Rationalizable Partition-Confirmed Equilibrium

Drew Fudenberg and Yuichiro Kamada*

First Version: January 29, 2011; This Version: July 30, 2014

Abstract

Rationalizable partition-confirmed equilibrium (RPCE) describes the steady state outcomes of rational learning in extensive-form games, when rationality is 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, 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: Haas School of Business, University of California, Berkeley, Berkeley, CA 94720, e-mail: ykamada@haas.berkeley.edu; We are grateful to Pierpaolo Battigalli, Eddie Dekel, Ignacio Esponda, Faruk Gul, Yuhta Ishii, David Kreps, David K. Levine, Alex Peysakhovich, Jérôme Renault, Tomasz Strzalecki, the two referees, and seminar participants at Harvard University, Princeton University, Tel Aviv University, Toulouse School of Economics, and University of Tokyo for helpful comments. NSF grants SES-0646816 and SES-0951462 provided financial assistance.

1 Introduction

Most applications of game theory suppose that the observed outcomes will correspond to equilibria, so it is important to consider which sorts of equilibrium concepts are applicable to various situations. The most compelling general explanation for equilibrium is that it arises as the long-run outcome of some sort of non-equilibrium process of learning and adjustment. If the game in question is simply one round of simultaneous moves, and participants observe the outcome each time the game is played, then if play converges we expect the long-run outcomes to correspond to Nash equilibria.[1] However, when the game has a non-trivial extensive form, observed play need not reveal the actions that would be taken at information sets that have never been reached, so even if play converges incorrect beliefs may persist, and play may not converge to a Nash equilibrium.[2]

Self-confirming equilibrium (SCE) formalizes the idea that incorrect off-path beliefs can persist for settings 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 equilibrium distribution on terminal nodes. However, because SCE places no *a priori* restrictions on the players' beliefs, it does not capture the idea that players use prior information about opponents' payoff functions to predict the opponents' play. To capture such predictions, Dekel, Fudenberg, and Levine (1999) (hereafter "DFL") define "rationalizable self-confirming equilibrium,"or "RSCE," which requires that players make certain inferences based on their knowledge of the other players' payoff functions and observation structure. For example, RSCE requires that player 1's conjecture about how player 2 thinks player 3 is playing be consistent with player 1's information about what player 2 observes.

Both SCE and RSCE apply to situations where all participants see the realized terminal node at the end of each play of the game. In some cases, though, players do not observe the exact terminal node that is reached. For example, in a sealed-bid uniform-price k-unit auction for a good of known value, the terminal node is the entire vector of submitted bids, but agents might only observe the winning price and the identity of the

[1] For example this is true for processes that are "asymptotically myopic" and "asymptotically empirical" in the sense of Fudenberg and Kreps (1993).

[2] We do not explicitly study dynamics here, but one motivation for the solution concepts we propose here is the idea that a large population of agents play the game repeatedly, with anonymous random matching, and no strategic links between repetitions. See Fudenberg and Levine (1993b) and Fudenberg and Takahashi (2011) for examples of the sorts of learning models we have in mind, and Fudenberg and Levine (2009) for a survey of related work. As this literature shows, incorrect beliefs are more of an issue when players are relatively impatient and so have less incentive to "experiment" with off-path actions; very patient players will experiment enough to rule out non-Nash outcomes (although not necessarily enough to justify backwards induction, see Fudenberg and Levine (2006).)

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. Terminal node partitions are also natural when there are many agents in each player role: If we model each agent as a distinct player then a given agent in the role of player i need not observe the play of other agents in that role.

The rationalizable partition-confirmed equilibrium (RPCE) defined in this paper generalizes RSCE by supposing that each player has a partition over terminal nodes, and that players' beliefs are consistent with the observed distribution over the the partition but not necessarily consistent with the true distribution on terminal nodes. We should stress that both of these implicitly suppose that equilibrium play corresponds to an objective distribution; 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 and supposes that each player sees the objective distribution over the cells of their own partition. 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. 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. RPCE 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.

Before proceeding to the formal part of the paper, we provide an informal illustration of RPCE in the two extensive-form games in Figure 1 (Example 1). 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. Player 1's payoffs are the amount that player 2 gets plus an "extra" of 0.1, if player 1 plays *In*. 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.

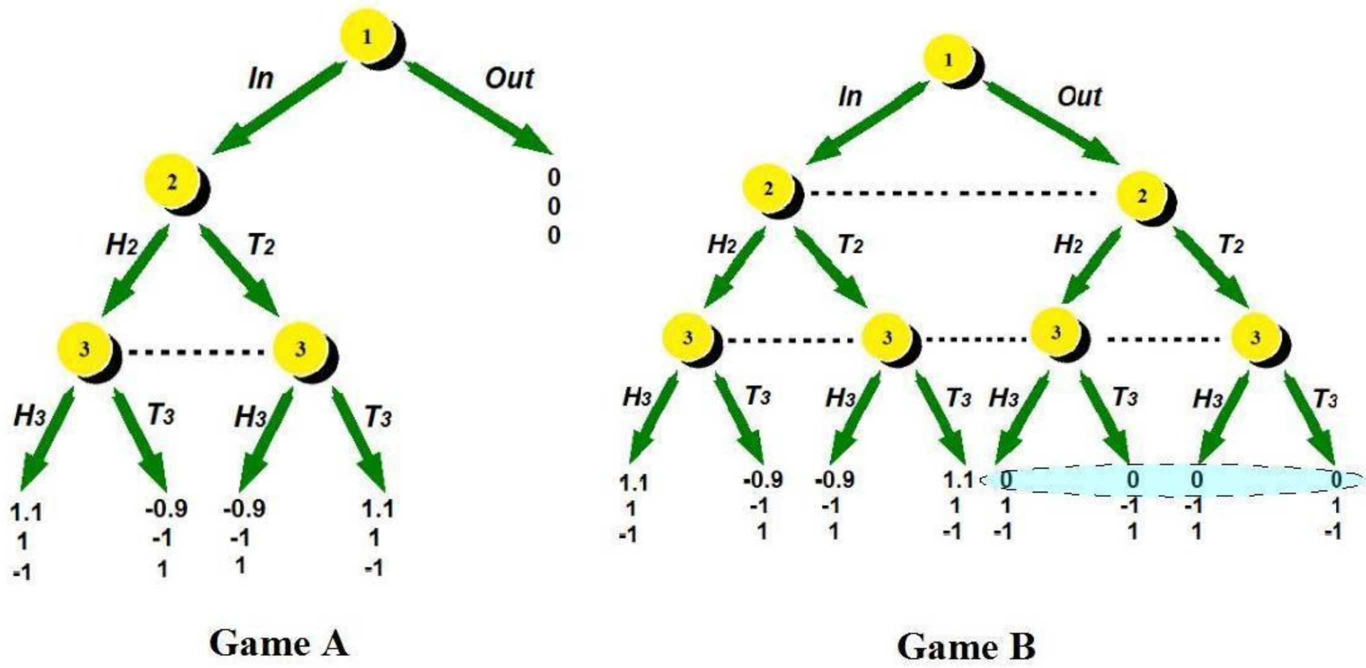

Figure 1: The dots connecting payoffs denote terminal node partitions.

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. 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. 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.

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 what follows we give a formal definition of RPCE, and provide results to show that RPCE behaves as expected and to relate it to past work, but much of our contribution comes from examples that illustrate various implications of RPCE. Many (but not all) of these examples use what we call "participation games;" we explore the impact of different terminal node partitions in these games, and also compare them to closely related games with sequential moves. The distinguishing feature of participation games is that some 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, 1's belief about their play must be concentrated on the set of Nash equilibria of the "subgame" between them. In 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 she 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 unobserved play of their opponents, as advocated by Fudenberg and Kreps (1988). As we argue in Example 8, terminal node partitions make the restriction to independent beliefs less compelling, even as a simplifying assumption: 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.

Hahn's (1977) conjectural equilibrium is a forerunner of SCE in a specific setting, as it allows firms to misperceive demand at out-of-equilibrium prices. Battigalli (1987) defines what we call self-confirming equilibrium with independent, unitary beliefs, where "unitary" means that every action in the support of a player's mixed strategy is a best response to the same belief about play of the opponents, and "independent" means that each player's subjective uncertainty about the play of the others corresponds to a product distribution. Fudenberg and Kreps (1988) give the first example where this sort of SCE has an outcome that cannot arise in Nash equilibrium. In the large-population learning models used to provide foundations for SCE, it is natural (though not necessary) to allow different agents to have different beliefs. The general definition of SCE, due to Fudenberg and Levine (1993a), allows beliefs to be heterogeneous as well as correlated.[3]

[3]Kalai and Lehrer (1993) give a version that corresponds to independent, unitary beliefs. Lehrer's (2012) "partially-specified equilibrium" is similar as it also allows players to only partially know the opponents' strategies. Ryall (2003) and Dekel, Fudenberg, and Levine (2004) develop SCE variants that in our terminology have specific sorts of partitions over terminal nodes.

Allowing for heterogenous beliefs about play when players use payoff information to make predictions is more complicated, so DFL restrict attention to unitary beliefs. This paper too restricts attention to unitary beliefs, to cut down on the number of new issues that need to be addressed at one time; note that unitary beliefs correspond to steady states of large-population learning systems when all agents in a given player role pool their information. Alternatively one can view our solution concept as providing predictions as a result of repeated interactions among a fixed set of players when the discount factor is small. In the companion paper (Fudenberg and Kamada, 2012) we allow for heterogeneous beliefs.

The paper is organized as follows. Section 2 defines a model of extensive-form games with terminal node partitions. Section 3 revisits Example 1, and analyzes other examples to show the implications of RPCE. Section 4 further motivates the RPCE definition by exploring the consequences of alternative specifications. Section 5 explains the connection between RPCE and other concepts from the literature, notably the rationalizable conjectural equilibrium (RCE) of Rubinstein and Wolinsky (1994).

2 The Model

2.1 Extensive-Form Games with Terminal Node Partitions

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

In the main text we restrict attention to "one-move games," in which for any path of play each player moves at most once, and there are no moves by Nature. In addition we assume that, for every h, h', if there is $x \in h$ and $x' \in h'$ such that $x < x'$ (where $<$ is the precedence order on nodes), then there is no $x'' \in h$ and $x''' \in h'$ such that $x''' < x''$. We then say that h' is after h if some $x' \in h'$ is after some $x \in h$, and we assume that this partial order on information sets is transitive.

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 and $\mathbf{P} = (\mathbf{P}_1, \dots, \mathbf{P}_n)$. 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 can be reached by different sequence of pure 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. 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.

Because we want to model equilibrium as an objective, steady-state distribution, while maintaining the simplicity of "unitary" beliefs (defined below) we need to allow for mixed strategies as outcomes of play. Here we adopt the simplest method, namely to let the players use mixed strategies, as in Rubinstein and Wolinsky (1993) and DFL.[4] 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 $h \in H_i$, where and subsequently, for any set X we let $\Delta(X)$ denote the set of probability distributions on X with finite support. 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. Say that an information set $h \in H_i$ is reachable under π_{-i} if there exists π_i such that h has a positive probability under (π_i, π_{-i}).

