Rationalizable Partition-Confirmed Equilibrium with Heterogeneous Beliefs[*]

Drew Fudenberg and Yuichiro Kamada[†]

Abstract

Many models of learning in games implicitly or explicitly assume there are many agents in the role of each player. In principle this allows different agents in the same player role to have different beliefs and play differently, and this is known to occur in laboratory experiments. To explore the impact of this heterogeneity, along with the idea that subjects use their information about other players' payoffs, we define rationalizable partition-confirmed equilibrium (RPCE). We provide several examples to highlight the impact of heterogeneous beliefs, and show how mixed strategies can correspond to heterogeneous play in a large population. We also show that every heterogeneous-belief RPCE can be approximated by a RPCE in a model where every agent in a large pool is a separate player.

Keywords: Rationalizability, extensive-form games, self-confirming equilibrium, heterogeneous beliefs, purification, random matching

1 Introduction

Fudenberg and Levine (1993b) and Fudenberg and Kreps (1995) showed how learning from repeated observation of the realized terminal nodes in each play of a game can allow the long run outcome to approximate a self-confirming equilibrium (Fudenberg and Levine, 1993a), in which the strategies used are best responses to possibly incorrect beliefs about play that are not disconfirmed by the players' observations. This paper defines and analyzes a solution concept that makes three modifications to the self-confirming concept, inspired by the following three considerations. First, the set of beliefs that is consistent with the players' observations depends on what they observe when the game is played, and in some cases of interest players do not observe the exact terminal node, but only a coarser *terminal node partition*, such as when bidders in an auction do not observe the losing bids. Second, both in the lab and in the field, there are often many agents in each player role, so that different agents in a given player role can have different beliefs and play differently, and experimental data frequently suggests that subjects' beliefs and play are indeed heterogeneous.[1] Third, experimental data shows that subjects play differently when they are informed of opponents' payoff functions than when they are not.[2] To model these three facts, we develop and analyze *heterogeneous rationalizable partition-confirmed equilibrium* (heterogeneous RPCE).

We do not develop an explicit learning theory here, but the model we develop is motivated by the idea that there is a large number of ex-ante identical agents in each player role, who are rematched each period to play an extensive form game and interact anonymously, so that they are strategically myopic and do not try to influence play in future matches. Such random matching is implicit in many learning models, and is explicitly modeled in the Bayesian learning models of Fudenberg and Levine (1993b) and Fudenberg and He (2016).

The long-run implications of learning with random matching depend on

---

[1]For example, Fudenberg and Levine (1997) relate heterogeneous beliefs to data from experiments on the best-shot, centipede, and ultimatum games.

[2]See for example Prasnikar and Roth (1992).

what information is revealed at the end of each round of play. In the *information-sharing model*, all agents in the same player role pool their information about what they observe after each round of play, which leads to rationalizable partition-confirmed equilibrium with unitary beliefs, which we studied in Fudenberg and Kamada (2015) (hereafter "FK").[3] In the *personal-information model*, each agent observes and learns only the play in her own match, and no information sharing takes place. This is the treatment most frequently used in game theory experiments. It allows different agents in the same player role to maintain different beliefs, even after many iterations of the game, and even when the agents are identical ex ante.

The large-population learning models described above assume personal information, and so their steady states can have heterogeneous beliefs, which is why Fudenberg and Levine (1993a) defined and analyzed heterogeneous self-confirming equilibrium. Dekel et al. (2004) argue that in Bayesian games it may be appropriate to allow different types of the same player to have different beliefs, and Battigalli et al. (2015) allow heterogeneous beliefs in their extension of self-confirming equilibrium to cases of "model uncertainty." Their model also allows for players to observe a "message" that is generated by a "feedback function" at the end of each play as opposed to the realized terminal node. In Section 2 we say more about the relationship between feedback functions and terminal node partitions.

1.1   Illustrative Examples

To motivate our extension of RPCE to heterogeneous beliefs, we will give informal descriptions of RSCE and its extension to unitary RPCE, and then discuss a few examples to show how allowing for heterogeneous beliefs can make a difference. We will return to these examples later after we have introduced our formal definitions.

---

[3]FK extends an earlier literature on equilibrium concepts that combine restrictions based on the agents' observations with restrictions based in their knowledge of opponents' payoffs, including (Rubinstein and Wolinsky (1994), Battigalli and Guaitoli (1997), Dekel et al. (1999), and Esponda (2013)).

Roughly speaking, RSCE (Dekel et al. 1999) combines the idea that there is a commonly known equilibrium outcome with that of "rationalizability at reachable nodes," which means that any strategy of player i that some other player j thinks i might be playing maximizes player i's payoff at information sets where player i has not yet been observed to deviate from the equilibrium path, and that each player believes that other players believe this, and so on. As Dekel et al. (1999) show, the combination of these two considerations leads to stronger conclusions than the intersection of the implications of each alone.

Unitary RPCE relaxes RSCE by replacing the assumption of a commonly known distribution on terminal nodes with the assumption of a commonly known partition structure. Both of these concepts are "unitary" in the sense that it is common knowledge that there is only a single strategy and belief that is really held by any agent in the role of player i. However, both solution concepts allow for "hypothetical versions" of each player's strategy and belief, which correspond not to what that player is doing but to things that other players might reasonably think she is doing. The heterogeneous RPCE studied in this paper allows for different agents in the role of a given player to have different beliefs, and moreover for agents to believe that other agents in the same player role have different beliefs than they do.

We now informally discuss some of the implications of heterogeneous beliefs. In those examples we claim that some outcome is inconsistent with unitary beliefs but is consistent with heterogenous beliefs, given our background assumption that all players have a common belief in rationality and in confirmed beliefs.

Perhaps the most immediate implication is that it allows the play of one or more players to be strictly mixed in cases where unitary beliefs require the outcome to be a single terminal node, as in the game in Figure 1.

Figure 1: The broker, seller, and buyer are denoted players 1, 2, and 3, respectively.

Here there are potential brokers who facilitate trades between sellers and buyers. Each potential broker decides whether she enters the market or stays out; if she stays out, the game ends with no trade. If the broker enters, the involved seller and buyer play a coordination game, where the efficient and inefficient outcomes correspond to the possible outcomes of an unmodeled process of negotiation. The personal-information model has a steady state in which some potential brokers stay out of the market and the others enter, while all of sellers and buyers play efficient negotiations. Although staying out is not a best response, the brokers can choose it if they believe the negotiations would be inefficient, and this belief will not be falsified by their observations. This is a heterogeneous self-confirming equilibrium, but it is not the outcome of a self-confirming equilibrium with unitary beliefs. This is because the aggregate play of the brokers corresponds to a mixed distribution, yet if the brokers pooled their information they would not be indifferent between the actions in the distribution's support.

Note that the distribution is a convex combination of the outcomes of subgame perfect equilibria.[4] The game in Figure 2 illustrates another way that heterogeneous beliefs can matter. Consider the following situation: Investors decide whether to attend a business event, and entrepreneurs simultaneously decide whether to prepare materials to solicit investments. This preparation must be done before the meeting, and any entrepreneur who does prepare will then make a solicitation. Each investor derives a positive benefit from coming to the event, but this is outweighed by the cost if she is approached by an entrepreneur who solicits money, while entrepreneurs only want to solicit if it is sufficiently likely they can talk with an investor. Specifically, entrepreneurs who do not solicit get 0; those who do solicit get 1 if the investor attends, and −1 if the investor stays *Out*. Similarly, entrepreneurs who don't attend get 0; those who attend get −1 if the entrepreneur solicits and 1 if it does not.


Figure 2: The investor and the entrepreneur are denoted players 1 and 2, respectively.

---

[4]We believe we could construct a more complicated game where this is not the case, just as Fudenberg and Levine (1993a) give an example of a SCE with heterogeneous beliefs whose outcome is not in the convex hull of the Nash equilibrium outcomes.

Note that the unique Nash equilibrium here is for both players to randomize $(1/2, 1/2)$. The dotted line indicates a *terminal node partition*; it shows that an investor who stays *Out* does not observe if the entrepreneur solicits, even though the entrepreneur's action is then on the path of play. Here, the profile $(Out, Don't)$ cannot be supported with unitary beliefs. To see this, note that for the investor to play *Out*, she has to expect a positive probability that the entrepreneur solicits, but soliciting is not a best response for the entrepreneur against *Out*. This contradicts the assumption that the investor knows the entrepreneur is rational.

With heterogeneous beliefs, on the other hand, it is not obvious why the outcome $(Out, Don't)$ should be rejected. To see why, note that if all agents in the role of investors think that the overall distribution of play corresponds to the Nash equilibrium of the game, these agents will be indifferent, and absent any information to the contrary, all of the investors could stay home even though none of the entrepreneurs solicit. In order to rationalize this belief, the investors must conjecture that some other agents in the same player role *Attend* because they believe some entrepreneurs *Solicit*, and *Solicit* is a best response only if there are investors who *Attend*. To capture this in our formal model, we will need to allow for "hypothetical versions" of the investors who in fact do *Attend*, so that the investors who stay out can assign positive (and sufficiently large) probability to these hypothetical versions.

The effect of heterogeneous beliefs is even starker in the game in Figure 3.

7

Figure 3: The tax attorney, IRS agent, and tax evader are denoted players 1, 2, and 3, respectively.

Here, Nature first chooses Good or Bad. If it chooses Bad, then a tax attorney (player 1) prepares a tax return, which can be either *Safe* or *Risky*. *Risky* results in auditing by an IRS agent (player 2), and depending on the agent's effort level, the attorney is either rewarded by the tax evader (player 3) or punished. If the attorney chooses *Safe*, the return will not be audited, and then the tax evader has a choice of staying with the attorney (*Stay*) or firing him (*Fire*). Nature's choice of Good represents the situation in which the person who is audited has filed her tax return sincerely. The IRS agent, who does not know if the return is good or bad, would like to exert effort (*E*) in auditing if and only if it faces the evader, and otherwise prefers to not exert effort (*N*). If the file is good, then the attorney does not observe what the IRS agent has chosen, as irrespective of the agent's effort, the auditing would not result in punishment.

With unitary beliefs, it is not possible for the attorney to play *Safe* with probability 1. To see this, note that the attorney knows that the IRS agent

observes the exact terminal node reached. This implies that if the attorney played $Safe$ with probability 1, then she would know that whenever the IRS agent moves the IRS agent would know that Nature gave him the move, which would then imply that the attorney expects that the agent should play N. Thus the attorney should expect the payoff of 2 from playing $Risky$, which dominates $Safe$.

