# GENERIC UNIQUENESS OF RATIONALIZABLE ACTIONS

MUHAMET YILDIZ

Abstract. For a finite set of actions and a rich set of fundamentals, consider the interim rationalizable actions on the universal type space, endowed with the usual product topology. (1) Generically, there exists a unique rationalizable action profile. (2) Every model can be approximately embedded in a dominance-solvable model. (3) For any given rationalizable strategy of any finite model, there exists a nearby finite model with common prior such that the given rationalizable strategy is uniquely rationalizable for nearby types.

*Key words*: higher-order uncertainty, rationalizability, universal type space, multiplicity

*JEL Numbers*: C72, C73.

## 1. Introduction

In Game Theoretical applications there are often many rationalizable strategies. In this paper, I try to understand when and why this multiplicity occurs and how we should address it when it occurs.

I start with the observation that Game Theoretical models of a given situation necessarily make strong simplifying assumptions, such as common knowledge assumptions, which idealize its true underlying features. These assumptions are meant to be satisfied only approximately in the actual situation. From the point of view of the modeler as outside observer, a model will look similar to the actual situation, but will be an idealization of it. Such idealizations may nevertheless have significant impact on the conclusions, as demonstrated by Kreps, Milgrom, Roberts, and Wilson (1982).

---

*Date*: October 2004.

I thank Jonathan Weinstein for long collaborations on the topic; this work is partly built on our joint work, and we had discussed some closely related ideas. I thank Stephen Morris for extensive discussions on the topic while I visited Cowles Foundation; some of the ideas in this paper occurred to us during a lunch discussion. I thank Daron Acemoglu, Glenn Ellison, Bart Lipman, and Casey Rothschild for invaluable comments, and Dov Samet and Aviad Heifetz for earlier discussions.

This paper unearths a deep connection between these idealizations and the presence of multiplicity of rationalizable outcomes. This connection was partially exposed in the seminal work of Carlsson and van Damme (1993), who illustrated that multiplicity may sometimes be a direct result of the implicit simplifying assumptions of our models. To be concrete, consider their well-known example.

Example 1 (Carlsson and van Damme (1993)). Consider the payoff matrix

| | $\alpha_2$ | $\beta_2$ |
|---|---|---|
| $\alpha_1$ | $\theta, \theta$ | $\theta - 1, 0$ |
| $\beta_1$ | $0, \theta - 1$ | $0, 0$ |

where $\theta$ is a real number. Assume that $\theta$ is unknown but each player $i \in \{1, 2\}$ observes a noisy signal $x_i = \theta + \varepsilon\eta_i$, where $(\eta_1, \eta_2)$ is independently distributed from $\theta$, and the support of $\theta$ contains an interval $[a, b]$ where $a < 0 < 1 < b$. When $\varepsilon = 0$, $\theta$ is common knowledge. If it is also the case that $\theta \in (0, 1)$, there exist two Nash equilibria in pure strategies and one Nash equilibrium in mixed strategies. Under complete information, the players are able to "coordinate" on different equilibria. With incomplete information, this is no longer possible. Under mild conditions, Carlsson and van Damme show that when $\varepsilon$ is small but positive, multiplicity disappears: for each signal value $x_i \neq 1/2$, there exists a unique rationalizable action. The rationalizable action is $\beta_i$ whenever $x_i < 1/2$, and it is $\alpha_i$ whenever $x_i > 1/2$.

The model in which $\theta = \bar{\theta} \neq 1/2$ is common knowledge (i.e. $\varepsilon = 0$) idealizes the situation in which incomplete information is small (i.e. $\varepsilon$ is small but positive) and players believe that $\theta$ is close to $\bar{\theta}$. We, modelers, use the idealized model of $\varepsilon = 0$ as an approximation of small incomplete information, in order to simplify our model. The simplification weakens our ability to make predictions, as there are now new outcomes that would have been ruled out in the relaxed model. It also makes it more difficult to generate insights about the actual situation. For example, if we take $\alpha$ and $\beta$ as attacking and not attacking a currency, respectively, and $\theta$ as the vulnerability of the economy, then Example 1 predicts that attack becomes likelier when the economy is more vulnerable (Morris and Shin (1998)), which is not clear in the complete information case.