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, and 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.

2.2 Beliefs, Consistency, and Best Responses

We will impose some restrictions on beliefs about off-path play, so we will need to specify assessments at off-path information sets: Player i's assessment at $h \in H_i$ is a probability distribution over nodes in h, so that the assessment at h is an element of $\Delta(h)$. For any $h \in H_i$, i's assessment at h and her opponents' behavioral strategies π_{-i} completely determine i's expected payoff for playing any strategy π_i, conditional on h. Denote by $\mu_i \in \Delta(\Pi_{-i}) \times [\times_{h\in H_i}\Delta(\Delta(h) \times \Pi_{-i})]$ the **belief** held by player i. That is, player i's belief consists of two terms. The first is a finite-support probability distribution over the opponents' strategy profiles, and the second is a vector that specifies, at each information set h of player i, a probability distribution over the product space of pairs of the form (assessments at that information set, opponents' strategy profiles). We denote by $b(\mu_i)$ the marginal of the belief μ_i on the first coordinate, and by $(\mu_i)_h$ the marginal of μ_i on the coordinate for information set h. Note that the belief has sufficient information to calculate conditional expected payoffs at each information set.

We allow $b(\mu_i)$ to be any distribution on Π_{-i} (with finite support), as opposed to a product of independent mixed strategies. Example 8 explains why this is desirable. We

[4] In Fudenberg and Kamada (2012) we show that we can replace mixed strategies with a distribution of players each of whom uses a pure strategy.

allow assessments to be correlated with the beliefs over opponents' strategies; Example 9 explains why.

Definition 1. Belief μ_i is an **independent belief** if the following conditions hold:

1. For each $\hat{\pi}_{-i}$ in the support of $b(\mu_i)$, we require

$$b(\mu_i)(\hat{\pi}_{-i}) = \prod_{j \neq i} \left(\sum_{\pi_{-i,j} \text{ s.t. } \exists \pi_j \text{ s.t. } (\pi_j, \pi_{-i,j}) \in \text{supp}(b(\mu_i))} b(\mu_i)(\hat{\pi}_j, \pi_{-i,j}) \right).$$

2. For each h and each $(\hat{a}_i, \hat{\pi}_{-i})$ in the support of $(\mu_i)_h$, we require

$$(\mu_i)_h(\hat{a}_i, \hat{\pi}_{-i}) = \left(\sum_{\pi_{-i} \text{ s.t. } \exists a_i \text{ s.t. } (a_i, \pi_{-i}) \in \text{supp}((\mu_i)_h)} (\mu_i)_h(\hat{a}_i, \pi_{-i})) \right)$$

$$\cdot \prod_{j \neq i} \left(\sum_{(a_i, \pi_{-i,j}) \text{ s.t. } \exists \pi_j \text{ s.t. } (a_i, (\pi_j, \pi_{-i,j})) \in \text{supp}((\mu_i)_h)} (\mu_i)_h(a_i, (\hat{\pi}_j, \pi_{-i,j})) \right).$$

That is, μ_i is independent if $b(\mu_i)$ and the $(\mu_i)_h$ are all product measures. We allow i's belief to vary with i's information sets, because the posterior belief about which element in the support of $b(\mu_i)$ has been used may be different from the prior belief. Example 10 explains why such variability is desirable.

Definition 2. A belief μ_i satisfies **accordance** if it satisfies the following.

1. $(\mu_i)_h$ is derived by Bayes rule if there exists π_{-i} in the support of $b(\mu_i)$ such that h is reachable under π_{-i}.[5]

2. For all $h \in H_i$, if $(\mu_i)_h$ assigns positive probability to $\hat{\pi}_{-i}$, then there exists $\tilde{\pi}_{-i} \in \text{supp}(b(\mu_i))$ such that $\hat{\pi}_{-i}(h') = \tilde{\pi}_{-i}(h')$ for each h' after h.

[5]For each π_{-i} in the support of $b(\mu_i)$ such that h is reachable under π_{-i}, let $a(x|h, \pi_{-i})$ be the probability that node $x \in h \in H_i$ is reached conditional on the event that π_{-i} is used and h is reached. Since h is reachable under π_{-i}, this conditional probability is well-defined. If π_{-i} is in the support of $b(\mu_i)$, define

$$(\mu_i)_h(a, \pi_{-i}) = \frac{b(\mu_i)(\pi_{-i}) \cdot \text{Prob}(h|\pi_{-i})}{\sum_{\pi'_{-i} \in \text{supp}(b(\mu_i))} b(\mu_i)(\pi'_{-i}) \cdot \text{Prob}(h|\pi'_{-i})}$$

where $a = a(\cdot|h, \pi_{-i})$. Otherwise we define $(\mu_i)_h(a, \pi_{-i}) = 0$.

The first part of the definition restricts the belief at on-path information sets, and the second part does so for off-path information sets. Given our restriction to one-move games, the information sets h' referred to in part 2 belong to players who did not move before h; part 2 imposes a form of consistency between player i's "initial" beliefs $b(\mu_i)$ about what these players will do and player i's beliefs conditional on unexpectedly arriving at h.

There are other reasonable alternatives for off-path restrictions on beliefs, both to weaker conditions that allow for the sort of correlation we discuss in the next example, and to conditions that impose additional restrictions in games where some players can act multiple times on a path of play. We do not examine these alternatives here, because refining off-path beliefs is not our focus. Instead, we assume accordance throughout the paper. Accordance is an easy-to-check condition, and in particular implies that if $b(\mu_i)$ has a singleton support then $b(\mu_i)$ and $(\mu_i)_h$ coincide. For example, consider the extensive form in Figure 2.

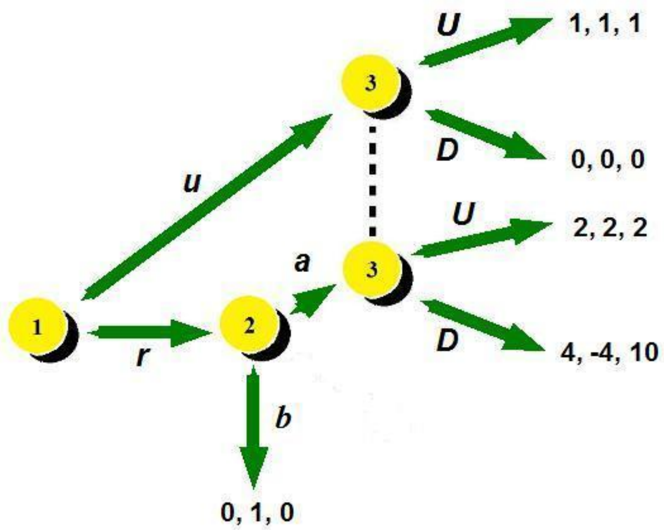

Figure 2

If $b(\mu_2)$ assigns probability 1 to (u, U) then, under accordance, player 2's belief $(\mu_2)_{h_2}$ at his information set h_2 must assign probability 1 to player 3 playing U. On the other hand, if we did not impose part 2 of Definition 2, then when player 2 unexpectedly sees player 1 play r, she could change her belief about the future play of player 3 from U to D, which would make her want to play b. This sort of change in beliefs can arise if deviations occur as the result of correlated trembles or payoff shocks. In particular, the outcome (u, U) is a c-perfect equilibrium (Fudenberg, Kreps, and Levine (1988)), because b for player 2 is a best response to the correlated distribution $((1-\epsilon)(u, U), \epsilon(r, D))$. However, this cannot occur when the trembles are required to be independent across players as in trembling-hand perfection and sequential equilibrium, and (u, U) is not a sequential equilibrium outcome. We implicitly impose "independent trembles" in the accordance condition for simplicity but as noted above alternative conditions may be reasonable as well.

The following result is immediate and is stated without a proof, and together with Theorem 2 below will establish that a RPCE exists.

Theorem 1. *Suppose that an assessment-strategy pair $(\tilde{a}, \tilde{\pi})$ satisfies Kreps and Wilson's (1982) consistency, $b(\mu_i)(\tilde{\pi}_{-i}) = 1$ for all i, and $(\mu_i)_h(\tilde{a}, \tilde{\pi}) = 1$ for each h. Then μ_i satisfies accordance.*

We say that $\pi_i \in \Pi_i$ is a **best response to a belief μ_i at $h \in H_i$** if the restriction of π_i to the subtree starting at h maximizes player i's expected payoff against $(\mu_i)_h$ in that subtree.[6]

2.3 Versions, Conjectures, and Belief Models

To facilitate comparison with DFL, we model the beliefs of the players about the beliefs and play of others- their "interactive beliefs"- in the same way as DFL, using the idea of "*versions*" v_i of each player i. Only one of these versions represents the way player i actually behaves; the other versions v_i of player i are descriptions of player i that some player j thinks is possible.[7] 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

[6] Note that for the purpose of computing this best response, the relevant part of $(\mu_i)_h$ is the distribution of play at successors of h. Formally, suppose that there are K points in the support of $(\mu_i)_h$, and index them by superscript k to write (a_i^k, π_{-i}^k). An assessment a_i at h and a strategy profile π together induce a unique probability distribution over terminal nodes, denoted $f(a_i(h), \pi))$. The restriction of a strategy π_i^* to the subtree starting at h maximizes player i's expected payoff against $(\mu_i)_h$ if

$$\pi_i^* \in \arg\max_{\pi_i \in \Pi_i} \sum_{k=1}^{K} \left((\mu_i)_h(a_i^k, \pi_{-i}^k) \cdot \sum_{z \in Z} \left[f(a_i^k, (\pi_i, \pi_{-i}^k))(z) \cdot u_i(z) \right] \right).$$

[7] An alternative approach would be to use the notion of an "epistemic structure," as in Ben-Porath (1997), Battigali and Siniscalchi (2002), and Battigali and Friedenberg (2011). That approach would facilitate comparison with some of the literature on rationalizability, but complicate the comparison with DFL.

will be slightly different, as instead of specifying beliefs we associate with each version a probability distribution over opponents' versions that we call a "conjecture." We use these conjectures below to formalize an analog of the usual belief-closed condition—the idea that the play that player i expects to see is generated by the versions he expects are present. To introduce the notion of conjectures formally, we first need to specify a profile of sets of versions.

A *belief model* is a collection $V = (V_1, \ldots, V_n)$ where each V_i is a finite set of player i's versions. In our setting, version v_i of player i is denoted by $v_i = (\pi_i, p_i)$, where the first element is version v_i's strategy $\pi_i \in \Pi_i$, and the second is her conjecture $p_i \in \Delta(\times_{j \neq i} V_j)$.