However, with heterogeneous beliefs the attorney can play $Safe$ with probability 1. Roughly, if each attorney thinks that all other attorneys play $Risky$, she has to infer each IRS agent assigns probability about .5 to each node, and this implies that the IRS agent must play E. But then if the evaders $Stay$, playing $Safe$ is a best response. And an attorney can believe that other attorneys think $Risky$ is a best response by supposing these other agents believe all the evaders play $Fire$.

We first lay out our model, then use it to analyze these and other examples with more rigor. We then show how heterogeneous play by a continuum of agents permits the "purification" of mixed strategy equilibria. That is, any outcome of a heterogeneous RPCE is the outcome of a heterogeneous RPCE in which all agents use pure strategies and believe that all other agents use pure strategies as well. Finally, we relate the heterogenous RPCE of a given game to the unitary RPCE of a larger *"anonymous-matching game"* in which each of the agents in a given player role is viewed as a distinct player. Because we assume that this larger game has a finite number of agents in each role, the result here is not quite an equivalence, but involves some approximations which vanish as the number of agents grows large.

2   Model and Notation

The game tree consists of a finite set X, with terminal nodes denoted by $z \in Z \subseteq X$. The initial node corresponds to Nature's move, if any. The set of Nature's actions is A_N. The set of players is $I = \{1, \ldots, |I|\}$. H_i is the collection of player i's information sets, with $H = \bigcup_{i \in I} H_i$ and $H_{-i} = H \setminus H_i$. Let $A(h)$ be the set of available actions at $h \in H$.

For each $z \in Z$, player i's payoff is $g_i(z)$. The information each player i observes at the end of each round of play is captured by a terminal node partition $\mathbf{P}_i$ that is a partition of Z, where we require that $g_i(z) = g_i(z')$ if terminal nodes z and z' are in the same cell of $\mathbf{P}_i$. We let $\mathbf{P} = (\mathbf{P}_i)_{i \in I}$ denote the collection of the partitions.[5] The formal definitions and assumptions we make below are motivated by the assumption that the structure of the game, including the payoff functions, terminal node partitions, and distributions over Nature's move, are common knowledge in the informal sense.

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 $\Pi_{-i} = \times_{j \neq i} \Pi_j$ with typical element π_{-i}. For $\pi \in \Pi$ and $\pi_i \in \Pi_i$, $H(\pi)$ and $H(\pi_i)$ denote the information sets reached with positive probability given π and (π_i, π'_{-i}), respectively, where π'_{-i} is any completely mixed behavioral strategy.

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. 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 $\mathbf{P}_i$ depends only on i's actions, the partition is called trivial.

For most of the paper we restrict attention to "generalized one-move games," in which for any path of pure actions each player moves at most once, and for each i, if there exist $h \in H_i$, $a_N, a'_N \in A_N$ with $a_N \neq a'_N$, and $\tilde{\pi}_{-i} \in \Pi_{-i}$ such that h is reached with positive probability under both $(a_N, \tilde{\pi}_{-i})$ and $(a'_N, \tilde{\pi}_{-i})$, then there does not exist $\hat{\pi}_{-i} \in \Pi_{-i}$ such that, for all $a''_N \in A_N$, h is reached

---

[5]Battigalli et al. (2015) model the information that players receive at the end of each play of a dynamic game by assigning each player i a "feedback function" f_i that maps from the terminal nodes to a finite set of messages M_i. Feedback functions obviously include our terminal node partitions as a special case. Conversely, our framework nests feedback functions. To see this, given f_i, construct a terminal node partition such that terminal nodes z and z' are in the same cell of the partition if and only if $f_i(z) = f_i(z')$.

with probability zero under $(a_N'', \hat{\pi}_{-i})$.[6] This restriction lets us neglect conceptual complications that would arise in specifying assessments at off-path information sets.[7] We use a slightly more general class of games to relate the heterogeneous and unitary solution concepts, as in those games Nature's move determines which agents are selected to play.

2.1 Heterogeneous Rationalizable Partition-Confirmed Equilibrium

Player i's *belief* is denoted $\gamma_i \in [\times_{h \in H_i} \Delta(h)] \times \Pi$, which includes her assessment over nodes at her information sets as well as her belief about the overall distribution of strategies. We denote the second element of γ_i by $\pi(\gamma_i)$, and let $\pi_{-i}(\gamma_i)$ denote the corresponding strategies of players other than i. Note that we suppose that the belief about strategies is a point mass on a single behavior strategy profile, as opposed to a probability distribution over strategy profiles.[8]

To model the idea that players are reasoning about the beliefs and play of others, we follow Dekel et al. (1999) and FK and use *versions* of player i.[9] We let V_i denote the set of versions of player i. For simplicity we assume that each V_i is finite and index the elements of V_i with integers k. Each version v_i^k of player i consists of a strategy $\pi_i^k \in \Pi_i$ and a *conjecture* $q_i^k \in \times_{j \in I} \Delta(V_j)$ about the distribution of versions in the population, where a conjecture q_i^k is a

---

[6]This is a generalization of the "one-move games" defined in FK. It reduces to FK's definition of one-move games without a move by Nature.

[7]For example, suppose that Nature chooses L or R and then player 1 chooses between U and D without knowing Nature's action. U ends the game, and D leads to a single information set h_2 of player 2, so if 1 plays U, then Bayes rule does not pin down 2's assessment at h_2. Note that this is not a generalized one-move game as h_2 groups together the two moves by Nature. One issue that could come up on in such general games is whether to require that a player's deviations cannot convey information about things they do not know.

[8]We allowed for beliefs to be possibly correlated probability distributions on Π_{-i} in FK. Here we restrict to a single strategy profile to focus on some of the issues that arise with heterogeneity. This also lets us avoid the need to impose an analog of FK's "accordance" condition.

[9]Note that versions are not used in the definition of self-confirming equilibrium, because that concept does not model reasoning about the beliefs of other players.

point belief over version distributions. To lighten notation, we will sometimes suppress the indices of the versions and write $\pi_i(v_i)$ for the strategy of type v_i, and $\pi_{-i}(v_{-i})$ for the strategies of the other types v_{-i}.

These versions, and the players' conjectures about them, serve a similar role as the epistemic type structures used in e.g. Dekel and Siniscalchi (2014). We use "version" instead both for ease of comparison with Dekel et al. (1999) and because we feel this structure is easier for readers to understand. We believe that we could map heterogeneous RPCE into the language of epistemic type structures, and that no new insights would emerge, but as we have not done so this is an open question that may be of interest to specialists in epistemic game theory.

Not all of the versions need actually be present in the population. We track the shares of the versions that are objectively present with the *share function* $\phi = (\phi_i)_{i \in I}$, where each $\phi_i \in \Delta(V_i)$ specifies the fractions of the population of player i that are each v_i^k; version v_i^k is called an "actual version" if $\phi_i(v_i^k) > 0$, and a "hypothetical version" otherwise. Hypothetical versions are the ones that some players think might be present but are not. Let $V := (V_1, \ldots, V_{|I|})$. We call (V, ϕ) a *belief model*.

Next, we show how a belief model induces a behavior strategy profile π that describes the aggregate play of the actual versions, and also induces, for each version v_i^k, the strategy profile that the version thinks describes actual play.

For each player j, define $\psi_j(\phi_j)$ for each ϕ_j by $\psi_j(\phi_j)(\hat{\pi}_j) = \sum_{k:\pi_j^k = \hat{\pi}_j} \phi_j(v_j^k)$ for each $\hat{\pi}_j$; this is the share of agents who play $\hat{\pi}_j$ under the belief model (V, ϕ). Note that $\psi_j(\phi_j)$ has finite support.

Definition H1 A belief model (V, ϕ) **induces** actual play $\hat{\pi}_j$ if for all $h_j \in H_j$ and $a_j \in A(h_j)$,

$$\hat{\pi}_j(h_j)(a_j) = \sum_{\pi_j' \in \text{supp}(\psi_j(\phi_j))} \psi_j(\phi_j)(\pi_j') \cdot \pi_j'(h_j)(a_j).$$

We say that (V, ϕ) induces $\hat{\pi}_j$ for version $v_i^k \in V_i$ if $\hat{\pi}_j$ is constructed by

replacing ϕ_j above by the j'th coordinate of q_i^k.

Definition H2 Given a belief model (V, ϕ), we say v_i^k is **self-confirming with respect to** π^* if there exists $\tilde{\pi}_{-i} \in \Pi_{-i}$ such that (i) for each $j \neq i$, (V, ϕ) induces $\tilde{\pi}_j$ for version v_i^k and (ii) $D_i(\pi_i^k, \tilde{\pi}_{-i}) = D_i(\pi_i^k, \pi_{-i}^*)$.

Note that π_i^k can be different from π_i^*. This is because an agent in player role i does not get to observe what other agents in the same role play.

Definition H3 Given a belief model (V, ϕ), v_i^k is **observationally consistent** if $(q_i^k)_j \, (\tilde{v}_j) > 0$ implies that there exists $\hat{\pi}_{-j} \in \Pi_{-j}$ such that (i) for each $l \neq j$, (V, ϕ) induces $\hat{\pi}_l$ for v_i^k and (ii) $\tilde{v}_j$ is self-confirming with respect to $(\pi_j(\tilde{v}_j), \hat{\pi}_{-j})$.

Intuitively, the self-confirming condition requires that the agent's belief is not rejected by her observations. Observational consistency requires that, if agent A thinks agent B exists, then A should expect B's belief not to be rejected by B's observations.

It is important here to note that the observational consistency condition defined above restricts v_i^k's belief about i's strategies as well as her beliefs about the strategies of the other players. This is needed because other agents in the same player role may play differently from v_i^k.

We say that $\pi_i \in \Pi_i$ is a **best response to** γ_i at $h \in H_i$ if the restriction of π_i to the subtree starting at h maximizes player i's expected payoff given the assessment at h given by γ_i and the continuation strategy of the opponents given by $\pi_{-i}(\gamma_i)$ in that subtree.

Definition H4 π^* is a **heterogeneous rationalizable partition-confirmed equilibrium**, or a **heterogeneous RPCE**, if there exists a heterogeneous belief model (V, ϕ) such that the following four conditions hold for each i:

1. (V, ϕ) induces actual play π_i^*;

2. For all v_i^k, there exists γ_i such that (i) (V, ϕ) induces $\pi_j(\gamma_i)$ for v_i^k for each $j \neq i$ and (ii) π_i^k is a best response to γ_i at all $h \in H_i$;

3. For all v_i^k, $\phi_i(v_i^k) > 0$ implies v_i^k is self-confirming with respect to π^*;

4. Each $v_i^k \in V_i$ is observationally consistent.

One significant change from the definition of unitary RPCE is in the self-confirming condition: In the unitary case, the self-confirming condition is imposed for those versions who have share 1 according to ϕ. In our current context, multiple versions may exist with strictly positive shares, and in such a case we require that all such versions are self-confirming.

