives a precise definition, and Example 15 in the Online Supplementary Appendix illustrates the need for this complication.

3.2 Versions, Conjectures, and Belief Models

As in DFL, we define a *version* of player i , denoted v_i . In DFL v_i specifies player i 's strategy, her assessment, and her belief about the opponents' play. The definition of a version in our context will be different from that of DFL, as we include the assessment as part of beliefs, and also associate with each version a probability distribution over opponents' versions that we call a "conjecture." This is a convenient way to formalize the condition that a belief is generated by some probability distribution over opponents' versions. To introduce the notion of conjectures formally, we first need to specify a profile of sets of versions.

DFL used the notion of *belief model* to define solution concepts. A belief model in DFL was $V = (V_1, \dots, V_n)$ where each V_i is a set of player i 's versions. We also identify a belief model. Formally, a belief model is expressed by $V = (V_1, \dots, V_n)$ where each V_i is

⁷Recall that independent randomizations over behavior strategies is equivalent to a behavior strategy from Kuhn's theorem- see e.g. Fudenberg and Kreps (1995). DFL consider only independent beliefs; as we explain below that restriction is less appealing in the context we consider here. Fudenberg and Levine (1993) discuss the interpretation of correlated beliefs and its implications for self-confirming equilibrium.

the set of player i 's versions.⁸

Version k of player i is denoted by $v_i^k = (\pi_i^k, \mu^{i,k}, p^{i,k})$, where the first element is version k 's strategy $\pi_i^k \in \Pi_i$, the second is her belief $\mu^{i,k} \in \Delta(\mathcal{A}_i \times \Pi_{-i})$, and the third is her conjecture $p^{i,k} \in \Delta(\times_{j \neq i} V_j)$. Notice that the specification of conjectures takes into account the idea of correlated beliefs, as otherwise, $p^{i,k}$ must lie in the space $\times_{j \neq i} \Delta(V_j)$. We allow $p^{i,k}$ to take any number in $[0, 1]$: Even if player i is sure that there is only a single agent in player j 's player role, she may not be sure whether this single agent is of version v_j^1 or v_j^2 . We let $K_i < \infty$ be the number of elements in V_i ; that is, $V_i = \{v_i^1, \dots, v_i^{K_i}\}$.

3.3 Partition-Confirmed Equilibrium and Rationalizable Partition-Confirmed Equilibrium

Before stating our solution concept, we develop several regularity conditions.

Let $W_j(\pi_j)$ be the set of j 's versions v_j^l such that $\pi_j^l = \pi_j$, and let $W(\pi_{-i}) = \times_{j \neq i} W_j(\pi_j)$.

Definition 1. Given a belief model V , version v_i^k has a **coherent belief** if $\mu^{i,k}(\cdot, \pi_{-i}) = \sum_{v_{-i} \in W(\pi_{-i})} p^{i,k}(v_{-i})$.

This requires that the version's belief matches his own conjecture about the opponents' versions.

The requirement that all versions in a belief model have a coherent belief is analogous to requiring the belief model be *belief-closed*, as defined in DFL. Specifically, in Appendix B we prove the following:

Claim 1. $\pi \in \Pi$ is generated by a belief model in which all versions have coherent beliefs and satisfy the best response condition and all beliefs correspond to product measures if and only if it is generated by a DFL-style belief-closed model in which the best response condition is satisfied for all versions.⁹

We require that, for each version v_i^k of each player role i , any pair of an assessment and opponents' strategies in the support of $\mu^{i,k}$ satisfies *KW-consistency*. We say that $\mu^{i,k}$ is KW-consistent if any $(a_i, \pi_{-i}) \in \mathcal{A}_i \times \Pi_{-i}$ in the support of $\mu^{i,k}$ is consistent in the sense of Kreps and Wilson (1982): there exists a sequence of pairs of an assessment and

⁸The reason that each V_i may contain multiple versions is that in our solution concept we want to allow for the possibility that some version's conjecture is incorrect. For example, a version v_i of player i may conjecture that v_j of player j exists with a positive probability, while v_j may not actually exist.

⁹Recall that DFL restrict attention to independent beliefs, which is why we need to do the same for this equivalence.

a completely mixed behavioral strategy $\{a_i^m, \pi^m\}_{m=1}^\infty$ with the property that each a_i^m is generated by π^m and Bayes rule, and that $a_i^m \rightarrow a_i$ and $\pi_{-i}^m \rightarrow \pi_{-i}$ as $m \rightarrow \infty$.

Definition 2. Given a belief model V , version $v_i^k \in V_i$ is **self-confirming with respect to π^*** if for any π_{-i} in the support of $\mu^{i,k}$, $D_i(\pi_i^k, \pi_{-i}) = D_i(\pi_i^k, \pi_{-i}^*)$.

Let us explain the equality in this definition. The left hand side is the distribution over i 's terminal partition generated by version k 's strategy and her belief about the opponents' play. The right hand side is the distribution that she observes if the actual distribution of the play is π^* . That is, this equality says that v_i^k 's observation (the left hand side) is equal to the actual play (the right hand side).

The next definition incorporates the idea that players know the terminal node partitions of other players and the self-confirming condition. First, let us define the link between strategies and versions.

Definition 3. Given a belief model V , $\pi \in \Pi$ is **generated by** a version profile $v \in \times_{j \in I} V_j$ if for each i , π_i is a strategy prescribed by v_i .

For notational simplicity, let $\pi_j(v_j)$, $\pi(v)$ and $\pi_{-i}(v_{-i})$ denote the strategy (profile) generated by $v_j \in V_j$, $v \in \times_{j \in I} V_j$ and $v_{-i} \in \times_{j \neq i} V_j$, respectively.

Definition 4. Given a belief model V , v_i^k is **observationally consistent** if $p^{i,k}(v_{-i}) > 0$ implies, for each $j \neq i$, v_j is self-confirming with respect to $\pi(v_i^k, v_{-i})$.

If v_j is self-confirming with respect to $\pi(v_i^k, v_{-i})$ then by Definition 2, for all π_{-j} in the support of v_j 's belief, $D_j(\pi_j(v_j), \pi_{-j}) = D_j(\pi_j(v_j), \pi_{-j}(v_i^k, v_{-i,j})) = D_j(\pi_i^k, \pi_{-i}(v_{-i}))$. Hence the definition is equivalent to the following:

Definition 4'. Given a belief model V , v_i^k is **observationally consistent** if $p^{i,k}(v_{-i}) > 0$ implies, for each $j \neq i$, $D_j(\pi_j(v_j), \pi_{-j}) = D_j(\pi_i^k, \pi_{-i}(v_{-i}))$ for all π_{-j} in the support of v_j 's belief.

$p^{i,k}(v_{-i}) > 0$ means that v_i^k assigns positive probability to the event that the profile of opponents' versions is v_{-i} (which induces the opponents' strategy profile $\pi_{-i}(v_{-i})$). When v_i^k thinks such a profile is possible, observational consistency requires the equality in the definition to hold. The left hand side in this equality is what v_j expects to observe given his belief under the partition given by D_j . The right hand side describes what v_i^k thinks v_j is observing under the partition given by D_j . Thus the equality claims that v_i^k believes that v_j 's belief is consistent with what v_j observes. In other words, the condition

states that v_i^k thinks v_j 's belief matches the reality- that is, v_i^k believes v_j satisfies the self-confirming condition with respect to the strategy profile generated by what v_i^k thinks the actual version is.

To understand this condition better, consider the following example: Suppose that v_1^1 believes that (v_2^1, v_3^1) and (v_2^2, v_3^2) are possible and that no other profiles are possible. Then we require that v_1^1 thinks what v_2^1 would be observing is consistent with his play, v_1^1 's play, and also v_3^1 's play. It is important to note that we do not require v_1^1 thinks v_2^1 's belief is consistent with v_3^2 's play. This is because, even though v_1^1 thinks each of v_2^1 and v_3^2 is possible, she thinks (v_2^1, v_3^2) is impossible.

Note that the equality in Definition 4' only need hold when v_i^k thinks the profile v_{-i} exists: Otherwise, v_i^k need not believe that v_j 's observation is consistent with her belief. Relatedly, even if v_i^k thinks v_j exists and v_j thinks version v_m^n exists, v_i^k 's belief need not be consistent with what v_m^n observes. This is because v_i^k might think that "I'm sure that v_j exists, but this v_j incorrectly conjectures that version v_m^n exists." Note also that the equality is imposed for all π_{-j} in the support of v_j 's belief. Otherwise, v_i^k could be conjecturing that v_j who does not satisfy the self-confirming condition exists. Finally, if v_j is self-confirming with respect to π^* , then in the left hand side of the equation in Definition 4', $D_j(\pi_j(v_j), \pi_{-j})$ can be replaced with $D_j(\pi_j(v_j), \pi_{-j}^*)$.

The rationale for this condition is that player i knows j 's terminal node partition and knows that j 's belief is consistent with what j observes, but in the model developed so far this knowledge is informal. In Subsection 3.4 we make this interpretation of the observational consistency precise, by constructing an epistemic model with the state space being the set of all possible version-partition configurations. We will show that our observational consistency condition corresponds to the situation where the version in question knows the opponents' terminal node partitions and that they satisfy the self-confirming condition.

Now we are ready to define our solution concepts. First, as a benchmark definition, we define an analog of SCE that does not incorporate the idea of rationalizability.

Definition 5. π^* is a **partition-confirmed equilibrium (PCE)** if there exist a belief model V and an actual version profile v^* such that the following three conditions hold:

1. π^* is generated by v^* ;
2. For all i and k , π_i^k is a best response to $\mu^{i,k}$ at $H(\pi_i^k, \pi_{-i})$ for all $(a_i, \pi_{-i}) \in \text{supp}(\mu^{i,k})$.
3. v_i^* is self-confirming with respect to π^* ;

Condition (1) says that the equilibrium strategy profile is generated by the specified belief model. We call the version that generates the strategy profile the “actual version.”

Condition (2) ensures that players optimize against their beliefs, but the condition merely requires that they optimize only at the “on-path” information sets. This is one of the conditions that we strengthen in our main solution concept.

Condition (3) is the “self-confirming” part of the equilibrium concept. Notice that this condition is imposed only for the versions that actually exist. However, imposing the self-confirming condition for all versions does not restrict the set of equilibria (See Theorem 7). However, in the main concept that we define shortly, imposing the self-confirming condition for non-actual versions does rule out some seemingly sensible outcomes. We will discuss this point later in more detail.

Our main interest is in the concept of RPCE.

Definition 6. π^* is a **rationalizable partition-confirmed equilibrium (RPCE)** if there exist a belief model V and an actual version profile v^* such that the following five conditions hold:

1. π^* is generated by v^* ;
- 2'. For all i and k , π_i^k is a best response to $\mu^{i,k}$ at $H(\pi_i^k)$;
3. v_i^* is self-confirming with respect to π^* ;
4. For all i and k , v_i^k has a coherent belief;
5. For each i and k , v_i^k is observationally consistent.

Note that RPCE strengthens PCE by replacing condition (2) with condition (2') and requiring that versions have coherent beliefs and are observationally consistent.

Condition (2') says that each version plays a best response to his belief at the information sets that he himself does not preclude. This is the condition called *rationalizability at reachable nodes* in DFL. As discussed there, rationalizability at all information sets is not robust to small uncertainty about players' rationality, hence we require it to be true only at those information sets that are not precluded by the player's own strategy.¹⁰

As mentioned earlier, condition (3), the self-confirming condition, is required only for versions that are objectively present, i.e. the actual versions. In Example 7 we will

¹⁰DFL and Greenberg, Gupta and Luo (2003) define an analog of RSCE that requires all versions optimize at all information sets; this version is not robust to small uncertainties about the payoffs. See Fudenberg, Kreps, and Levine (1988) for an analysis of the robustness of equilibrium refinements to payoff uncertainty, and Dekel and Fudenberg (1990) and Börgers (1994) for the robustness of iterated weak dominance.

explain the reason for this restriction. There was no such restriction in RSCE. Roughly, this is because in DFL’s model, all players’ terminal node partitions are discrete, so the probability distribution over terminal nodes is common knowledge among players, which precludes the need for there being two versions with different observations in a set of versions. We clarify this point in Theorem 6, where we show that requiring the condition for all versions do not restrict the set of equilibria when the terminal nodes partitions are discrete. Also, in the (nonrationalizable) PCE concept, this distinction does not matter because the coherent-belief condition is not imposed. This point will be clarified in Theorem 7.

The “replacement” discussed in the third remark after the definition of observational consistency (replacing “ $D_j(\pi_j(v_j), \pi_{-j})$ ” with “ $D_j(\pi_j(v_j), \pi_{-j}^*)$ ”) works only for actual versions; for other versions, this replacement may not work.¹¹

We did not require observational consistency when defining PCE. This is because observational consistency is hard to interpret when there is no connection between beliefs and conjectures, and such a connection is described by the coherent belief condition, which we do not require in the PCE concept.

Finally, notice that if players do not get to observe any consequence of opponents’ actions, conditions (3) and (5) have no bite, hence RPCE reduces to (extensive-form) rationalizability.

3.4 Observational Consistency and Knowledge about Terminal Node Partitions and the Self-Confirming Condition

In this section we make our interpretation of observational consistency precise, using an epistemic model.

Let $\Omega = \mathcal{V} \times \mathcal{P}$, where \mathcal{V} is the set of all version profiles, and \mathcal{P} is the set of all terminal node partitions. The information set $h_i(\omega)$ is the set of states that i thinks possible. That is,

$$h_i((v, P)) = \{(v_i, v'_{-i}) \in \mathcal{V} \mid v_i \text{'s conjecture assigns a positive probability to } v_{-i}\} \times \bar{\mathcal{P}}_i(P),$$

where $\bar{\mathcal{P}}_i : \mathcal{P} \rightarrow 2^{\mathcal{P}} \setminus \{\emptyset\}$ is some given function.

Notice that it may be the case that $\omega \notin h_i(\omega)$, as agents’ conjectures can be incorrect.

To simplify the notation, let $\pi(v)$ be the strategy profile played by version profile v .

¹¹See for example the belief model provided in Example 7.

