

Common Certainty of Rationality Revisited

Stephen Morris*

Satoru Takahashi†

February 18, 2011

Abstract

In conventional epistemic analysis of solution concepts in complete information games, complete information is implicitly interpreted to mean common certainty of (i) a mapping from action profiles to outcomes; (ii) players' (unconditional) preferences over outcomes; and (iii) players' preferences over outcomes conditional on others' actions. We characterize a new solution concept—preference-correlated rationalizability—that captures common certainty of (i) and (ii) but not (iii). We show that it is badly behaved, with failures of upper hemicontinuity giving rise to counter-intuitive results. We discuss restrictions that restore well-behaved results.

1 Introduction

What solution concept in game theory embodies the implications of common certainty of payoffs (i.e., complete information) and common certainty of rationality?¹ Some formal approaches to this question were developed in the 1980s (Bernheim (1986), Brandenburger and Dekel (1987) and Tan and Werlang (1988)). With no additional assumptions beyond complete information and common certainty of rationality, one could conclude only that players would choose actions that survive iterated deletion of strictly dominated actions. This is equivalent to assuming that players choose correlated rationalizable actions—i.e., actions that survive iterated deletion of actions that are not a best response to any conjecture (perhaps correlated) over the remaining actions of the opponents.

Under correlated rationalizability, the actions of players 1 and 2 may be correlated in player 3's mind. This is not possible under the original definition of (independent) rationalizability in Bernheim (1984) and Pearce (1984). Common certainty of rationality *and* the fact that each player

*Princeton University, smorris@princeton.edu

†Princeton University, satorut@princeton.edu

¹We follow the convention of using “certainty” to mean “belief with probability one.” An older tradition in the economics literature referred to this as “knowledge,” although this is inconsistent with formal and informal usage elsewhere; for example, in philosophy, an event is said to be known when it is true and believed to be true.

treats other players' actions as stochastically independent implies that independent rationalizable actions are played.

How can players' correlated conjectures be interpreted? In an influential argument, Aumann (1987) wrote that this correlation (between the actions of players 1 and 2) in player 3's mind

...has no connection with any overt or even covert collusion between 1 and 2; they may be acting entirely independently. Thus it may be common knowledge that both 1 and 2 went to business school, or perhaps to the same business school; but 3 may not know what is taught there. In that case 3 would think it quite likely that they would take similar actions, without being able to guess what those actions might be.

Put differently, Aumann is taking Savage's (1954) "small world" view of state spaces. State spaces should not be understood to include everything, and therefore we should not assume that the state space includes everything that plays a role in determining players' beliefs.²

In this paper, we want to re-visit this debate, relaxing a key but implicit maintained assumption. In the literature on epistemic foundations of game theoretic solution concepts in complete information games, "complete information" means the assumption that there is common certainty of players' "payoff functions." The payoff that a player gets from a certain action profile is the von Neumann-Morgenstern index of the physical outcome that results from the action profile, under some representation of that player's expected utility preferences. Thus it is (implicitly) assumed that there is common certainty of both a mapping from action profiles to outcomes and players' preferences over (lotteries over) those outcomes. But, more than this, it is assumed that players' preferences over outcomes remain the same conditional on the realized actions/types of opponents. This may be a natural, useful and insightful assumption. However, it is inconsistent with the maintained "small worlds" assumption in arguments justifying correlated rationalizability as the standard implication of common certainty of rationality. In particular, to pursue Aumann's story, there may be common certainty of the (unconditional) preferences of all players over lotteries; but player 3—understanding that players 1 and 2 went to business school and he didn't—understands that the actions of 1 or 2 may end up being correlated with how he (player 3) feels about outcomes, and thus he may have different preferences conditional on the realized actions of the other players than his unconditional preferences.

We can illustrate this with an example. Consider a game where there are two players, Ann and

²A related debate arises in deriving the "right" definition of rationalizability in two player games with incomplete information. The definition of interim correlated rationalizability in Dekel, Fudenberg and Morris (2007) allows a player to believe that the other player's action is correlated with an unknown state, even when that player does not believe she has any information about the state.

Bob, and two outcomes, party and no party. Each player has two possible actions and the following table describes the probability of a party for each action profile, where Ann chooses the row and Bob chooses the column:

	<i>L</i>	<i>R</i>
<i>U</i>	0	$\frac{1}{3}$
<i>D</i>	$\frac{1}{3}$	1

There is common certainty that each player would strictly prefer a party to no party. Clearly actions *D* and *R* are rationalizable for Ann and Bob respectively. But we will argue why actions *U* and *L* are also rationalizable, if we do not assume that conditional preferences must equal unconditional preferences. Note that players in this game are not choosing among outcomes, or lotteries over outcomes, because their choices do not generate a lottery over outcomes. By choosing an action in the game, a player is choosing a mapping from the actions of the opponent to lotteries over outcomes, i.e., an Anscombe-Aumann act over the opponent's actions, and it is players' expected utility preferences over acts that will determine their choices. Suppose Ann assigned probability $\frac{1}{2}$ to Bob choosing *L* and her own utilities from no party and party being 0 and 6 respectively; and assigned probability $\frac{1}{2}$ to player 2 choosing *R* and her own utilities from outcomes no party and party being 0 and -4 respectively. Given a choice between no party or party, she would choose party, since it gives expected utility $\frac{1}{2} \times 6 + \frac{1}{2} \times (-4) = 1 > 0$. Thus her unconditional preferences do have her strictly preferring a party. But given a choice between playing action *U* or *D* in the above game, Ann would choose *U* because it gives expected utility $\frac{1}{2} \times \frac{1}{3} \times (-4) = -\frac{2}{3}$, while *D* gives expected utility $\frac{1}{2} \times \frac{1}{3} \times 6 + \frac{1}{2} \times (-4) = -1$. Thus *U* is rationalizable for Ann, who strictly prefers a party to no party.

We describe a general framework where agents' beliefs and von Neumann-Morgenstern indices are treated symmetrically. If we assume common certainty of (i) rationality, (ii) the mapping from action profile to outcomes and (iii) unconditional preferences over outcomes, then we get a solution concept—weaker than correlated rationalizability—that we label preference-correlated rationalizability. As the above example illustrates this solution concept is very permissive. We can capture the standard approach, and correlated rationalizability, by adding common certainty that conditional preferences are equal to unconditional preferences. More generally, we can examine a variety of solution concepts where we make intermediate assumptions about how unconditional preferences relate to conditional preferences.