Notice that the specification of conjectures allows correlated beliefs, as otherwise, p_i must lie in the space $\times_{j \neq i} \Delta(V_j)$. We do not require that p_i assigns probability 1 to a single version profile of the opponents: 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 or v''_j.

Finally, we will associate with each V_i in a belief model an **actual version** $v_i^* \in V_i$, which is the version that is objectively present. Any other versions of player i are called **hypothetical versions**, as they exist only in the minds of the other players.

2.4 Rationalizable Partition-Confirmed Equilibrium

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 3. A belief μ_i is **coherent** with a conjecture p_i if $b(\mu_i)$ assigns probability $\sum_{\pi_{-i}(v_{-i}) = \tilde{\pi}_{-i}} p_i(v_{-i})$ to each $\tilde{\pi}_{-i} \in \Pi_{-i}$.

In the definition of RPCE, we require that all versions in a belief model have a coherent belief; this is analogous to requiring the belief model be *belief-closed*, as defined in DFL.

Definition 4. Given a belief model V, version $v_i = (\pi_i, p_i) \in V_i$ is **self-confirming with respect to** π^* if for all v_{-i} in the support of p_i, $D_i(\pi_i, \pi_{-i}(v_{-i})) = D_i(\pi_i, \pi^*_{-i})$.[8]

In the defining equality, the left hand side is the distribution over i's terminal node partition generated by version v_i's strategy and the belief about the opponents' play that is induced by v_i's conjecture. The right hand side is the distribution that version v_i observes if the actual distribution of the play is π^*. That is, this equality says that v_i's observation (the left hand side) is equal to the actual play (the right hand side).

[8]Because each version $v_i = (\pi_i, p_i)$ views $\pi_{-i}(v_{-i})$ as possible for each v_{-i} in the support of p_i, every such $\pi_{-i}(v_{-i})$ must be consistent with the version's observations.

Definition 5. Given a belief model V, $\pi^* \in \Pi$ **is generated by** a version profile $v = (v_i)_{i\in I} = (\pi_i, p_i)_{i\in I} \in \times_{j\in I} V$ if for each i, $\pi_i = \pi_i^*$.

Definition 6. Given a belief model V, $v_i = (\pi_i, p_i)$ is **observationally consistent** if $p_i(\tilde{v}_{-i}) > 0$ implies, for each $j \neq i$, $\tilde{v}_j$ is self-confirming with respect to $\pi(v_i, \tilde{v}_{-i})$.

Remark 1.

(a) If $\tilde{v}_j$ is self-confirming with respect to $\pi(v_i, \tilde{v}_{-i})$ then by Definition 4, for all $\hat{v}_{-j}$ in the support of $\tilde{p}_j$, $D_j(\pi_j(\tilde{v}_j), \pi_{-j}(\hat{v}_{-j})) = D_j(\pi_j(\tilde{v}_j), \pi_{-j}(v_i, \tilde{v}_{-i,j})) = D_j(\pi_i, \pi_{-i}(\tilde{v}_{-i}))$. Hence Definition 6 is equivalent to the following: "Given a belief model V, "$v_i = (\pi_i, p_i)$ is observationally consistent if $p_i(\tilde{v}_{-i}) > 0$ implies, for each $j \neq i$, $D_j(\pi_j(\tilde{v}_j), \pi_{-j}(\hat{v}_{-j})) = D_j(\pi_i, \pi_{-i}(\tilde{v}_{-i}))$ for all $\hat{v}_{-j}$ in the support of $\tilde{p}_j$." The left hand side in this equality is what $\tilde{v}_j$ expects to observe given his belief under the partition given by D_j. The right hand side describes what v_i thinks $\tilde{v}_j$ is observing under the partition given by D_j. Thus the equality requires that v_i believes that $\tilde{v}_j$'s belief is consistent with what $\tilde{v}_j$ observes. Thus this definition incorporates the idea that players know (i) the terminal node partitions of other players and (ii) that the opponents satisfy the self-confirming condition.

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

(c) Note that the condition in Definition 6 only need hold when v_i thinks the profile $\tilde{v}_{-i}$ has positive probability: Otherwise, v_i need not believe that $\tilde{v}_j$'s observation is consistent with her belief. Relatedly, even if v_i thinks $\tilde{v}_j$ has positive probability and $\tilde{v}_j$ thinks version v_k has positive probability, v_i's belief need not be consistent with what v_k observes. This is because v_i might think that $\tilde{v}_j$ has positive probability, and $\tilde{v}_j$ incorrectly conjectures that v_k has positive probability. Finally, if $\tilde{v}_j$ is self-confirming with respect to π^*, then in the left hand side of the equation of the alternative definition in Remark 1(a), $D_j(\pi_j(\tilde{v}_j), \pi(\hat{v}_{-j}))$ can be replaced with $D_j(\pi_j(\tilde{v}_j), \pi^*_{-j})$.

(d) The rationale for requiring observational consistency is that player i knows j's terminal node partition and knows that j's belief is consistent with what j observes. In the model developed so far this knowledge is informal. In the Online Supplementary Appendix we make this interpretation precise by constructing an epistemic model.

Definition 7. 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 conditions hold:

1. π^* is generated by v^*.
2. For each i and $v_i = (\pi_i, p_i)$, there exists μ_i such that (i) μ_i is coherent with p_i and (ii) π_i is a best response to μ_i at all $h \in H_i$.
3. For all i, v_i^* is self-confirming with respect to π^*.
4. For all i and v_i, v_i is observationally consistent.

One consequence of the definition of RPCE is that the set of RPCE shrinks if terminal node partitions become coarser.[9] Since the belief model that supports a strategy profile as a RPCE under finer partitions can support the same strategy profile under coarser partitions, the following comparative statics is immediate.

Theorem 2. *If the terminal node partitions* $\mathbf{P}$ *are coarser than* $\mathbf{P}'$ *then any strategy profile that is RPCE under partition* $\mathbf{P}'$ *is also a RPCE under* $\mathbf{P}$*.*

Theorem 1 implies that with discrete terminal node partitions for each player, every sequential equilibrium is an RPCE with a single version for each player and correct beliefs.[10] Hence all sequential equilibria are RPCE under the discrete terminal node partitions. Combining this observation with Theorem 2 and the fact that every game has a sequential equilibrium (Kreps and Wilson, 1982) yields the following result.

Corollary 1. *In any extensive-form game, a RPCE exists.*

[9]The Online Supplementary Appendix provides further examples to illustrate the effects of changes of terminal node partitions.

[10]Note that this applies even to sequential equilibria that are ruled out by strategic stability (Kohlberg and Mertens (1986)). Thus RPCE corresponds to sequential equilibrium's assumption that players view deviations as trembles, as opposed to the "forward induction" view that whenever possible deviations should be viewed as a deliberate choice.

3 Implications of RPCE

In this section we consider several examples to illustrate the implications of RPCE. One theme 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 assign probability one to actions that are not only rationalizable but also Nash. The Appendix generalizes this result to a class of "participation games." Example 4 provides an example that shows that some RPCE outcomes can only be sustained with belief models in which multiple versions of a given player play the same strategy. In that example, players 1 and 2 each has a single version, and their beliefs involve differing implicit models of the beliefs of player 4. However, there is a strong restriction on the versions that actual versions can assign positive probability: Lemma 1 in the Appendix shows that in a RPCE, any version profile to which an actual version assigns positive probability is the actual version profile of some RPCE.

Terminal node partitions have various effects on the set of strategies that a player can play in a RPCE. Example 5 demonstrates that a player need not expect unobservable play by the opponents to resemble a Nash equilibrium if their terminal node partitions are not discrete. Example 6 shows how giving a player a more refined terminal node partition can change his RPCE play even though that player's beliefs were correct in the RPCE for the coarser partition: The effect comes from the fact that with the finer partition other players know that the player's beliefs are correct. The Online Supplementary Appendix also provides examples to illustrate how the terminal node partitions change the set of strategies in RPCE.

Example 1 Revisited.

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

Consider game A, in which players 2 and 3 play matching pennies if and only if player 1 plays *In*. We argue that player 1 can play *Out* in a RPCE with the following belief model and actual versions[11]:

$$V_1 = \{v_1'\}, \quad v_1' = (Out, (v_2', v_3'));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (H_2, (v_1', v_3'')), v_2'' = (T_2, (v_1', v_3'));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (T_3, (v_1', v_2')), v_3'' = (H_3, (v_1', v_2''));$$

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

Here, v_2' can believe that player 3 plays H_3 because she never gets to observe 3's play, while v_3'' 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' 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.

Now we turn to game B, where players 2 and 3 play matching pennies regardless of player 1's action but 1 only observes their play when she chooses In. Fix a RPCE π^*, with an associated belief model V. Suppose that some version of player 1's conjecture assigns positive probability to a version profile $(\tilde{v}_2, \tilde{v}_3)$ such that $\pi(\tilde{v}_2)$ and $\pi(\tilde{v}_3)$ are not best responses to each other. Suppose without loss of generality that $\pi(\tilde{v}_2)$ is not a best response to $\pi(\tilde{v}_3)$. Notice that by the observational consistency condition, we have $D_2(\tilde{\pi}_2, \pi_{-2}(v_{-2})) = D_2(\tilde{\pi}_2, \cdot, \pi(\tilde{v}_3))$ for all v_{-2} in the support of $\tilde{p}_2$. Since player 2 observes the exact terminal node reached, this implies that $\tilde{p}_2$ assigns probability 1 to v_3 such that $\pi_3(v_3) = \pi_3(\tilde{v}_3)$. But this means that any belief $\tilde{\mu_2}$ coherent with $\tilde{p}_2$ has a property that $b(\mu_2)$ assigns probability one to $\pi_3(\tilde{v}_3)$, so the best response condition is violated for player 2.

Therefore, it must be the case that, for any $v_1 = (\pi_1, p_1)$, any belief μ_1 coherent with p_1 assigns probability $\frac{1}{2}$ to each of H_2 and H_3. The best-response condition then implies that π_1 assigns probability 1 to In, 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 of player 1 and π^* is generated by the actual versions, we conclude that π_1^* assigns probability 1 to In, that is, player 1 plays In with probability 1.

Example 2.

Consider a modification of game B, where we add "player 0" at the top of the extensive-

[11] The notation that we use when presenting belief models in examples involves a slight abuse of notation. In particular, when a player's conjecture is a point mass on a particular version profile v_{-i} we write that profile in place of the Dirac measure concentrated on v_{-i}.

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 game B of 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.

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

Example 3.