Define the following:

$$E(P) = \mathcal{V} \times \{P\};$$

$$E_i^{SC} = \{(v, P) \in \Omega \mid v_i \text{ is self-confirming with respect to } \pi(v) \text{ under } P_i\};$$

$$E_{-i}^{SC} = \bigcap_{j \neq i} E_j^{SC}, \quad E^{SC} = \bigcap_{j \in I} E_j^{SC};$$

$$E_i^{OC} = \{(v, P) \in \Omega \mid v_i \text{ is observationally consistent under } P_{-i}\};$$

To interpret, $E(P)$ is the set of all states with partition P ; E_i^{SC} and E_i^{OC} are the sets of all states ω such that the self-confirming condition and observational consistency condition, respectively, are satisfied if ω is the reality.

Note that whether v_i is self-confirming or not depends on P_i while whether it is observationally consistent or not depends on P_{-i} .

We define knowledge, mutual knowledge, and common knowledge as in the usual model. That is, for any given $E \subseteq \Omega$, define $K_i(E) = \{\omega \mid h_i(\omega) \subseteq E\}$ and $K(E) = \bigcap_{i \in I} K_i(E)$. Let $K^n(E) = K(K^{n-1}(E))$ with $K^0(E) = E$, and let $K^\infty(E) = \bigcap_{n=0}^\infty K^n(E)$.

Now we show that if i knows the terminal node partitions and she believes that the self-confirming condition is satisfied for all players, then she must be observationally consistent.

Theorem 1. *Suppose that the definition of information sets are given by h_i above with some \bar{P}_i for each player i . Then, for each i ,*

$$\left(\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)] \right) \cap K_i(E_{-i}^{SC}) \subseteq E_i^{OC}.$$

Proof. Fix $i \in I$ and suppose $(\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)]) \cap K_i(E_{-i}^{SC}) \neq \emptyset$, because if it were empty then we would be done. Take an arbitrary $\omega = (v, P)$ in this set.

First, $\omega \in (\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)])$ implies that $P' = P$ for all $(v', P') \in h_i(\omega)$.

Second, $\omega \in K_i(E_{-i}^{SC})$ implies that, for any $j \neq i$, v'_j is self-confirming with respect to $\pi(v')$ under P'_j for all $(v', P') \in h_i(\omega)$.

These two mean that v'_j is self-confirming with respect to $\pi(v')$ under P_j for all $(v', P') \in h_i(\omega)$.

By the definition of h_i , this means that v'_j is self-confirming with respect to $\pi(v_i, v'_{-i})$ under P_j for all v'_{-i} to which v_i 's conjecture assigns a positive probability. Hence, by Definition 4, v_i is observationally consistent under P_{-i} . Hence $\omega \in E_i^{OC}$. \square

Note that the opposite direction of set inclusion is generally false, as the right hand side is always nonempty, while $(\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)])$ may be empty depending on

how we specify \bar{P}_i in the definition of h_i .

We say that players have correct beliefs about the partitions if

$$h_i((v, P)) = \{(v_i, v'_{-i}) \in \mathcal{V} | v_i \text{'s conjecture assigns a positive probability to } v_{-i}\} \times \{P\}.$$

Theorem 2. *Suppose that players have correct beliefs about the partitions. Then,*

$$\left(\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)] \right) \cap K_i(E_{-i}^{SC}) = E_i^{OC}.$$

Proof. Suppose that players have correct beliefs about the partitions. We have already proven one direction. So we prove the opposite direction.

Notice that when players have correct beliefs about the partitions, $(\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)]) = \Omega$. Hence all we need to prove is $K_i(E_{-i}^{SC}) \supseteq E_i^{OC}$.

Suppose that $\omega = (v, P) \in E_i^{OC}$. Then, by Definition 4, v'_j is self-confirming with respect to $\pi(v_i, v'_{-i})$ under P_j for all v'_{-i} to which v_i assigns a positive probability. By the definition of h_i and the assumption that players have correct beliefs about the partitions, this means that v'_j is self-confirming with respect to $\pi(v_i, v'_{-i})$ under P'_j for all $(v', P') \in h_i^*((v, P))$. Since whenever $(v', P') \in h_i(v, P)$ we must have $v'_i = v_i$ by definition, we have that v'_j is self-confirming with respect to $\pi(v')$ under P'_j for all $(v', P') \in h_i((v, P))$. Hence $h_i((v, P)) \subseteq \bigcap_{j \neq i} E_j^{SC} = E_{-i}^{SC}$. By definition $K_i(E_{-i}^{SC}) = \{w' | h_i(w') \subseteq E_{-i}^{SC}\}$, so we have that $\omega \in K_i(E_{-i}^{SC})$. \square

Now we consider higher order knowledge. The next theorem states that RPCE implies common knowledge of the partition structure and the self-confirming condition.

Theorem 3. *Given partitions P^* , if π^* is an RPCE then there exists v such that $\pi(v) = \pi^*$ and $(v, P^*) \in (\bigcup_{P \in \mathcal{P}} K^\infty(E(P))) \cap K^\infty(E^{SC}) = K^\infty(E^{OC})$ where players have correct beliefs about the partitions.*

Proof. Fix P^* . Take a belief model that rationalizes π^* . Let the actual version profile be v . Let $\omega = (v, P^*)$. Suppose that players have correct beliefs about the partitions.

First we show $\omega \in E^{OC}$. To see this, suppose that for some $n = 0, 1, 2, \dots$, $\omega \in K^n(E^{SC})$ but $\omega \notin K^{n+1}(E^{SC})$. Then there must exist $\omega' = (v', P^*) \in E^{SC} \cap \left[\bigcup_{\omega^0 = \omega, \{\omega^k\}_{k=0}^n, \omega^{k+1} \in h_i(\omega^k)} \omega^n \right]$ such that $h_i(\omega') \not\subseteq E^{SC}$. But by the definition of h_i , v' must be an element of the belief model that rationalizes π^* . Hence if $\omega'' = (v'', P^*) \in h_i^*(\omega')$ then v'' must also be an element of the belief model that rationalizes π^* , so in particular v''_i must be observationally consistent under P^*_{-i} . This means $\omega'' \in E^{OC}$, so $h_i(\omega') \subseteq E^{SC}$. Contradiction.

Now, we show that $(\bigcup_{P \in \mathcal{P}} K^\infty(E(P))) \cap K^\infty(E^{SC}) = K^\infty(E^{OC})$, which completes the proof. Notice that $(\bigcup_{P \in \mathcal{P}} [K^n(E(P))]) = \Omega$ for any $n = 0, 1, 2, \dots$. Hence

$$\begin{aligned} K^\infty(E^{OC}) &= K^\infty \left(\left(\bigcup_{P \in \mathcal{P}} [K_i(E(P)) \cap E(P)] \right) \cap K_i(E_{-i}^{SC}) \right) \\ &= K^\infty(E^{SC}) = \left(\bigcup_{P \in \mathcal{P}} [K^\infty(E(P))] \right) \cap K^\infty(E^{SC}). \end{aligned}$$

Thus the proof is complete. □

4 Implications of RPCE

In this section we revisit Example 1 and consider several examples to illustrate the implications of RPCE. One of the themes will be the difference between situations where player 1 (say) prevents other players from acting (and thus from learning) and situations where the other players do act but player 1 does not observe their play. First we revisit Example 1 to show how the RPCE definition delivers the desired conclusion there. Example 2 adds a player to game B to study the assumption of higher order knowledge of rationality. In Example 3, RPCE implies that belief about unobservable play should correspond not only to rationalizable actions, but also a Nash equilibrium. We then generalize this result to a class of “participation games.” Example 4 points out the implication of the general result presented in Example 3, and argues that the RPCE concept may be too restrictive in some cases. Example 5 demonstrates that a player need not expect the unobservable play by the opponents to resemble a Nash equilibrium if their terminal node partitions are not discrete. Finally, Example 6 demonstrates that a player can learn from an opponent’s play the information about a third player’s play that she does not directly observe.

Example 1 Revisited.

Here we show that in games A and C it is possible for player 1 to play *Out* in RPCE, but this is not possible in game B.

Consider game A. We argue that player 1 can play *Out* in an RPCE with the following belief model and actual versions¹²:

$$V_1 = \{v_1^1\}, \quad v_1^1 = (Out, (H_2, T_3), p^{1,1}(v_2^1, v_3^1) = 1);$$

¹²The notation that we use when presenting belief models in examples involves slight abuse of notation. In particular, we do not specify the assessments as they are clear in the examples.

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (H_2, (Out, H_3), p^{2,1}(v_1^1, v_3^2) = 1), v_2^2 = (T_2, (Out, T_3), p^{2,2}(v_1^1, v_3^1) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (T_3, (Out, H_2), p^{3,1}(v_1^1, v_2^1) = 1), v_3^2 = (H_3, (Out, T_2), p^{3,2}(v_1^1, v_2^2) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold. For example, v_2^1 can believe that player 3 plays H_3 because she never gets to observe 3's play, while v_3^1 plays H_3 because he believes that 2 plays T_2 , which again is justified by the fact that he is not observing 2's play. Since v_1^1 never observes 2 and 3's play, and she knows that they do not get to play on the path so do not observe each other's play, she can believe that they can have such mutually inconsistent beliefs, hence can entertain a belief that the opponents play (H_2, T_3) , which is consistent with the self-confirming condition. In game C, exactly the same belief model (and the same argument) can be used to justify player 1 choosing *Out*.

Now we turn to game B. Fix an RPCE π^* , with an associated belief model V . Suppose that some version of player 1's belief assigns a positive probability to a version profile (v_2^k, v_3^l) such that π_2^k and π_3^l are not best responses to each other. Suppose without loss of generality that π_2^k is not a best response to π_3^l . Notice that by the observational consistency condition, we have $D_2(\pi_2^k, \pi_{-2}) = D_2(\pi_2^k, \cdot, \pi_3^l)$ for all π_{-2} in the support of $\mu^{2,k}$. Since player 2 observes the exact terminal node reached, this implies that $\mu^{2,k}$ assigns probability 1 to π_3^l . But this means that the best response condition is violated for player 2.

Therefore, it must be the case that $\mu^{1,k}$ assigns the weight exactly equal to $\frac{1}{2}$ to each of H_2 and H_3 . The best-response condition then implies that $\pi_1^k(h_1)(In) = 1$, as playing *In* gives her the expected payoff of 0.1 while playing *Out* gives her 0. Because this is true for any version v_1^k of player 1 and π^* is generated by the actual versions, we conclude that $\pi_1^*(h_1)(In) = 1$, that is, player 1 plays *In* with probability 1. \square

Example 2.




Figure 2

Now consider a modification of game B, depicted in Figure 2, where we add “player 0” at the top of the extensive-form game. Specifically, player 0 moves first, choosing between *In* and *Out*. whatever action is played, the game goes on and game B is played, where only player 1 knows the action taken by player 0. The map from the action profile for players 1, 2, and 3 to their payoffs are exactly the same as in game B, while player 0 gets 0 if he plays *Out*, 1 if he plays *In* and player 1 also plays *In*, and -1 if he plays *In* and player 1 plays *Out*. The terminal node partitions are the same as in game B, where everyone knows the move by player 0, and player 0 observes everything if he plays *In* and does not observe anything if he plays *Out*.

In any RPCE of this game, player 0 must play *In*, because player 0 must infer that player 1 plays *In*. Remember that in Example 1 all versions of player 1 must play *In*; the coherent belief condition ensures that player 0 believes that 1 plays *In* with probability 1.¹³

This example shows that RPCE assumes that a player not only knows that the play by the opponents has converged, but she also knows that an opponent knows that the play by these opponents has converged. \square

¹³Player 0 can play *Out* if players 2 and 3 observe whether 0 played *In* or *Out* before they move. In that game, if 0 usually plays *Out*, then when 0 unexpectedly plays *In* 1 has no reason to think that 2 and 3 have converged to equilibrium play, so 1 can play *Out* in the belief that 2 would lose, so 0 can play *Out*.

Example 3.




Figure 3

Consider the game in Figure 3. Three players move simultaneously. Player 1 chooses *In* or *Out*. Players 2 and 3 choose H_i or T_i or N_i . The terminal node partitions are such that everyone observes the exact terminal node reached, except that player 1 cannot distinguish between the opponents' action profiles if she plays *Out*.

Notice that H_2 is a best response to H_3 , which is a best response to T_2 , which is a best response to T_3 , which in turn is a best response to H_2 . The game has a unique Nash equilibrium, namely (N_2, N_3) . Hence all actions are rationalizable, while only N_i is in the support of Nash equilibrium strategy.

In this game, RPCE requires not only that 1 expects 2 and 3 to play rationalizable actions, but also that she expects their play to correspond to a Nash equilibrium. Hence 1 should expect the payoff of 1 from playing *In*, so 1 should play *In*. The proof of this is exactly the same as in Example 1: if player 1's belief assigns a positive probability to a version profile such that player 2 is not best responding to player 3, observational consistency condition for player 1 implies that the best response condition for player 2 should be violated.

This example shows that an action being rationalizable does not necessarily imply that there exists an RPCE in which some player expects it to be played with a positive probability. Again, it is important in this example that 2 and 3 do not know 1's action when they move; otherwise 1 can play *Out*, believing that 2 and 3 play H_i or T_i , which are rationalizable if she plays *Out*.

Both in game B and this example, RPCE requires that player 1 expects a Nash play by players 2 and 3. This point can be generalized in the next part. \square

A Theorem for Participation Games

A **player 1 participation game** Γ is a game with a set of players I , action profiles A , and payoff function u , in which all players move simultaneously (so $Z = A$), where player 1 has an option to play an action “Out” which gives her a constant payoff of 0. The terminal node partition is such that everyone observes the exact terminal node reached, except that 1 does not observe anything if she plays Out. Γ_{-1} is a game with a set of players $I \setminus \{1\}$, action profiles A_{-1} , and payoff function v_j ’s defined by $v_j(a_{-1}) = u_j(Out, a_{-1})$.

Theorem 4. Fix a player 1 participation game Γ . If player 1 plays Out with probability 1 in an RPCE, then there is a convex combination of Nash equilibria of game Γ_{-1} such that no action of player 1 has a strictly positive payoff.

Corollary 1. Fix a player 1 participation game Γ such that there is a unique Nash equilibrium in Γ_{-1} . If player 1 plays Out with probability 1 in an RPCE, then no action of player 1 gives her a positive payoff against this unique Nash equilibrium.

To sum up, in player 1 participation game, 1 must expect not only that the opponents’ actions are rationalizable, but also that their play resembles a Nash equilibrium of their game. This restriction rules out the possibility of 1’s playing Out in some cases, for instance in game B and in the game in Example 3. \square

Example 4 (Shapley Example).