2.2 Brief Review of Unitary RPCE

Here we briefly review the definition of unitary RPCE. In this solution concept, a belief is $\mu_i \in [\times_{h \in H_i} \Delta(\Delta(h) \times \Pi_{-i})] \times \Delta(\Pi_{-i})$. The coordinate for information set h is denoted $(\mu_i)_h$ which is assumed to have finite support. The second coordinate describes the strategy distribution of the opponents player i believes she is facing and is denoted $b(\mu_i)$. Note that the belief μ_i in the unitary model is in a different space than that for the belief γ_i for the heterogeneous model. In particular, μ_i specifies an element in Π_{-i} for each information set, which means that what i thinks about the continuation play can vary with the information set. Each μ_i is required to satisfy accordance, meaning the following:

Definition U0 A belief μ_i satisfies **accordance** if (i) $(\mu_i)_h$ 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} and (ii) 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.

As Claim 1 of Appendix A explains, we do not need to assume accordance in heterogenous RPCE because we have assumed each player's belief about the opponents' continuation strategies is the same at every information set. We say that $\pi_i \in \Pi_i$ is a best response to μ_i at $h \in H_i$ if the restriction of π_i to the subtree starting at h is optimal against the probability distribution over assessments and continuation strategies given by μ_i.[10]

---

[10]This is essentially the same definition as above, the only difference is that the domain of the best responses has been changed.

A belief model $U := (U_j)_{j \in I}$ is a profile of finite sets, where $U_j = \{u_j^1, \ldots, u_j^{K_j}\}$ with K_j being the number of elements in U_j. Each element in U_j is called a version. For each $j \in I$ and $k \in \{1, \ldots, K_j\}$ associate to u_j^k a pair (π_j^k, p_j^k), where $\pi_j^k \in \Pi_j$ and $p_j^k \in \Delta(\times_{j' \neq j} U_{j'})$, and write $u_j^k = (\pi_j^k, p_j^k)$. When there is no room for confusion, we omit superscripts that distinguish different versions in the same player role.

There are two main differences in the definitions of versions in the unitary and heterogeneous belief models. First, in unitary belief models, conjectures do not specify a probability measure over the player's own versions. Second, in the unitary model players are sure that only one actual version exists for each player role, but unsure which one is actual. In the heterogeneous model they assign probability one to a single version distribution for each player role. Thus, in the heterogenous model, a conjecture of $(\frac{1}{2} v_2^k + \frac{1}{2} v_2^l, \frac{1}{2} v_3^k + \frac{1}{2} v_3^l)$ means that $1/2$ of the player 2's are v_2^k and not that there is probability $1/2$ that all of them are v_2^k, which is allowed in the unitary model.[11]

Definition U1   (a) Given a belief model U, π^* is **generated** by a version profile $(\pi_i, p_i)_{i \in I} \in \times_{i \in I} U_i$ if for each i, $\pi_i = \pi_i^*$.

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

Definition U2 Given a belief model U, version $u_i = (\pi_i, p_i)$ is **self-confirming** with respect to π^* if $D_i(\pi_i, \pi_{-i}(u_{-i})) = D_i(\pi_i, \pi_{-i}^*)$ for all u_{-i} in the support of p_i.

Definition U3 Given a belief model U, version $u_i = (\pi_i, p_i)$ is **observationally consistent** if $p_i(\tilde{u}_{-i}) > 0$ implies, for each $j \neq i$, $\tilde{u}_j$ is self-confirming with respect to $\pi(u_i, \tilde{u}_{-i})$.

Using these notions, we define unitary rationalizable partition-confirmed equilibrium as follows:

---

[11]Similarly, $1/2$ of the player 3's are v_3^k.

Definition U4 π^* is a **unitary rationalizable partition-confirmed equilibrium** if there exist a belief model U and an actual version profile u^* such that the following conditions hold:

1. π^* is generated by u^*.

2. For each i and $u_i = (\pi_i, p_i)$, 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, u_i^* is self-confirming with respect to π^*.

4. For all i and u_i, u_i is observationally consistent.

3 Examples

In this section we illustrate heterogeneous RPCE with several examples. We first revisit Examples 1-3 to formalize the arguments provided there.

Example 1 (Mixed Equilibrium and Heterogeneous Beliefs)

We revisit the game of Figure 1 to explain why $\left(\left(\frac{1}{2}In_1, \frac{1}{2}Out_1\right), U_2, U_3\right)$ is not a unitary RPCE but is a heterogeneous RPCE.

To see that it is not a unitary RPCE, note that if it were, then by the self-confirming condition the actual version of player 1 must believe that player 2 and player 3 play (U_2, U_3) with probability one. But given this belief the only best response is to play action In_1 with probability one, so 1's strategy contradicts the best response condition.

However the profile is a heterogeneous RPCE. To see this, consider the following belief model[12]:

$$V_1 = \{v_1^1, v_1^2\} \quad \text{with} \quad v_1^1 = (Out_1, (v_1^1, v_2^2, v_3^2)), \ v_1^2 = (In_1, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, v_2^1, v_3^1));$$

---

[12]As in Fudenberg and Kamada (2015), 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}.

$$V_2 = \{v_2^1, v_2^2\} \quad \text{with} \quad v_2^1 = (U_2, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, v_2^1, v_3^1)), \ v_2^2 = (D_2, (v_1^1, v_2^2, v_3^2));$$

$$V_3 = \{v_3^1, v_3^2\} \quad \text{with} \quad v_3^1 = (U_3, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, v_2^1, v_3^1)), \ v_3^2 = (D_3, (v_1^1, v_2^2, v_3^2));$$

$$\phi_1(v_1^1) = \phi_1(v_1^2) = \frac{1}{2}, \ \phi_2(v_2^1) = 1, \ \phi_3(v_3^1) = 1.$$

It is easy to check that the RPCE conditions hold (note that v_2^2 and v_3^2 must believe that all player 1's play Out_1, because otherwise v_1^1's observational consistency would be violated).

Example 2 (Investor-Entrepreneur)

Here we revisit the investor-entrepreneur game of Figure 2. We first show that $(Out, Don't)$ cannot be a unitary RPCE. To see this, suppose the contrary. Note that the best response condition implies that the actual version of player 1 has to assign a strictly positive probability to a version v_2' of player 2 that plays $Solicit$ with strictly positive probability. But then observational consistency applied to the actual version of player 1 implies that the belief of v_2' assigns probability 1 to Out, which would make Solicit strictly suboptimal.

To show that $(Out, Don't)$ is a heterogeneous RPCE, consider the following belief model:

$$V_1 = \{v_1^1, v_1^2\} \quad \text{with} \quad v_1^1 = (Out, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, \tfrac{1}{2}v_2^2 + \tfrac{1}{2}v_2^3)),$$

$$v_1^2 = (Attend, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, \tfrac{1}{2}v_2^2 + \tfrac{1}{2}v_2^3));$$

$$V_2 = \{v_2^1, v_2^2, v_2^3\} \quad \text{with} \quad v_2^1 = (Don't, (v_1^1, v_2^2))$$

$$v_2^2 = (Don't, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, \tfrac{1}{2}v_2^2 + \tfrac{1}{2}v_2^3)), \ v_2^3 = (Solicit, (\tfrac{1}{2}v_1^1 + \tfrac{1}{2}v_1^2, \tfrac{1}{2}v_2^2 + \tfrac{1}{2}v_2^3));$$

$$\phi_1(v_1^1) = 1, \ \phi_2(v_2^1) = 1.$$

Fudenberg and Levine (1993a) show by example that there are heterogeneous self-confirming equilibria that are not unitary self-confirming equilibrium. The profile they construct uses mixed strategies, and the mixing is

necessary: If a strategy profile is a heterogeneous self-confirming equilibrium but is not a unitary self-confirming equilibrium, then it uses mixed strategies.

In contrast, in this example there is a heterogeneous RPCE in which the distribution of strategies generated by ϕ is pure, yet the observed play cannot be the outcome of a unitary RPCE.[13]

Example 3 (Heterogeneous RPCE with Pure Strategies)

Here we revisit the tax evasion example of Figure 3. To see that the profile $(Safe, N, Stay)$ cannot be a unitary RPCE, suppose the contrary. Then the actual version of the attorney (player 1) must play $Safe$, so by observational consistency her conjecture assigns probability 1 to versions of the IRS agent (player 2) whose assessment assigns probability 1 to the left node in 2's information set. By the best response condition these versions must play N. Then the coherent belief condition implies that the actual version of player 1 believes that 2 plays N, and by the best response condition she has to play $Risky$ instead of $Safe$ irrespective of her belief about the play by the tax evader (player 3).

To see that the profile is a heterogeneous RPCE, consider the following belief model:

$$V_1 = \{v_1^1, v_1^2\} \quad \text{with} \quad v_1^1 = (Safe, (v_1^2, v_2^2, v_3^2)), \ v_1^2 = (Risky, (v_1^2, v_2^2, v_3^3));$$

$$V_2 = \{v_2^1, v_2^2\} \quad \text{with} \quad v_2^1 = (N, (v_1^1, v_2^1, v_3^1)), \ v_2^2 = (E, (v_1^2, v_2^2, v_3^3));$$

$$V_3 = \{v_3^1, v_3^2, v_3^3\} \quad \text{with} \quad v_3^1 = (Stay, (v_1^1, v_2^2, v_3^1)),$$

$$v_3^2 = (Stay, (v_1^2, v_2^2, v_3^2)), \qquad v_3^3 = (Fire, (v_1^2, v_2^2, v_3^3));$$

$$\phi_1(v_1^1) = 1, \ \phi_2(v_2^1) = 1, \ \phi_3(v_3^1) = 1.$$

Notice that in this belief model, version v_1^2 has share 0, but each agent of

---

[13]Note that in a simultaneous-move game with a discrete terminal node partition, each agent's belief must be correct, so every heterogeneous RPCE is a Nash equilibrium.

version v_1^1 thinks that all other agents in his player role are version v_1^2.[14] This is possible, because if version v_1^2 was an actual version, he could not observe 3's play, so his belief about 3's play can be arbitrary. Given that all other agents are playing $Risky$, v_1^1 infers 2 should be playing E, and such a belief is "self-confirming" because v_1^1 does not observe 2's choice due to the terminal node partition.[15]

Note that players 1 and 2 have strict incentives to play the equilibrium actions, unlike in the heterogeneous RPCE in Example 2. Player 3 is indifferent, but one can replace his move with a simultaneous-move game by two players to avoid ties.