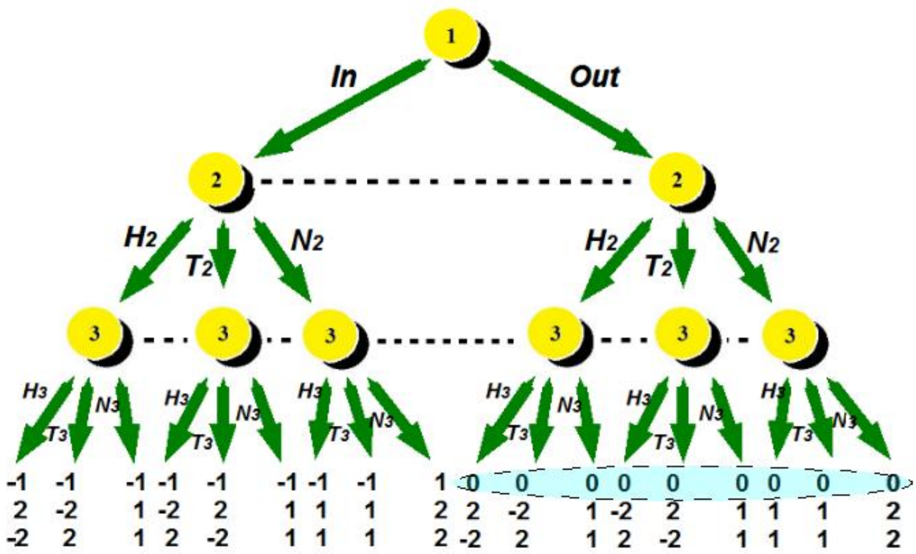

Figure 3

Consider the game in Figure 3. Everyone observes the exact terminal node reached, except that player 1 cannot distinguish between the opponents' action profiles if she plays Out.

Player 1's action has no effect on the information or payoffs of players 2 and 3, so it makes sense to talk of the subgame involving just those two players. 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 so all actions in the subgame are rationalizable, but it has a unique Nash equilibrium, namely (N_2, N_3).

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 be a Nash equilibrium of the subgame. 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 conjecture 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.

It is important here that 2 and 3 do not observe 1's action before they move, as otherwise 1 can play Out, believing that 2 and 3 play H_i or T_i after In. This example shows that in RPCE, beliefs about unobserved actions on the path of play should assign probability one to actions that are not only rationalizable but also Nash. We generalize this in Theorem 5 in the Appendix. □

Example 4 (Need for Duplicate Versions).

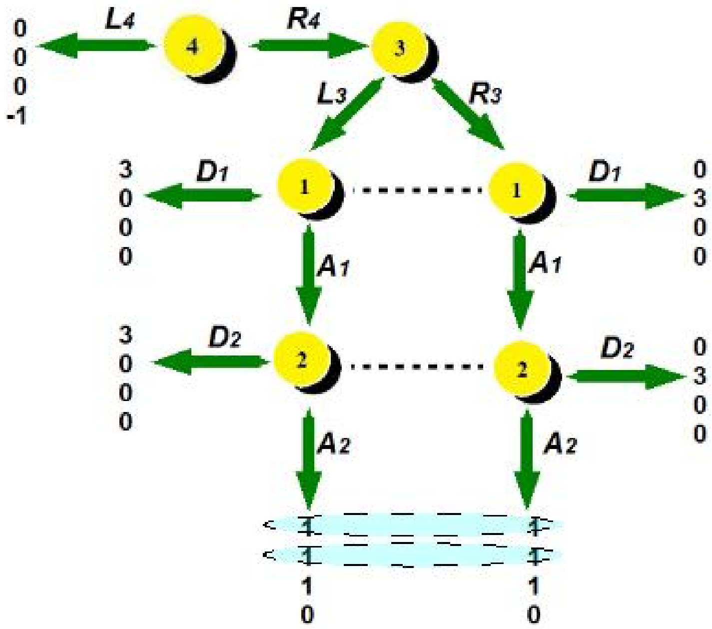

Figure 4

The game depicted in Figure 4 is a modification of the "horse" example in Fudenberg and Kreps (1988). Instead of having player 3 move only when 1 or 2 plays down, we now suppose that 3 moves whenever 4 plays a dominant action at the root node, and 1 and 2 do not know 3's play as long as both play "across." The terminal node partitions are such that everyone observes the terminal node reached, except that if (A_1, A_2) is taken then 1 and 2's partitions do not reveal 3's choice.

This game has a RPCE in which (A_1, A_2) is chosen. Specifically, consider the following belief model:

$$V_1 = \{v_1'\}, \quad v_1' = (A_1, (v_2', v_3', v_4'));$$

$$V_2 = \{v_2'\}, \quad v_2' = (A_2, (v_1', v_3'', v_4''));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (R_3, (v_1', v_2', v_4')),\ v_3'' = (L_3, (v_1', v_2', v_4''));$$

$$V_4 = \{v_4', v_4''\}, \quad v_4' = (R_4, (v_1', v_2', v_3')),\ v_4'' = (R_4, (v_1', v_2', v_3''));$$

The actual version profile is (v_1', v_2', v_3', v_4').

Notice that V_4 has two versions, both of which play the same strategy. This is a necessary feature of any belief model that supports the outcome involving (A_1, A_2). This is because this (A_1, A_2) can happen only when 1 and 2 disagree about 3's play, and know that 4 observes 3's play. This means 1 and 2 must also disagree about what 4 believes, which requires there be (at least) two versions of player 4, and both versions need to play R_4 as it is a dominant action.[12]

This need for two versions that play the same strategy is a new feature that arises with nondiscrete terminal node partitions; such duplicate versions do not enlarge the set of RSCE, because in RSCE players can only disagree about play off of the equilibrium path.[13] □

[12] A formal proof goes as follows: Suppose that there is only one version $\hat{v}_4$ in V_4, and that $\hat{v}_4$ believes that L_3 is played with probability $p \in [0,1]$. By coherency, all versions of players 1 and 2 must have a conjecture that assigns probability 1 to $\hat{v}_4$. Then observational consistency implies that all versions of players 1 and 2 must believe that L_3 is played with probability $p \in [0,1]$. But since $p > \frac{1}{3}$ implies that D_1 is strictly better than A_1 and $p < \frac{2}{3}$ implies that D_2 is strictly better than A_2, (A_1, A_2) cannot be played.

[13] Fix a belief model used to justify a RSCE π^* in the DFL model, and suppose that it has m versions $(v_i^{(1)}, \ldots, v_i^{(m)})$ that use the same strategy in a single player role i. Now consider a new belief model formed by eliminating $(v_i^{(2)}, \ldots, v_i^{(m)})$. If a version of some opponent player role j has a mixture over $(v_i^{(2)}, \ldots, v_i^{(m)})$ in DFL's belief-closed condition, then the belief-closed condition will be satisfied in the new belief model by assigning the sum of probabilities on $(v_i^{(1)}, \ldots, v_i^{(m)})$ in the original mixture to $v_i^{(1)}$ in the new mixture. Thus the new belief model supports the RPCE π^*.

Example 5 (Participation Game with Unobservable Actions).

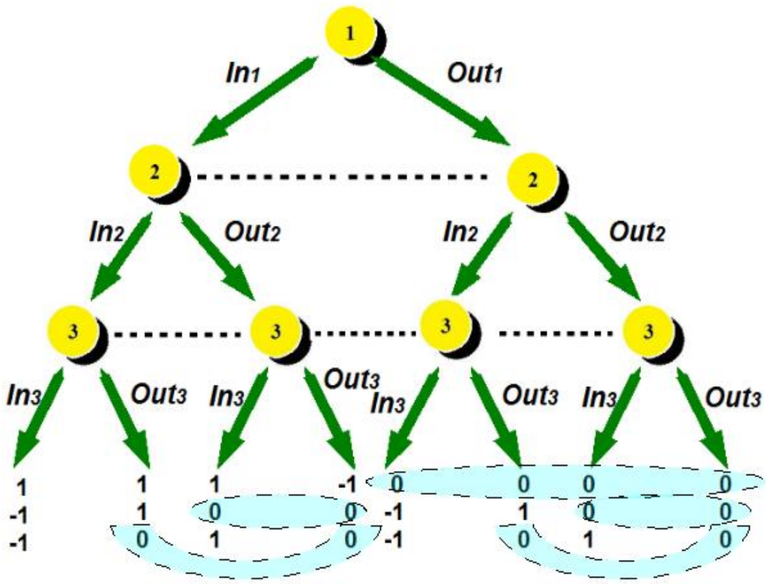

Figure 5

Consider the game in Figure 5. Player 1 does not observe the exact terminal node if she plays Out_1, and she observes the exact terminal node reached if she plays In_1. The other players' terminal node partitions always reveal 1's move but only reveal the exact terminal node if they play 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 play a Nash equilibrium of the subgame, 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:

$$V_1 = \{v_1'\}, \quad v_1' = (Out_1, (v_2', v_3'));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (Out_2, (v_1', v_3'')), v_2'' = (In_2, (v_1', v_3'));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (Out_3, (v_1', v_2'')), v_3'' = (In_3, (v_1', v_2'));$$

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

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.

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 player 1's payoff is $\frac{1}{2}$ in every Nash equilibrium of the game between players 2 and 3, so by Theorem 5 in the Appendix 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. In other words, the conclusion of Theorem 5 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 to be 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 otherwise. □

Example 6 (**Terminal Node Partitions and Learning One Player's Actions from those of Another**).

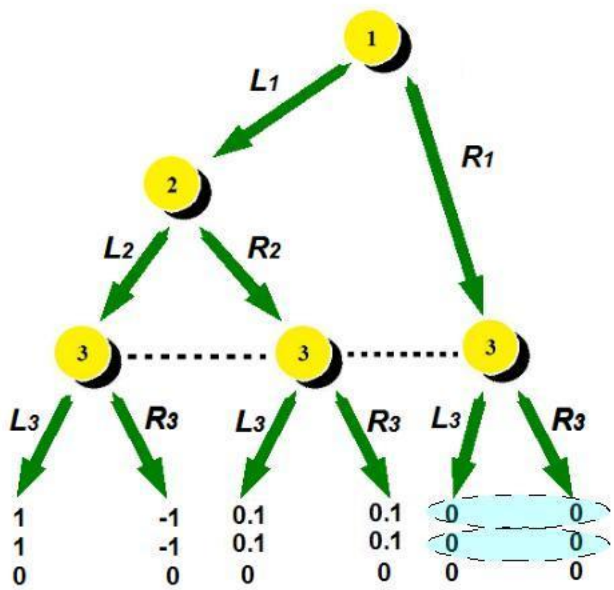

Figure 6

Here we provide an example in which the difference in terminal node partitions affects learning. If one player's terminal node partition is finer than another player's then the latter can learn by observing the play of the former and can respond accordingly, while if the partitions are the same then there is nothing to learn.

In the game in Figure 6, all players observe the exact terminal node reached, except that player 1's and 2's partitions do not reveal 3's action if 1 plays R_1.