Figure 4

The game in Figure 4 replaces the matching pennies game in game B with the “Shapley game” (Shapley, 1964) . Specifically, player 1 moves first, choosing between In and Out . Whatever 1 chooses, 2 and 3 get their moves, choosing between A_i , B_i , and C_i . The terminal node partitions are such that everyone observes the exact terminal node reached, except that player 1 cannot distinguish between the opponents’ action profiles if she plays Out .

As shown in Theorem 4 (and Corollary 1), in RPCE player 1 expects players 2 and 3’s play resembles that of the Nash equilibrium of their game, hence plays In . Note that in the Shapley game fictitious play (Brown, 1951) does not converge.¹⁴ RPCE requires that 1 expects 2 and 3 to play as in Nash equilibrium even in such a case.

If the Shapley game is played by two players then the only RSCE corresponds to the unique Nash equilibrium of the game. RPCE further requires that a player involved in the game expects the unobserved play to resemble to a Nash equilibrium, which may be too strong an assumption as a description of long-run outcome of learning processes. In the Online Supplementary Appendix, we define a less restrictive solution concept that allows player 1 to play Out ; as we will see, though, the less restrictive concept may in some settings be too inclusive. \square

Example 5 (Participation Game with Unobservable Actions).



¹⁴Brown (1951) introduced fictitious play as a way to compute Nash equilibria. Fudenberg and Kreps (1993) give fictitious play a descriptive interpretation in strategic form games, and point out some problems with that interpretation when the process cycles as instead of converging to constant play of a fixed pure action profile.

Figure 5

In Figure 5, players $i = 1, 2, 3$ move simultaneously, choosing between In_i and Out_i . The terminal node partitions are such that player 1 does not observe anything if she plays Out_1 , while she observes the exact terminal node reached if she plays In_1 . Each of players 2 and 3 does not observe anything if he plays Out_i , while he observes the exact terminal node reached if he plays In_i .

Notice that for any Nash equilibrium of 2 and 3's simultaneous move game, player 1 expects a payoff of at least $\frac{1}{2}$ from playing In_1 . Thus if 1 believes that 2 and 3's play corresponds to a Nash profile, she must play In_1 . We argue, however, that in RPCE it is possible for player 1 to play Out_1 . Specifically, consider the following belief model and actual versions:

$$V_1 = \{v_1^1\}, \quad v_1^1 = (Out_1, (Out_2, Out_3), p^{1,1}(v_2^1, v_3^1) = 1);$$

$$V_2 = \{v_2^1, v_2^2\},$$

$$v_2^1 = (Out_2, (Out_1, In_3), p^{2,1}(v_1^1, v_3^2) = 1), v_2^2 = (In_2, (Out_1, Out_2), p^{2,2}(v_1^1, v_3^1) = 1);$$

$$V_3 = \{v_3^1, v_3^2\},$$

$$v_3^1 = (Out_3, (Out_1, In_2), p^{3,1}(v_1^1, v_2^2) = 1), v_3^2 = (In_3, (Out_1, Out_2), p^{3,2}(v_1^1, v_2^1) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold. In this belief model, player 1 believes that both players 2 and 3 play Out_i . Although Out_2 is not a best response against Out_3 , player 2 does not observe 3's play when he is playing Out_2 , and so he can believe that 3 plays In_3 . Likewise, player 3 can play Out_3 , believing that 2 plays In_2 . Player 1 plays Out_1 because she believes that (Out_2, Out_3) is played as a result of such mutually inconsistent beliefs. The point is that player 1 knows the terminal node partitions of players 2 and 3, and this knowledge leads her to believe in the non-Nash play by the opponents.

We note that Out_1 could not be played in any RPCE if the terminal node partitions for players 2 and 3 were discrete. This is because in any Nash equilibrium of the game by players 2 and 3, player 1 should expect the payoff of $\frac{1}{2}$, so by Theorem 4 she should play In . Hence, nondiscrete terminal node partitions allow an action to be played even if the action is outside the support of equilibria under finer partitions.¹⁵ In other words,

¹⁵The key here is that 2 and 3 act not knowing 1's action: Otherwise their play is only required to be rationalizable in the subgame that follows 1's off-path action. Thus even if they have discrete partitions,

the conclusion of Theorem 4 may fail if the hypothesis that player 1's opponents have discrete partitions is weakened.

To sum up, this example shows that a player need not expect unobserved play by the opponents to resemble a Nash equilibrium if these opponents do not observe the exact terminal nodes, and as a consequence she may play an action that she would not play if the opponents could observe the exact terminal nodes.¹⁶ \square

Example 6 (Learning from an Opponent's Play).




Figure 6

In the game in Figure 6, there are three players, $i = 1, 2, 3$, each choosing between L_i and R_i . Player 1 moves first, and player 2 moves only when 1 chooses L_1 . Whatever has happened, 3 then moves, terminating the game. All players observe the exact terminal node reached, except that player 2 does not observe the consequence of 3's action if 1 plays R_1 .

We show that, if player 1's partition is discrete, she cannot play R_1 . To see this, suppose, to the contrary, that she plays R_1 in an RPCE. The self-confirming condition and the best response condition for player 1 imply that 3 is playing R_3 with probability strictly greater than $\frac{1}{2}$. But then, the best response for player 2 is to play R_2 with

¹⁶1 can still play Out_1 , expecting both the opponents to play Out_i .

¹⁶In Example 8 we will discuss a game in which if a player's own terminal node partition becomes finer (with other players' partitions held fixed), she no longer plays an action that she plays when her partition is coarser.

probability 1. However, this implies that 1's payoff from playing L_1 is $0.1 > 0$, so she cannot play R_1 .

What happens in this example is that, with a discrete terminal node partition for player 1, player 2 can and should learn 3's play by observing 1's play. \square

5 Justification of the RPCE Definition

In this section we consider several examples to justify the details of the definitions of RPCE and PCE. In each of these examples, we consider a game and its outcome that we think is plausible as a consequence of rational learning, and argue that such an outcome is possible in RPCE, while it would be impossible under alternative versions of the definition.

Example 7 argues that the self-confirming condition should not be imposed on a non-actual version, and Example 8 argues that we should allow for correlated beliefs in our model. Example 9 justifies our specification of the space that beliefs lie in, and Example 10 shows that adding the coherent belief condition to the PCE concept may rule out some of sensible outcomes.

Example 7 (Self-Confirming Condition for Zero-Share Versions).




Figure 7

In this example we argue that we should not require the self-confirming condition for non-actual versions.

In the game in Figure 7, player 1 moves first, choosing between In and Out . Regardless of 1's choice, players 2 and 3 play a coordination game, choosing between L_2 and R_2 , and L_3 and R_3 , respectively, not knowing 1's action. The terminal node partitions are such

that everyone observes the exact terminal node reached, except that 1 does not observe 2 and 3's play if she plays *Out*.

For the solution concept to correspond only to the implications of common knowledge of rationality and (almost) common knowledge of the payoff functions, it should allow for the outcome (Out, L_2, L_3) . Intuitively, if 1 thinks that 2 and 3 coordinate on the (R_2, R_3) equilibrium, she has an incentive to play *Out*, which makes her unable to observe what 2 and 3 actually play. Given what 1 is observing, 1's belief that 2 and 3 coordinate on the (R_2, R_3) equilibrium does not contradict the assumption of the common knowledge of rationality and the observation structure.

Indeed, the outcome (Out, L_2, L_3) is possible in RPCE. To see this, consider the belief model and actual versions:

$$V_1 = \{v_1^1\}, \quad v_1^1 = (Out, (R_2, R_3), p^{1,1}(v_2^2, v_3^2) = 1);$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (L_2, (Out, L_3), p^{2,1}(v_1^1, v_3^1) = 1), v_2^2 = (R_2, (Out, R_3), p^{2,2}(v_1^1, v_3^2) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (L_3, (Out, L_2), p^{3,1}(v_1^1, v_2^1) = 1), v_3^2 = (R_3, (Out, R_2), p^{3,2}(v_1^1, v_2^2) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

It is clear from inspection that (Out, L_2, L_3) is an RPCE with this belief model. Notice v_2^2 and v_3^2 are non-actual versions, and they do not satisfy the self-confirming condition. v_1^1 plays *Out* because she conjectures that these non-actual versions exist, and her conjecture is never falsified because she plays *Out*.

Now we show that the outcome in which (Out, L_2, L_3) is played is impossible if we further require the self-confirming condition for non-actual versions. To see this, suppose that we strengthen Definition 6 by replacing condition (3) with the condition that for all v_i^k , v_i^k is self-confirming with respect to π^* .

First note that, since 2 is playing L_2 , by this modified condition and the best response condition all versions of player 3 should play L_3 . Similarly, all versions of player 2 should play L_2 . Then, since we assume coherent belief, any version of player 1 must believe that players 2 and 3 play (L_2, L_3) . But by the best response condition 1 must then play *In* (which gives her payoff 1) as opposed to *Out* (which gives her payoff 0).

Therefore, the outcome in which (Out, L_2, L_3) is played is impossible if we use the modified self-confirming condition.¹⁷ □

Example 8 (Correlated Beliefs).

¹⁷Notice that this argument does not go through with the notion of PCE: Since we do not require beliefs to be coherent, 1 may still believe that 2 and 3 to play R_2 and R_3 , respectively.




Figure 8

Our formulation of beliefs is more complicated than DFL, because we allow for correlated beliefs, while DFL restricted attention to independent beliefs. In this example we motivate correlated beliefs by arguing that a player can play an action only when she believes in correlation at such information sets.

Consider the game depicted in Figure 8. This game is similar to Example 7, but player 1 has two actions that make the terminal nodes observable for her. Specifically, 1 moves first, choosing among three alternatives: A , B , and Out . Regardless of her action, players 2 and 3 play the simultaneous-move coordination game choosing between L_i and R_i for each player i , not knowing 1's action. The terminal node partitions are such that everyone observes the exact terminal node reached except that player 1 cannot distinguish among four terminal nodes that are caused by the action Out .

To capture the long-run consequences of rational learning, RPCE should allow for the possibility that 1 plays Out . Intuitively, since players 2 and 3 get to play on the path, they should play as in a Nash equilibrium of their coordination game. Hence it makes sense for player 1 to believe that players 2 and 3 coordinate on either (L_2, L_3) or (R_2, R_3) , but that they are equally likely. Given this belief, the expected payoff from playing action A is the average of 1 and -2 , which is $-\frac{1}{2}$, and the payoff for action B is also $-\frac{1}{2}$ in the same way. Hence, with this belief, playing Out , which gives her the payoff of 0, is indeed an optimal action to take.

Note that RPCE allows for correlated beliefs, so the exact argument that we have just made implies that player 1 can play Out in RPCE. Indeed, it is clear from inspection that all the conditions in the definition of RPCE hold in the following belief model and actual

versions:

$$V_1 = \{v_1^1\}, \quad v_1^1 = (Out, (\frac{1}{2}(L_2, L_3) + \frac{1}{2}(R_2, R_3)), p(v_2^1, v_3^1) = p(v_2^2, v_3^2) = \frac{1}{2});$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (L_2, (Out, L_3), p(v_1^1, v_3^1) = 1), v_2^2 = (R_2, (Out, R_3), p(v_1^1, v_3^2) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (L_3, (Out, L_2), p(v_1^1, v_2^1) = 1), v_3^2 = (R_3, (Out, R_2), p(v_1^1, v_2^2) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

However, if player 1 is restricted to hold an independent belief, the action *Out* is impossible. To see this, notice that for *Out* to be at least as good as playing *A* for a version of player 1, her belief has to assign probability at least $\frac{1}{3}$ to (L_2, L_3) . In the same way, for *Out* to be at least as good as playing *B* for a version of player 1, her belief has to assign probability at least $\frac{1}{3}$ to (R_2, R_3) . However, any independent randomization by players 2 and 3 leads to the situation where the minimum of the probabilities assigned to (L_2, L_3) and (R_2, R_3) is no more than $\frac{1}{4}$. Hence for any independent beliefs, *Out* cannot be a best response.

We note that, as in Example 5, if the terminal node partitions were discrete, player 1 could not play *Out*. However, the reason behind this effect of terminal node partitions is different: It is now that player 1 can entertain a correlated belief, which she would be unable to have if she actually observes 2 and 3's play.¹⁸ \square

Example 9 (Assessment-Strategies Correlation).

¹⁸A similar argument can be made in Example 9 to show that player 1 cannot play *Out* and so player 4 cannot play *R*₄.




Figure 9

We have allowed v_i^k 's belief $\mu^{i,k}$ to lie in the space $\Delta(\mathcal{A}_i \times \Pi_{-i})$ and not necessarily in $\mathcal{A}_i \times \Delta(\Pi_{-i})$. Here we provide an example that justifies this specification.

Consider the extensive-form game depicted in Figure 9. In this example, player 4 moves first, choosing between L_4 and R_4 . In either case player 2 moves next, choosing between L_2 and R_2 , not knowing 4's choice. If 4 played L_4 , player 1 moves next, knowing that 4 chose L_4 but not knowing 2's choice. She has three choices, A , B , and Out . Unless 1 gets her move and plays Out , 3 gets to play without knowing anything about the past play other than the fact that 1 has not played Out . 3 chooses between L_3 and R_3 . The terminal node partition is such that everyone observes the exact consequence of any sequence of actions, while player 1 does not distinguish among those terminal nodes that are caused by R_4 .

We argue that it is plausible that 4 plays R_4 . To support this outcome, it must be possible that 4 believes 1 to play Out once her information set is reached. For this play to satisfy the best response condition at this information set, we should allow for player 1 to believe that players 2 and 3's play is correlated, just as in Example 8. Notice that player 1 knows that 2 and 3 are actually playing the coordination game on the path of play because 4 plays R_4 , thus this correlated belief seems to be one of the plausible beliefs (and indeed it is possible in our RPCE concept). However, to allow for this correlation, 1 should expect 3 to play L_3 given the node caused by L_2 and R_3 given the node caused by R_2 . This is possible when each profile of opponents' strategies is associated with a

different assessment, but is impossible if only a single assessment is used for a distribution of the opponents' strategies. Indeed, for any single assessment at 1's information set, 1's expected payoff from playing either A or B is at least $\frac{1}{4}$, so playing Out can never be a best response. Hence player 4 should expect the payoff of 1 by playing L_4 , which means that 4 cannot play R_4 . \square

Example 10 (PCE with and without Coherent Beliefs).




Figure 10

As mentioned earlier, adding the coherent belief condition to the PCE concept rules out some profiles which seem plausible when players do not know the opponents' payoff functions.