Note also that the construction here is different from that of Example 1, where each agent thinks that all other agents in the same role are playing in the same way as she does, while here each agent thinks that other agents in the same player role behave differently than herself. Example 7 in the Online Supplementary Appendix extends this idea to show that a heterogeneous RPCE can be different from a unitary RPCE because an actual version of one player role can conjecture that different versions in *another* player role play differently.

Example 4 (Inferring the Play of Other Agents in the Same Role)

[14]Hence each agent in v_1^1 thinks that he has measure zero. This is not necessary here: the same conclusion applies if v_1^1 believes the share of v_1^1's is strictly less than $1/2$. In Example 5, which has a weakly dominated strategy, it does matter that a version can think it has share 0.

[15]Player 3 has three versions although she has two actions because we need two versions who play $Stay$: the actual version who observes $Safe$, and a hypothetical version who observes $Risky$, which is needed so that v_1^1 is observationally consistent.

Figure 4

Here we show how knowledge of the payoff functions and the observation structure can rule out heterogeneous beliefs when neither of these forces would do so on its own, as agents in a given player role may be able to use their observations to make inferences about the play of other agents in their own role. In the game in Figure 4, the terminal node partitions are discrete. One might conjecture that some player 2's can play Out_2 while some play In_2 and some player 1's play In_1, as Out_2 prevents player 2 from observing 3 and 4's play. However, we claim that whenever a heterogeneous RPCE assigns strictly positive probability to In_1, player 2 plays In_2 with probability 1.

To see this, consider a heterogeneous RPCE such that 1 plays In_1 with a strictly positive probability. Fix a belief model that rationalizes this hetero-

geneous RPCE and fix an actual version v_2 of player 2. We show that v_2 must play In_2 with probability 1 in this heterogenous RPCE.

First, by the self-confirming condition, v_2's conjecture assigns positive probability to a version $\bar{v}_1$ of player 1 that plays In_1 with positive probability. Suppose that v_2's conjecture assigns probability zero to versions of player 2 that play In_2 with positive probability. Then, by observational consistency applied to v_2, $\bar{v}_1$ believes 2 plays Out_2 with probability 1. But this contradicts the best response condition for $\bar{v}_1$. Hence v_2's conjecture must assign positive probability to versions who play In_2 with positive probability. Pick one such version who plays In_2, and call it $\bar{v}_2$.

Since $\bar{v}_2$ must satisfy the best response condition, she must assign probability at least $\frac{5}{6}$ to (L_3, L_4). This in particular implies that $\bar{v}_2$'s belief must assign probability at least $\frac{5}{6}$ to L_3. But L_4 is the unique best response to a strategy that plays L_3 with probability at least $\frac{5}{6}$ (given that player 4 is on the path), so observational consistency applied to $\bar{v}_2$ and the best response condition for player 4 imply that $\bar{v}_2$'s belief must assign probability 1 to L_4, and by a symmetric argument it must assign 1 to L_3. But since $\bar{v}_2$ observes play by players 3 and 4, observational consistency applied to v_2 implies that v_2 assigns probability 1 to (L_3, L_4), and so the best response condition for v_2 implies that she must play In_2 with probability 1.

Note that extensive-form rationalizability alone does not preclude heterogeneity, as all actions of all players are extensive-form rationalizable. Now we show that common knowledge of the observation structure alone does not rule out heterogeneity either.

Consider requiring conditions 1 and 3 in the definition of heterogeneous RPCE and a weaker version of condition 2 where we require optimality only at the on-path information sets (all $h \in H(\pi_i^k, \pi_{-i})$ for all π_{-i} is in the support of γ_i^k). This concept would correspond to relaxing the unitary assumption of partition-confirmed equilibrium defined in FK. With this definition, player 2 is not assumed to know the payoff function of player 1, so he can believe that 3 and 4 play (R_3, R_4). Specifically, it is easy to check by inspection that all

the above conditions are satisfied in the following belief model:

$$V_1 = \{v_1^1\} \quad \text{with} \ \ v_1^1 = (In_1, (v_1^1, \frac{1}{2}v_2^1 + \frac{1}{2}v_2^2, v_3^1, v_4^1));$$

$$V_2 = \{v_2^1, v_2^2\} \quad \text{with} \ \ v_2^1 = (In_2, (v_1^1, \frac{1}{2}v_2^1 + \frac{1}{2}v_2^2, v_3^1, v_4^1)), \ \ v_2^2 = (Out_2, (v_1^1, \frac{1}{2}v_2^1 + \frac{1}{2}v_2^2, v_3^1, v_4^1));$$

$$V_3 = \{v_3^1\} \quad \text{with} \ \ v_3^1 = (L_3, (v_1^1, \frac{1}{2}v_2^1 + \frac{1}{2}v_2^2, v_3^1, v_4^1)),$$

$$V_4 = \{v_4^1\} \quad \text{with} \ \ v_4^1 = (L_4, (v_1^1, \frac{1}{2}v_2^1 + \frac{1}{2}v_2^2, v_3^1, v_4^1)),$$

$$\phi_1(v_1^1) = 1, \ \phi_2(v_2^1) = \phi_2(v_2^2) = \frac{1}{2}, \ \phi_3(v_3^1) = 1, \phi_4(v_4^1) = 1.$$

4 A "Purification" Result

In the heterogeneous model, the aggregate play of agents in each player role can correspond to a mixed (behavior) strategy. One standard interpretation of mixed strategies in equilibrium is that the mixing describes the aggregate play of a large population, with different agents in the same player role using different pure strategies. Here we show that this interpretation of mixed-strategy equilibrium also applies to heterogeneous RPCE with a continuum of agents in each player role, so that there exists a belief model with the share functions ϕ_i that describe the mass of each population i whose play and conjectures are generated by various versions that use pure strategies. The continuum of agents allows ϕ_i to take on any value between 0 and 1. In the next section we relate this continuum model to one with a large but finite population.

We say that a belief model (V, ϕ) rationalizes a heterogenous RPCE π^* if, given π^*, the four conditions in Definition H4 hold for (V, ϕ) for each i.

Remark 1 *Any heterogeneous RPCE can be rationalized with a belief model in which all versions use pure strategies.*

The proof of this remark is provided in the appendix.

5 Anonymous-Matching Games in Large Finite Populations

The interpretation of heterogeneous beliefs and play is that there are many agents in each player role. An alternative way of thinking of such situations is that every agent is a "player," but each period only a subset of the agents get to actually play; the agents who are not playing do not receive any feedback on what happened that period. In this way, we can identify an *anonymous-matching game* with any given extensive form. To do this, we view each of the agent k's in the role of player i as distinct players, so the anonymous-matching game has as many players as the original model has agents. Each period Nature picks $|I|$ players to anonymously participate in the game, where $|I|$ is the number of player roles in the original extensive form, and only one player is picked from each of the respective groups.

We will show that each heterogeneous RPCE of a given extensive-form game is an "approximate" unitary RPCE in the corresponding anonymous-matching game, where the approximation becomes arbitrarily close as the population of the anonymous-matching game becomes large.

5.1 Anonymous-Matching Games

Formally, given an extensive-form game Γ with a set of players $I = \{1, ..., |I|\}$ and the terminal node partitions $\mathbf{P} = (\mathbf{P}_1, \ldots, \mathbf{P}_{|I|})$, we define an **anonymous-matching game of** Γ parameterized by a positive integer T defined below, denoted $Y(\Gamma, T)$, as follows[16]:

1. The set of players is $J := \bigcup_{i \in I} J_i$, with $J_i := \{(i, 1), \ldots, (i, T)\}$, where T is a positive integer.

2. Nature N moves at the initial node, choosing $|I|$ players who will move at subsequent nodes. For each $i \in I$, a unique player is chosen from J_i

---

[16]We assume there is the same number of players in each player role; none of our results hinges on this assumption.

independently, according to the uniform distribution over J_i. Let the chosen player for each $i \in I$ be (i, r_i).

3. The chosen players, $((1, r_1), (2, r_2), \ldots, (n, r_n))$, play Γ, without knowing the identity of the opponents. Unchosen players receive the constant payoffs of 0. Formally,

   (a) Each node of $Y(\Gamma, T)$ is denoted $(x, (i, w_i)_{i \in I})$, where x is an element of X of Γ and w_i is the index of the agent in player role i who "plays" in the game that contains the node.[17]

   (b) For each player $(i, w_i) \in J_i$, nodes $(x, (j, w_j)_{j \in I})$ and $(x', ((i, w_i), (j, w_j')_{j \neq i}))$ are in the same information set if and only if x and x' are in the same information set of player i in Γ (This formalizes the idea that the identity of the matched agents cannot be observed).

   (c) For any $(i, w_i)_{i \in I}$, actions available at an information set that includes $(x, (i, w_i)_{i \in I})$ are the same as the actions available at an information set that includes x in Γ.

   (d) The payoff function is such that if a player in J_i is chosen and an action profile a of the chosen players (which lies in A) is realized, she receives a payoff identical to $g_i(z)$ where the action profile a leads to the terminal node z in Γ. If a player is not chosen, she receives the payoff of 0.

   (e) The terminal node partition is such that if a player is not chosen then she does not observe anything (except the fact that she was not chosen). If a player $(i, r_i) \in J_i$ is chosen and a terminal node $(z, (j, r_j)_{j \in I})$ is reached, all she knows is that some node $(z', ((i, r_i), (j, w_j')_{j \neq i}))$ for some z' and $w'_{-i} \in \times_{j \neq i} J_j$ is reached, where z' and z are in the same partition cell of $\mathbf{P}_i$ in Γ (In particular, she does not know the identity of the opponents).

---

[17]In this section, we use (i, w_i) to denote a generic agent in the role of player i.

5.2 Motivating ϵ-Unitary RPCE

Here we use two examples to motivate our use of an approximate version of unitary RPCE in the equivalence result.

**Example 5 (Heterogeneous RPCE with Dominated Strategies)**

Figure 5

Our definition allows each version to believe that the aggregate play of her player role does not assign positive mass to her own strategy. For example, even if v_i^k plays L_i, her belief may assign probability 1 to R_i. This reflects the premise that there is a continuum of agents in each player role and no one agent can change the aggregate distribution of play. This continuum model is meant to be an approximation of a large but finite population model. In Section 5, we formalize this idea of approximation by using ϵ-self-confirming and ϵ-observational consistency conditions, as opposed to the exact self-confirming and the observational consistency conditions. This example shows why some sort of approximate equilibrium notion is needed.