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

$$V_1 = \{v_1'\}, \quad v_1' = (R_1, (v_2', v_3'));$$

$$V_2 = \{v_2'\}, \quad v_2' = (L_2, (v_1', v_3''));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (R_3, (v_1', v_2')), v_3'' = (L_3, (v_1', v_2'));$$

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

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 we show that if player 1's partition is discrete, she can no longer play R_1 in a RPCE. In such a situation, because player 1 has a discrete terminal node partition and player 2 does not, player 2 can learn 3's play by observing 1's play. For this reason, player 1 cannot play R_1 in a RPCE, while 1 could play R_1 if players are not required to believe that other players act rationally.[14] To see that R_1 cannot be played, suppose the contrary. The best response condition for player 1 and observational consistency applied to player 2 imply that $b(\mu_2)$ assigns probability at least $\frac{1}{2}$ to R_3. By accordance, player 2's belief $(\mu_2)_{h_2}$ at his information set h_2 assigns... probability at least $\frac{1}{2}$ to R_3. Then the best response for player 2 is to play R_2 with probability 1. However, this implies that 1's payoff from playing L_1 is $0.1 > 0$, so she cannot play R_1. On the other hand, if player 1 does not know 2's payoff function, the fact that 2's behavior reflects her belief about 3's play doesn't convey any information to player 1. So 1 can believe (L_2, R_3) is played with probability 1, making R_1 possible. The key is the observational consistency condition: player 2 knows player 1 observes 3's play, so 2's belief about 3's play must match with what 2 thinks 1 is best-responding against.

Notice that player 1's belief in the RPCE we constructed for the original terminal node partitions is in fact correct. However, when 1's terminal node partition is discrete, 1 can no longer play R_1: With a discrete terminal node partition for player 1, player 1 knows 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. □

[14]The Online Supplementary Appendix develops the concept of "partition-confirmed equilibrium" or PCE, which extends SCE to games with non-discrete terminal node partitions. Roughly speaking PCE weakens condition (2) of RPCE and also drops the coherency condition in condition (2).

4 Justification of the RPCE Definition

In this section we discuss several examples of a game and a RPCE outcome that we think is a plausible consequence of rational learning, and study whether the outcome would still be a RPCE under alternative definitions that might seem natural to some readers.

Specifically, Example 7 explains why the self-confirming condition should not be imposed on hypothetical versions, Example 8 argues that we should allow for correlated beliefs in our model, Example 9 justifies our specification of the space of beliefs, and Example 10 discusses the role of accordance.

Example 7 (Self-Confirming Condition for Hypothetical Versions).

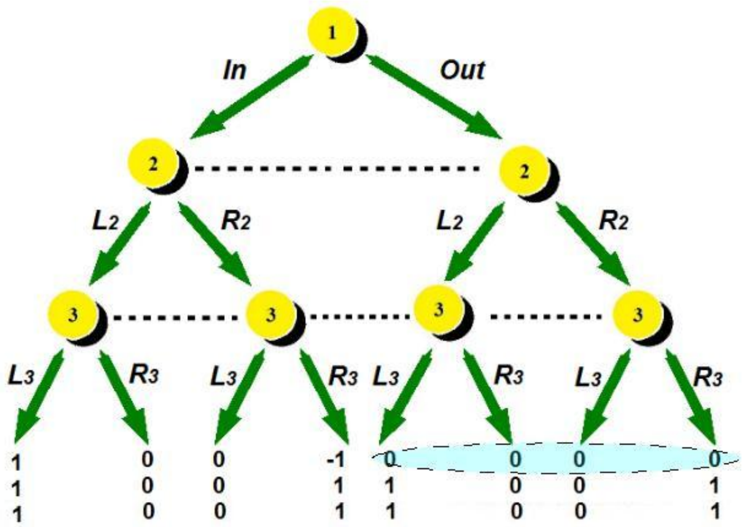

Figure 7

Consider the game in Figure 7. 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.

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 how 2 and 3 play. Indeed, the outcome (Out, L_2, L_3) is possible in RPCE. To see this, consider the belief model and actual versions:

$$V_1 = \{v_1'\}, \quad v_1' = (Out, (v_2'', v_3''));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (L_2, (v_1', v_3')), \ v_2'' = (R_2, (v_1', v_3''));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (L_3, (v_1', v_2')), \ v_3'' = (R_3, (v_1', v_2''));$$

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

Notice v_2'' and v_3'' are hypothetical versions, and they do not satisfy the self-confirming condition. Version v_1' plays Out because she conjectures that these hypothetical versions exist, and her conjecture is never falsified because she plays Out.

Now we show that the outcome (Out, L_2, L_3) is impossible if we require the self-confirming condition with respect to the equilibrium strategy profile for hypothetical versions. To see this, suppose that we strengthen Definition 7 by replacing condition (3) with the condition that for all i and v_i, v_i is self-confirming with respect to π^*. If (Out, L_2, L_3) is a RPCE under this stronger condition, the best response condition implies that all versions of player 2 should play L_2 and that all versions of player 3 should play L_3, so player 1 must believe that players 2 and 3 play (L_2, L_3). But then by the best response condition player 1 must play In.[15] $\square$

Example 8 (Correlated Beliefs).

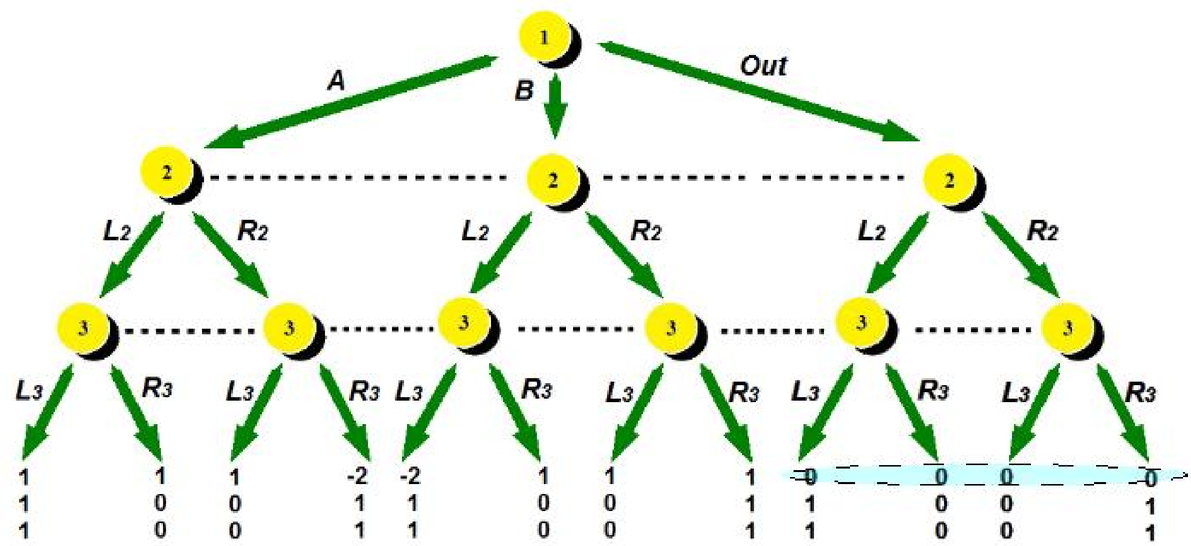

Figure 8

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 player

[15]Without knowledge of opponents' payoff functions (as in the partition-confirmed equilibrium concept in the Online Supplementary Appendix), 1 may still play Out, believing that 2 and 3 play R_2 and R_3, respectively.

1 can play an action only when she has correlated beliefs about the play at information sets that she does not observe.

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; her decision amounts to either betting on the action that players 2 and 3 will coordinate on, or declining to bet. 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), and each is equally likely.[16] 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* is optimal, as it leads to the payoff of 0.

Player 1 can play *Out* in RPCE as shown by the system[17]:

$$V_1 = \{v_1'\}, \quad v_1' = (Out, (\frac{1}{2}(v_2', v_3'), \frac{1}{2}(v_2'', v_3'')));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (L_2, (v_1', v_3')),\ v_2'' = (R_2, (v_1', v_3''));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (L_3, (v_1', v_2')),\ v_3'' = (R_3, (v_1', v_2''));$$

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

Just as with SCE, one can refine the set of RPCE by requiring independent beliefs. In some cases this might be viewed as an innocuous simplifying assumption, but we think the restriction would be problematic here, because the fact that players 2 and 3 observe each other's play means that the extensive form and terminal node partitions provide them with a particular sort of correlating device.[18]

[16]Note that with this belief, player 1 is certain that the actual play of 2 and 3 is deterministic and hence independent; the correlation here is in player 1's subjective uncertainty about which pure strategies players 2 and 3 are using. This is the same sort of subjective correlation that SCE allows for in beliefs about off-path play.

[17]Version v_i's conjecture $(pv'_{-i}, (1-p)v''_{-i})$ denotes the probability distribution with probability p on v'_{-i} and $1-p$ on v''_{-i}.

[18]If players 2 and 3 have trivial terminal node partitions (and so do not observe their own ex-post payoffs) then there is no reason for player 1 to think their play has converged. In this case too RPCE would allow player 1 to have correlated beliefs about the actions of 2 and 3, but absent the explicit correlating device of own past moves the restriction to independent beliefs strikes us as less problematic.

Moreover, if player 1 is restricted to hold an independent belief, the action *Out* cannot be played in a RPCE. 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 (R_2, R_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 (L_2, L_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: Here it is that player 1 can entertain a correlated belief, which she would be unable to have if she actually observes 2 and 3's play.[19] □

Example 9 (Assessment-Strategies Correlation).

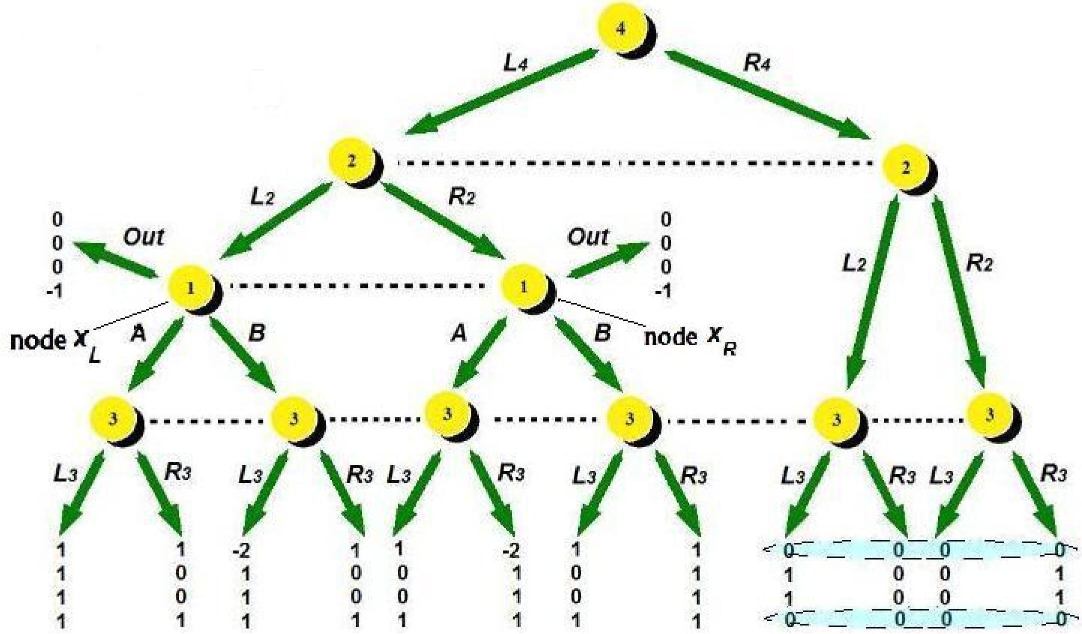

Figure 9

We have allowed v_i's belief at h, $(\mu_i)_h$, to lie in the space $\Delta(\Delta(h) \times \Pi_{-i})$ and not necessarily in $\Delta(h) \times \Delta(\Pi_{-i})$. Here we provide an example that justifies this specification.