Figure 10 displays a simultaneous-move game between players 1 and 2, where player 1 chooses between In and Out , while player 2 chooses between L and R . Note that L is the dominant strategy for player 2. The terminal node partition is that both players observe the exact terminal node reached except that player 1 does not observe the consequence of player 2's play if she plays Out .

First we argue that (Out, L) is a sensible outcome in this game if players do not know the opponents' payoff functions. To see this, note that L is a best response against Out , which player 2 indeed observes. Out is not a best response against L , but player 1 does not observe player 2's play when she plays Out , so she may well believe that player 2 is playing R . In this case the expected payoff from playing In is -1 , so playing Out is indeed a best response against such a belief. Thus PCE allows for this outcome.

However, if we add the coherent belief condition, this outcome is no longer supported in PCE. To see this, notice that the best response condition ensures that all versions of player 2 play L , as it is the dominant strategy. If we impose the coherent belief condition, player 1's belief has to be a convex combination of player 2's strategies specified in player

2's possible versions. Hence player 1 must believe that player 2 will play L with probability 1. But then the best response against this belief is In , invalidating the candidate outcome of (Out, L) .

This is in contrast with the DFL's Theorem 2.1, which shows that adding the belief-closed condition (which corresponds to our coherent belief condition) to the SCE concept does not restrict the set of possible outcomes. In their context players know opponents' play on the equilibrium path. Thus if a player's belief about an opponent's play at an information set h corresponds to a dominated strategy then h must lie off the path of play. This conclusion fails if players do not necessarily observe all on-path play, which is why adding the coherent belief condition matters for PCE but not SCE.¹⁹

We note that, if the terminal node partitions were discrete, player 1 could not play In in any PCE. So terminal node partitions allow extra actions not only under RPCE but also under PCE. \square

6 The Effect of Changes in Terminal Node Partitions

In this section we discuss the effect of changing terminal node partitions on RPCE outcomes. The focus of this section is mostly on how the set of an individual player's RPCE strategies depends on terminal node partitions. Before discussing this, in Subsection 6.1, we briefly discuss how the set of RPCE strategy profiles depends on terminal node partitions.

In Subsection 6.2, we ask how the set of an individual player's RPCE strategies depends on the terminal node partitions.²⁰ Specifically, we identify four reasons that the set of an individual player's RPCE strategies is affected by terminal node partitions, and show that in a special class of games the set is invariant with respect to changes in terminal node partitions.

6.1 The Effect of Terminal Node Partitions on RPCE Strategy Profiles

First, consider how the set of RPCE strategy *profiles* (not an individual player's strategies) changes with the terminal node partitions. If the terminal node partitions \mathbf{P} are coarser

¹⁹One may wonder if this result hinges on the fact that we do not impose observational consistency. This is not true, as a belief model to support the (Out, L) outcome can satisfy the observational consistency condition (for example, specify $V_1 = \{(Out, R)\}$ and $V_2 = \{(L, Out)\}$).

²⁰One motivation for this is that the analyst may only know the terminal node partitions of some of the players and/or may only observe some players' moves.

than \mathbf{P}' then any strategy profile that is RPCE under \mathbf{P}' is also an RPCE under \mathbf{P} : if a belief model rationalizes a strategy profile under \mathbf{P}' then it can also be used to rationalize the same strategy profile under \mathbf{P} .

On the other hand, versions V in the belief model that rationalizes a strategy profile under \mathbf{P} may not rationalize it under a finer partition \mathbf{P}' . Perhaps the most obvious reason is that a player may not want to play a particular action once she learns the unobserved play by the opponents. For example, the strategy profile discussed in Example 7 ((Out, L_2, L_3)) would not be an RPCE if player 1's terminal node partition were discrete: If she observes that the equilibrium that the opponents are coordinating on is different from the one that she was expecting, she wants to play In .

These examples show that not only the set of RPCE strategies but also the RPCE outcome of these games (the distribution over terminal nodes) can depend on the terminal node partitions.

6.2 The Effect of Terminal Node Partitions on an Individual Player's RPCE Strategies

Now we ask how the set of an individual player's RPCE strategies depends on the terminal node partitions. To begin with, note that if terminal node partitions \mathbf{P} are coarser than \mathbf{P}' then any strategy that i can play in RPCE under \mathbf{P}' can also be played by i in an RPCE under \mathbf{P} . The main purpose of this section is to identify four reasons that versions V in the belief model that rationalizes a strategy under \mathbf{P} may not rationalize it under a finer partition \mathbf{P}' . As we will see, these reasons are that:

- (i) Player i 's opponents' terminal node partitions have changed (Example 5);
- (ii) Player i has a correlated belief (Example 8);
- (iii) Player i 's belief is coherent with the conjecture that assigns strictly positive probabilities to multiple versions of the opponents (Example 11);
- (iv) Some player j believes player i has an incorrect belief (Example 11 Revisited, Example 12, and Example 13).

In the absence of all these reasons, we can specify a new actual versions for i 's opponent as in i 's belief, and such an actual version profile defined in the original belief model would rationalize a strategy under \mathbf{P}' .²¹ Hence, when these conditions are absent, the set of player i 's RPCE strategies is invariant with respect to her terminal node partition.

²¹This is proven in Theorem 5. For example, in Example 7, the new actual version profile for the opponents would assign share 1 to (R_1, R_2) , which corresponds to 1's actual version's belief.

The plan of the rest of this subsection is as follows. First we briefly review reasons (i) and (ii). Then we present an example for reason (iii) (Example 11), and discuss how it is related to the characterization of SCE outcomes in Fudenberg and Levine (1993a) and Kamada (2010). We then use the example and a related similar one to illustrate (iv). In response, we define a notion of “own-action independence,” which is violated in all of these examples. We show that reason (iv) can still occur under the “own-action independence” condition (Example 13), but not in simultaneous move games. This leads us to Theorem 5, which shows that in simultaneous move with own-action independence for player i 's opponents and measurable payoffs (with respect to terminal node partitions), the set of i 's RPCE strategies is invariant with respect to her terminal node partition.

We have already discussed the first two reasons. For reason (i), as discussed in Example 5, if the opponents' terminal node partitions were different in two games, a player may expect different opponents' plays by playing In , so she may play differently. For reason (ii), as discussed in Example 8, if a player plays an action entertaining a correlated belief about the opponents' unobserved play under one terminal node partition, she may not play that action if the partition were discrete, because the observed play should correspond to an independent belief. In both examples, the strategies that are ruled out by refining the partitions are played on the path of play, thus changing the partitions alters not only the set of strategies that the player in question can play in RPCE, but also the outcomes that she can induce in RPCE. The same comment applies to all examples that follow when we change terminal node partitions.

The next example explains reason (iii).

Example 11 (Reason (iii)).




Figure 11

This example considers a 2-player game, so reason (ii) does not apply. Here if player 1's terminal node partition is changed (keeping the opponents' partitions unchanged) she can no longer play an action that she plays under the original terminal node partition.

Consider the game in Figure 11. This game is a two-player simultaneous-move game. Player 1 chooses among A_1 , B_1 , and Out_1 ; player 2 chooses among In_2 and Out_2 . The terminal node partitions are such that both players observe the exact terminal node reached except that player 1 does not observe the consequence of player 2's action if she plays Out_1 , and player 2 does not observe the consequence of player 1's action if he plays Out_2 .

First we show that player 1 can play Out_1 . To see this, consider the following belief model and actual versions:

$$V_1 = \{v_1^1, v_1^2\}, \quad v_1^1 = (Out_1, \frac{1}{2}In_2 + \frac{1}{2}Out_2, p(v_2^1) = p^{1,1}(v_2^2) = \frac{1}{2}), v_1^2 = (B_1, Out_2, p^{1,2}(v_2^1) = 1);$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (In_2, Out_1, p^{2,1}(v_1^1) = 1), v_2^2 = (Out_2, B_1, p^{2,2}(v_1^2) = 1);$$

The actual version profile is (v_1^1, v_2^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold. Note that, since each player faces only one opponent, all players have independent beliefs.

Notice that, although player 1's action is rationalized by a belief that corresponds to 2's mixed strategies, she is sure that 2 is playing a pure strategy: Both versions v_2^1 and v_2^2 play pure strategies. If 1's conjecture assigns probability 1 to either of these versions, 1 cannot play Out_1 : If 1 expects In_2 with probability 1 then she wants play A_1 ; if she expects Out_2 with probability 1 then she wants play B_1 . Thus, the action Out_1 is possible only when 1's belief corresponds to 2's mixed strategy.

Player 1 can be unsure which of v_2^1 and v_2^2 is present, because she plays Out_1 and this makes it unable for her to observe the exact terminal node reached, so she never knows which action player 2 is playing.

Now we argue that, if 1's terminal node partition is discrete, she can never play a strategy that assigns probability 1 to Out_1 . To see this, we first note that no version of player 2 can play a mixed strategy if 1 plays Out_1 . This is because if player 1 plays Out_1 with probability 1 and player 2 assigns a positive probability to In_2 , then 2's terminal node partition allows him to observe 1's choice of Out_1 , and hence he expects the payoff of 1 from playing In_2 and 0 from playing Out_2 . This means he is not indifferent, so he cannot mix.

Thus, whenever 1 plays Out_1 with probability 1, 2 should not play a mixed strategy. But this implies that, if 1's terminal node partition is discrete, 1 is observing either (a) 2

is playing In_2 with probability 1 or (b) 2 is playing Out_2 with probability 1. However, as we have explained above, player 1 would be better off by playing A_1 than Out_1 in case (a), and B_1 than Out_1 in case (b). Hence, she cannot play a strategy that assigns probability 1 to Out_1 if her terminal node partition is discrete, although this action could be played if the partition were not discrete.

The key here is that player 1's belief is coherent with the conjecture that assigns strictly positive probabilities to multiple versions of player 2, but the corresponding "mixed strategy" by player 2 cannot be played in RPCE. \square

A Remark on Example 11.

Fudenberg and Levine (1993a) and Kamada (2010) identify the conditions that guarantee that the outcome of an SCE is identical to a Nash outcome, in games with discrete terminal node partitions. Roughly, players are required to have independent and unitary beliefs, and the equilibrium has to be "strongly consistent," as defined in Kamada (2010). To prove this theorem, they explicitly construct a Nash equilibrium from an SCE that satisfies these conditions: While for an off-path information set h_j that player i can deviate to reach, they set player j to play as in i 's belief, strategies at other information sets are unchanged. The three conditions ensure that this modification is well-defined. In particular, the independent beliefs condition guarantees that the modification can be done information set by information set.

Given this, it might seem natural to conjecture that given independent beliefs, if π^* is an RPCE under partitions $(\mathbf{P}_i, \mathbf{P}_{-i})$ then in a game under $(\bar{\mathbf{P}}_i, \mathbf{P}_{-i})$ with $\bar{\mathbf{P}}_i$ being the discrete partition, we can let i 's opponents play "as in i 's belief" (while we do not change i 's strategy) and the modified strategy profile constitutes an RPCE of the game under $(\bar{\mathbf{P}}_i, \mathbf{P}_{-i})$, because of common knowledge of rationality. Reason (iii) above shows why this argument fails: The problem is that we cannot replace i 's opponents' strategies "as in i 's belief" even if we impose independent beliefs. This is what happens in Example 11. In Example 11, two versions of player 2 that player 1 assigns positive probabilities play different strategies which are rationalized by different beliefs, and it is not necessary the case that we can rationalize a convex combination of these pure strategies by some single belief. The intuition is similar to the idea behind the need of unitary beliefs to establish the outcome equivalence between SCE and Nash: If heterogeneous beliefs are allowed, a player's two pure strategies in the support of her mixed strategy may not be rationalized by a single belief, so the mixed action may not be played in a Nash equilibrium. \square

So far we have identified three reasons that versions V in the belief model that rationalizes a strategy under \mathbf{P} may not rationalize it under a finer partition \mathbf{P}' : (i) player i 's opponents' terminal node partitions have changed; (ii) i has a correlated belief; (iii) i 's belief is coherent with the conjecture that assigns strictly positive probabilities to multiple versions of the opponents.

Further investigation of Example 11 shows that there is yet another reason that we may expect a different RPCE actions for different terminal node partitions, which is reason (iv):

Example 11 Revisited.

Now we analyze the effect of making player 2's terminal node partition discrete in Example 11.

Now, player 2 is no longer able to play a strategy that plays In_2 with probability 1. To see this, suppose that In_2 is played with probability 1 when 2's terminal node partition is discrete. Notice that In_2 is a best response only when player 1 assigns probability no less than $\frac{1}{2}$ to Out_1 . However, as we have discussed, Out_1 can be a best response only when 1's belief correspond to a mixed strategy of player 2. The version of player 1 who plays Out_1 should expect that the versions that she assigns a positive probability should be observing Out_1 regardless of these versions' strategies because 2's terminal node partition is discrete. But the best response condition implies that these versions should play In_2 with probability 1, contradicting our earlier conclusion that 1's belief should correspond to a mixed strategy.

What happens here is as follows: When player 2's terminal node partition was not discrete, player 1 could believe that player 2 has an incorrect belief about player 1's action. However, if 2's terminal node partition is discrete, player 1 cannot entertain such a belief, which rules out some of player 1's actions, and it in turn rules out some actions of player 2.

Notice that, although the example involves "reason (iii)" that we discussed above, the logic is independent of it. The next simple example clarifies this point. The example also shows that reason (i) can be present in two-player games without involving reason (iii). □

Example 12 (Reason (iv)).




Figure 12

Consider the game depicted in Figure 12. In this game, players 1 and 2 move simultaneously, choosing between In_i and Out_i , $i = 1, 2$. Notice that in this game, player 2 is indifferent between In_2 and Out_2 when 1 plays Out_1 . As usual, the terminal node partition is such that player i observes the consequence of the opponent's action when she plays In_i , while she cannot do so when she plays Out_i .

We first show that player 1 can play In_1 given these terminal node partitions. To see this, consider the following belief model and actual versions:

$$V_1 = \{v_1^1, v_1^2\}, \quad v_1^1 = (In_1, Out_2, p(v_2^1) = 1), v_1^2 = (Out_1, In_2, p(v_2^2) = 1);$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (Out_2, Out_1, p(v_1^1) = 1), v_2^2 = (In_2, Out_1, p(v_1^2) = 1);$$

The actual version profile is (v_1^1, v_2^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold. Notice that player 2 plays Out_2 because he believes player 1 is playing Out_1 . Such a belief is justified because given Out_1 , 1 does not observe 2's play, so 1 can *incorrectly* believe that 2 is playing In_2 . However, such an "incorrect belief" is not possible if player 1's terminal node partition is discrete, so player 2 cannot believe that 1 plays Out_1 when he plays Out_2 . This in turn rules out the possibility of the strategy that assigns probability 1 to In_1 .