While we argue that preference-correlated rationalizability captures the consequences of common certainty of rationality in a natural way, it turns out to be “badly behaved.” Consider the following variant of the above game, where we replace the probability of party if Ann choose *D* and

Bob chooses R with a number $p \in [0, 1]$:

	L	R
U	0	$\frac{1}{3}$
D	$\frac{1}{3}$	p

A surprising result is that U is preference-correlated rationalizable for any value of p other than $\frac{2}{3}$. Suppose Ann assigned probability $\frac{1}{2}$ to Bob choosing L and her own utilities from outcomes no party and party being 0 and $\frac{6p}{3p-2}$ respectively; and assigned probability $\frac{1}{2}$ to player 2 choosing R and her own utilities from outcomes no party and party being 0 and $-\frac{4}{3p-2}$ respectively. Since $\frac{1}{2} \times \frac{6p}{3p-2} + \frac{1}{2} \times (-\frac{4}{3p-2}) = 1 > 0$, Ann strictly prefers a party to no party. But the expected utility from action U is $\frac{1}{2} \times \frac{1}{3} \times (-\frac{4}{3p-2}) = -\frac{2}{3(3p-2)}$, while the expected utility from action D is $\frac{1}{2} \times \frac{1}{3} \times \frac{6p}{3p-2} + \frac{1}{2} \times p \times (-\frac{4}{3p-2}) = -\frac{p}{3p-2} < -\frac{2}{3(3p-2)}$. However, if $p = \frac{2}{3}$, then the difference in the probability of party going from action U to action D is the same ($\frac{1}{3}$) whatever Bob's action is, so D is strictly preferred to U whenever she strictly prefers a party to no party unconditionally.³

This example illustrates a failure of upper hemicontinuity of the solution concept with respect to the game. This failure reflects a failure of “compactness”: the preferences of Ann that we described above to make U preference-correlated rationalizable have the feature that a percentage point increase in having a party conditional on that Bob chooses L is worth $\frac{3p}{3p-2}$ percentage point increase in having a party unconditionally, and that the “marginal rate of substitution” of a conditional party for an unconditional party is unbounded as p becomes close to $\frac{2}{3}$. The main results in this paper involve identifying restrictions on expected utility preferences that are sufficient to restore upper hemicontinuity of rationalizability. The failure of compactness also has implications for dual dominance characterizations of rationalizability. Recall that in the standard framework, an action is never a best response to any conjecture over the opponents' actions if and only if it is strictly dominated, i.e., there exists a mixed action strictly preferred to it for every conjecture over the opponents' action profiles. We report sufficient conditions for an analogous dual dominance characterization of preference-correlated rationalizability. Such a dual dominance characterization exists in the above example, but, because of the failure of compactness, it takes a somewhat counter-intuitive form. As a by-product of our analysis of alternative restrictions on expected utility preferences in games, we are able to explore some interesting connections with the work of Ledyard (1986) and Börgers (1993) which also allow richer than standard expected utility preferences in games.

While our motivation concerns solution concepts in games, the key questions about the upper hemicontinuity and dual characterizations of best responses already arise in understanding the op-

³This example and argument are re-stated more formally in Section 2.4

timal actions of an expected utility maximizer in single person decision problems, when we have partial information about that agent’s preferences. For example, the leading example of interest from the discussion above is “unconditional cardinal” preferences, where the agent’s unconditional preferences over lotteries are known, but the agent’s preferences over acts (and the implied conditional preferences) are not known. In Section 2, we formally analyze restrictions on sets of expected utility preferences and their consequences for best responses. We identify a class of “inequality restrictions” on preferences, i.e., sets of expected utility preferences characterized by a finite set of weak and strict preference inequalities. We show that any set of preferences within this class has a dual dominance characterization, where an action is not a best response to any preference in the set if and only if there is a mixed action which is strictly better than that action under all preferences in the set. Börgers (1993) analyzed games where a strict ordinal ranking of outcomes was known and allowed any expected utility preferences with conditional preferences consistent with that ordinal ranking. These preferences do not satisfy our definition of inequality restrictions and do not have a dual dominance characterization. In the analysis of independent rationalizability, expected utility preferences consistent with fixed (unconditional and conditional) cardinal preferences *and* independent beliefs over the opponents’ actions are studied. These preferences also fail our definition of inequality restrictions and—as is well known—do not have a dual dominance characterization. It turns out that “unconditional cardinal” preferences—our leading example—can be represented by inequality restrictions and thus best responses have a dual dominance characterization, but the characterization is complex because it is necessary to check that the dominating mixed action does better for an infinite set of preferences. We also identify a stronger class of “polytopic” restrictions on preferences, requiring that the set of preferences is the convex hull of a finite set of preferences in an appropriate sense. In this case, the dual dominance characterization can be sharpened to require checking against only a finite set of preferences (the extreme points), and upper hemicontinuity of best responses can be established. Our leading “unconditional cardinal” preferences are not polytopic; as noted above, they fail upper hemicontinuity and do not have a dominance characterization using extreme points.

We use these findings in analyzing and characterizing solution concepts for complete and incomplete information games in Sections 3 and 4 respectively. For each profile of preference restrictions, we prove that preference-correlated rationalizability captures the implications of common certainty of rationality and the preference restriction profile. Epstein (1997) and Di Tillio (2008) prove analogous results for more general non-expected utility preferences. Our marginal contribution here is to highlight that when we restrict back to expected utility preferences in the settings of Epstein (1997) and Di Tillio (2008), there are many subtleties about which exact classes of expected utility preferences are allowed. In particular, if we impose no restrictions on correlation between

each player’s conditional preferences and his opponents’ actions, then we show that preference-correlated rationalizability is equivalent to one round of elimination of never best responses both in the complete information setting (Claim 2) and the incomplete information setting (Claim 3).

In a companion paper, Bergemann, Morris and Takahashi (2011), we examine when different interdependent types can be “strategically distinguished” in the sense that there exists a mechanism where they are guaranteed to choose different actions. In that paper, we restrict attention to solution concepts and type spaces where each agent has a worst outcome that remains the worst outcome conditional on the behavior and types of other agents. This restriction is polytopic as defined in this paper. Under this restriction, we show that under the solution concept of equilibrium or under the incomplete information version of preference-correlated rationalizability defined in this paper, two types can be strategically distinguished if and only if they map to different points in a universal space of expected utility interdependent preferences that we construct. The current paper then supplies an epistemic foundation for the rationalizability concept studied in the companion paper. The polytopic worst outcome restriction ensures that the rationalizable outcomes are upper hemicontinuous. If we did not impose such a restriction, it is an implication of Claim 3 in this paper that it is not possible to strategically distinguish any two types that have the same first order preferences, i.e., preferences over acts contingent on states but not on opponents’ types. Thus a contribution of this paper is to show why some form of compactness restriction is necessary to derive the main results in Bergemann, Morris and Takahashi (2011).

2 Decision Problems

2.1 Expected Utility Preferences

Let Z be a finite set of outcomes and X be a finite set of states. Let $F(X)$ be the set of Anscombe-Aumann acts over X , $f: X \rightarrow \Delta(Z)$. We identify each outcome with the lottery that assigns probability 1 on the outcome, and each lottery with the constant act that yields the lottery. A preference on $F(X)$ is a binary relation satisfying completeness and transitivity. A preference is a (state dependent) expected utility preference if it has a linear representation as follows.

Definition 1. A preference \succeq on $F(X)$ is an *expected utility preference* if there exists a state dependent utility index $u: X \times Z \rightarrow \mathbb{R}$ such that

$$f \succeq f' \Leftrightarrow \sum_{x \in X} \sum_{z \in Z} f(x)(z)u(x, z) \geq \sum_{x \in X} \sum_{z \in Z} f'(x)(z)u(x, z)$$

for all $f, f' \in F(X)$.

A preference \succeq is an expected utility preference if and only if it satisfies *independence* (for every $f, f', f'' \in F(X)$ and $\lambda \in (0, 1]$, $f \succeq f'$ if and only if $\lambda f + (1 - \lambda)f'' \succeq \lambda f' + (1 - \lambda)f''$) and *continuity* (for every $f, f', f'' \in F(X)$, if $f \succ f' \succ f''$, then there exists $\varepsilon \in (0, 1)$ such that $(1 - \varepsilon)f + \varepsilon f'' \succ f' \succ (1 - \varepsilon)f'' + \varepsilon f$). See, for example, Fishburn (1970, Theorem 13.1).