[19] A similar argument can be made in Example 9 below to show that, with the discrete terminal node partition player 1 cannot play *Out* and so player 4 cannot play R_4.

Consider the extensive-form game depicted in Figure 9. All players observe the exact terminal node reached, except that players 1 and 4 do not distinguish among those terminal nodes that are caused by R_4.

We first show that R_4 is compatible with RPCE. For example, consider the following belief model and actual versions.

$$V_1 = \{v_1'\}, \quad v_1' = (Out, (\frac{1}{2}(v_2', v_3', v_4'), \frac{1}{2}(v_2'', v_3'', v_4')));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (L_2, (v_1', v_3', v_4')),\ v_2'' = (R_2, (v_1', v_3'', v_4'));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (L_3, (v_1', v_2', v_4')),\ v_3'' = (R_3, (v_1', v_2'', v_4'));$$

$$V_4 = \{v_4'\}, \quad v_4' = (R_4, (v_1', v_2', v_3'));$$

The actual version profile is (v_1', v_2', v_3', v_4').

To support R_4, it must be possible that 4 believes 1 plays *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. Specifically, suppose a belief μ_1' of v_1' is such that $b(\mu_1')$ assigns equal probabilities to (L_2, L_3, R_4) and (R_2, R_3, R_4), and $(\mu_1')_h$ assigns equal probabilities to $(x_L, (L_2, L_3, R_4))$ and $(x_R, (R_2, R_3, R_4))$, where h is 1's information set. This belief satisfies coherency and accordance, and it makes 1 playing *Out* a best response. 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 plausible, and it is possible in RPCE when each profile of opponents' strategies is associated with a different assessment. However, it 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.

Because the belief model underlying the play of R_4 seems sensible, we would not want to refine the set of RPCE by insisting that each version has a point distribution on assessments. The definition of RPCE allows each version to have a non-point distribution on assessments, and in particular it enables player 4 to play R_4 in this example. □

Example 10 (Accordance).

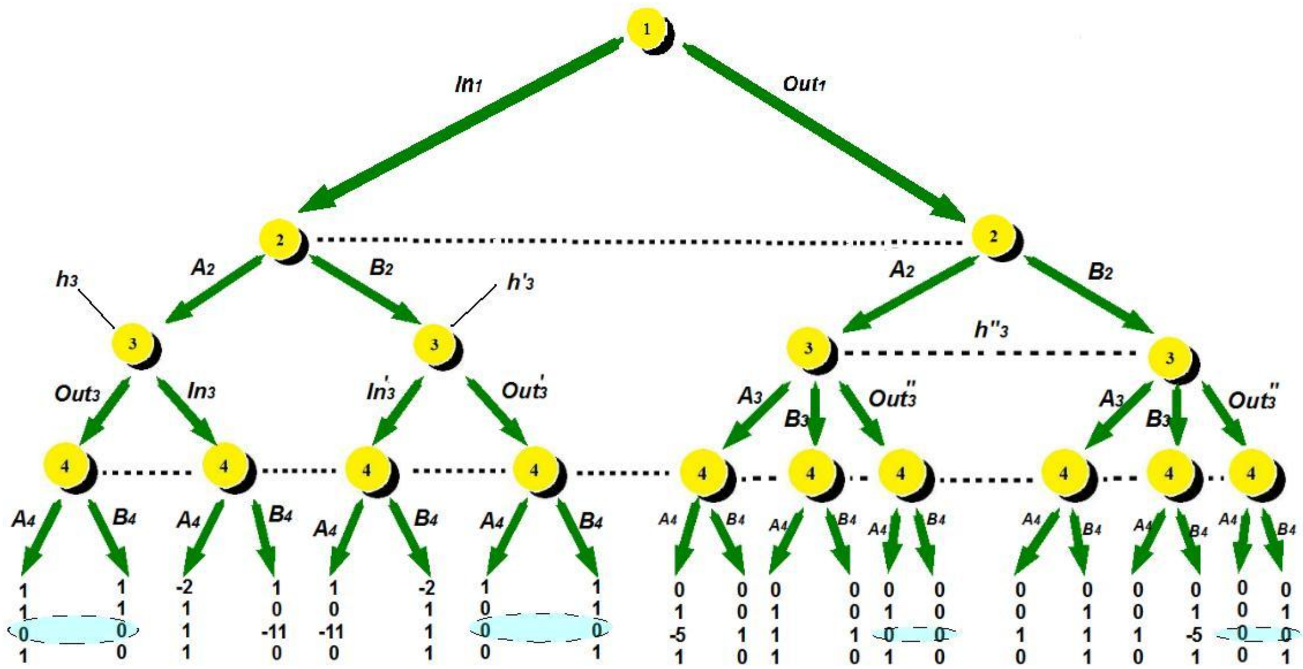

Figure 10

We use Figure 10 to explain why the accordance condition allows for a belief $(\mu_i)_h$ to have different marginals over continuation strategies across different h's. All players observe the exact terminal node reached, except that player 3 does not observe 4's choice when 3 plays Out. Intuitively, if 1 plays Out_1 and 3 thinks that 2 and 4 play either the (A_2, A_4) or (B_2, B_4) equilibria regardless of 1's play, then 1's deviation to In_1 would inform player 3 of which equilibrium 2 and 4 are coordinating on. To model this inference, 3's belief about the continuation play has to vary across information sets, which the definition of accordance allows.

We show here that $(Out_1, A_2, Out_3'', A_4)$ is a RPCE outcome while it would not be if we strengthened the definition by replacing part 2 to the following: For all $h \in H_i$,

$$(\mu_i)_h(\hat{\pi}_{-i}) = \sum_{\tilde{\pi}_{-i}(h')=\hat{\pi}_{-i}(h') \text{ for all } h' \text{ after } h} b(\mu_i)(\tilde{\pi}_{-i}). \tag{1}$$

This stronger condition requires that the continuation play has to agree with $b(\mu)$ as opposed to merely having a weakly smaller support.

First we show that $(Out_1, A_2, Out_3'', A_4)$ is a RPCE. In particular, it satisfies accordance. To see this, consider the following belief model:

$$V_1 = \{v_1', v_1''\}, \quad v_1' = (Out_1, ((v_2', v_3', v_4'))), \; v_1'' = (Out_1, (v_2'', v_3', v_4'')));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (A_2, (v_1', v_3', v_4')),\ v_2'' = (B_2, (v_1'', v_3', v_4''));$$

$$V_3 = \{v_3'\}, \quad v_3' = ((In_3, In_3', Out_3''), (\frac{1}{2}(v_1', v_2', v_4'), \frac{1}{2}(v_1'', v_2'', v_4''));$$

$$V_4 = \{v_4', v_4''\}, \quad v_4' = (A_4, (v_1', v_2', v_3')),\ v_2'' = (B_4, (v_1'', v_2'', v_3'));$$

The actual version profile is (v_1', v_2', v_3', v_4').

In this belief model, player 3 assigns probability 1/2 to each of (A_2, A_4) and (B_2, B_4). Then, when 1 plays Out_1, RPCE allows for a belief in which (i) player 3 at h_3 thinks that 4 will play A_4; (ii) player 3 at h_3' thinks that 4 will play B_4; (iii) player 3 at h_3'' thinks that 4 will play A_4 after A_2, and B_4 after B_2. Given this belief, (In_3, In_3', Out_3'') is a best response for player 3.

If 1 thinks that 3's version is as above, then 1 would expect payoff -2 from playing In_1, and 0 from playing Out_1. So Out_1 is a best response.

Now we show that $(Out_1, A_2, Out_3'', A_4)$ is not a RPCE outcome with the stronger version of accordance that imposes (1).

To see this, note that player 3's belief about 2 and 4's play can only assign positive probability to the strategy profiles (A_2, A_4), (B_2, B_4), or $((1/2A_2, 1/2A_4), (1/2B_2, 1/2B_4))$. Some observations about player 3's incentives are in order: First, for A_3'' not to be strictly better than Out_3'' at h_3'', player 3 cannot assign probability more than 5/6 to (B_2, B_4) while for B_3'' not to be strictly better than Out_3'' at h_3'', player 3 cannot assign probability more than 5/6 to (A_2, A_4).[20]

Next, for Out_1 to be a best response for player 1, 1 has to think that 3 plays In_3 or In_3' with a positive probability. This means that player 1 thinks that either player 3 at h_3 thinks that 4 would play A_4 with probability at least 11/12,[21] or player 3 at h_3' thinks that 4 would play B_4 with probability at least 11/12. But this is impossible under the strengthened definition of accordance because, given the conclusion above, player 3's belief about player 4's continuation strategy can assign probability at most $\max_p [(5/6) \cdot (1-p) + (1/2) \cdot p] = 5/6$ to each of A_4 and B_4, where p in the maximand denotes the probability that player 3 attaches to the mixed equilibrium play by players 2 and 4.

[20] If (B_2, B_4) is assigned probability 5/6, then playing A_3'' ensures the payoff of $1 \times (5/6) + (-5) \times (1/6) = 0$ because -2 is the worst payoff that player 3 can get given Out_1. A similar computation applies to the play of B_3'' at h_3''.

[21] If A_4 is assigned probability 11/12, then playing In_3 ensures the payoff of $1 \times (11/12) + (-11) \times (1/12) = 0$. A similar computation applies to case (v) as well.

5 RPCE, RSCE, and RCE

In this section we compare RPCE with other concepts from the literature. In Subsection 5.1 we compare RPCE with RSCE, and show that RPCE "reduces" to RSCE if the terminal node partitions are discrete and beliefs are independent. In Subsection 5.2 we compare RPCE with RCE (Rubinstein and Wolinsky, 1994), and show that when the signal function specified in the definition of RCE gives the same information as the partitions of the terminal nodes, RPCE is equivalent to RCE if moves are simultaneous.

5.1 Rationalizable Self-Confirming Equilibrium

In this subsection we show that RPCE is implied by RSCE if we require independent beliefs. One part of this argument is that any independent beliefs can be reduced to a single behavior strategy profile for the opponents, as shown by Fudenberg and Kreps (1995); the idea is that Kuhn's theorem allows us to associate a behavior strategy to any probability distribution on strategies, and that with independence the profile of these associated behavior strategies is equivalent to the original belief.