To see this formally, suppose that player 1's terminal node partition is discrete and she plays In_1 with probability 1. Then, for the best response condition for player 2 to be satisfied, player 2 must be playing In_2 with probability 1, or Out_2 with probability 1. However for the best response condition for player 1 to hold, it must be the case that

Out_2 is played with probability 1. For Out_2 to be a best response for player 2, his belief must assign probability 1 to Out_1 . But then the observational consistency condition and the assumption that player 1's terminal node partition is discrete imply that there exists a version of player 1 who plays Out_1 with a belief that assigns probability 1 to Out_2 . However such a version violates the best response condition, as In_1 gives a strictly higher payoff than Out_1 against Out_2 .

Note that the example hinges on the assumption that player 2 is indifferent between In_2 and Out_2 when 1 plays Out_1 , as otherwise either one of (In_2, Out_1) or (Out_2, Out_1) cannot satisfy the best response condition. However the logic behind reason (iv) is independent of ties, as shown by Example 16 in the Online Supplementary Appendix.

Finally we note that player 1 cannot play In_1 if player 2's terminal node partition becomes discrete. This is easy to check: If it were discrete, player 2 must play In_2 with probability 1 if player 1 plays In_1 . However, then, player 1 would be better off by playing Out_1 than In_1 . This shows that reason (i) can be present in two-player games without involving reason (iii). \square

Note that, in Examples 11, 11 Revisited, and 12, it is important that an opponent's observation about other players' strategies depends on his own action. In these examples, this dependence is captured by the terminal node partitions. Notice that terminal node partitions are not the only ways to capture this dependence— if whether a player's opponent gets a move depends on the player's action, then the player's observation will depend on her own action. In Example 17 in the Online Supplementary Appendix, we present one such example.

To formalize this dependence, we introduce a notion of “own-action independence”:

Let $\zeta : A \rightarrow Z$ be the map that assigns to each action profile the terminal node caused by that action profile.

Definition 7. A game with player i 's terminal node partition \mathbf{P}_i is **own-action independent for i** if, $\zeta(a_i, a_{-i})$ and $\zeta(a_i, a'_{-i})$ are in the same cell of \mathbf{P}_i if and only if $\zeta(a'_i, a_{-i})$ and $\zeta(a'_i, a'_{-i})$ are in the same cell of \mathbf{P}_i .

That is, if i 's action does not affect what she observes, the game is own-action independent for i . Notice that, for instance in Example 1, own-action independence for player 1 not only rules out games B and C which involve nondiscrete terminal node partitions, but also game A in which the terminal node partition is discrete. The reason is that in game A, the terminal node the cell of terminal node partition caused by Out is independent of 2 and 3's actions, while it does depend on their actions if she plays In . The condition

is satisfied, for example, in simultaneous-move games with discrete partitions, but it is slightly more general. For example, game A is own-action independent for players 2 and 3.

Even with own-action independence for players other than i , reason (iv) still has bite, as seen in the following example.

Example 13 (Learning from an Opponent’s Play II).




Figure 13

Consider the game depicted in Figure 13. This game is exactly the same as the one in Example 6, except that now the terminal node partition for player 1 is such that she does not observe the consequence of 3’s action if 1 plays R_1 . Notice that the game is own-action independent for players 2 and 3.

First, we show that player 1 can play R_1 in an RPCE. To see this, consider the following belief model and actual versions:

$$V_1 = \{v_1^1\}, \quad v_1^1 = (R_1, (L_2, R_3), p(v_2^1, v_3^1) = 1);$$

$$V_2 = \{v_2^1\}, \quad v_2^1 = (L_2, (R_1, L_3), p(v_1^1, v_3^2) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (R_3, (R_1, L_2), p(v_1^1, v_2^1) = 1), v_3^2 = (L_3, (R_1, L_2), p(v_1^1, v_2^1) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold. Notice that players 1 and 2 disagree about player 3's action, which neither of them observe when 1 plays R_1 , which is why 2 can play L_2 even though 1 is playing R_1 .

Now, remember that in Example 6 we showed that, if player 1's partition is discrete, she can no longer play R_1 : With a discrete terminal node partition for player 1, player 2 can and should learn 3's play by observing 1's play. But this is impossible when 1 and 2's terminal node partitions coincide. \square

Notice that in the above example, it is important that, with nondiscrete partition, some player believes another player has an incorrect belief. The difference from the logic in Examples 11 and 12 is that, in these examples i 's opponent j believes that i is not best responding to j 's play with a nondiscrete partition, so if j knows i observes j 's play then j should expect i is best responding to j , so j should play differently. In Example 13, on the other hand, when the partition is discrete, j learns a third player m 's strategy from the fact that i is observing m 's play and best responding to it, and this information that j gets from observing i 's play changes how j should act. This learning from player i 's play was not an issue in Examples 11 and 12.

The key of Example 13 is that player 2's information set lies at an off-path information set, but we require him to play optimally there. A simple condition that rules out this possibility is to restrict ourselves to simultaneous-move games.

Now we have explained all four reasons that versions V in the belief model that rationalizes a strategy under \mathbf{P} may not rationalize it under a finer partition \mathbf{P}' . We are now ready to state a theorem.

Let $\Pi_i^*(\mathbf{P}_i, \mathbf{P}_{-i}) \subseteq \Pi_i$ be the set of i 's strategies that i can play in RPCE under terminal node partitions $(\mathbf{P}_i, \mathbf{P}_{-i})$. Also, let $\Pi_i^{IND}(\mathbf{P}_i, \mathbf{P}_{-i}) \subseteq \Pi_i^*(\mathbf{P}_i, \mathbf{P}_{-i})$ be the set of i 's strategies that i can play in RPCE in which the actual version of i is restricted to entertain an independent belief, under terminal node partitions $(\mathbf{P}_i, \mathbf{P}_{-i})$.

Theorem 5. *Suppose that a game is simultaneous-moves and own-action independent for all $j \neq i$, and that payoffs are measurable with respect to terminal node partitions. Then, for all pair of i 's partitions, \mathbf{P}_i and \mathbf{P}'_i , and all \mathbf{P}_{-i} , $\Pi_i^{IND}(\mathbf{P}_i, \mathbf{P}_{-i}) = \Pi_i^{IND}(\mathbf{P}'_i, \mathbf{P}_{-i})$.²²*

²²The key feature of Example 13 is that there is a path of actions from the off-path information set of player 2 to reach the on-path information set of player 3. We conjecture that the theorem would hold even if we relaxed simultaneous moves to the assumption that player i "induces a single hypothesis," that is, for all information set of player i , $h_i \in H_i$, and for every information set h that is a successor of h_i , there is only a single sequence of information sets $\{h^{(\iota)}\}_{\iota=1}^{\kappa}$ such that $h^{(1)} = h_i$, for each $\iota = 1, \dots, \kappa - 1$ there is an action from the node in $h^{(\iota)}$ which leads to a node in information set $h^{(\iota+1)}$, and $h^{(\kappa)} = h$. This condition is satisfied, for example, in game A and all two-player games (with perfect recall).

This theorem restricts beliefs to be independent. Since the distinction between independent and correlated beliefs has bite only for games with three or more players, we have the following corollary:

Corollary 2. *In two-player simultaneous-move games, if a game is own-action independent for $j \neq i$ and payoffs are measurable with respect to terminal node partitions, then for all pair of i 's partitions, \mathbf{P}_i and \mathbf{P}'_i , and all \mathbf{P}_{-i} , $\Pi_i^*(\mathbf{P}_i, \mathbf{P}_{-i}) = \Pi_i^*(\mathbf{P}'_i, \mathbf{P}_{-i})$.²³*

We note that the theorem and the corollary rely on the measurability of payoffs with respect to terminal node partitions. Indeed, the corollary (and hence the theorem) fails if the payoffs are not measurable with respect to terminal node partitions. In the Online Supplementary Appendix, we provide a counterexample for such a case (Example 18).

To sum up, the set of strategies a player can use in equilibrium is typically sensitive to the details of her terminal node partition.

As a final remark in this section, we note that the set of an individual player's PCE strategies are unaffected by reasons (i), (iii), and (iv) as they have bite only when a player knows the opponent's payoff functions. However, even without reason (ii), the set of an individual player's PCE strategies depends on terminal node partitions because, as we have seen in Example 10, PCE does not require the knowledge about the opponents' payoff functions.

7 RPCE, RSCE, and RCE

In this section we compare RPCE with other concepts from the literature. In Subsection 7.1 we compare RPCE with RSCE, and show that RPCE “reduces” to RSCE if the terminal node partition is discrete and beliefs are unitary and independent. In Subsection 7.2 we compare RPCE with RCE, and show that when the “signal function” specified in the definition of RCE takes an appropriate form, RPCE is equivalent to RCE if moves are simultaneous.

²³Notice that, in Example 11, player 1 cannot assign probability 1 to Out_1 when player 2's terminal node partition is discrete, so reason (i) still has bite even in two-player games. Hence the equivalence in this corollary would fail if we changed the opponent's terminal node partition. To see that player 1 cannot assign probability 1 to Out_1 if player 2's partition becomes discrete, suppose, to the contrary, that player 1 plays Out_1 with probability 1 in an RPCE when player 2's partition becomes discrete. Since player 1 must expect player 2 to best respond to Out_1 and the unique best response to Out_1 is In_2 , player 1's belief must assign probability 1 to In_2 . However, player 1 would then be better off by playing A_1 rather than Out_1 , contradicting the best response condition for player 1.

7.1 Comparison with Rationalizable Self-Confirming Equilibrium

In this subsection we show that RPCE reduces to RSCE if we require independent beliefs when the terminal node partitions are discrete. Given these restrictions, superficially there are two differences between RSCE and RPCE: First, in RSCE the self-confirming condition is required for all versions but in RPCE it is required only for actual versions. Second, observational consistency is not directly imposed in RSCE but it is imposed in RPCE. We show that these two differences do not affect the set of equilibria when unitary beliefs are required and the terminal node partitions are discrete.

To see this formally, let us first define RSCE (notations are adjusted to accord with ours). This concept is defined for games with discrete terminal node partitions. First, we define SCE:

Definition 8. π^* is a **self-confirming equilibrium** if there exist a belief model V and an actual version profile v^* such that the following three conditions hold:

1. π^* is generated by v^* .
2. For all i and k , π_i^k is a best response to $\mu^{i,k}$ at $H(\pi_i^k, \pi_{-i}^{i,k})$.
- 3'. For any versions v_i^k , $d(\pi_i^k, \pi_{-i}) = d(\pi^*)$ for all π_{-i} in the support of $\mu^{i,k}$.
6. For all i and k , v_i^k has an independent belief.

As in the difference between PCE and RPCE, RSCE strengthens condition (2), and adds the coherent belief condition.

Definition 9. π^* is a **rationalizable self-confirming equilibrium** if there exist a belief model V and an actual version profile v^* such that the following three conditions hold:²⁴

1. π^* is generated by v^* .
- 2'. For all i and k , π_i^k is a best response to $\mu^{i,k}$ at $H(\pi_i^k)$.
- 3'. For any versions v_i^k , $d(\pi_i^k, \pi_{-i}) = d(\pi^*)$ for all π_{-i} in the support of $\mu^{i,k}$.
4. For all i and k , v_i^k has a coherent belief;
6. For all i and k , v_i^k has an independent belief.

²⁴DFL allows all $\hat{\pi}$ that have the same distribution over terminal nodes as π^* to be RSCE, but this difference is not important for our purpose.

As noted earlier, there are two differences between these definitions and those of PCE and RPCE, namely that condition (3') (*every* version expects the same distribution over terminal nodes) is stronger than condition (3), and that PCE and RPCE do not impose the independence condition (6). Even with a discrete terminal node partition the way condition (3) is stated is somewhat different than condition (3'), but as the next result shows this difference is irrelevant.

Theorem 6. *Fix a game with discrete terminal node partitions.*

1. *Condition (3') implies condition (5).*
2. *If a strategy profile is generated by a belief model that satisfies conditions (3), (4), and (5) then it can also be generated by a belief model that satisfies conditions (3'), (4), and (5).*

Part 1 is not surprising: Since the terminal node partitions are discrete, condition (3') essentially requires that the terminal node reached is common knowledge, so observational consistency holds. Part 2 says that in the presence of the observational consistency condition, requiring the self-confirming condition for non-actual versions does not further restrict the set of equilibria. Notice that this conclusion was not true when we considered RPCE with nondiscrete terminal node partitions (See Example 7).

Corollary 3. *RPCE with independent beliefs in games with discrete terminal node partitions is equivalent to RSCE.*

In the next example, which is taken from DFL's Example 3.2, we show that replacing condition (3) by condition (3') without requiring observational consistency would change the set of possible outcomes even when the terminal node partitions are discrete.

Example 14 (DFL Example).




Figure 14

Consider the game depicted in Figure 14. In this game, there are three players, 1, 2, and 3. Player 1 moves first, choosing between u and r . If 1 chooses r , 2 gets the move, choosing between a and d , where the latter ends the game. If u or (r, a) are chosen then 3 gets the move, without knowing which causes his move. 3 chooses between U and D . All players' terminal node partitions are discrete.

DFL argue that the outcome (u, U) is impossible in RSCE, because if 3 chooses U then 2 should play a since he observes the terminal node, and then 1 should take r . However, if we replace condition (3') by condition (3) in Definition 9 where observational consistency is not imposed, this outcome becomes possible. To see this, consider the following belief model and actual versions:

$$V_1 = \{v_1^1, v_1^2\}, \quad v_1^1 = (u, (d, U), p^{1,1}(v_2^2, v_3^1) = 1), v_1^2 = (r, (a, D), p^{1,1}(v_2^1, v_3^2) = 1);$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (a, (u, U), p^{2,1}(v_1^1, v_3^1) = 1), v_2^2 = (d, (u, D), p^{2,2}(v_1^1, v_3^2) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (U, (u, a), p^{3,1}(v_1^1, v_2^1) = 1), v_3^2 = (D, (r, a), p^{3,2}(v_1^2, v_2^2) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

It is clear from inspection that all the conditions in the definition of RSCE other than condition (3') hold and that condition (3) holds. Notice that v_1^2 , v_2^2 , and v_3^2 are not self-confirming, and they are non-actual versions.