25

The game in Figure 5 has the same extensive form as in the game in Example 3, with a different payoff function for player 2. Notice that R_2 is weakly dominated.

We first show that 1 can play R_1 in a heterogeneous RPCE. To see this, consider the following belief model:

$$V_1 = \{v_1^1, v_1^2\} \quad \text{with} \quad v_1^1 = (R_1, (v_1^2, v_2^2, v_3^2)), \ v_1^2 = (L_1, (v_1^2, v_2^2, v_3^3));$$

$$V_2 = \{v_2^1, v_2^2\} \quad \text{with} \quad v_2^1 = (L_2, (v_1^1, v_2^1, v_3^1)), \ v_2^2 = (R_2, (v_1^2, v_2^2, v_3^2));$$

$$V_3 = \{v_3^1, v_3^2\} \quad \text{with} \quad v_3^1 = (L_3, (v_1^1, v_2^1, v_3^1)),$$

$$v_3^2 = (L_3, (v_1^2, v_2^2, v_3^2)), \ v_3^3 = (R_3, (v_1^2, v_2^2, v_3^3));$$

$$\phi_1(v_1^1) = 1, \ \phi_2(v_2^1) = 1, \ \phi_3(v_3^1) = 1.$$

Notice that in this belief model, the actual version v_1^1 of player 1 conjectures that version v_1^2 has share 1, which justifies his belief that player 2 is indifferent between two actions so can play a weakly dominated action R_2. One can check by inspection that all conditions in the definition of heterogeneous RPCE are met.

However, player 1 cannot play R_1 if we require that each version's belief has to assign a positive weight to her own strategy. To see this, notice that if this condition were imposed, observational consistency and Bayes rule would imply that each version v_1^k of player 1 assigns probability 1 to versions of player 2 whose assessments assign probability strictly less than $\frac{1}{2}$ to the node that follows L_1. By the best response condition, these versions must play L_2. But then from condition 2(i) of Definition H4 v_1^k would need to believe that 2 will play L_2 with probability 1, so that R_1 would give her a strictly smaller payoff than the maximal possible payoff from L_1, which contradicts the best response condition.

The point is that in a finite population model an agent's own actions can give her information about an opponent's belief and hence about their strategy, but such an inference is not captured by heterogeneous RPCE.

**Example 6 (Heterogeneous RPCE with Irrational-Valued Beliefs)**

Figure 6: Terminal node partitions are not depicted in the figure but are described in the text.

Another reason that approximation is needed is that the number of agents in our anonymous-matching games are integers, so all agents have rational probabilities of being chosen to play, but payoffs and the probability distribution over Nature's move need not take only rational values.[18] In the extensive-form game in Figure 6, each player only observes their own payoffs, except that player 1 observes the exact terminal node reached. One can check by inspection that $\left(M_1, \frac{1}{\sqrt{2}}L_2 + \left(1 - \frac{1}{\sqrt{2}}\right)R_2, L_3, L_4\right)$ is a heterogenous RPCE by verifying that the following belief model meets the necessary conditions:

$$V_1 = \{v_1^1, v_1^2, v_1^3\} \quad \text{with} \quad v_1^1 = \left(L_1, \left(v_1^1, v_2^1, v_3^1, v_4^1\right)\right), \ v_1^2 = \left(R_1, \left(v_1^2, v_2^2, v_3^2, v_4^2\right)\right),$$

$$v_1^3 = \left(M_1, \left(v_1^3, \frac{1}{\sqrt{2}}v_2^1 + \left(1 - \frac{1}{\sqrt{2}}\right)v_2^2, v_3^1, v_4^1\right)\right);$$

$$V_2 = \{v_2^1, v_2^2\} \quad \text{with} \quad v_2^1 = (L_2, (v_1^1, v_2^1, v_3^1, v_4^1)), \ v_2^2 = (R_2, (v_1^2, v_2^2, v_3^2, v_4^2));$$

$$V_3 = \{v_3^1, v_3^2\} \quad \text{with} \quad v_3^1 = (L_3, (v_1^1, v_2^1, v_3^1, v_4^1)), \ v_3^2 = (R_3, (v_1^2, v_2^2, v_3^2, v_4^2));$$

$$V_4 = \{v_4^1, v_4^2\} \quad \text{with} \quad v_4^1 = (L_4, (v_1^1, v_2^1, v_3^1, v_4^1)), \ v_4^2 = (R_4, (v_1^2, v_2^2, v_3^2, v_4^2));$$

---

[18]The point we make here would disappear if payoffs and the probabilities of Nature's moves are restricted to rational numbers.

$$\phi_1(v_1^3) = 1, \ \phi_2(v_2^1) = \frac{1}{\sqrt{2}}, \ \phi_3(v_3^1) = 1, \ \phi_4(v_4^1) = 1.$$

In this belief model, all (actual and hypothetical) versions of players 2, 3, and 4 think that either it is common knowledge that everyone plays left, or it is common knowledge that everyone plays right.

Now consider generating this outcome using an anonymous-matching game. First, notice that M_1 is a best response for player 1 only if she believes L_2 is chosen with probability $\frac{1}{\sqrt{2}}$. Second, if an agent represented by version u_2 of player 2 is indifferent between L_2 and R_2, then his belief has to assign probability $\frac{1}{2}$ to the profile (L_3, L_4). At the same time, since players 3 and 4 play a symmetric coordination game, u_2's belief can assign probability either 0, $\frac{1}{4}$, or 1 to (L_3, L_4).[19] None of these probabilities are even close to $\frac{1}{2}$, which is a contradiction. Hence, u_2 cannot be indifferent between L_2 and R_2. Since the agent we chose in player role 2 was arbitrary, we conclude that no agent of player 2, whether actual or hypothetical, is indifferent between L_2 and R_2. This has two implications.

First, if an agent of player 1 plays M_1, then she cannot assign probability 1 to a single realization of the version profiles of the agents. To see this, note that every agent of player 1 can only assign positive probability to agents playing a pure strategy. Hence, if an agent of player 1 assigns probability 1 to a single realization of the version profiles of agents, conjectures that each agent of player 2 is chosen with equal probability, and the number of agents is finite, then coherency implies that it is impossible for the agent of player 1 to believe that the probability of L_2 is $\frac{1}{\sqrt{2}}$.

The second implication is again about the need for an approximation. To see this, consider the agents corresponding to actual versions of player 1. These agents observe the true distribution of the play by the agents of player 2, which must take rational values since all agents of player 2 play pure strategies. The self-confirming condition then implies that it is impossible for the actual agents in player role 1 to believe that the probability of L_2 is $\frac{1}{\sqrt{2}}$. However, the belief

---

[19]In this game, observation of their own payoff is sufficient for players 3 and 4 to coordinate on a Nash equilibrium in any RPCE.

can become arbitrarily close to $\frac{1}{\sqrt{2}}$ when the number of agents becomes large.

5.3 Unitary ϵ-RPCE

Before Section 5, we have restricted attention to generalized one-move games. Anonymous-matching games do not fall into this class, because they allow two different actions of nature to lead to the same information set under some strategy profiles but not under others. For this reason, we will now need to impose a more general form of accordance.

Hereafter, we assume that any belief $\mu_i \in [\times_{h \in H_i} \Delta(\Delta(h) \times \Pi_{-i})] \times \Pi_{-i}$ in the unitary belief model satisfies the following condition:

Definition U5 A belief μ_i satisfies **convex structurally-consistent accordance** if it satisfies the following two conditions:

1. μ_i satisfies accordance.

2. For each $h \in H_i$ and each $(\alpha_i, \pi_{-i}) \in \text{supp}((\mu_i)_h)$, there exists a probability distribution $\gamma \in \Delta(\Pi_{-i})$ such that (i) there exists $\hat{\pi}_{-i} \in \text{supp}(\gamma)$ such that h is reachable under $\hat{\pi}_{-i}$, and (ii) α_i at h is derived by Bayes rule under γ.

In anonymous-matching games, Nature's move determines which agents get to play. The second condition in the definition imposes a restriction on the assessment regarding the probability distribution over the agents chosen by Nature at each off-path information set h: we require that this probability distribution does not change even after a deviation that leads to h (note that the deviator does not move at h because we still assume that each player moves only once at each path of the extensive form). The reason we allow for a distribution γ of strategies here is analogous to the argument for convex structural consistency in Kreps and Ramey (1987): the stronger condition that one single profile generates the beliefs is not compatible with Kreps-Wilson's (1982) consistency. Note that the second condition is moot in the generalized one-move games we have considered before Section 5.

29

To allow an approximation motivated by Examples 5 and 6 in a model with a finite number of agents, we relax the definitions of "self-confirming" and "observationally consistent."

Let $||\cdot||$ denote the supremum norm. We say that π_i represents $(\pi_{(i,1)}, \ldots, \pi_{(i,T)})$ if, for each $h_i \in H_i$ and $a \in A(h_i)$, $\pi(h_i)(a_i) = \frac{1}{T} \sum_{j=1}^{T} \pi_{(i,j)}(h_i)(a_i)$. We say that π_i ϵ-**represents** $(\pi_{(i,1)}, \ldots, \pi_{(i,T)})$ if there is π_i' such that π_i' represents $(\pi_{(i,1)}, \ldots, \pi_{(i,T)})$ and $||\pi_i - \pi_i'|| < \epsilon$.[20]

Definition U2(ϵ) Given a belief model U, version $u_i = (\pi_i, p_i)$ is ϵ-**self-confirming** with respect to π^* if $||D_i(\pi_i, \pi_{-i}(u_{-i})) - D_i(\pi_i, \pi_{-i}^*)|| < \epsilon$ for all u_{-i} in the support of p_i.

Definition U3(ϵ) Given a belief model U, version $u_i = (\pi_i, p_i)$ is ϵ-**observationally consistent** if $p_i(\tilde{u}_{-i}) > 0$ implies, for each $j \neq i$, $\tilde{u}_j$ is ϵ-self-confirming with respect to $\pi(u_i, \tilde{u}_{-i})$.

Using these notions, we define unitary ϵ-rationalizable partition-confirmed equilibrium.

Definition U4(ϵ) π^* is a **unitary ϵ-rationalizable partition-confirmed equilibrium** (**unitary ϵ-RPCE**) if there exist a belief model U and an actual version profile u^* such that the following conditions hold:

1. π^* is generated by u^*.