In the rest of the paper, all preferences will be expected utility preferences. The trivial preference is a preference that is indifferent between all acts. Write $P(X)$ for the set of expected utility preferences on $F(X)$, and $P^*(X) \subsetneq P(X)$ excludes the trivial preference.

A decision problem consists of a finite set of actions A and an outcome function $g: A \rightarrow F(X)$. The domain of g is extended to mixed actions $\Delta(A)$ in the usual way. For any restriction $\Phi \subseteq P(X)$ on preferences, an action $a \in A$ is a Φ -best response in (A, g) if it is optimal according to some preference in Φ .

Definition 2. An action $a \in A$ is a Φ -best response in (A, g) if there exists $\succeq \in \Phi$ such that $g(a) \succeq g(a')$ for all $a' \in A$.

Definition 3. An action $a \in A$ is Φ -dominated in (A, g) if there exists a mixed action $\alpha \in \Delta(A)$ such that $g(\alpha) \succ g(a)$ for all $\succeq \in \Phi$.

2.2 Examples of Preference Restrictions

Some subsets of $P^*(X)$ we will be interested in are as follows:

Example 1 (conditional cardinal). A state independent “payoff function” over outcomes is known but there is complete uncertainty about the state. Thus we fix a non-constant von Neumann-Morgenstern index $v: Z \rightarrow \mathbb{R}$, and we are interested in preferences that are represented by

$$u(x, z) = \lambda(x)v(z)$$

for some $\lambda \in \Delta(X)$.

Example 2 (conditional ordinal). There is a known state independent ordinal ranking of outcomes but there is complete uncertainty about cardinal payoffs and the state. Thus we fix a strict total order $>$ on Z and we are interested in non-trivial preferences that are represented by $u: X \times Z \rightarrow \mathbb{R}$ satisfying

$$z > z' \Rightarrow u(x, z) \geq u(x, z').$$

Example 3 (conditional worst outcome). There is a known worst outcome but nothing else is known. Thus we fix an outcome $w \in Z$ and we are interested in non-trivial preferences that are represented by $u: X \times Z \rightarrow \mathbb{R}$ satisfying

$$u(x, z) \geq u(x, w).$$

Example 4 (unconditional cardinal). A “payoff function” over outcomes pins down preferences over state independent lotteries but nothing else is known. Thus we fix a nonconstant von Neumann-Morgenstern index $v: Z \rightarrow \mathbb{R}$ and we are interested in preferences that are represented by $u: X \times Z \rightarrow \mathbb{R}$ satisfying

$$\sum_{x \in X} u(x, z) = v(z).$$

In what follows, we generalize these examples and place the “bad behavior” of the unconditional cardinal case (Example 4), discussed in the introduction, in a wider context. In particular, we will consider three classes of restrictions on preferences, inequality restrictions, polytopic restrictions and conditional and unconditional preference restrictions.

2.3 Inequality Restrictions

We first consider the set of preferences where some weak and strict rankings are known to hold. Let W and S be finite subsets of $F(X) \times F(X)$.

Definition 4. A preference $\succeq \in P(X)$ satisfies *inequality restrictions* (W, S) if $f \succeq f'$ for all $(f, f') \in W$ and $f \succ f'$ for all $(f, f') \in S$.

Each weak or strict ranking of acts corresponds to a “change of direction” in $F(X)$ going from one act to another. The inequality restrictions W and S characterize a set of preferences such that an act is preferred to another if and only if the direction of change is a nonnegative linear combination of the directions of change in the sets W and S .

Lemma 1. *Let Φ be the set of preferences satisfying inequality restrictions (W, S) . For $f, f' \in F(X)$, we have $f \succeq f'$ for all $\succeq \in \Phi$ if and only if there exists $\lambda \in \mathbb{R}_+^{W \cup S}$ such that*

$$f - f' = \sum_{(g, g') \in W \cup S} \lambda_{g, g'} (g - g'). \quad (1)$$

Moreover, $f \succ f'$ for all $\succeq \in \Phi$ if and only if there exists $\lambda \in \mathbb{R}_+^{W \cup S}$ with $\lambda_{g, g'} > 0$ for some $(g, g') \in S$ such that (1) holds.

We omit the proof, which is a straightforward application of Farkas’ lemma.

If a restriction Φ is generated by weak and strict inequalities, then never Φ -best response is characterized by Φ -dominance in any decision problem.

Proposition 1. *Let Φ be the set of preferences satisfying inequality restrictions (W, S) . Then an action is not a Φ -best response if and only if it is Φ -dominated.*

Proof. The if direction is immediate. To show the only-if direction, suppose that a is not a Φ -best response. Assume $\Phi \neq \emptyset$ without loss of generality. Then there exists no $u: X \times Z \rightarrow \mathbb{R}$ such that

$$\sum_{x,z} f(x)(z)u(x,z) \geq \sum_{x,z} f'(x)(z)u(x,z)$$

for any $(f, f') \in W$,

$$\sum_{x,z} f(x)(z)u(x,z) > \sum_{x,z} f'(x)(z)u(x,z)$$

for any $(f, f') \in S$, and

$$\sum_{x,z} g(a)(x)(z)u(x,z) \geq \sum_{x,z} g(a')(x)(z)u(x,z)$$

for any $a' \in A$. By Farkas' lemma, there exist $\lambda \in \mathbb{R}_+^{W \cup S}$ and $\alpha \in \mathbb{R}_+^A$ such that $\lambda_{f,f'} > 0$ for some $(f, f') \in S$ and

$$\sum_{(f,f') \in W \cup S} \lambda_{f,f'}(f - f') + \sum_{a'} \alpha(a')(g(a) - g(a')) = 0.$$

Since $\Phi \neq \emptyset$, we have $\alpha \neq 0$, thus α can be normalized to be in $\Delta(A)$. The rest follows from Lemma 1. ■

This Proposition generalizes the classical Farkas' lemma on the “dual dominance” of best responses: an action is never a best response to any conjecture over states if and only if there exists a mixed action that gives a strictly higher payoff for all conjectures over the states (Pearce (1984) Lemma 3).

All the examples in Section 2.2 fit the formulation of inequality restrictions. For example, consider the unconditional cardinal case (Example 4). Let b and w be the best and the worst outcomes according to v . (If there are multiple best (or worst) outcomes, choose one arbitrarily.) Then the restriction is characterized by requiring $|Z| - 2$ indifference equalities (or $2|Z| - 4$ weak inequalities) between constant acts

$$z \sim \frac{v(z) - v(w)}{v(b) - v(w)}b + \frac{v(b) - v(z)}{v(b) - v(w)}w$$

for $z \in Z \setminus \{b, w\}$ and a single strict inequality $b \succ w$.

It is instructive to consider two examples of preference restrictions which do not have the equivalence between Φ -dominance and never Φ -best response and thus (by Proposition 1) cannot be characterized by inequality constraints. As will discuss in Section 3, each of these example corresponds to a well-known version of complete information rationalizability.