The key is that the actual version of player 1, v_1^1 , believes that player 2 plays d , while such a belief is impossible if observational consistency is imposed: the equation in the definition of observational consistency applied to v_1^1 's belief is $d(\pi_2^2, \pi_{-2}^{2,2}) = d(\pi_2^2, \pi_1^1, \pi_{-1,2}^{1,1})$, which is $d(d, (u, D)) = d(d, u, U)$. But this equation is false. \square

Notice that the distinction between conditions (3) and (3') described in the above example relies on the fact that RSCE requires common knowledge of rationality (at reachable nodes). Indeed, this type of examples does not exist if we consider (non-rationalizable) SCE. The next theorem generalizes this point, considering the correlated PCE concept in games with terminal node partitions that are not necessarily discrete.

Theorem 7. *The set of PCE does not change if we replace condition (3) with the following:*

For all v_i^k , v_i^k is self-confirming with respect to π^ .*

The intuition for this result is simple: Since we do not require the coherent belief condition in the PCE concept, conjectures do not mean much, so eliminating the non-actual versions (who may not satisfy the self-confirming condition) does not invalidate the belief model as a justification of a PCE.

7.2 Comparison with Rationalizable Conjectural Equilibrium

In this subsection we compare RPCE to RCE. Rubinstein and Wolinsky describe RCE as corresponding to situations where “each player chooses an action which maximizes his payoff given a conjecture regarding the actions of the others; each agent’s conjectures are consistent with his signal and own choice, (and) the conjectures are also consistent with the understanding that everyone rationalizes his action in this manner.” One obvious difference is that RPCE, like RSCE, requires players believe others will play rationally (maximize the presumed payoff functions) as long as they have not behaved irrationally in the past, while RCE is designed to model normal form games and places no restrictions on play at off-path information sets.^{25,26} Because of this difference, RPCE makes stronger predictions than RCE in most extensive-form games. If all players move on every path, this distinction becomes moot, and the two concepts become very similar. In particular, like RPCE, RCE can require a player to believe play in an unobserved subgame has converged to a Nash equilibrium, as in Example 4, which might not be apparent from the Rubinstein and Wolinsky paper.

²⁵See the example in Figure 2.1 of DFL.

²⁶Gilli (1999) proposes a related solution concept; Battigalli (1999) shows it is equivalent to RCE.

In one-shot simultaneous-move games, we can make a precise connection between RCE and RPCE. To state the comparison, let us first define the RCE concept.

Here we follow the notation used in Rubinstein and Wolinsky (1994), whenever appropriate.²⁷

Consider a normal-form game with players $I = \{1, \dots, n\}$, the action set A_i , $A = \times_{i \in I} A_i$, and $A_{-i} = \times_{j \neq i} A_j$, the payoff function $u_i : A_i \rightarrow \mathbb{R}$. The set of mixed strategies are $M_i = \Delta(A_i)$, $M = \times_{i \in I} M_i$, and $M_{-i} = \times_{j \neq i} M_j$. There is a set of private signals S_i , and a signal function $g_i : A \rightarrow S_i$. $g_i(a)$ is the signal that i privately observes when the action profile is $a \in A$. With an abuse of notation we write $g_i(m)$ for a probability distribution over S_i given the mixed profile $m \in M$, called a random signal. Let $\sigma_i \in \Delta(S_i)$ be the general element of the set of random signals.

The strategy-signal pair (m_i, σ_i) is said to be *g-rationalized* by $\mu \in \Delta(M_{-i})$ if (i) $g_i(m_i, m_{-i}) = \sigma_i$ for all $m_{-i} \in \text{supp}(\mu)$, and (ii) m_i is a best response against μ . (i) says that whatever i thinks is possible must be consistent with her strategy and the signal.

The sets of strategy-signal pairs B_1, \dots, B_n are *g-rationalizable* if for all i , every $(m_i, \sigma_i) \in B_i$ is *g-rationalized* by some μ such that for all $m_{-i} \in \text{supp}(\mu)$ and all j , $(m_j, g_j(m_i, m_{-i})) \in B_j$. What this says is as follows: Fix i 's version who plays m_i and gets the signal σ_i , with a belief μ that is consistent with this strategy-signal pair. Then, if she thinks that m_{-i} is possible, an opponent j who plays m_j must exist in B_j , and he receives a signal that is consistent with i 's play (m_i) and what she thinks is possible (m_{-i}).

An *RCE* is $m^* \in M$ such that there exists *g-rationalizable* sets B_1, \dots, B_n such that $(m_i^*, g_i(m^*)) \in B_i$ for each i .

For an extensive-form game Γ with terminal node partitions $\mathbf{P} = (\mathbf{P}_1, \dots, \mathbf{P}_n)$, let $(A^\Gamma, g^\mathbf{P})$ be the pair of normal-form representation of Γ and the profile of signal functions (denoted by $g^\mathbf{P} := (g_1^\mathbf{P}, \dots, g_n^\mathbf{P})$) that corresponds to the map from action profiles to the terminal nodes according to partition \mathbf{P} . Conversely, given any (A, g) such that $g_i(m) = g_i(m')$ implies $m_i = m'_i$ (so that the (extended notion of) perfect recall assumption is satisfied), we define the related simultaneous-move extensive form game Γ^A , and endow it with the terminal node partition that corresponds to (g_1, \dots, g_n) .

Finally, we say that a behavioral strategy π is equivalent to a mixed strategy profile m or a mixed strategy profile m is equivalent to a behavioral strategy π if π is generated by m according to the Kuhn's theorem.

Now we are ready to state the formal connection between the two concepts.²⁸ We omit

²⁷The formal part their paper considers only pure actions. The "Comments" section informally describes two extensions to mixed strategies; we adopt the first of those here.

²⁸An analogous relationship can be made between PCE and conjectural equilibrium (CE), where CE

the proof.

Theorem 8.

1. Any RPCE in (Γ, \mathbf{P}) is equivalent to some RCE in $(A^\Gamma, g^{\mathbf{P}})$.
2. Any RCE in (A, g) is equivalent to some RPCE in (Γ^A, \mathbf{P}^g) .

References

- Aumann, R. J. [1987]: “Correlated equilibrium as an expression of Bayesian rationality” *Econometrica*, 55, 1-18.
- Battigalli, P. [1987]: “Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali,” unpublished dissertation, Università Commerciale “L. Bocconi”, Milano.
- Battigalli, P. [1999]: “A Comment on Non-Nash Equilibria,” mimeo.
- Battigalli, P. and D. Guatoli [1997]: “Conjectural Equilibria and Rationalizability in a Game with Incomplete Information” in *Decisions, Games and Markets*, Kluwer Academic Publishers, Norwell, MA.
- Bernheim, D. B. [1984]: “Rationalizable Strategic Behavior,” *Econometrica*, 52, 1007-1028.
- Börgers, T. [1994]: “Weak Dominance and Approximate Common Knowledge,” *Journal of Economic Theory*, 64, 265-276.
- Brown, G. W. [1951]: “Iterative Solutions of Games by Fictitious Play” in *Analysis of Production and Allocation*, ed. by T. C. Koopmans. New York: Wiley.
- Dekel, E. and D. Fudenberg [1990]: “Rational Behavior with Payoff Uncertainty,” *Journal of Economic Theory*, 52, 243-267.
- Dekel, E., D. Fudenberg and D. Levine [1999]: “Payoff Information and Self-Confirming Equilibrium,” *Journal of Economic Theory*, 89, 165-185.
- Dekel, E., D. Fudenberg and D. Levine [2002]: “Subjective Uncertainty over Behavior Strategies: A Correction,” *Journal of Economic Theory*, 104, 473-478.
- Dekel, E., D. Fudenberg and D. Levine [2004]: “Learning to Play Bayesian Games,” *Games and Economic Behavior*, 46, 282-303.
- Esponda, I. [2011] “Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions,” mimeo.

is $m^* \in M$ such that $(m_i^*, g_i(m^*))$ is g -rationalized by some $\mu \in \Delta(M_{-i})$ for all i .

- Forges, F. [1986]: “An Approach to Communication Equilibria,” *Econometrica*, 54, 1375-1386.
- Foster, D. P. and R. V. Vohra [1997]: “Calibrated Learning and Correlated Equilibrium,” *Games and Economic Behavior*, 21, 40-55.
- Fudenberg, D. and D. M. Kreps [1993]: “Learning Mixed Equilibria,” *Games and Economic Behavior*, 5, 320-367.
- Fudenberg, D. and D. M. Kreps [1995]: “Learning in Extensive Games, I: Self-Confirming Equilibrium,” *Games and Economic Behavior*, 8, 20-55.
- Fudenberg, D., D. M. Kreps, and D. K. Levine [1988]: “On The Robustness of Equilibrium Refinements,” *Journal of Economic Theory*, 44, 354-380.
- Fudenberg, D. and D.K. Levine [1993a]: “Self-Confirming Equilibrium” *Econometrica*, 61: 523-546.
- Fudenberg, D. and D. K. Levine [1993b]: “Steady State Learning and Nash Equilibrium” *Econometrica*, 61, 547-574.
- Fudenberg, D. and D.K. Levine [2009]: “Learning and Equilibrium” *Annual Review of Economics* 1, 385-419.
- Gilli, M. [1999] “On Non-Nash Equilibria,” *Games and Economic Behavior* , 27, 184-203.
- Greenberg, J., S. Gupta, and X. Luo [2003]: “Towering Over Babel: Worlds Apart-But Acting Together,” *mimeo*, ” McGill Univ.
- Kamada, Y. [2010]: “Strongly Consistent Self-Confirming Equilibrium,” *Econometrica*, 78, 823-832.
- Kreps, D. and F. Wilson [1982] “Sequential Equilibria,” *Econometrica* 50, 863-894.
- Pearce, D. [1984]: “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 52, 1029-1050.
- Rubinstein, A. and A. Wolinsky [1994]: “Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability,” *Games and Economic Behavior*, 6, 299- 311.
- Shapley, L. [1964]: “Some Topics in Two-Person Games,” in *Advances in Game Theory*, ed. by M. Drescher, L.S. Shapley, and A. W. Tucker. Princeton University Press.

A Definition of Best Response

Given $\mu \in \Delta(\mathcal{A}_i \times \Pi_{-i})$, consider a sequence of beliefs $\{\mu^m\}_{m=1}^\infty$ such that

- For each point (a_i, π_{-i}) in the support of μ , and for each m , μ^m assigns probability $\mu(a_i, \pi_{-i})$ to some point in the interior of $\mathcal{A}_i \times \Pi_{-i}$.²⁹ Denote that point by (a_i^m, π_{-i}^m) .
- $\{(a_i^m, \pi_{-i}^m)\}_{m=1}^\infty$ converges to (a_i, π_{-i}) .
- a_i^m is generated by π_{-i}^m by Bayes rule, so in particular KW-consistency holds for each point in the support of μ .

Notice that each $h \in H_i$ is reached with a positive probability given any point in the support of μ^m for any m , provided that i 's own strategy does not rule it out. Thus, given belief μ^m and each point (a_i^m, π_{-i}^m) in the support of μ^m , we can compute by Bayes rule the probability (conditional on reaching h) that the opponents' strategy profile is π_{-i}^m . Given a point (a_i^m, π_{-i}^m) in the support of μ^m , let $r_h(a_i^m, \pi_{-i}^m)$ be this conditional probability. Let $r_h(a_i, \pi_{-i}) = \lim_{m \rightarrow \infty} r_h(a_i^m, \pi_{-i}^m)$ be its limit if it exists.

Definition 10. Given $\bar{H}_i \subseteq H_i$, π_i^k is a **best response at \bar{H}_i** if there exists a sequence of beliefs $\{\mu^m\}_{m=1}^\infty$ that converges to $\mu^{i,k}$ as above such that for all $h \in \bar{H}_i$, the restriction on π_i^k to the subtree starting at h is optimal against the probability distribution $r_h(\cdot, \cdot)$ generated by the sequence $\{\mu^m\}_{m=1}^\infty$ and $\mu^{i,k}$ in that subtree.

Note that a single sequence applies to all $h \in \bar{H}_i$, as opposed to an alternative specification in which for each h , one sequence is constructed. However it can be the case that $r_h(a_i, \pi_{-i}) \neq r_{h'}(a_i, \pi_{-i})$ if $h \neq h'$.

B Proofs

In this section we provide the proofs of the results stated in the main text.

B.1 Proof of Claim 1

Proof.

Fix a strategy profile that is generated by a belief model V and an actual version profile v^* in which all versions have coherent beliefs and satisfy the best response condition and all beliefs correspond to product measures. Construct a belief model \hat{V} in which everything is the same as in V except that each belief $\mu^{i,k}$ in V is replaced with belief $\hat{\mu}^{i,k}$ that has a unit

²⁹Remember that μ 's support is finite.

mass on a single point that corresponding to the weighted average of the original belief $\mu^{i,k}$. Specify the same actual versions as in the original model. Since two beliefs $\mu^{i,k}$ and $\hat{\mu}^{i,k}$ induce the same probability distribution over terminal nodes given any i 's strategies by construction, the best response condition is still satisfied in the new belief model. Also, since the point that $\hat{\mu}^{i,k}$ assigns the unit mass can be generated by a “mixture” (explained in footnote 11 of DFL; Dekel et al. (2002) corrects the definition of “mixture” but for our analysis both the old and the corrected definitions have the same implications) over strategies in the support of $\mu^{i,k}$ by the assumption of product measures, the belief-closed condition is also satisfied. The converse direction is analogous: for any strategy profile which is generated by a belief-closed belief model in which the best response condition is satisfied for all versions, we can replace each belief by a probability distribution over the support of the “mixture” in the belief-closed condition, which guarantees the coherent beliefs condition and the product measure requirement, and the best response property is preserved. \square

B.2 Proof of Theorem 4

Proof.

Fix a player 1 participation game Γ . Suppose that 1 plays *Out* with probability 1 in an RPCE. Let the actual version of player 1 be v_1^k . By the best response condition, *Out* is a best response to $\mu^{1,k}$. Thus it suffices to show that all the point in the support of $\mu^{i,k}$ corresponds to a Nash equilibrium of Γ_{-1} . That is, it suffices to show that each profile π_{-1} in the support of $\mu^{1,k}$ is a Nash equilibrium of Γ_{-1} .