2. For each i and $u_i = (\pi_i, p_i)$, 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, u_i^* is ϵ-self-confirming with respect to π^*.

4. For all i and u_i, u_i is ϵ-observationally consistent.

---

[20]When $||\cdot||$ applies to the distance between two behavioral strategies π_i and π_i', it is given by $||\pi_i - \pi_i'|| = \max_{h \in H_i, a \in A(h)} |\pi_i(h)(a) - \pi_i'(h)(a)|$.

5.4 The Equivalence Theorem

Theorem 1 *For any $\epsilon > 0$, Γ, and a heterogeneous RPCE π^* of Γ, there exist T and a pure unitary ϵ-RPCE π^{**} of $Y(\Gamma, T)$ such that for each $i \in I$, π_i^* ϵ-represents $(\pi_{(i,1)}^{**}, \ldots, \pi_{(i,T)}^{**})$.*

The proof is provided in the appendix. In outline, the way the proof handles heterogeneous RPCE with irrational probabilities is as follows. Let $\tilde{u}_{(1,j)}^3$ be the version in the constructed unitary belief model who corresponds to v_1^3 in the heterogeneous belief model in Example 6. We have to construct a conjecture and an associated belief so that coherency holds and the strategy of $\tilde{u}_{(1,j)}^3$ is an exact best response. To do so, we suppose that $\tilde{u}_{(1,j)}^3$ is not certain about what the share functions are. For example, if $T = 100$ then we let the conjecture of $\tilde{u}_{(1,j)}^3$ assign probability 1 to the event that 70 agents are the versions who play L_2 and 29 agents are the versions who play R_2, assign probability $\frac{1}{\sqrt{2}} - \frac{70}{100}$ to the event that the remaining one agent is the version who plays L_2, and assign the remaining probability to the event that the remaining one agent is the version who plays R_2.[21] With this construction any point in the support of the belief of $\tilde{u}_{(1,j)}^3$ is close to the corresponding point in the support of v_1^3 (because 99 agents play deterministically) and the belief is essentially unchanged so playing M_1 is still a best response (because we allow the remaining one agent to be either one of the two possible versions with the probability computed from the original mixing probability of player 2). There is more subtlety in making sure the best-response condition holds also at zero-probability information sets, which we will detail in the proof.

Theorem 1 relaxes the self-confirming and observational consistency conditions to approximate heterogeneous RPCE with unitary RPCE played by a finite number of agents. The conclusion of the theorem holds with the exact (e.g. $\epsilon = 0$) versions of self-confirming and observational consistency if instead we specify that the probability distribution over agents in the anonymous-

---

[21] If $\tilde{u}_{(1,j)}^3$ conjectures that each agent of player 2 plays either L_2 or R_2 with probabilities $\frac{1}{\sqrt{2}}$ and $1 - \frac{1}{\sqrt{2}}$, respectively, then the ϵ-self-confirming condition would be violated because this conjecture would assign strictly positive probability to the event that no agent of player 2 plays L_2.

matching game depends on the unitary RPCE that is being replicated. We do not state this version of the result formally, as we do not find it satisfactory to vary the probability distribution over agents to match the target equilibrium.

In the Online Appendix, we discuss our choice of approximation criterion used in defining ϵ-observational consistency, and explain the implications of an alternative.

6   Conclusion

This paper has developed an extension of RPCE to allow for heterogeneous beliefs, both on the part of the agents who are objectively present, and also in the "versions" that represent mental states agents think other agents can have. This extension allows the model to fit the heterogeneity that naturally arises when there are many agents in the role of each player, as implicitly assumed by most learning theories and implemented in the random-matching protocols of most game theory experiments. It also permits RPCE to be restricted to pure strategies without loss of generality. The paper explored the impact of heterogeneous beliefs in various examples. It also showed how heterogeneous RPCE relates to the unitary RPCE of a larger anonymous-matching game with many agents in each player role.

This paper is only the first look at the new issues posed by heterogeneous beliefs. It would be interesting to explore some of the complications that we have avoided here, such as the possibility of a player's beliefs being a correlated distribution over the opponents' strategies, defining heterogeneous RPCE for a class of games larger than generalized one-move games, and a full dynamic model of learning. It would also be interesting to know more about the relationship between unitary and heterogeneous RPCE. The examples that we provided in this paper are only a first step in this direction.

32

References

[1] Battigalli, P., S. Cerreia-Viglio, F. Maccheroni and M. Marinacci. (2015): "Selfconfirming Equilibrium and Model Uncertainty" *American Economic Review*, 105, 646-677.

[2] 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.

[3] Dekel, E., D. Fudenberg and D. Levine (1999): "Payoff Information and Self-Confirming Equilibrium," *Journal of Economic Theory*, 89, 165-185.

[4] Dekel, E., D. Fudenberg and D. Levine (2004): "Learning to Play Bayesian Games," *Games and Economic Behavior*, 46, 282-303.

[5] Dekel, E. and M. Siniscalchi (2014): "Epistemic Game Theory," Forthcoming in the *Handbook of Game Theory*, 4.

[6] Esponda, I. (2013): "Rationalizable Conjectural Equilibrium: A Framework for Robust Predictions," *Theoretical Economics*, 8: 467-501.

[7] Fudenberg, D. and K. He (2016): "Gittins Equilibrium in Signaling Games," mimeo.

[8] Fudenberg, D. and Y. Kamada (2015): "Rationalizable Partition-Confirmed Equilibrium," *Theoretical Economics*, 10: 775-806.

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

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

[11] Fudenberg, D. and D. K. Levine (1993b) "Steady State Learning and Nash Equilibrium," *Econometrica*, 61: 547-573.

[12] Fudenberg, D. and D. K. Levine (1997) "Measuring Player's Losses in Experimental Games," *Quarterly Journal of Economics*, 112: 479-506.

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

[14] Kreps, D. and G. Ramey (1987): "Structural Consistency, Consistency, and Sequential Rationality" Econometrica 55, 1331-1348.

[15] Kreps, D. and F. Wilson (1982): "Sequential Equilibria," *Econometrica* 50, 863-894.

[16] Prasnikar, V. and A. Roth (1992): "Considerations of Fairness and Strategy: Experimental Data from Sequential Games," *Quarterly Journal of Economics*, 107: 865-888.

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

A Independence and Accordance

Recall that unitary RPCE is only defined for 1-move games. In that solution concept, a belief μ_i belongs to the space $[\times_{h \in H_i} \Delta(\Delta(h) \times \Pi_{-i})] \times \Delta(\Pi_{-i})$. We denote $(\mu_i)_h$ the coordinate of μ_i that corresponds to h, and $b(\mu_i)$ the last coordinate that does not correspond to any particular information sets.

Although we do not use it in the next claim, keep in mind that the space of the beliefs in the heterogeneous model is $[\times_{h \in H_i} \Delta(h)] \times \Delta(\Pi)$.

Claim 1 *Suppose that under a belief μ_i, $b(\mu_i)$ assigns probability one to a single strategy profile for i's opponents, π^*_{-i}. Suppose also that for every $h \in H_i$, $(\mu_i)_h$ assigns probability one to a single assessment-strategy profile pair such that the strategy profile to which probability one is assigned is π^*_{-i}. Then, accordance holds.*

Proof. Accordance requires two conditions. We check them one by one.

The first condition of accordance requires that $(\mu_i)_h$ 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}. Since $b(\mu_i)$ and $(\mu_i)_h$ for each h assigns probability 1 to the same strategy profile for $-i$, this part holds.

The second condition requires that for each $h \in H_i$, 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 all h' after h.

Now, $(\mu_i)_h$ assigns positive probability only to π^*_{-i}. Also, π^*_{-i} is in the support of $b(\mu_i)$. Thus we can always take $\tilde{\pi}_{-i} = \pi^*_{-i}$ to satisfy the equality. ∎

B   Proof of Remark 1

Proof.

Fix a heterogeneous RPCE π^* and a belief model (V, ϕ) that rationalizes it. Pick any version $v_i^k = (\pi_i^k, q_i^k)$ in V, let σ_i^k be a mixed strategy that induces π_i^k,[22] and suppose that σ_i assigns positive probability to K distinct pure strategies. We construct copies of version v_i^k, each playing a distinct pure strategy in the support of σ_i^k. The copy corresponding to pure strategy s_i, denoted $v_i^k(s_i)$, plays s_i and has the same belief as v_i^k.

To construct the conjectures in the new belief model from the conjectures in the old one, we suppose that all of the copies corresponding to v_i^k have the same conjecture $\bar{q}_i^k$, where $\bar{q}_i^k(v_j^l(s_j)) = (q_i^k)_j(v_j^l) \cdot \sigma_j^l(s_j)$ for all $v_j^l \in V_j$ and all $s_j \in \text{supp}(\sigma_j^l)$.[23] Finally, denoting the share function in the new belief model by $\bar{\phi}$, we let $\bar{\phi}_i(v_i^k(s_i)) = \phi_i(v_i^k) \cdot \sigma_i^k(s_i)$ for each $s_i \in \text{supp}(\sigma_i^k)$.

It is straightforward to check that with this construction the new belief model rationalizes the original heterogeneous RPCE. ∎

---

[22] This mixed strategy exists from Kuhn's theorem; see for example the proof of Theorem 4 in Fudenberg and Levine (1993a).

[23] With a slight abuse of notation, we denote by $q_i^k(v_j^l)$ the weight on v_j^l of the j'th coordinate of q_i^k.

C   Proof of Theorem 1

Proof.