A first example is a strict version of the conditional ordinal case (Example 3) studied in Börgers (1993). Given a strict total order $>$ on Z , he considered the set of preferences that are represented by

$$u(x, z) = \lambda(x)v(z)$$

for some $\lambda \in \Delta(X)$ and some utility index $v: Z \rightarrow \mathbb{R}$ that is *strictly* increasing with respect to $>$, i.e.,

$$z > z' \Leftrightarrow v(z) > v(z').$$

This restriction does not fit the formulation of inequality restrictions if $|X| \geq 2$ and $|Z| \geq 3$. To see this, it is sufficient to show that Proposition 1 does not hold in this case. Suppose for simplicity that $X = \{x, x'\}$ and $Z = \{z, z', z''\}$ with $z > z' > z''$, and let Φ be the restriction à la Börgers. Consider the following decision problem (A, g) with $A = \{a, a', a''\}$ and

$$g = \begin{array}{c|cc} & x & x' \\ \hline a & z & z'' \\ a' & z' & z' \\ a'' & z' & z \end{array}.$$

Then a' is not a Φ -best response, but there is no $\alpha \in \Delta(A)$ such that $g(\alpha) \succ g(a')$ for any $\succeq \in \Phi$.⁴

Another example that does not fit this formulation is a variant of the conditional cardinal case (Example 1) where beliefs are required to be independent over a product set of states. Suppose $X = X_1 \times X_2$ with $|X_1|, |X_2| \geq 2$. Given a nonconstant von Neumann-Morgenstern index $v: Z \rightarrow \mathbb{R}$, we require that preferences be represented by

$$u(x_1, x_2, z) = \lambda_1(x_1)\lambda_2(x_2)v(z)$$

for some $\lambda_1 \in \Delta(X_1)$ and $\lambda_2 \in \Delta(X_2)$. This set of preferences cannot be characterized by linear inequalities, and Proposition 1 does not hold.

2.4 Polytopic Restrictions

We now consider another formulation of restrictions on expected utility preferences.

Definition 5. Let $V = \{\succeq_1, \dots, \succeq_K\} \subseteq P(X)$ be a finite set of preferences. A preference $\succeq \in P(X)$ is a *mixture of V* if

$$f \succ_1 f', \dots, f \succ_K f' \Rightarrow f \succ f' \tag{2}$$

⁴Ledyard (1986, Section 4.3) allowed for state dependent cardinal payoffs, i.e., $u: X \times Z \rightarrow \mathbb{R}$ satisfying $z > z' \Leftrightarrow u(x, z) > u(x, z')$ for any $x \in X$. This restriction is not characterized by finitely many inequalities, either. Bogomolnaia and Moulin (2001) introduced a related notion of ordinal efficiency in the random assignment problem.

for all $f, f' \in F(X)$.

The following lemma follows from Farkas' lemma.

Lemma 2. *Suppose that \succeq is represented by $u: X \times Z \rightarrow \mathbb{R}$ and, for each $k = 1, \dots, K$, \succeq_k is represented by $u_k: X \times Z \rightarrow \mathbb{R}$. Then \succeq is a mixture of $\{\succeq_1, \dots, \succeq_K\}$ if and only if there exist $\lambda \in \mathbb{R}_+^K \setminus \{0\}$ and $\mu: X \rightarrow \mathbb{R}$ such that*

$$u(x, z) = \sum_{k=1}^K \lambda_k u_k(x, z) + \mu(x).$$

Definition 6. A set of non-trivial preferences $\Phi \subseteq P^*(X)$ is *polytopic* if Φ is equal to the set of all mixtures of some finite set V of preferences. In this case, we say that V is the set of *vertices* of Φ .

Note that our definition requires that a polytopic restriction Φ must exclude the trivial preference. This implies that (2) does not hold vacuously, i.e., there exists a pair of acts $f, f' \in F(X)$ such that $f \succ f'$ for all $\succeq \in V$, and hence for all $\succeq \in \Phi$.

Examples 1–3 are polytopic restrictions as we will see in Section 2.5. Generally, any polytopic restriction is characterized by finitely many weak inequalities and a single strict inequality (just to exclude the trivial preference), but not every restriction characterized by weak and strict inequalities is polytopic. For example, the unconditional cardinal case (Example 4) is not polytopic if $|X| \geq 2$ and $|Z| \geq 2$. To see this, suppose for simplicity that $X = \{x, x'\}$, $Z = \{z, z'\}$ and $v(z) > v(z')$. Then, for a state dependent utility index $u: X \times Z \rightarrow \mathbb{R}$ that represents a preference such that $z \succ z'$, the “relative weight” on state x

$$\lambda(x; u) = \frac{u(x, z) - u(x, z')}{u(x, z) + u(x', z) - u(x, z') - u(x', z')}$$

does not necessarily belong to $[0, 1]$, and can take any real value, unlike in the conditional cardinal case (Example 1).⁵ However, if \succeq were a mixture of $\{\succeq_1, \dots, \succeq_K\}$ with each \succeq_k represented by some $u_k: X \times Z \rightarrow \mathbb{R}$, then, by Lemma 2

$$\lambda(x; u) = \frac{\sum_{k=1}^K \lambda_k (u_k(x, z) - u_k(x, z'))}{\sum_{k=1}^K \lambda_k (u_k(x, z) + u_k(x', z) - u_k(x, z') - u_k(x', z'))},$$

which would be bounded by $\min_k \lambda(x; u_k)$ and $\max_k \lambda(x; u_k)$.⁶ This is a contradiction.

If Φ is polytopic, then we can restate Proposition 1 in terms of vertex preferences.

Proposition 2. *Let Φ be the polytopic restriction with vertices V . Then the following conditions are equivalent.*

⁵Note that $\lambda(x; u)$ is well-defined since the denominator is strictly positive.

⁶Note that $(a + b)/(c + d)$ is between a/c and b/d if $c, d > 0$.

1. a is not a Φ -best response;
2. a is Φ -dominated;
3. a is V -dominated.

Using only finitely many vertex preferences greatly simplifies checking the dominance condition. We noted above that the unconditional cardinal case (Example 4) is not polytopic, and such a simplification is not available. Indeed, we will show in Claim 1 that the characterization of Φ -dominance is not simple in this case.

The next proposition shows that the set of Φ -best responses depends on underlying decision problems in the upper hemicontinuous manner. This is an extension of the well-known result for the conditional cardinal case (Example 1).

Proposition 3. *Suppose that Φ is polytopic. Let $\{g^n\}$ be a sequence of functions $g^n: A \rightarrow F(X)$ such that $g^n \rightarrow g$ as $n \rightarrow \infty$. If a is a Φ -best response in (A, g^n) for each n , then a is also a Φ -best response in (A, g) .*

Proof. Let $V = \{\succeq_1, \dots, \succeq_K\}$ be the set of vertices of Φ . For each $k = 1, \dots, K$, $u_k: X \times Z \rightarrow \mathbb{R}$ represents \succeq_k . If a is a Φ -best response in (A, g^n) , then, by Lemma 2, there exists $\lambda^n \in \mathbb{R}_+^K \setminus \{0\}$ such that

$$\sum_{k,x,z} g^n(a)(x)(z) \lambda_k^n u_k(x, z) \geq \sum_{k,x,z} g^n(a')(x)(z) \lambda_k^n u_k(x, z)$$

for any $a' \in A$. By normalizing λ^n , we can assume that $\lambda^n \in \Delta(\{1, \dots, K\})$. By the compactness of $\Delta(\{1, \dots, K\})$, there exists an accumulation point λ of $\{\lambda^n\}$, where we have

$$\sum_{k,x,z} g(a)(x)(z) \lambda_k u_k(x, z) \geq \sum_{k,x,z} g(a')(x)(z) \lambda_k u_k(x, z)$$

for any $a' \in A$. ■

The set of Φ -best responses may not be upper hemicontinuous in g if Φ is not polytopic. We can illustrate this using the example discussed in the introduction, stated more formally as a single agent decision problem for Ann (with Bob's action treated as an unknown state), and with outcome "party" labelled z and outcome "no party" labelled z' . Thus consider the unconditional cardinal case (Example 4) with $X = \{L, R\}$, $Z = \{z, z'\}$ and $(v(z), v(z')) = (1, 0)$. Consider decision problems parameterized by $p \in [0, 1]$ with $A = \{U, D\}$ and $g^p: A \times X \rightarrow \Delta(Z)$ given by