Fix $\tilde{\pi}_{-1}$ in the support of $\mu^{1,k}$, and let the corresponding version profile be \tilde{v}_{-1} . First, observational consistency for player 1 implies that for any player $j \neq 1$, $d(\tilde{\pi}_j, \pi_{-j}) = d(\tilde{\pi}_j, \text{Out}, \tilde{\pi}_{-1,j})$ for all π_{-j} in the support of \tilde{v}_j 's belief, because $D_j = d$ as player j 's terminal node partition is discrete. This means that \tilde{v}_j 's belief assigns probability 1 to $(\text{Out}, \tilde{\pi}_{-1,j})$ because player j 's terminal node partition is discrete and Γ is a simultaneous-move game. From the best response condition $\tilde{\pi}_j$ is a best response against \tilde{v}_j 's belief, so $\tilde{\pi}_j$ is a best response against $(\text{Out}, \tilde{\pi}_{-1,j})$. Hence $\tilde{\pi}_j$ is a maximizer of $v_j(\cdot, \tilde{\pi}_{-1,j})$. Since this is true for all $j \neq 1$, this means that $\tilde{\pi}_{-1}$ is a Nash equilibrium of Γ_{-1} . \square

B.3 Proof of Theorem 5

Proof.

Fix a simultaneous-move game with terminal node partitions $(\mathbf{P}_i, \mathbf{P}_{-i})$ that is own-action independent for all players $j \neq i$. Also, fix an RPCE of this game π^* with an

associated belief model V and an actual version profile v^* in which the actual version of i has an independent belief. It suffices to show that player i can play π_i^* in an RPCE of the game under the discrete terminal node partitions.

Denote player i 's actual version by $v_i^k = v_i^*$, and consider a strategy profile $(\pi_i^k, \pi_{-i}^{i,k})$. We show that this is a Nash equilibrium of the game (hence it is an RPCE in the game with discrete terminal node partitions), which proves our claim because $\pi_i^k = \pi_i^*$ as π^* has to be generated by v^* , and all the information sets are reached with positive probability in simultaneous-move games.

First, since the best response condition is satisfied in the original belief model V and v_i^k has an independent belief, π_i^k is a best response to $\pi_{-i}^{i,k}$.

Now we show that $\pi_j^{i,k}$ is a best response against $(\pi_i^k, \pi_{-i,j}^{i,k})$. Suppose to the contrary that $\pi_j^{i,k}$ is not a best response against $(\pi_i^k, \pi_{-i,j}^{i,k})$. Then, by the linearity of the payoff functions, there must exist a version $v_j^l \in V_j$ such that $p^{i,k}(v_j^l) > 0$ and π_j^l is not a best response against $(\pi_i^k, \pi_{-i,j}^{i,k})$.

Observational consistency and the assumption that i has an independent belief imply that $D_j(\pi_j^l, \pi_{-j}) = D_j(\pi_j^l, \pi_i^k, \pi_{-i,j}^{i,k})$ for all π_{-j} in the support of $\mu^{j,l}$. This implies, by the own-action independence for j , that for any $\pi_j^l \in \Pi_j$, $D_j(\pi_j^l, \pi_{-j}) = D_j(\pi_j^l, \pi_i^k, \pi_{-i,j}^{i,k})$ must hold for all π_{-j} in the support of $\mu^{j,l}$.

This and the measurability of payoffs with respect to the terminal node partitions imply that, for any $\pi_j^l \in \Pi_j$, the payoff from playing π_j^l against $\mu^{j,l}$ is the same as that of playing π_j^l against any π_{-j} in the support of $\mu^{j,l}$, which in turn is the same as that of playing π_j^l against $(\pi_i^k, \pi_{-i,j}^{i,k})$. Since the payoff from playing π_j^l against $\mu^{j,l}$ is no less than that of playing any other π_j^l against $\mu^{j,l}$, we have that the payoff from playing π_j^l against $(\pi_i^k, \pi_{-i,j}^{i,k})$ is no less than that of playing any other π_j^l against $(\pi_i^k, \pi_{-i,j}^{i,k})$. In other words, π_j^l is a best response to $(\pi_i^k, \pi_{-i,j}^{i,k})$. \square

B.4 Proof of Theorem 6

Proof.

Part 1: Since terminal node partitions are discrete, the observational consistency condition for version v_i^k reduces to the requirement that $p^{i,k}(v_{-i}) > 0$ implies, for each $j \neq i$, $d(\pi_j(v_j), \pi_{-j}) = d(\pi_i^k, \pi_{-i}(v_{-i}))$ for all π_{-j} in the support of v_j 's belief. But the conclusion of this requirement is automatically satisfied by condition (3'), as both sides of the equality are equal to $d(\pi^*)$.

Part 2: Fix a strategy profile π^* and an associated belief model V and an actual version profile v^* that satisfies conditions (3) and (5). Construct a new belief model \hat{V} that is identical to the original ones, except that all versions that do not satisfy the self-confirming

condition in the original belief model are eliminated and each version's conjecture assigns the same weight to the versions that are still in \hat{V} . Specify the same actual version profile as in the original model (such versions are not eliminated because of condition (3)). By definition, v^* generates π^* . By construction, condition (3') holds. Hence by part 1, condition (5) holds as well. Finally, the coherence belief condition holds because condition (3) implies that the actual version $v_i^k = v_i^*$ must satisfy $d(\pi_i^k, \pi_{-i}^{i,k}) = d(\pi^*)$, and whenever $d(\pi_j^l, \pi_{-j}) = d(\pi^*)$ for all π_{-j} in the support of $\mu^{j,l}$, observational consistency implies that for any version of j 's opponent v_m^n in the support of $p^{j,l}$, $d(\pi_m^n, \pi_{-m}) = d(\pi^*)$ for all π_{-m} in the support of $\mu^{m,n}$. This means that the elimination of versions that we have done does not invalidate the coherent belief condition. \square

B.5 Proof of Theorem 7

Proof.

Fix a PCE π^* with an associated belief model V and an actual version profile v^* in which some non-actual version v_i^k 's self-confirming condition is violated. Construct a new belief model \hat{V} that is identical to the original one, except that all versions that do not satisfy the self-confirming condition in the original belief model are eliminated and each version's conjecture is arbitrary. Specify the same actual versions as in the original model (such versions are not eliminated because of condition (3)). By definition, v^* generates π^* . Since we did not change π_i^k and $\mu^{i,k}$ for each v_i^k that still exists in \hat{V}_i for all player i , the best response condition still holds for all $v_i^k \in \hat{V}_i$. Finally, condition (3') holds by construction. \square

Online Supplementary Appendix to: Rationalizable Partition-Confirmed Equilibrium

Drew Fudenberg and Yuichiro Kamada*

August 11, 2011

Abstract

This note is an online supplementary appendix to “Rationalizable Partition-Confirmed Equilibrium.” We discuss additional examples (Section 1 of this note) and an extension of the main solution concept (Section 2).

*Fudenberg: Department of Economics, Harvard University, Cambridge, MA 02138, e-mail: dfudenberg@harvard.edu; Kamada: Department of Economics, Harvard University, Cambridge, MA 02138, e-mail: ykamada@fas.harvard.edu.

1 Additional Examples

In this appendix, we provide three examples that are omitted in the main text.

Example 15 (Complication in the Definition of Best Response).




Figure 15

Consider the game depicted in Figure 15. This game is similar to a part of Example 9, but now player 2 has an option to play Out_2 , which ends the game. Specifically, player 2 moves first, choosing between L_2 , R_2 , and Out_2 . Playing Out_2 terminates the game, while playing the other actions lead to player 1's information set, where 1 does not know which of L_2 and R_2 are played. 1 has three actions, A , B , and Out_1 . Out_1 terminates the game, while A and B lead to player 2's information set, where 3 does not know which of A and B are played. 3 chooses between L_3 and R_3 , both of which end the game. The terminal node partitions are discrete.

We first show that player 2 can play Out_2 in RPCE. Notice that this requires 2 to believe 1 to play Out_1 with a high probability. Consider the following belief model:

$$V_1 = \{v_1^1\}, \quad v_1^1 =$$

$$(Out_1, q(a^{1,1}(x_L) = 1, (Out_2, L_3)) + (1-q)(a^{1,1}(x_R) = 1, (Out_2, R_3)), p^{1,1}(v_2^1, v_3^1) = p^{1,1}(v_2^2, v_3^2) = \frac{1}{2});$$

$$V_2 = \{v_2^1\}, \quad v_2^1 = (Out_2, (Out_1, L_3), p^{2,1}(v_1^1, v_3^1) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (L_3, (a^{3,1}(L_2) = 1, (Out_1, Out_2)), p^{3,1}(v_1^1, v_2^1) = 1),$$

$$v_3^2 = (R_3, (a^{3,2}(R_2) = 1, (Out_1, Out_2))), p^{3,2}(v_1^1, v_2^1) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) ,

where $q = \frac{1}{2}$, and $a^{3,1}(L_2)$ and $a^{3,2}(R_2)$ denote the probabilities that 3's assessment assigns to the nodes corresponding to actions L_2 and R_2 , respectively. It is clear from inspection that all the conditions in the definition of RPCE hold.

In this belief model, player 1's information set lies off the path of play. Hence, to calculate the best response for player 1, there must be a way to specify the weights for each assessment-strategy pair, which is given by q . This is why we need a sequence of beliefs in the definition of best response in Appendix A of the paper. In this example, we consider a sequence of beliefs $\{\mu^m\}_{m=1}^\infty$ such that μ^m assigns probability $\frac{1}{2}$ to $(a^{1,1}(x_L) = 1, ((1 - \frac{\epsilon}{m})Out_2 + \frac{\epsilon}{m}L_2, L_3))$ and the remaining probability to $(a^{1,1}(x_R) = 1, ((1 - \frac{\epsilon}{m})Out_2 + \frac{\epsilon}{m}R_2, R_3))$.

Example 16 (Reason (iv) without Relevant Ties).




Figure 16

Consider the game depicted in Figure 16. There are four players, 1, 2, 3, and 4. Players $i = 1, 2$ move first, to choose between In_i and Out_i . The game ends except when (Out_1, In_2) is chosen. If it is chosen, then players $i = 3, 4$ move simultaneously, choosing between L_i and R_i . The terminal node partitions are such that players observe the exact consequence of the opponents' actions unless they do not move on the path of play or they play Out_i , while in such cases they do not observe the exact consequence of the opponents' play.¹

Note that this example is the same as Example 12, except that now (Out_1, In_2) is followed by a simultaneous-move game by players 3 and 4. In Example 12, it is important that there is a tie in payoffs. The point of this example is to show that reason (iv) does not hinge on the ties in payoffs.

We first show that player 1 can play In_1 in an RPCE. To see this, consider the following belief model:

$$V_1 = \{v_1^1, v_1^2\}, \quad v_1^1 = (In_1, Out_2, \cdot, \cdot), p^{1,1}(v_2^1, \cdot, \cdot) = 1, v_1^2 = (Out_1, (In_2, L_3, L_4), p^{1,2}(v_2^2, v_3^1, v_4^1) = 1);$$

$$V_2 = \{v_2^1, v_2^2\},$$

$$v_2^1 = (Out_2, (Out_1, R_3, R_4), p^{2,1}(v_3^2, v_4^2) = 1), v_2^2 = (In_2, (Out_1, L_3, L_4), p^{2,2}(v_3^1, v_4^1) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (L_3, (Out_1, In_2, L_4), p^{3,1}(v_4^1) = 1), v_3^2 = (R_3, (\cdot, \cdot, R_4)), p^{3,2}(\cdot, \cdot, v_4^2) = 1);$$

$$V_4 = \{v_4^1, v_4^2\}, \quad v_4^1 = (L_4, (Out_1, In_2, L_3), p^{4,1}(v_3^1) = 1), v_4^2 = (R_4, (\cdot, \cdot, R_3)), p^{4,2}(\cdot, \cdot, v_3^2) = 1);$$

The actual version profile is $(v_1^1, v_2^1, \cdot, \cdot)$.

It is clear from inspection that all the conditions in the definition of RPCE hold. Notice that version v_1^1 believes that version v_2^1 is incorrectly believing that player 1's actual version is v_1^2 , who, from the viewpoint of v_2^1 , incorrectly believes that player 2 is version v_2^2 .

Now we show that In_1 cannot be played with probability 1 if player 1's terminal node partition is discrete. To see this, suppose to the contrary that player 1 plays In_1 with probability 1. If 2 plays In_2 with positive probability, then by the self-confirming condition, 2 expects 1 to play In_1 with probability 1, so by the best response condition 2 should play In_2 with probability 1. But then player 1 is not best-responding, so player 2 must play Out_2 with probability 1, and from observational consistency when 1's partition is discrete, 2 must expect 1 to believe 2 plays Out_2 with probability 1. Since 1's best

¹The assumption that 3 and 4 do not observe the opponents' play when they do not get to play simplifies the belief model but is not essential to the result.

response to such a belief is to play In_1 with probability 1, player 2 must believe 1 plays In_1 by the coherent belief condition. But the best response to such a belief cannot assign a positive probability to Out_2 . This means that player 1 cannot play In_1 with probability 1. \square

Example 17 (Reason (iii) with Discrete Terminal Node Partitions for Opponents).




Figure 17

Consider the game depicted in Figure 17. There are three players, 1, 2, and 3. Players 1 and 2 move first, where 1 chooses between A , B , and Out , and 2 chooses between L_2 and R_2 . If R_2 is chosen, the game ends. If L_2 is chosen, player 3 gets his move, choosing between L_3 and R_3 . The terminal node partitions are such that everyone observes the exact terminal node reached, except that player 1 does not observe the consequence of the opponents' actions if she plays Out .

We first show that player 1 can play Out in an RPCE. To see this, consider the following belief model:

$$V_1 = \{v_1^1\}, \quad v_1^1 = (Out, (\frac{1}{2}L_2 + \frac{1}{2}R_2, L_3), p^{1,1}(v_2^1, v_3^1) = p^{1,1}(v_2^2, v_3^1) = \frac{1}{2});$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (L_2, (Out_1, L_3), p^{2,1}(v_1^1, v_3^1) = 1), v_2^2 = (R_2, (Out_1, R_3), p^{2,2}(v_1^1, v_3^2) = 1);$$

$$V_3 = \{v_3^1, v_3^2\}, \quad v_3^1 = (L_3, (Out_1, L_2), p^{3,1}(v_1^1, v_2^1) = 1), v_3^2 = (R_3, (Out_1, R_2), p^{3,2}(v_1^1, v_2^2) = 1);$$

The actual version profile is (v_1^1, v_2^1, v_3^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold.