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.

Definition 8. π^* is a **rationalizable self-confirming equilibrium** if there exist a belief model V and an actual version profile v^* such that the following five conditions hold:[22]

1. π^* is generated by v^*.

2' For each i and $v_i = (\pi_i, p_i)$, there exists μ_i such that (i) μ_i is coherent with p_i, (ii) π_i is a best response to μ_i at all $h \in H_i$, and (iii) μ_i is an independent belief.[23]

3'. For all i and $v_i = (\pi_i, p_i)$, $d(\pi_i, \pi_{-i}(v_{-i})) = d(\pi^*)$ for all v_{-i} in the support of p_i.

There are two main differences between this definition and that of RPCE, namely that condition (3') (every version expects the same distribution over terminal nodes) is stronger than condition (3), and that observational consistency (4) is not directly imposed in RSCE. Even with a discrete terminal node partition the way condition (3) is stated is somewhat different than condition (3'),[24] but as the next result shows this difference is irrelevant.

[22]DFL 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.

[23]DFL required optimality only at the information sets that have positive probability under π_i, but the difference is immaterial in one-move games.

[24]If v_i is self-confirming then $d(\pi_i, \pi_{-i}(v_{-i}))$ equals $d(\pi_i, \pi^*_{-i})$ for all v_{-i} in the support of p_i, but condition (3') states that it is equal to $d(\pi^*_i, \pi^*_{-i})$.

Theorem 3. *Fix a game with discrete terminal node partitions.*

1. *If an actual version profile v^* and a belief model V satisfy conditions (3) and (4) then there exists a belief model $\hat{V}$ such that v^* and $\hat{V}$ satisfy conditions (3') and (4).*

2. *Condition (3') implies condition (4).*

The proof of this result is in the Appendix. Part 2 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 1 says that in the presence of the observational consistency condition, requiring the self-confirming condition for hypothetical 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 2. *In games with discrete terminal node partitions, any outcome of a RSCE is the outcome of a RPCE with independent beliefs.*

In the next example, which is taken from DFL's Example 3.2, we show that the set of possible outcomes can expand if we relax the definition of RSCE equilibrium by replacing condition (3') with condition (3).

Example 11 (DFL).

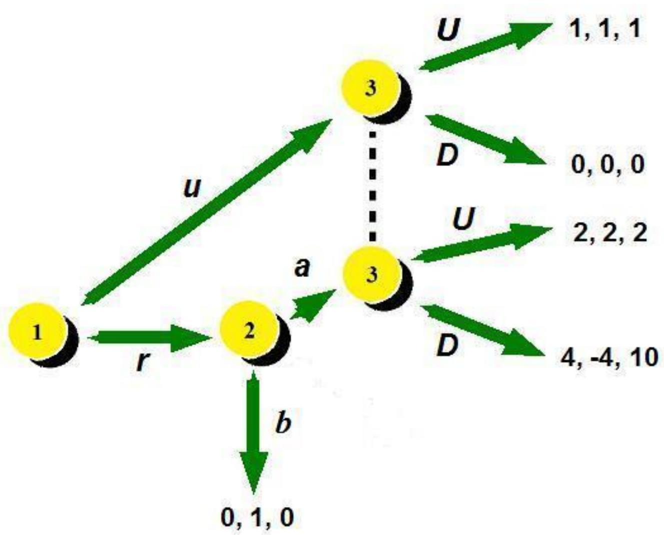

Figure 11

Consider the game depicted in Figure 11, where 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 8 where observational consistency is not imposed, this outcome becomes possible. To see this, consider the following belief model and actual versions:

$$V_1 = \{v_1', v_1''\}, \quad v_1' = (u, (v_2'', v_3')), v_1'' = (r, (v_2', v_3''));$$

$$V_2 = \{v_2', v_2''\}, \quad v_2' = (a, (v_1', v_3')), v_2'' = (b, (v_1', v_3''));$$

$$V_3 = \{v_3', v_3''\}, \quad v_3' = (U, (v_1', v_2')), v_3'' = (D, (v_1'', v_2''));$$

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

Here all the conditions in the definition of RSCE other than condition (3') hold, as does condition (3). Notice that v_1'', v_2'', and v_3'' are not self-confirming with respect to the actual distribution $\pi(v_1', v_2', v_3')$, and they are hypothetical versions and not actual ones.

The key is that the actual version of player 1, v_1', conjectures that 2 believes that 3 plays D, and this conjecture is ruled out by observational consistency: The equation in Remark 1(a) of observational consistency applied to v_1''s belief is $d(b, (u, D)) = d(b, (u, U))$. But this equation is false. □

Notice that the set of SCE is the same with (3) or (3'), thus requiring optimality at off-path information sets is the key to this example.

Finally, we have shown in Section 2 that Kreps and Wilson's (1982) consistency implies our restriction on beliefs. Since our restrictions do not imply consistency, the converse of Corollary 2 need not hold.

5.2 Rationalizable Conjectural Equilibrium

The main difference between RPCE and RCE 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

[25] See the example in Figure 2.1 of DFL.

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

of this difference, RPCE makes stronger predictions than RCE in most extensive-form games. If all information sets are on every path, this distinction becomes moot, and the two concepts become equivalent. In particular, in one-shot simultaneous-move games, we can state the precise connection between RCE and RPCE. To do so we first define RCE.

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 $u_i : A \to \mathbb{R}$. 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 $g_i : A \to S_i$. $g_i(a)$ 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 $\gamma \in \Delta(M_{-i})$ if (i) $g_i(m_i, m_{-i}) = \sigma_i$ for all $m_{-i} \in \text{supp}(\gamma)$, and (ii) m_i is a best response against γ.

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 γ such that for all $m_{-i} \in \text{supp}(\gamma)$ and all j, $(m_j, g_j(m_i, m_{-i})) \in B_j$.

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 $g^\mathbf{P} := (g_1^\mathbf{P}, \dots, g_n^\mathbf{P})$) such that $g_i^\mathbf{P}(a) = \mathbf{P}_i(a)$ for each action profile $a \in A^\Gamma$. Conversely, given any (A, g) with $g = (g_1, \dots, g_n)$ 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 $\mathbf{P}^g$ such that $\mathbf{P}_i^g(a) = g_i(a)$ for each action profile $a \in A^\Gamma$.

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. We omit the proof.

Theorem 4.

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

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

One consequence of this equivalence is that RCE, like RPCE, requires that in games like Example 3 when player 1 plays Out_1 she believes the play of the others is a Nash equilibrium of the subgame.[27] In particular this is true even in a three-player game where players 2 and 3 play the game of Shapley (1964), where fictitious play and smooth fictitious play do not converge.[28] Because the long-run joint distribution over actions in the Shapley cycle is a correlated equilibrium, this example may suggest an alternative equilibrium concept in which players expect that the empirical distribution of unobserved on-path play is a correlated equilibrium in the subgame. We do not define this alternative here because it is typically too inclusive.

6 Conclusion

Like RCE and RSCE, RPCE combines the idea that players have partial but objective information about equilibrium play with the idea that players reason about the observations and incentives of others. RSCE applies to extensive-form games where players see the realized terminal node at the end of each play of the game; RPCE generalizes this to situations where players see only a partition of the terminal nodes. In addition, RPCE relaxes the independent-beliefs condition of RSCE to allow for correlation.

The examples show that (1) under RPCE a player's belief about the actions of others can depend on whether those others get to act along the equilibrium path, (2) unobserved on-path play provides a natural form of correlating device, (3) a player can learn about the unobserved actions of a second player from the actions of a third, and finally, (4) the precise implications of all of the above depend on the nature of the terminal node partitions. In general, coarsening a player's terminal node partition cannot restrict the set of that player's RPCE strategies, but it can enlarge it. We identified four reasons that this enlargement can occur, and provided a sufficient condition under which coarsening a player's terminal partition has no effect on his RPCE strategies. We also showed how RPCE reduces to RCE and RSCE in the appropriate special cases.

The Online Supplementary Appendix discusses three additional topics: The definition of partition-confirmed equilibrium or PCE, the epistemic interpretation of observational consistency, and the effect of changes in terminal node partitions on the outcomes under RPCE.

[27]The Appendix gives a formal definition of the class of "player-1 participation games" and proves this claim.

[28]Brown (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.

References

[1] Aumann, Robert J. (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality" *Econometrica*, 55. 1-18.

[2] Battigalli, Pierpaolo (1987): "Comportamento razionale ed equilibrio nei giochi e nelle situazioni sociali," unpublished dissertation, Università Commerciale "L. Bocconi", Milano.

[3] Battigalli, Pierpaolo (1999): "A Comment on Non-Nash Equilibria," mimeo.

[4] Battigalli, Pierpaolo, Simone Cerreia-Vioglio, Fabio Maccheroni, and Massimo Marinacci (2012): "Selfconfirming Equilibrium and Model Uncertainty," mimeo.

[5] Battigalli, Pierpaolo and Amanda Friedenberg (2011): "Forward Induction Reasoning Revisited," *Theoretical Economics*, forthcoming.

[6] Battigalli Pierpaolo, Mario Gilli and M. Christina Molinari (1992): "Learning and convergence to equilibrium in repeated strategic interaction" *Research in Economics*, 46, 335-378.

[7] Battigalli, Pierpaolo and Danilo Guatoli (1997): "Conjectural Equilibria and Rationalizability in a Game with Incomplete Information" in *Decisions, Games and Markets*, Kluwer Academic Publishers, Norwell, MA.

[8] Battigalli, Pierpaolo and Marciano Siniscalchi, (2002): "Strong Belief and Forward-Induction Reasoning," *Journal of Economic Theory*, 106, 356-391.

[9] Ben-Porath, Elchanan (1997): "Rationality, Nash Equilibrium and Backward Induction in Perfect Information Games" *Review of Economic Studies* 64, 1997, 23-46.

[10] Bernheim, B. Douglas (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007- 1028.

[11] Brandenburger, Adam and Eddie Dekel (1987): "Common Knowledge with Probability 1," *Journal of Mathematical Economics*, 16, 237-245.

[12] Brandenburger, Adam and Amanda Friedenberg (2008): "Intrinsic Correlation in Games," *Journal of Economic Theory*, 141, 28-67.

[13] Brown, George W. (1951): "Iterative Solutions of Games by Fictitious Play" in *Analysis of Production and Allocation*, ed. by T. C. Koopmans. New York: Wiley.

[14] Dekel, Eddie, Drew Fudenberg and David K. Levine (1999): "Payoff Information and Self-Confirming Equilibrium," *Journal of Economic Theory*, 89, 165-185.

[15] Dekel, Eddie, Drew Fudenberg and David K. Levine (2002): "Subjective Uncertainty over Behavior Strategies: A Correction," *Journal of Economic Theory*, 104, 473-478.