	L	R	
U	z'	$\frac{1}{3}z + \frac{2}{3}z'$.
D	$\frac{1}{3}z + \frac{2}{3}z'$	$pz + (1-p)z'$	

As we saw in the introduction, if $p \neq \frac{2}{3}$, then U is a Φ -best response; this can be seen by considering the preference represented by

$$(u(L, z), u(R, z), u(L, z'), u(R, z')) = \left(\frac{3p}{3p-2}, -\frac{2}{3p-2}, 0, 0 \right),$$

but, if $p = \frac{2}{3}$, then D is a unique Φ -best response. Thus as $p \rightarrow \frac{2}{3}$, there is a failure of upper hemicontinuity of best responses.

2.5 Conditional and Unconditional Preference Restrictions

All the examples we introduced in Section 2.2 took the form of imposing restrictions on the decision maker's conditional and unconditional preferences. In this section, we describe such restrictions in general and relate them back to the examples and to inequality and polytopic restrictions.

For each $\succeq \in P(X)$ and $x \in X$, the *conditional preference* \succeq_x of \succeq at state x is the preference over lotteries defined by

$$y \succeq_x y' \Leftrightarrow \begin{pmatrix} y \text{ on } x \\ y'' \text{ otherwise} \end{pmatrix} \succeq \begin{pmatrix} y' \text{ on } x \\ y'' \text{ otherwise} \end{pmatrix}$$

for all $y, y' \in \Delta(Z)$. By the independence axiom, the definition of \succeq_x does not depend on the choice of y'' . A state $x \in X$ is \succeq -null (in the sense of Savage) if \succeq_x is the trivial preference. The *unconditional preference* of \succeq is the restriction of \succeq to constant acts. A preference \succeq is *certain about event* $E \subseteq X$ if every $x \in X \setminus E$ is \succeq -null.

Let U and C be subsets of $P(\{*\})$, to be interpreted as sets of expected utility preferences over lotteries. ($\{*\}$ is an arbitrary singleton set.)

Definition 7. A preference $\succeq \in P(X)$ satisfies *conditional and unconditional preference restrictions* (C, U) if the conditional preference of \succeq at every non- \succeq -null state belongs to C and the unconditional preference of \succeq belongs to U .

Let $\Phi^{C,U}$ be the set of non-trivial preferences that satisfy conditional and unconditional restrictions (U, C) . The following proposition shows that properties on conditional and unconditional preference restrictions carry over to the restriction on entire preferences.

Proposition 4. 1. *If C is characterized by weak inequalities (without strict inequalities) and U is characterized by weak and strict inequalities, then $\Phi^{C,U}$ is characterized by weak and strict inequalities.*

2. *If C and U are polytopic, then $\Phi^{C,U}$ is also polytopic.*

Proof. Part 1 is obvious. For part 2, since U is polytopic, there exists a pair $y, y' \in \Delta(Z)$ such that $y \succ y'$ for any $\succeq \in \Phi^{C,U}$. Thus, by the usual normalization, without loss of generality, we can work with representations $u: X \times Z \rightarrow \mathbb{R}$ that satisfy $\sum_{x,z} y(z)u(x, z) = 1$ and $\sum_z y'(z)u(x, z) = 0$ for all x . Let \mathcal{U} be the set of such u functions, which is an affine subspace of $\mathbb{R}^{X \times Z}$ that does not contain the origin. By Lemma 2, the set of non-trivial preferences that satisfy the conditional restriction C is represented by a convex polytope (a convex hull of finitely many points) in \mathcal{U} ; the set of preferences that satisfy the unconditional restriction U is represented by a convex polyhedron (an intersection of finitely many closed half-spaces, possibly unbounded) in \mathcal{U} . Thus, the intersection of these two sets is a convex polytope, which represents a polytopic restriction. ■

In part 2 of Proposition 4, it is important to impose polytopic restrictions on both conditional and unconditional preferences. If conditional preferences are completely unrestricted, as in the unconditional cardinal case (Example 4), the restriction as a subset of $P^*(X)$ may not be polytopic even if the unconditional preference restriction is.

Now let us return to the examples in Section 2.2. By Proposition 4, Examples 1–3 are polytopic since they are characterized by polytopic conditional and unconditional restrictions $C = U$. In each case, the vertices consist of, for each state $x \in X$ and each vertex preference of $C = U$ over lotteries, the preference that is certain about $\{x\}$ and the preference conditional on x is given by that lottery preference. Thus we can apply Proposition 2 to characterize Φ -best responses as follows.

For the conditional cardinal case (Example 1), an action a is not a Φ -best response in (A, g) if and only if it is strictly dominated in the usual sense, i.e., there exists $\alpha \in \Delta(A)$ such that

$$\sum_{z \in Z} g(\alpha)(x)(z)v(z) > \sum_{z \in Z} g(a)(x)(z)v(z)$$

for any $x \in X$.

For the conditional ordinal case (Example 2), let w be the worst outcome according to $>$. Then a is not a Φ -best response in (A, g) if and only if it is *stochastically dominated*, i.e., there exists $\alpha \in \Delta(A)$ such that $g(\alpha)(x)$ stochastically dominates $g(a)(x)$, i.e.,

$$\sum_{z' > z} g(\alpha)(x)(z')v(z') > \sum_{z' > z} g(a)(x)(z')v(z')$$

for any $x \in X$ and $z \in Z \setminus \{w\}$.

For the conditional worst outcome case (Example 3), a is not a Φ -best response in (A, g) if and only if it is *worst outcome dominated*, e.g., there exists $\alpha \in \Delta(A)$ such that

$$g(\alpha)(x)(z) > g(a)(x)(z')$$

for any $x \in X$ and $z \in Z \setminus \{w\}$.

Unlike those examples, the unconditional cardinal case (Example 4) is not polytopic. We can nevertheless provide a complete characterization of Φ -best responses as below.

Claim 1. *In the unconditional cardinal case, an action a is not a Φ -best response in (A, g) only if there exists $\alpha \in \Delta(A)$ such that*

1. $h := g(\alpha)(x) - g(a)(x)$ is independent of x , and
2. $\sum_{z \in Z} h(z)v(z) > 0$.

Proof. The if direction is immediate. To show the only-if direction, recall that the unconditional cardinal case is characterized by inequalities between various constant acts $(y, y') \in W \cup S$.⁷ Thus, if a is not a Φ -best response in (A, g) , then it follows from Lemma 1 and Proposition 1 that there exist $\alpha \in \Delta(A)$ and $\lambda \in \mathbb{R}_+^{W \cup S}$ such that

$$h := g(\alpha)(x) - g(a)(x) = \sum_{(y, y') \in W \cup S} \lambda_{y, y'} (y - y')$$

with $\lambda_{y, y'} > 0$ for some $(y, y') \in S$. Then h is independent of x . Since $\sum_z (y(z) - y'(z))v(z) \geq 0$ for $(y, y') \in W \cup S$ with strict inequalities for $(y, y') \in S$, we have $\sum_z h(z)v(z) > 0$. ■

This Claim is related to Theorem 4 and Corollary 4.3 in Ledyard (1986). He asked when an action is not a strict best response given a set of expected utility preferences without the unconditional preference restriction and obtained a characterization which essentially requires the first condition of the claim to hold.