Notice that player 1 is unsure which action player 2 is going to play, which leads 1 to play *Out*. Even though the terminal node partition for player 2 is discrete, 2 does not observe player 3's action if 1 plays *Out*, hence 2 can play either L_2 with probability 1 or R_2 with probability 1, while he cannot play a mixed strategy unless he is exactly indifferent.

Now we show that, if player 1's terminal node partition is discrete, player 1 cannot play *Out*. To see this, suppose to the contrary that player 1 plays *Out*. Then, for *Out* to be at least as good as A , the probability of L_2 must be no greater than $\frac{2}{3}$. Similarly, for *Out* to be at least as good as B , the probability of R_2 must be no greater than $\frac{2}{3}$. These two imply that player 2 must be playing a mixed strategy, and hence, as player 2 observes player 3's play, the payoff from playing L_2 must be exactly equal to 0. This means that player 3 is playing R_3 with probability $\frac{1}{2}$.² So the probability of (L_2, R_3) must be no less than $\frac{1}{3} \times \frac{1}{2} = \frac{1}{6}$. However, if this is the case, either the payoff from playing A or that from playing B is strictly greater than $\frac{1}{6} \times 10 + \frac{5}{6}(-2) = 0$, so player 1 cannot play *Out*.

Remember that in Example 11 (which illustrates reason (iii)), it was important that the opponent of the player whose terminal node partitions are changed has nondiscrete terminal node partition. This example shows that reason (iii) can be present even if the opponents' terminal node partitions are discrete. The key is that what the opponent j observes depends on j 's action. \square

Example 18 (Measurability of Payoffs with Respect to Terminal Node Partition).

²In this example player 3 is indifferent among any terminal nodes. We can avoid ties by replacing player 3's move by, say, a simultaneous-move coordination game by players 3 and 4.

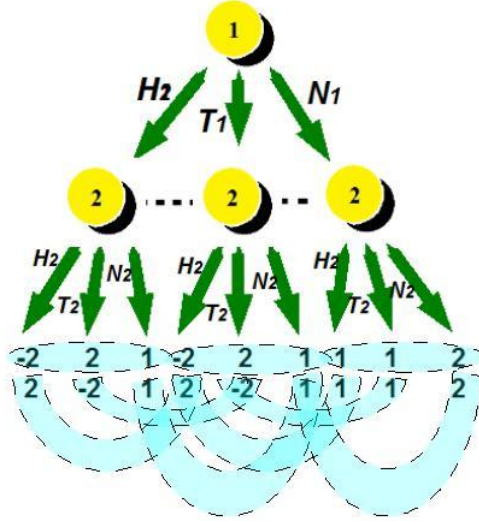


Figure 18

In this example we point out that the conclusion of Theorem 2 may fail if payoffs are not measurable with respect to terminal node partitions.

Consider the game depicted in Figure 18. Players $i = 1, 2$ move simultaneously, choosing between H_i , T_i , and N_i . Notice that this is the game played by players 2 and 3 in Example 3. The terminal node partitions are such that each player does not observe anything about the consequence of the opponent's play. This implies that the payoffs are not measurable with respect to terminal node partitions. For example, player 1 with H_1 gets 2 if player 2 plays H_2 and -2 if player 2 plays T_2 , while she does not observe the consequence of 2's action.

Since no one gets to see anything, RPCE allows for all rationalizable strategies. In particular, (H_1, H_2) is an RPCE, with the following belief model:

$$V_1 = \{v_1^1, v_1^2\}, \quad v_1^1 = (H_1, H_2, p^{1,1}(v_2^1) = 1), v_1^2 = (T_1, T_2, p^{1,2}(v_2^2) = 1);$$

$$V_2 = \{v_2^1, v_2^2\}, \quad v_2^1 = (H_2, T_1, p^{2,1}(v_1^2) = 1), v_2^2 = (T_2, H_1, p^{2,2}(v_1^1) = 1);$$

The actual version profile is (v_1^1, v_2^1) .

It is clear from inspection that all the conditions in the definition of RPCE hold.

Now consider the case where we modify player 1's terminal node partition to the discrete partition (with player 2's terminal node partition held fixed). Note that in the above belief model, player 2 who plays H_2 believes that 1 is incorrectly believing that 2 is playing T_2 . However, when player 1's terminal node partition is discrete, that belief

model cannot satisfy observational consistency of player 2, as now player 2 can no longer believe that 1 has an incorrect belief.

Moreover, H_1 cannot be played in any RPCE. To see this, note that player 2 must expect that player 1 is best responding to whatever 2 plays. Since 2 has to play a best response to his belief, this means that 2 has to play a Nash equilibrium strategy in any RPCE. However there is only one Nash equilibrium action, which is N_2 . As a unique best response to N_2 is N_1 , the only possible 1's RPCE strategy is to play N_1 with probability 1. Hence H_1 cannot be played in any RPCE. \square

2 Correlated Equilibrium

In the main sections of this paper we considered a model in which it is common knowledge that the play has converged. However, as pointed out in Example 4, this may sometimes be too strong an assumption, as under many learning dynamics play need not converge to a stationary distribution in all games. This suggests we use a weaker notion to capture the effect of long run learning; one such notion is that of correlated equilibrium.³ In this section we present a model that uses the idea of correlated equilibrium to develop a less restrictive alternative to RPCE. We use the same specification of extensive-form and terminal node partitions.

To introduce the idea of correlated equilibrium, we postulate a state space and a probability distribution over it. Specifically, consider a state space Ω with a typical element ω , and probability distribution over Ω , denoted q . For each information set of player i , h_i , there is a partition over Ω , $Q_i(h_i) = (Q_i^1(h_i), \dots, Q_i^{J_i}(h_i))$, where if h_i precedes h'_i then for each $Q_i^j(h'_i)$ there exists some $Q_i^k(h_i) \supseteq Q_i^j(h'_i)$. That is, $Q_i(h'_i)$ is a weakly finer partition than $Q_i(h_i)$. Denote $Q = (Q_1, \dots, Q_n)$. **Player i 's strategy at h_i , $\pi_i(h_i)$** , is a map from the cells of the partition $Q_i(h_i)$ to distributions over actions available at h_i . Then, **player i 's strategy** is defined to be a mapping from her information sets to her strategies at her information sets. Let Π_i be the set of all i 's strategies, and denote $\Pi = \times_i \Pi_i$ and $\Pi_{-i} = \times_j \Pi_j$. The probability distribution over terminal nodes and over cells of the partition for player i , d and D_i , are defined in an analogous manner as in Section 3 of the paper. Note that d (and hence D_i) depends on the probability distribution q . $H(\pi)$ and $H(\pi_i)$ are defined analogously to the ones in Section 3 of the paper.

³See Aumann (1987) for the definition and discussion of (normal-form) correlated equilibrium. Foster and Vohra (1997) shows that if two players use learning procedures with calibrated forecasts and play best responses to such forecasts, the time average of the frequency of play must converge to the set of correlated equilibria.

Assessment a_i is again a mapping from her information sets to her assessments at her information sets. i 's assessment at h_i , $a_i(h_i)$, is a map from the cells of the partition $Q_i(h_i)$ to distributions over nodes in h_i . Let α_i be the set of all i 's assessments.

As in Section 3 of the paper, a belief μ^i is defined as a probability distribution over $\alpha_i \times \Pi_{-i}$.

Given (Ω, \mathbf{Q}, q) , best response is defined analogously to the one in Section 3 of the paper.

We postulate a belief model again. Now it is written as $(V, (\Omega, \mathbf{Q}, q))$. Each version v_i^k in V_i is specified as

$$v_i^k = (\pi_i^k, \mu^{i,k}, p^{i,k}),$$

where $\pi_i^k \in \Pi_i$, $\mu^{i,k} \in \Delta(\alpha_i \times \Pi_{-i})$, and $p^{i,k} \in \Delta(\times_{j \neq i} V_j)$.⁴

We note that the way we model the idea of correlated equilibrium is analogous to but slightly different from the definition of *extensive-form correlated equilibrium* by Forges (1985).⁵

With this set of notation in hand, Definitions 1-3 can be used literally except that we replace ‘‘Given a belief model V ’’ with ‘‘Given a belief model $(V, (\Omega, \mathbf{Q}, q))$.’’

Definitions of partition-confirmed equilibrium and rationalizable partition-confirmed equilibrium can be used literally except that we replace ‘‘there exist a belief model V ’’ with ‘‘there exist a belief model $(V, (\Omega, \mathbf{Q}, q))$.’’

The Shapley Example Revisited.

Consider the game discussed in Example 4. In Example 4, we concluded that player 1 has to expect the equilibrium play by players 2 and 3, so 1 should play *In*. However,

⁴Here we follow the way we modeled correlated beliefs in Section 3 of the paper. We note that there is an alternative way to model correlated beliefs in the current context. Namely, to have an enlarged state space. Then the belief simply becomes a point in $\alpha_i \times \Pi_{-i}$ and a conjecture becomes a point in $\times_{j \neq i} \Delta(V_j)$. For example, suppose that there are three players, 1, 2, and 3, and two states, ω and ω' . Player 1 does not observe the state while 2 and 3 do. Then 1's belief about 2 and 3's play may be ‘‘subjectively correlated’’ if 1 believes that 2 and 3 play one action (denote L) given ω and another action (denote R) given ω' . Strictly speaking, the correlation that we considered in the main text and the correlation considered in this footnote are different: In the former, 1 believes that 2 and 3 play L with probability 1, or R with probability 1. In the latter, 1 believes that 2 and 3 both play L with probability a half, and R with probability a half. We employed the current specification to make the comparison with the model in Section 3 of the paper transparent.

⁵One small difference worth noting is that in Forges the partition $Q_i(h_i)$ is weakly finer than partition $Q_j(h_j)$ even for two different players i and j if h_j precedes h_i , and hence the notion was defined only for games in which information sets are ordered by precedence. In our case this requirement is imposed only for partitions associated with the same player i , because in our context a player who does not observe the opponents' action profile may move after their moves, but in such a case we do not want the player to observe the realization of the information that the opponents obtain. As a consequence of this generalization, our concept is defined also for games in which information sets are not ordered by precedence.

in the solution concept that we defined in this section, this is not necessary. To see this, notice that there is a correlated equilibrium of the simultaneous-move game by players 2 and 3, which assigns an equal weight to each of the profiles (A_2, B_3) , (A_2, C_3) , (B_2, A_3) , (B_2, C_3) , (C_2, A_3) , and (C_2, B_3) . If player 1 expects that 2 and 3's play is as in this correlated equilibrium, she expects the payoff of -1 from playing *In*, so she chooses *Out* which gives her the payoff of 0. Formally, consider the state space $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$ and partitions $\mathbf{Q} = (\mathbf{Q}_1, \mathbf{Q}_2, \mathbf{Q}_3)$ such that $\mathbf{Q}_1(h_1) = (Q_1^1(h_1))$, $\mathbf{Q}_2(h_2) = (Q_2^1(h_2), Q_2^2(h_2), Q_2^3(h_2))$, and $\mathbf{Q}_3(h_3) = (Q_3^1(h_3), Q_3^2(h_3), Q_3^3(h_3))$, where

$$Q_1^1(h_1) = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\};$$

$$Q_2^1(h_2) = \{\omega_1, \omega_2\}, \quad Q_2^2(h_2) = \{\omega_3, \omega_4\}, \quad \text{and} \quad Q_2^3(h_2) = \{\omega_5, \omega_6\};$$

$$Q_3^1(h_3) = \{\omega_3, \omega_5\}, \quad Q_3^2(h_3) = \{\omega_1, \omega_6\}, \quad \text{and} \quad Q_3^3(h_3) = \{\omega_2, \omega_4\},$$

and a distribution $q(\omega_i) = \frac{1}{6}$ for all $i = 1, \dots, 6$.

The belief model $(V, (\Omega, \mathbf{Q}, q))$ is

$$V_1 = \{v_1^1\}, \quad V_2 = \{v_2^1\}, \quad \text{and} \quad V_3 = \{v_3^1\},$$

where the own strategies are

$$\pi_1^1(h_1)(Q_1^1(h_1))(Out) = 1;$$

$$\pi_2^1(h_2)(Q_2^1(h_2))(A_2) = 1, \quad \pi_2^1(h_2)(Q_2^2(h_2))(B_2) = 1, \quad \text{and} \quad \pi_2^1(h_2)(Q_2^3(h_2))(C_2) = 1;$$

$$\pi_3^1(h_3)(Q_3^1(h_3))(A_3) = 1, \quad \pi_3^1(h_3)(Q_3^2(h_3))(B_3) = 1, \quad \text{and} \quad \pi_3^1(h_3)(Q_3^3(h_3))(C_3) = 1,$$

the beliefs are

$$\mu^{1,1}(\cdot, (\pi_2^1, \pi_3^1)) = 1, \quad \mu^{2,1}(\cdot, (\pi_1^1, \pi_3^1)) = 1 \quad \text{and} \quad \mu^{3,1}(\cdot, (\pi_1^1, \pi_2^1)) = 1,$$

and the conjectures are

$$p^{1,1}(v_2^1, v_3^1) = 1, \quad p^{2,1}(v_1^1, v_3^1) = 1, \quad \text{and} \quad p^{3,1}(v_1^1, v_2^1) = 1.$$

Notice that, in this belief model, if player $i = 2, 3$ finds himself in the partition cell Q_i^1 (resp. Q_i^2 ; resp. Q_i^3), he plays action A_i (resp. B_i ; resp. C_i). Each state corresponds to an action profile in the support of correlated equilibrium described above, and the partition implies that at each cell a player knows that one of the two states have realized

but does not know which. By exactly the same reasoning as the one that proves the profile is indeed a correlated equilibrium, all versions in the above belief model satisfy the best response condition. It is clear from inspection that all other conditions in the (modified) definition of RPCE defined in this section hold, and hence player 1 can play *Out* in (modified) RPCE. \square

Although the specification in this section resolves the issue associated with the Shapley cycle, it allows for perhaps too many outcomes. For example in a 2-player 2×2 coordination game, there is a correlated equilibrium that assigns probability $\frac{1}{2}$ to each of pure equilibria. Our formulation allows for such a distribution of play as a solution, but this does not seem to be a sensible outcome of rational learning process, as typically the dynamic converges to one of these equilibria.