[16] Dekel, Eddie, Drew Fudenberg and David K. Levine (2004): "Learning to Play Bayesian Games," *Games and Economic Behavior*, 46, 282-303.

[17] Dekel, Eddie, and Marciano Siniscalchi (2012): "An Epistemic Analysis of Games," mimeo.

[18] Esponda, Ignacio (2011): "Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions," mimeo.

[19] Forges, Francoise (1986): "An Approach to Communication Equilibria," *Econometrica*, 54, 1375-1386.

[20] Foster, Dean P. and Rrakesh V. Vohra (1997): "Calibrated Learning and Correlated Equilibrium," *Games and Economic Behavior*, 21, 40-55.

[21] Fudenberg, Drew and Yuichiro Kamada (2012): "Rationalizable Partition-Confirmed Equilibrium, II: Heterogeneous Beliefs," mimeo.

[22] Fudenberg, Drew and David M. Kreps (1988): "A Theory of Learning, Experimentation and Equilibrium in Games," mimeo, Stanford.

[23] Fudenberg, Drew and David M. Kreps (1993): "Learning Mixed Equilibria," *Games and Economic Behavior*, 5, 320-367.

[24] Fudenberg, Drew and David M. Kreps (1995): "Learning in Extensive Games, I: Self-Confirming Equilibrium," *Games and Economic Behavior*, 8, 20-55.

[25] Fudenberg, Drew, David M. Kreps, and David K. Levine (1988): "On The Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44, 354–380.

[26] Fudenberg, Drew and David K. Levine (1993a): "Self-Confirming Equilibrium" *Econometrica*, 61, 523-546.

[27] Fudenberg, Drew and David K. Levine (1993b): "Steady State Learning and Nash Equilibrium" *Econometrica*, 61, 547-574.

[28] Fudenberg, Drew and David K. Levine (2006): "Superstition and Rational Learning" *American Economic Review* 96, 630-651.

[29] Fudenberg, Drew and David K. Levine (2009): "Learning and Equilibrium" *Annual Review of Economics* 1, 385-419.

[30] Fudenberg, Drew and Satoru Takahashi (2011): "Heterogeneous Beliefs and Local Information in Stochastic Fictitious Play," *Games and Economic Behavior*, 71, 100-120.

[31] Gilli, Mario (1999): "On Non-Nash Equilibria," *Games and Economic Behavior* , 27, 184-203.

[32] Hahn, Frank H. (1977): "Exercises in Conjectural Equilibria" *The Scandinavian Journal of Economics*, 79, 210-226.

[33] Kalai, Ehud and Ehud Lehrer (1993): "Subjective Equilibrium in Repeated Games," *Econometrica*, 61, 1231-1240.

[34] Kamada, Yuichiro (2010): "Strongly Consistent Self-Confirming Equilibrium," *Econometrica*, 78, 823-832.

[35] Kohlberg, Elon and Jean-Francois Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003-1037.

[36] Kreps, David M. and Garey Ramey (1987): "Structural Consistency, Consistency, and Sequential Rationality," *Econometrica*, 55, 1331-1348.

[37] Kreps, David M. and Robert Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863-894.

[38] Lehrer, Ehud (2012): "Partially-Specified Probabilities: Decisions and Games," *American Economic Journal: Microeconomics*, 4, 1-33.

[39] Pearce, David (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029-1050.

[40] Rubinstein, Ariel and Asher Wolinsky (1994): "Rationalizable Conjectural Equilibrium: Between Nash and Rationalizability," *Games and Economic Behavior*, 6, 299-311.

[41] Ryall, Michael D. (2003): "Subjective Rationality, Self-Confirming Equilibrium, and Corporate Strategy," *Management Science*, 49, 936-49.

[42] Shapley, Lloyd (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 Lemma 1

The next lemma shows that in a RPCE, any version profile to which an actual version assigns positive probability is the actual version profile of some RPCE.

Lemma 1. *Fix a RPCE π^*, a belief model V, and actual versions v^* that support it. For every player i, if v_i^*'s conjecture assigns positive probability to $\tilde{v}_{-i}$, then $\pi(v_i^*, \tilde{v}_{-i})$ is also a RPCE.*

Proof. Pick player i and $\tilde{v}_{-i}$ to which v_i^* assigns positive probability. We will use the belief model V to support a $\pi(v_i^*, \tilde{v}_{-i})$ as a RPCE. Conditions 2 and 4 of the definition of RPCE hold for all versions in V, so they hold for strategy $\pi(v_i^*, \tilde{v}_{-i})$ and the belief model V as well. So it remains to show that, if $(v_i^*, \tilde{v}_{-i})$ is the actual version profile, then v_i^* satisfies the self-confirming condition with respect to $\pi(v_i^*, \tilde{v}_{-i})$, and that $\tilde{v}_j$ for each player $j \neq i$ does as well. First, since v_i^* satisfies the self-confirming condition in the original RPCE, and $\tilde{v}_{-i}$ is in the support of the conjecture of v_i^* by assumption, v_i^* is self-confirming with respect to $\pi(v_i^*, \tilde{v}_{-i})$. Second, since v_i^* satisfies observational consistency in the original RPCE, $D_j(\pi_j(\tilde{v}_j), \pi_{-j}(\hat{v}_{-j})) = D_j(\pi(v_i^*, \tilde{v}_{-i}))$ for all $\hat{v}_{-j}$ in the support of the conjecture of $\tilde{v}_j$. Thus $\tilde{v}_j$ satisfies the self-confirming condition with respect to $\pi(v_i^*, \tilde{v}_{-i})$. $\square$

Remark 2. The proof of Lemma 1 shows that even the hypothetical version $\tilde{v}_j$ must satisfy the self-confirming condition with respect to $\pi(v_i^*, \tilde{v}_{-i})$ if v_i^* assigns positive probability to $\tilde{v}_{-i}$. This does not imply that $\tilde{v}_j$ is self-confirming with respect to π^*. Indeed, imposing that condition (i.e. imposing condition (3) for all versions and not just the actual ones) would be unduly restrictive, as we show in Example 7.

B A Theorem for Participation Games

The class of participation games generalizes some of the examples from the text. Intuitively, this is a game in which player 1 has an option to play *Out* at the root node that prevents her from observing the consequence of the opponents' actions at the terminal nodes, and other players play a game, not knowing player 1's action. Formally, a **player-1 participation game** Γ (with a payoff function u and the set of players I) is an extensive-form game with the following properties: Fix player 1's set of actions A_1 such that one of its element is *Out*, and another extensive-form game Γ' with a payoff function v and the set of players $I \setminus \{1\}$. Denote by $n(x, a)$ the node in Γ that corresponds to x in Γ' after action $a \in A_1$ is taken.

- At the root node player 1 moves, choosing between *In* and *Out*.
- Whichever action is taken, Γ' is played after player 1's decision.
- Nodes $n(x, a)$ and $n(x', a')$ for $x, x' \in X \setminus Z$ are in the same information set in Γ if and only if x and x' are in the same information set in Γ'.
- Terminal nodes $n(z, a)$ and $n(z', a')$ for $z, z' \in Z$ are in the same cell of the terminal node partition of player i if and only if z and z' are in the same cell for i and $a = a'$, except for the following exception.
- Terminal nodes $n(z, Out)$ and $n(z', Out)$ for $z, z' \in Z$ are in the same cell of the terminal node partition of player 1.

- $u_i(n(z, In)) = v_i(z)$ for all i, $u_i(n(z, Out)) = v_i(z)$ for all $i \neq 1$, and $u_1(n(z, Out)) = 0$.

Theorem 5. *Fix a player-1 participation game* Γ. *If player 1 plays Out with probability 1 in a RPCE, then there is a convex combination of RPCE of* Γ' *such that no action of player 1 has a strictly positive payoff.*

Proof. Fix a RPCE in which player 1 plays *Out* with probability 1. Pick a version profile for player 1's opponents $\tilde{v}_{-1}$ to which the conjecture of v_1^* assigns positive probability. By Lemma 1, $\pi(v_1^*, \tilde{v}_{-1})$ is a RPCE of Γ, and it is clear from the proof of the lemma that the same belief model V that supports the original RPCE can be used to support $\pi(v_1^*, \tilde{v}_{-1})$ as a RPCE, where the actual version profile is $(v_1^*, \tilde{v}_{-1})$. By the definition of a player-1 participation game, player 1's action does not affect any opponent's payoff or observation. Thus $\pi_{-1}(\tilde{v}_{-1})$ is trivially a RPCE of Γ', with the belief model simply deleting player 1. Since this is true for any $\tilde{v}_{-1}$ in the support of the conjecture of v_1^*, and the strategy of the actual version of player 1 is a best response to her belief in the original RPCE, the proof is complete. □

Corollary 3. *Fix a player-1 participation game* Γ *such that* Γ' *is a simultaneous-move game with discrete terminal node partitions and a unique Nash equilibrium. If player 1 plays Out with probability 1 in a RPCE of* Γ, *then no action of player* 1 *gives her a positive payoff against this unique Nash equilibrium.*

C Proof of Theorem 3

Proof.

Part 1: Fix an actual version v^* and a belief model V that satisfies conditions (3) and (4). Construct a new belief model $\hat{V}$ that is identical to the original one, except that all versions that do not satisfy the equality in (3') in V 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 V (such versions are not eliminated because of condition (3)). By construction, condition (3') holds. Hence by part 2 that we prove below, condition (4) holds as well. Finally, we check that the sum of probabilities assigned by the conjecture of any remaining version is unity. To see this, note first that condition (3) implies that the actual version $v_i^* = (\pi_i^*, p_i^*)$ must satisfy $(\pi_i^*, \pi_{-i}(v_{-i})) = d(\pi(v^*))$ for all v_{-i} in the support of p_i^*. Also, for $\tilde{v}_j = (\tilde{\pi}_j, \tilde{p}_j)$, whenever $d(\tilde{\pi}_j, \pi_{-j}(v_{-j})) = d(\pi(v^*))$ for all v_{-j} in the support of $\tilde{p}_j$, observational consistency implies that for any version of j's opponent $\hat{v}_k = (\hat{\pi}_k, \hat{p}_k)$ in the support of $\tilde{p}_j$, $d(\hat{\pi}_k, \pi_{-k}(v_{-k})) = d(\pi(v^*))$ for all v_{-k} in the support of

$\hat{v}_k$. This means that no version who is assigned a positive probability by any remaining version is eliminated, implying that the sum of probabilities is still unity.

Part 2: Since terminal node partitions are discrete, the observational consistency condition for version $v_i = (\pi_i, p_i)$ reduces to the requirement that $p_i(v_{-i}) > 0$ implies, for each $j \neq i$, $d(\pi_j(v_j), \pi_{-j}(v_{-j})) = d(\pi_i, \pi_{-i}(v_{-i}))$ for all v_{-j} in the support of v_j's conjecture. But the conclusion of this requirement is implied by condition (3'), as (3') implies that both sides of the equality are equal to $d(\pi^*)$. □