Note that, unlike other dominance conditions, this condition involves equalities rather than strict inequalities, and thus the set of outcome functions g that satisfy this condition is not open. This is consistent with the failure of upper hemicontinuity we saw in the previous subsection.

3 Complete Information Games

Here we propose solution concepts that capture the implications of common certainty of rationality under various restrictions on preferences.

For $E \subseteq X$, let $P(X | E) \subseteq P(X)$ be the set of preferences that are certain about E .

⁷See the paragraph below Proposition 1 for exact forms of W and S

3.1 Preference-Correlated Rationalizability

Fix a game $G = (I, (A_i)_{i \in I}, g)$, where I is a finite set of players, A_i is a finite set of player i 's actions and $g: A = \prod_i A_i \rightarrow \Delta(Z)$. Strictly speaking, G is not a game but a game form or mechanism, which does not describe real-valued payoffs. Note that, if player i chooses an action $a_i \in A_i$, then he faces an act $g(a_i, \cdot)$ contingent on the opponents' actions. Thus, from player i 's perspective, a game is a decision problem (A_i, g) over states $A_{-i} = \prod_{j \neq i} A_j$.

Each player i has a restriction on his preferences $\Phi_i \subseteq P^*(A_{-i})$. Each profile of preference restrictions $\Phi = (\Phi_i)_{i \in I}$ will give rise to a difference rationalizability concept. Namely, for a product set $B_{-i} = \prod_{j \neq i} B_j$, with each $B_j \subseteq A_j$, we say that a_i is a Φ_i -best response to B_{-i} for player i if a_i is a $\Phi_i \cap P(A_{-i} | B_{-i})$ -best response in player i 's decision problem (A_i, g) , i.e., there exists $\succeq_i \in \Phi_i \cap P(A_{-i} | B_{-i})$ such that

$$g(a_i, \cdot) \succeq_i g(a'_i, \cdot)$$

for any $a'_i \in A_i$. Φ -rationalizability is defined by iteration of Φ -best responses: let $R_{i,0}^\Phi = A_i$, and for $n \geq 1$, $R_{i,n}^\Phi$ be the set of Φ_i -best responses to $R_{-i,n-1}^\Phi = \prod_{j \neq i} R_{j,n-1}^\Phi$. Then an action in $R_i^\Phi = \bigcap_n R_{i,n}^\Phi$ is called Φ -rationalizable for player i .

If each Φ_i is characterized by weak and strict inequalities (polytopic), then $\Phi_i \cap P(A_{-i} | B_{-i})$ is also characterized by weak and strict inequalities (resp. polytopic). Thus, by Propositions 1 and 2, we can compute each step of iteration $R_{i,n}^\Phi$ by eliminating Φ_i -dominated actions with respect to $R_{-i,n-1}^\Phi$. Also, if each Φ_i is polytopic, then, by Proposition 3, $R_{i,n}^\Phi$ and R_i^Φ depend upper hemicontinuously on outcome function g .

Now by considering this definition of Φ -rationalizability letting the profile of preference restrictions $\Phi = (\Phi_i)$ take various forms, we can generate various important notions of rationality. If every player has conditional cardinal preference restriction Φ_i with nonconstant von Neumann-Morgenstern index $v_i: Z \rightarrow \mathbb{R}$ (i.e., Example 1) then Φ -rationalizability collapses to correlated rationalizability in game (not game form) $\tilde{G} = (I, (A_i, \tilde{g}_i)_{i \in I})$ with $\tilde{g}_i(a) = \sum_z g(a)(z)v_i(z)$. Conditional ordinal restrictions (Example 2) give an upper hemicontinuous solution concept that is equivalent to iterated deletion of stochastically dominated actions. Conditional worst outcome restrictions (Example 3) give an upper hemicontinuous solution concept that is equivalent to iterated deletion of worst outcome dominated actions. But in the extreme case where Φ_i restricts only unconditional preferences (Example 4), we can show that Φ -rationalizability is "badly behaved" in that the iteration of eliminating never Φ_i -best responses stops in at most one round.

Claim 2. *In the unconditional cardinal case, we have $R_i^\Phi = R_{i,1}^\Phi$.*

See the Appendix for the proof. This is a consequence of characterization of best responses given in Claim 1.

In the discussion following Proposition 1, we gave two examples of preference restrictions which could not be characterized by inequality restrictions and did not have dominance characterizations. The first one was a strict version of the conditional ordinal case. The definition of rationalizability in Börgers (1993) is equivalent to imposing this restriction as Φ_i in the definition of Φ -rationalizability. The second one was a version of the conditional cardinal case where beliefs over a product set of states were required to be independent. The original definition of (independent) rationalizability in Bernheim (1984) and Pearce (1984) is equivalent to imposing this restriction as Φ_i in the definition of Φ -rationalizability, where the independence is over the action profiles of the opponents.

3.2 Epistemic Foundations

Let $\mathcal{T} = (T_i, s_i, \pi_i)_{i \in I}$ be an epistemic type space of players' actions and preferences, where, for each $i \in I$, T_i is a finite set of player i 's types, $s_i: T_i \rightarrow A_i$, and $\pi_i: T_i \rightarrow P^*(T_{-i})$. Type t_i of player i knows his own action $s_i(t_i)$ and has a non-trivial preference $\pi_i(t_i)$ over $F(T_{-i})$, regarding other players' types as payoff-relevant states.

Given $s_{-i} = (s_j)_{j \neq i}: T_{-i} \rightarrow A_{-i}$, type t_i 's preference $\pi_i(t_i)$ induces another preference $\hat{\pi}_i(t_i) \in P(A_{-i})$ as follows:

$$f \hat{\pi}_i(t_i) f' \Leftrightarrow f \circ s_{-i} \pi_i(t_i) f' \circ s_{-i}$$

for all $f, f' \in F(A_{-i})$. (Note that $f \circ s_{-i}, f' \circ s_{-i} \in F(T_{-i})$.) We say that type t_i is *rational* if he plays a best response with respect to his preference:

$$g(s_i(t_i), \cdot) \hat{\pi}_i(t_i) g(a_i, \cdot)$$

for any $a_i \in A_i$. Let Rat_i be the set of player i 's rational types. Let

$$\hat{\Phi}_i = \{t_i \in T_i \mid \hat{\pi}_i(t_i) \in \Phi_i\}$$

be the set of player i 's types whose induced preferences over $F(A_{-i})$ satisfy the restriction Φ_i .

For each product event $E = \prod_{i \in I} E_i$, with each $E_i \subseteq T_i$, let $C_i(E_{-i})$ be the set of player i 's types $t_i \in T_i$ whose preferences $\pi_i(t_i)$ are certain about E_{-i} . Let $C(E) = \prod_i C_i(E_{-i})$. Let $CC(E) = \bigcap_{n=0}^{\infty} C^n(E)$ be the set of states at which E holds and there is common certainty about E .

The next proposition formalizes the idea that the implications of common certainty of rationality and preference restrictions Φ are captured by Φ -rationalizability.