Fix $\epsilon > 0$ and a heterogeneous RPCE of Γ, π^*. By Remark 1 there exists a belief model (V, ϕ) that rationalizes π^* such that all versions in the belief model use pure strategies. Fix one such belief model. For each i and k, let $\gamma_i^k = (\alpha_i^{[i,k]}, \pi^{[i,k]})$ be the belief of v_i^k used in condition 2 of the definition of heterogeneous RPCE. Pick an integer T such that $T > \max\{\frac{2(\max_{i \in I} |V_i|)(\#A)^2}{\epsilon}, \frac{1}{\underline{G}}\}$, where

$$\underline{G} = \min_{i \in I} \left(\min_{k \in \{1,\dots,|V_i|\}} \left(\max_{\pi_i \in \Pi_i} \min_{z \in Z(\pi_i, \pi_{-i}^{[i,k]})} p(\pi_i, \pi_{-i}^{[i,k]})(z) \right) \right),$$

$p(\pi)(z)$ is the probability that $z \in Z$ is reached under π, and $Z(\pi)$ is the set of terminal nodes z such that $p(\pi)(z) > 0$. To prove the claim we need to construct a belief model U and an actual version profile u^* for the game $Y(\Gamma, T)$ such that there is a pure unitary ϵ-RPCE π^{**} of $Y(\Gamma, T)$ where for each $i \in I$, π_i^* ϵ-represents $(\pi_{(i,1)}^{**}, \dots, \pi_{(i,T)}^{**})$.

a) Constructing the belief model

For each $i \in I$ and each $(i, j) \in J_i$, define $U_{(i,j)} = \{\tilde{u}_{(i,j)}(v_i^k) | v_i^k \in V_i\}$, where $\tilde{u}_{(i,j)}(v_i^k) = (\tilde{\pi}_{(i,j)}^k, \tilde{p}_{(i,j)}^k)$ and we define $\tilde{\pi}_{(i,j)}^k$ and $\tilde{p}_{(i,j)}^k$ in what follows. Below we simply denote $\tilde{u}_{(i,j)}(v_i^k)$ by $u_{(i,j)}^k$.

First, $\tilde{\pi}_{(i,j)}^k = \pi_i^k$;[24] note this is a pure strategy.

Second, we let $\tilde{p}_{(i,j)}^k$ be independent, and abuse notation to denote by $\tilde{p}_{(i,j)}^k(u_{(n,m)})$ the probability assigned to $u_{(n,m)}$ by the conjecture of $v_{(i,j)}^k$. That is, $\tilde{p}_{(i,j)}^k((\bar{u}_{(n,m)})_{(n,m) \neq (i,j)}) = \prod_{(n,m) \neq (i,j)} \tilde{p}_{(i,j)}^k(\bar{u}_{(n,m)})$ for each $(\bar{u}_{(n,m)})_{(n,m) \neq (i,j)} \in U_{-(i,j)}$. Similarly, we abuse notation to write $q_i^k(v_n^m)$ (recall that q_i^k is necessarily independent by definition).

Below we specify $\tilde{p}_{(i,j)}^k$ in the way we described in the example of 100 agents before this proof. In that method, there are 70 agents for whom $\tilde{p}_{(i,j)}^3$ assigns

---

[24]Recall that each player j_i in $Y(\Gamma, T)$ has the same number of information sets as player i in Γ; here we abuse notation to use the same notation for an information set h in Γ and the information set in $Y(\Gamma, T)$ of player j_i that includes the nodes corresponding to the nodes included in h.

probability one to a version who plays action L_2, 29 agents for whom it assigns probability zero to such a version, and 1 agent for whom the probability is in $(0, 1)$. The cases (i), (ii), and (iii) below correspond to these three cases, respectively. The way we specify probabilities for case (iii) is clarified in (iii)-(a), (iii)-(b), and (iii)-(c).

For all $(n, m) \in \left(\bigcup_{n' \in I} J_{n'} \right) \setminus \{(i, j)\}$ and all $l \in \{1, \ldots, |V_n|\}$, we set

(i) $\quad \tilde{p}_{(i,j)}^k(u_{(n,m)}^l) = 1$ if $\sum_{l' < l} \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor < m \leq \sum_{l' \leq l} \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor$,

(ii) $\quad \tilde{p}_{(i,j)}^k(u_{(n,m)}^l) = 0$ if $m \leq \sum_{l' < l} \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor$

$$\text{or} \ \sum_{l' \leq l} \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor < m \leq \sum_{l' \leq |V_n|} \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor,$$

(iii) $\quad \tilde{p}_{(i,j)}^k(u_{(n,m)}^l) \in [0, 1]$ if $\sum_{l' \leq |V_n|} \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor < m \leq T.$

To define $\tilde{p}_{(i,j)}^k$ for case (iii) concretely, for each $n \in I$ and $l' \in \{1, \ldots, |V_n|\}$, let

$$f(l'; n, q_i^k) = T \cdot q_i^k(v_n^{l'}) - \left\lfloor T \cdot q_i^k(v_n^{l'}) \right\rfloor.$$

That is, $f(l'; n, q_i^k)$ is the error that the approximation in (i) and (ii) above miss out. More specifically, (i) and (ii) assign too small a weight for each possible version in the support of the conjecture, and $f(l'; n, q_i^k)$ is the probability that needs to be added to make the conjecture exactly in line with the original conjecture q_i^k. Now we allocate these probabilities to remaining agents considered in (iii). To do this, we define $l(w; n, q_i^k)$ for each $w \in \mathbb{N}$ with $w \leq \sum_{l' \leq |V_n|} f(l'; n, q_i^k)$ as follows:

$$\sum_{l' < l(w; n, q_i^k)} f(l'; n, q_i^k) < w \leq \sum_{l' \leq l(w; n, q_i^k)} f(l'; n, q_i^k).$$

That is, $l(w; n, q_i^k)$ is the maximum number of versions such that the sum of the error probabilities can be no more than w, when we add these errors in the order of the indices of the versions.

Then we define

(iii)-(a) $\quad \tilde{p}^k_{(i,j)}(u^l_{(n,m)}) = \left(\sum_{l' \leq l} f(l'; n, q^k_i) \right) - (w-1)$

$$\text{if } m = \sum_{l' \leq |V_n|} \left\lfloor T \cdot q^k_i(v^{l'}_n) \right\rfloor + w \text{ and } l = l(w-1; n, q^k_i),$$

(iii)-(b) $\quad \tilde{p}^k_{(i,j)}(u^l_{(n,m)}) = f(l; n, q^k_i)$

$$\text{if } m = \sum_{l' \leq |V_n|} \left\lfloor T \cdot q^k_i(v^{l'}_n) \right\rfloor + w \text{ and } l(w-1; n, q^k_i) < l < l(w; n, q^k_i),$$

(iii)-(c) $\quad \tilde{p}^k_{(i,j)}(u^l_{(n,m)}) = w - \left(\sum_{l' < l} f(l'; n, q^k_i) \right)$

$$\text{if } m = \sum_{l' \leq |V_n|} \left\lfloor T \cdot q^k_i(v^{l'}_n) \right\rfloor + w \text{ and } l = l(w; n, q^k_i).$$

Note that (i,j) has only $T-1$ opponents in player role i, so the above specification of the belief may not give rise to the conjecture that exactly corresponds to the one in the heterogenous model. However it will not lead to violation of best response condition, as beliefs about the strategy of agents of player i do not affect the expected payoff of an agent in player role i.

Last, we construct a belief of $u^k_{(i,j)}$, denoted $\tilde{\mu}^k_{(i,j)}$, that is used to satisfy the best response condition. We let $\tilde{\mu}^k_{(i,j)}$ to be defined by the following rule. First,

$$b_{(i,j)}(\tilde{\mu}^k_{(i,j)})(\pi_{-(i,j)}) = \sum_{\substack{u_{-(i.j)} \in \times_{(n,m) \neq (i,j)} U_{(n,m)}: \\ \pi_{-(i,j)} = \pi_{-(i,j)}(u_{-(i,j)})}} \tilde{p}^k_{(i,j)}(u_{-(i,j)}).$$

Second, $\left(\tilde{\mu}^k_{(i,j)} \right)_h$ is computed by Bayes rule under $b_{(i,j)}(\tilde{\mu}^k_{(i,j)})$ if $h \in H(\hat{\pi}_{-i})$ (note that Bayes rule induces a well-defined probability distribution at such h under $b_{(i,j)}$ because $T > \frac{1}{G}$). For $h \notin H(\hat{\pi}_{-i})$, we set

$$\left(\tilde{\mu}^k_{(i,j)} \right)_h \left(\hat{\hat{\alpha}}_i(h), (\hat{\hat{\pi}}_{(n,m)})_{(n,m) \in \bigcup_{w \neq i} J_w} \right) = 1$$

where $\hat{\hat{\pi}}_{(n,m)}(h') = \hat{\pi}_n(h')$ holds for all $h' \in H_n$ for all $n \in I$ and

$$\hat{\hat{\alpha}}_i(h)(x, ((i,j), (n, r_n)_{n \neq i})) = \frac{1}{T^{|I|-1}} \hat{\alpha}_i(h)(x)$$

for each $(n, r_n)_{n \neq i} \in \times_{n \neq i} J_n$ and $(x, ((i,j), (n, r_n)_{n \neq i})) \in h$.

b) Constructing the actual versions

We specify the actual versions u^* as follows: For each $i \in I$ and each (i,j), we set

$$u^*_{(i,j)} = u^k_{(i,j)} \text{ if } \sum_{k' < k} \left\lfloor T \cdot \phi_i(v_i^{k'}) \right\rfloor < m \leq \sum_{k' \leq k} \left\lfloor T \cdot \phi_i(v_i^{k'}) \right\rfloor,$$

$$u^*_{(i,j)} = u^l_{(i,j)} \text{ if } \sum_{k' \leq |V_i|} \left\lfloor T \cdot \phi_i(v_i^{k'}) \right\rfloor < m \leq T, \quad \phi_i(v_i^l) > 0 \text{ and } \phi_i(v_i^{l'}) = 0 \text{ for all } l' < l.$$

Let $\pi^{**} = \pi(u^*)$.

c) Checking that the conditions of unitary ϵ-RPCE hold

Since $\frac{T - \sum_{k' \in \mathbb{N}, k' \leq |V_i|} \left\lfloor T \cdot \phi_i(v_i^{k'}) \right\rfloor}{T} \leq \frac{|V_i|}{T} < \epsilon$, it is straightforward that π_i^* $\frac{1}{T}$-represents $(\pi^{**}_{(i,1)}, \ldots, \pi^{**}_{(i,T)})$ for each i. Also, by definition π^{**} is generated by u^*. Coherency holds for each $i \in I$, each $(i,j) \in J_i$ and each $k \in \{1, \ldots, |V_i|\}$ by the construction of $\tilde{\mu}^k_{(i,j)}$. Accordance holds by construction. Moreover, the best response condition holds by construction (recall that randomization is conducted independently across players in the construction of $\tilde{p}^k_{(i,j)}$). Thus it remains to check the self-confirming condition and the observational consistency condition. To this end, we first note that, for any Γ and T, $D_{(i,j)}$ in the model $Y(\Gamma, T)$ can be seen as an element in the same space as D_i in the model Γ by the construction of the terminal node partitions in $Y(\Gamma, T)$. Henceforth, we abuse notation to write $D_i = D_{(i,j)}$.