Proposition 5. *For any epistemic type space,*

$$s_i \left(\text{proj}_{T_i} CC \left(\prod_{j \in I} (\text{Rat}_j \cap \hat{\Phi}_j) \right) \right) \subseteq R_i^\Phi.$$

Moreover, there exists an epistemic type space in which the other direction of the set-inclusion also holds.

Proof. The first part follows since we can show the following by induction on n :

$$s_i \left(\text{proj}_{T_i} \bigcap_{k=0}^n C^k \left(\prod_{j \in I} (\text{Rat}_j \cap \hat{\Phi}_j) \right) \right) \subseteq R_{i,n+1}^\Phi.$$

For $n = 0$, any type $t_i \in \text{Rat}_i \cap \hat{\Phi}_i$ plays a best response with respect to his preference $\hat{\pi}_i(t_i) \in \Phi_i$, thus $s_i(t_i) \in R_{i,1}^\Phi$. For $n \geq 1$, pick any $t_i \in \text{proj}_{T_i} \bigcap_{k \leq n} C^k \left(\prod_j (\text{Rat}_j \cap \hat{\Phi}_j) \right) = \text{Rat}_i \cap \hat{\Phi}_i \cap C_i \left(\bigcap_{k \leq n-1} C^k \left(\prod_j (\text{Rat}_j \cap \hat{\Phi}_j) \right) \right)$. Since $\pi_i(t_i)$ is certain about $\text{proj}_{T_{-i}} \bigcap_{k \leq n-1} C^k \left(\prod_j (\text{Rat}_j \cap \hat{\Phi}_j) \right)$, it follows from the induction hypothesis that $\hat{\pi}_i(t_i)$ is certain about $R_{-i,n}^\Phi$. Since t_i plays a best response with respect to $\hat{\pi}_i(t_i) \in \Phi_i \cap P(A_{-i} \mid R_{-i,n}^\Phi)$, we have $s_i(t_i) \in R_{i,n+1}^\Phi$.

The second part holds for $\mathcal{T} = (T_i, s_i, \pi_i)$, where $T_i = R_i^\Phi$, s_i is the identity function, and, for each $a_i \in T_i = R_i^\Phi$, $\pi_i(a_i)$ is a preference in $\Phi_i \cap P(A_{-i} \mid R_{-i}^\Phi)$ that rationalizes a_i . ■

In the case of unconditional cardinal restrictions Φ , the above epistemic characterization is still true, but Claim 2 has an unusual implication: Φ -rationalizability captures rationality and preference restrictions Φ , and imposing common certainty of rationality and preference restrictions Φ does not have further implications on actions: $s_i(\text{Rat}_i \cap \hat{\Phi}_i) \subseteq R_{i,1}^\Phi = R_i^\Phi$.

4 Incomplete Information Games

The formulation in the previous section can extend to incomplete information games to incorporate uncertainty about the outcome function and other players' preferences. We relax common certainty of the outcome function by introducing a finite set of states, Θ , that can affect the outcome. In addition, rather than assuming common certainty of preferences over outcomes, we allow players' preferences to depend on Θ and other players' (higher order) preferences. Thus we consider a type space $\mathcal{T} = (T_i, \pi_i)_{i \in I}$, where, for each $i \in I$, T_i is a finite set of player i 's types and $\pi_i: T_i \rightarrow P^*(\Theta \times T_{-i})$. Now a game is given by $G = (I, (A_i)_{i \in I}, g)$, where $g: A \times \Theta \rightarrow \Delta(Z)$ is an outcome function that depend on states. We impose preference restrictions $\Phi_i \subseteq P(A_{-i} \times \Theta \times T_{-i})$.⁸⁹

⁸With slightly heavier notations, we may allow for A_i and Φ_i to be dependent on player i 's type t_i .

⁹Strictly speaking, our incomplete information setup is not a generalization of the corresponding complete information setup since, even if $|\Theta| = |T| = 1$, π_i uniquely determines player i 's unconditional preference over lotteries.

We define an interim version of preference-correlated rationalizability as follows. For a correspondence $\Gamma = (\Gamma_i)_{i \in I}$ with $\Gamma_i: T_i \rightrightarrows A_i$, we say that a_i is a Φ_i -best response to Γ_{-i} for type t_i if there exists $\succeq_i \in P(A_{-i} \times \Theta \times T_{-i})$ such that

- $\succeq_i \in \Phi_i$
- the restriction of \succeq_i to acts $F(\Theta \times T_{-i})$ (constant over A_{-i}) is equal to $\pi_i(t_i)$,
- the restriction of \succeq_i to acts $F(A_{-i} \times T_{-i})$ (constant over Θ) is certain about the graph of Γ_{-i} , and
- a_i is a best response with respect to \succeq_i (or its restriction to acts $F(A_{-i} \times \Theta)$), i.e., $g(a_i, \cdot, \cdot) \succeq_i g(a'_i, \cdot, \cdot)$ for any $a'_i \in A_i$.

Let $R_{i,0}^\Phi(t_i) = A_i$, and for $n \geq 1$, $R_{i,n}^\Phi(t_i)$ be the set of Φ_i -best responses to $R_{-i,n-1}^\Phi$. Then an action in $R_i^\Phi(t_i) = \bigcap_n R_{i,n}^\Phi(t_i)$ is called *interim Φ -rationalizable for type t_i* .

For example, suppose that each player i has conditional cardinal preference restriction Φ_i with nonconstant von Neumann-Morgenstern index $v_i: Z \rightarrow \mathbb{R}$. With this restriction, each preference $\pi_i(t_i) \in P^*(\Theta \times T_{-i})$ is associated with a belief $\tilde{\pi}_i(t_i) \in \Delta(\Theta \times T_{-i})$, and the outcome function induces state dependent payoff functions $\tilde{g}_i(a, \theta) = \sum_z g(a, \theta)(z)v_i(z)$. Then interim Φ -rationalizability is equivalent to interim correlated rationalizability (Dekel, Fudenberg and Morris (2007)) applied to belief type space $\tilde{T} = (T_i, \tilde{\pi}_i)_{i \in I}$ and incomplete information game $\tilde{G} = (I, (A_i, \tilde{g}_i)_{i \in I})$.

In the case that players do not impose any preference restrictions, $\Phi_i = P(A_{-i} \times \Theta \times T_{-i})$, one can prove the following “bad behavior.” The proof is given in the Appendix.

Claim 3. *In the case of no preference restrictions, we have $R_i^\Phi(t_i) = R_{i,1}^\Phi(t_i)$. Especially, if $\pi_i(t_i)$ and $\pi_i(t'_i)$ agree on acts $F(\Theta)$ (constant over T_{-i}), then $R_i^\Phi(t_i) = R_i^\Phi(t'_i)$.*

Thus no preference relation $\pi_i(t_i)$ on T_{-i} -dependent acts affects the set of Φ -rationalizable actions. This is in stark contrast to the results in Dekel, Fudenberg and Morris (2006, 2007) and Bergemann, Morris and Takahashi (2011), who show that, if we impose conditional cardinal (respectively, worst outcome) restrictions, then the set of interim Φ -rationalizable actions for type t_i depend on higher order beliefs (respectively, preferences) of t_i .

To make these setups more consistent with each other, one could either impose unconditional preference restrictions $\pi_i \in P^*(\{*\})$ in the complete information setup, or use correspondence $\pi_i: T_i \rightrightarrows P^*(\Theta \times T_{-i})$ in the incomplete information setup. We maintain this slight inconsistency in notation because we believe the formalisms for the complete information and incomplete information cases, respectively, make the cleanest connections with the respective related literatures.

A Proof of Claims 2 and 3

Here we prove Claim 3. Claim 2 follows by setting $|\Theta| = |T| = 1$.

First, fix any player $i \in I$. For each $j \neq i$ and $t_j \in T_j$, if $a_j \in R_{j,1}^\Phi(t_j)$, then let $\sigma_j(\cdot | a_j, t_j)$ be the point mass on a_j . If $a_j \in A_j \setminus R_{j,1}^\Phi(t_j)$, then, as in Claim 1, it follows from Lemma 1 and Proposition 1 that there exists a mixed action $\sigma_j(\cdot | a_j, t_j) \in \Delta(A_j)$ such that $g(\sigma_j(\cdot | a_j, t_j), a_{-j}, \theta) - g(a_j, a_{-j}, \theta)$ is independent of a_{-j} (but may be dependent on θ). Without loss of generality, we assume that $\sigma_j(\cdot | a_j, t_j) \in \Delta(R_{j,1}^\Phi(t_j))$. For each $a_{-i} \in A_{-i}$ and $t_{-i} \in T_{-i}$, define $\sigma_{-i}(\cdot | a_{-i}, t_{-i}) \in \Delta(R_{-i,1}^\Phi(t_{-i}))$ by $\sigma_{-i}(a'_{-i} | a_{-i}, t_{-i}) = \prod_{j \neq i} \sigma_j(a'_j | a_j, t_j)$ for each $a'_{-i} \in R_{-i,1}^\Phi(t_{-i})$.

Pick any $t_i \in T_i$ and $a_i \in R_{i,1}^\Phi(t_i)$, which is a best response under state dependent utility index $u_i: A_{-i} \times \Theta \times T_{-i} \times Z \rightarrow \mathbb{R}$ and that $\sum_{a_{-i}} u_i(a_{-i}, \theta, t_{-i}, z)$ represents $\pi_i(t_i)$. We will show that a_i is a best response in the second round of iteration.

Define another state dependent utility index $u'_i: A_{-i} \times \Theta \times T_{-i} \times Z \rightarrow \mathbb{R}$ by

$$u'_i(a'_{-i}, \theta, t_{-i}, z) = \sum_{a_{-i} \in A_{-i}} \sigma_{-i}(a'_{-i} | a_{-i}, t_{-i}) u_i(a_{-i}, \theta, t_{-i}, z)$$

for $a'_{-i} \in A_{-i}$, $\theta \in \Theta$, $t_{-i} \in T_{-i}$ and $z \in Z$. First,

$$\begin{aligned} \sum_{a'_{-i} \in A_{-i}} u'_i(a'_{-i}, \theta, t_{-i}, z) &= \sum_{a_{-i}, a'_{-i} \in A_{-i}} \sigma_{-i}(a'_{-i} | a_{-i}, t_{-i}) u_i(a_{-i}, \theta, t_{-i}, z) \\ &= \sum_{a_{-i} \in A_{-i}} u_i(a_{-i}, \theta, t_{-i}, z), \end{aligned}$$

which represents $\pi_i(t_i)$. Second, since $\sigma_{-i}(\cdot | a_{-i}, t_{-i}) \in \Delta(R_{-i,1}^\Phi(t_{-i}))$ for any $a_{-i} \in A_{-i}$ and $t_{-i} \in T_{-i}$, u'_i represents a preference whose restriction to acts over $A_{-i} \times T_{-i}$ is certain about the graph of $R_{-i,1}^\Phi$. Third, for every $a'_i \in A_i$,

$$\begin{aligned} &\sum_{a'_{-i} \in A_{-i}} g(a'_i, a'_{-i}, \theta)(z) u'_i(a'_{-i}, \theta, t_{-i}, z) \\ &= \sum_{a_{-i}, a'_{-i} \in A_{-i}} g(a'_i, a'_{-i}, \theta)(z) \sigma_{-i}(a'_{-i} | a_{-i}, t_{-i}) u_i(a_{-i}, \theta, t_{-i}, z) \\ &= \sum_{a_{-i} \in A_{-i}} g(a'_i, \sigma_{-i}(\cdot | a_{-i}, t_{-i}), \theta)(z) u_i(a_{-i}, \theta, t_{-i}, z) \\ &= \sum_{a_{-i} \in A_{-i}} (g(a'_i, a_{-i}, \theta)(z) + h(a_{-i}, \theta, t_{-i})(z)) u_i(a_{-i}, \theta, t_{-i}, z) \\ &= \sum_{a_{-i} \in A_{-i}} g(a'_i, a_{-i}, \theta)(z) u_i(a_{-i}, \theta, t_{-i}, z) + \sum_{a_{-i} \in A_{-i}} h(a_{-i}, \theta, t_{-i})(z) u_i(a_{-i}, \theta, t_{-i}, z), \end{aligned}$$

for $\theta \in \Theta$, $t_{-i} \in T_{-i}$ and $z \in Z$, where $h(a_{-i}, \theta, t_{-i}) := g(a'_i, \sigma_{-i}(\cdot \mid a_{-i}, t_{-i}), \theta) - g(a_i, a_{-i}, \theta)$ is independent of a'_i by the definition of $\sigma_{-i}(\cdot \mid a_{-i}, t_{-i})$. Since a_i is a best response with respect to u_i , it is also a best response with respect to u'_i .

References

- [1] Aumann, R. (1987). “Correlated Equilibrium as an Expression of Bayesian Rationality,” *Econometrica* 55, 1-18.
- [2] Bergemann, B., S. Morris and S. Takahashi (2011). “Interdependent Preferences and Strategic Distinguishability.”
- [3] Bernheim, D. (1984). “Rationalizable Strategic Behavior,” *Econometrica* 52, 1007-1028.
- [4] Bernheim, D. (1986). “Axiomatic Characterizations of Rational Choice in Strategic Environments,” *Scandinavian Journal of Economics* 88, 473-488.
- [5] Bogomolnaia, A. and H. Moulin (2001). “A New Solution to the Random Assignment Problem,” *Journal of Economic Theory* 100, 295-328.
- [6] Börgers, T. (1993). “Pure Strategy Dominance,” *Econometrica* 61, 423-430.
- [7] Brandenburger, A. and E. Dekel (1987). “Rationalizability and Correlated Equilibria,” *Econometrica* 55, 1391-1402.
- [8] Dekel, E., D. Fudenberg and S. Morris (2006). “Topologies on Types,” *Theoretical Economics* 1, 275-309.
- [9] Dekel, E., D. Fudenberg and S. Morris (2007). “Interim Correlated Rationalizability,” *Theoretical Economics* 2, 15-40.
- [10] Di Tillio, A. (2008). “Subjective Expected Utility in Games,” *Theoretical Economics* 3, 287-323.
- [11] Epstein, L. G. (1997). “Preference, Rationalizability and Equilibrium,” *Journal of Economic Theory* 73, 1-29.
- [12] Fishburn, P. (1970). *Utility Theory for Decision Making*. New York, Wiley.
- [13] Ledyard, J. (1986). “The Scope of the Hypothesis of Bayesian Equilibrium,” *Journal of Economic Theory* 39, 59-82.

- [14] Pearce, D. (1984). "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica* 52, 1029-1050.
- [15] Savage, L. (1954, 1972). *The Foundations of Statistics*. New York, Dover.
- [16] Tan, T. and A. Werlang (1988). "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory* 45, 370-391.