The self-confirming condition is satisfied in the original heterogeneous RPCE, so for each v_i^k in the support of ϕ_i, there exists $\tilde{\pi}_{-i} \in \Pi_{-i}$ such that (i) (V, ϕ)

induces $\tilde{\pi}_j$ for version v_i^k for each $j \neq i$ and (ii) $D_i(\pi_i^k, \tilde{\pi}_{-i}) = D_i(\pi_i^k, \pi_{-i}^*)$.

First, by the construction of $\tilde{\pi}_{(i,j)}^k$ and $\tilde{p}_{(i,j)}^k$ and Claim 2 that we present below (i) implies:

$$||D_i(\tilde{\pi}_{(i,j)}^k, \pi_{-(i,j)}(u_{-(i,j)})) - D_i(\pi_i^k, \tilde{\pi}_{-i})|| \leq (\#A)^2 ||(\tilde{\pi}_{(i,j)}^k, \hat{\pi}_{-i}) - (\pi_i^k, \tilde{\pi}_{-i})|| \leq (\#A)^2 \frac{\max_{n \neq i} |V_n|}{T} < \frac{\epsilon}{2}$$

for each $u_{-(i,j)}$ in the support of $\tilde{p}_{(i,j)}^k$, where $\hat{\pi}_n$ represents $(\pi_{-(i,j)}(u_{-(i,j)}))_{(n,m) \in J_n}$ for each $n \neq i$, where $A = \times_{i \in I} \bigcup_{h \in H_i} A(h)$.

Second, by the construction of the actual versions u^* and Claim 2, we have that

$$||D_i(\pi_i^k, \pi_{-i}^*) - D_i(\pi_{(i,j)}^k, \pi_{-(i,j)}^{**})|| \leq (\#A)^2 ||(\pi_i^k, \pi_{-i}^*) - (\pi_{(i,j)}^k, \hat{\hat{\pi}}_{-i})|| \leq (\#A)^2 \frac{\max_{n \neq i} |V_n|}{T} < \frac{\epsilon}{2},$$

where $\hat{\hat{\pi}}_n$ represents $(\pi_{-(i,j)}^*)_{(n,m) \in J_n}$ for each $n \neq i$.

Thus, by the triangle inequality,

$$||D_i(\pi_{(i,j)}^k, \pi_{-(i,j)}(u_{-(i,j)})) - D_i(\pi_{(i,j)}^k, \pi_{-(i,j)}^{**})|| \leq$$

$$||D_i(\pi_{(i,j)}^k, \pi_{-(i,j)}(u_{-(i,j)})) - D_i(\pi_i^k, \tilde{\pi}_{-i})|| + ||D_i(\pi_i^k, \tilde{\pi}_{-i}) - D_i(\pi_i^k, \pi_{-i}^*)||$$

$$+ ||D_i(\pi_i^k, \pi_{-i}^*) - D_i(\pi_{(i,j)}^k, \pi_{-(i,j)}^{**})|| < \frac{\epsilon}{2} + 0 + \frac{\epsilon}{2} = \epsilon$$

for each $u_{-(i,j)}$ in the support of $\tilde{p}_{(i,j)}^k$.

Thus, $v_i = (\pi_i, p_i)$ is ϵ-self-confirming with respect to π^*.

The observational consistency condition is satisfied in the original heterogeneous RPCE, so for each v_i^k, $q_i^k(v_n^l) > 0$ implies that there exists $\hat{\pi}_{-n} \in \Pi_{-n}$ such that (i') (V, ϕ) induces $\hat{\pi}_w$ for v_i^k for each $w \neq n$ and (ii') there exists $\check{\pi}_{-n} \in \Pi_{-n}$ such that (ii')-(i) (V, ϕ) induces $\check{\pi}_w$ for version v_n^l for each $w \neq n$ and (ii')-(ii) $D_n(\pi_n(v_n^l), \check{\pi}_{-n}) = D_n(\pi_n(v_n^l), \hat{\pi}_{-n})$.

First, by the construction of $\tilde{\pi}_{(n,m)}^l$ and $\tilde{p}_{(n,m)}^l$ and Claim 2, we have that (ii')-(i) implies:

$$||D_n(\pi_{(n,m)}(u_{(n,m)}^l), \pi_{-(n,m)}(u_{-(n,m)})) - D_n(\pi_n(v_n^l), \check{\pi}_{-n})|| \leq$$

$$(\#A)^2||(\pi_{(n,m)}(u^l_{(n,m)}), \dot{\pi}_{-n}) - (\pi_n(v^l_n), \breve{\pi}_{-n})|| \leq (\#A)^2 \frac{\max_{w \neq n}|V_w|}{T} < \frac{\epsilon}{2}$$

for each $u_{-(n,m)}$ in the support of $\tilde{p}^l_{(n,m)}$, where $\dot{\pi}_w$ represents $(\pi_{-(n,m)}(u_{-(n,m)}))_{(w,r) \in J_w}$ for each $w \neq n$.

Second, by the construction of $\tilde{p}^k_{j_i}$ and Claim 2, we have that (i') implies:

$$||D_n(\pi_n(v^l_n), \hat{\pi}_{-n}) - D_n(\pi_{(n,m)}(u^l_{(n,m)}), (\pi(v_{(i,j)}, \tilde{v}_{-(i,j)}))_{-(n,m)})|| \leq$$

$$(\#A)^2||(\pi_n(v^l_n), \hat{\pi}_{-n}) - (\pi_{(n,m)}(u^l_{(n,m)}), \ddot{\pi}_{-n})|| \leq (\#A)^2 \frac{\max_{w \neq n}|V_w|}{T} < \frac{\epsilon}{2}$$

where $\ddot{\pi}_w$ represents $(\pi(u_{(i,j)}, \tilde{u}_{-(i,j)}))_{(w,r) \in J_w}$ for each $w \neq n$.

Thus, by the triangle inequality, for all $\tilde{u}_{(n,m)}$ in the support of $\tilde{p}^k_{(i,j)}$, it must be the case that for all $u_{-(n,m)}$ in the support of the conjecture of $\tilde{u}_{(n,m)}$,

$$||D_n(\pi_{(n,m)}(\tilde{u}_{(n,m)}), \pi_{-(n,m)}(u_{-(n,m)})) - D_n(\pi_{(n,m)}(\tilde{u}_{(n,m)}), (\pi(u_{(i,j)}, \tilde{u}_{-(i,j)}))_{-(n,m)})|| \leq$$

$$||D_n(\pi_{(n,m)}(\tilde{u}_{(n,m)}), \pi_{-(n,m)}(u_{-(n,m)})) - D_n(\pi_n(v^l_n), \breve{\pi}_{-n})|| + ||D_n(\pi_n(v^l_n), \breve{\pi}_{-n}) - D_n(\pi_n(v^l_n), \hat{\pi}_{-n})||$$

$$+ ||D_n(\pi_n(v^l_n), \hat{\pi}_{-n}) - D_n(\pi_{(n,m)}(\tilde{u}_{(n,m)}), (\pi(u_{(i,j)}, \tilde{u}_{-(i,j)}))_{-(n,m)})||$$

$$< \frac{\epsilon}{2} + 0 + \frac{\epsilon}{2} = \epsilon.$$

Thus, u_{j_i} is ϵ-observationally consistent. ∎

Finally, we provide the statement and the proof of the claim in the above proof.

Claim 2 *For all $\pi, \pi' \in \Pi$ and $i \in I$, $||D_i(\pi) - D_i(\pi')|| \leq (\#A)^2 \cdot ||\pi - \pi'||$ holds.*

Proof. First we show that $||d(\pi) - d(\pi')|| \leq |A| \cdot ||\pi - \pi'||$ for any $\pi, \pi' \in \Pi$. To see this, fix π and π', and let $||\pi - \pi'|| = \epsilon$. Let $\tilde{A}(z)$ be the set of actions that are taken to reach $z \in Z$, $h(a_j)$ be the information set such that action a_j can be taken, and $j(a_j)$ be the player such that $h(a_j) \in H_{j(a_j)}$. Note that

$\tilde{A}(z)$ is finite for each $z \in Z$. For any $z \in Z$,

$$|d(\pi)(z) - d(\pi')(z)| = \left| \prod_{a_j \in \tilde{A}(z)} \pi_{j(a_j)}(h(a_j))(a_j) - \prod_{a_j \in \tilde{A}(z)} \pi'_{j(a_j)}(h(a_j))(a_j) \right|$$

$$\leq \left| \prod_{a_j \in \tilde{A}(z)} \max\{\pi_{j(a_j)}(h(a_j))(a_j), \pi'_{j(a_j)}(h(a_j))(a_j)\} \right.$$

$$\left. - \prod_{a_j \in \tilde{A}(z)} \min\{\pi_{j(a_j)}(h(a_j))(a_j), \pi'_{j(a_j)}(h(a_j))(a_j)\} \right|$$

$$\leq \left| \prod_{a_j \in \tilde{A}(z)} \max\{\pi_{j(a_j)}(h(a_j))(a_j), \pi'_{j(a_j)}(h(a_j))(a_j)\} \right.$$

$$\left. - \prod_{a_j \in \tilde{A}(z)} (\max\{\pi_{j(a_j)}(h(a_j))(a_j), \pi'_{j(a_j)}(h(a_j))(a_j)\} - \varepsilon) \right|$$

$$\leq 1 - (1 - \varepsilon)^{\#A} \leq \#A \cdot \varepsilon = \#A \, \|\pi - \pi'\| \, .$$

Hence,

$$||d(\pi) - d(\pi')|| \leq \#A \cdot ||\pi - \pi'||. \tag{1}$$

Next, for any $i \in I$,

$$||D_i(\pi) - D_i(\pi')|| = \max_{P_i^l \in \mathbf{P}_i} |D_i(\pi)(P_i^l) - D_i(\pi')(P_i^l)| = \max_{P_i^l \in \mathbf{P}_i} \left| \sum_{z \in P_i^l} (d(\pi)(z) - d(\pi')(z)) \right|$$

$$\leq \max_{P_i^l \in \mathbf{P}_i} \sum_{z \in P_i^l} |d(\pi)(z) - d(\pi')(z)| \leq \sum_{z \in Z} |d(\pi)(z) - d(\pi')(z)|$$

$$\leq \#Z \cdot \max_{z \in Z} |d(\pi)(z) - d(\pi')(z)| \leq \#A \cdot ||d(\pi) - d(\pi')||. \tag{2}$$

Combining equations (1) and (2), we have that $||D_i(\pi) - D_i(\pi')|| \leq (\#A)^2 \cdot ||\pi - \pi'||.$ ∎

42