In this paper, I first show that this intuition of Carlsson and van Damme is quite general. Under their assumption that each action can be dominant at some parameter value, I show that, by introducing a small amount of incomplete information, we can *always* relax the implicit assumptions of a model and obtain an open set of situations in which there is a unique rationalizable outcome, which is the same across all these situations. Therefore, without a very precise knowledge of the situation, we cannot rule out the possibility that we could have predicted accurately what the rationalizable outcome is by learning more about the actual situation. In contrast, when there is a unique rationalizable outcome, slight relaxations of the assumptions will not have any effect.

My second result characterizes the situations with multiple rationalizable outcomes. To be concrete, consider the signal value $x_i = 1/2$. This value can be considered to be an idealization of two cases: (i) $x_i$ is close to $1/2$ but smaller, in which case $\beta_i$ is the unique outcome, and (ii) $x_i$ is close to $1/2$ but larger, in which case $\alpha_i$ is the unique outcome. The signal value $x_i = 1/2$ idealizes two strategically distinct situations, and both $\alpha_i$ and $\beta_i$ must be rationalizable at $x_i = 1/2$ due to upper-semicontinuity. I show that multiplicity occurs in a model precisely for this reason. Whenever there are multiple rationalizable actions for a type, the type simultaneously idealizes multiple strategically distinct situations, with distinct unique rationalizable outcomes. Each rationalizable action profile $a$ corresponds to an open set of situations where the assumptions are slightly relaxed and $a$ is the only rationalizable outcome.

Given our limitations in observing the actual situation, we will always idealize the situation to some degree. Sometimes, we may want to get insights about the situation by considering a more detailed information structure with unique rationalizable outcome. The first result shows that, in principal, this is always possible, and one can do it without having to rely on rigid assumptions on the information structure. The second result suggests that, when there are multiple rationalizable outcomes in the original model, the insights generated in this way may depend on the information structure one considers. Sometimes, we may want to predict what the players will do, using a refinement of rationalizability, such a particular equilibrium refinement. We would also want to know how our predictions would change if we used a model with slightly relaxed assumptions. In that case, by the second result, the set of rationalizable actions gives us invaluable information. It tells us which predictions we would necessarily make under various information structures where our assumptions are only approximately satisfied.

Formulating the above results for general games inherently requires topological notions on large spaces, and interpretation of such results requires great care. In the rest of this introduction, I will explain and justify my formulation, describe my formal results, and discuss their implications.

In real life, players' incomplete information is not of the form we describe in a usual type space. They have some (vague) beliefs about the payoffs, which are called the first-order beliefs, some (vaguer) beliefs about the other players' beliefs, which are called second order beliefs, and so on. Following Harsanyi (1967), we, modelers, use a type space and a type profile $t$ within this space to model these beliefs. The type profile $t$ is such that when we compute the beliefs of $t$ about payoffs and so on using the type space, they are similar to the beliefs of the players we try to model. In my formulation, I will directly consider these (coherent) belief hierarchies, each of which corresponds to a type in some type space. The set of all such hierarchies is called the universal type space (Mertens and Zamir (1985), Brandenburger and Dekel (1993)). My results will apply both to the universal type space and to smaller but sufficiently rich type spaces.

I consider a topology on belief hierarchies that mathematically captures (a) the usual continuity notion in usual models and (b) the idea that we modelers cannot observe the entire hierarchy of beliefs, as they may not be well-articulated in real life. For example, in Example 1, beliefs are continuous functions of $(\varepsilon, x_1, x_2)$,[1] and we would like to consider types corresponding to $(\varepsilon, x_1, x_2)$ and $(\varepsilon', x'_1, x'_2)$ close, when $(\varepsilon, x_1, x_2)$ and $(\varepsilon', x'_1, x'_2)$ are close to each other in the usual sense.[2] We would like to consider an open set $O$ of such parameters to be a relaxation of the assumptions in $(\varepsilon, x_1, x_2)$—about the parameters—when $(\varepsilon, x_1, x_2)$ is on the boundary of $O$.

Mertens and Zamir (1985) have shown that the product topology is the topology described above. Consider a usual model with a compact state space with any topology, and assume that the beliefs are continuous functions of states. (In Example 1, a state is $(\varepsilon, \theta, x_1, x_2)$ .) Assume that there are no two types with the same hierarchy. Then, the function that maps the states to the corresponding belief hierarchies and underlying payoff parameters is an isomorphism when we put the product topology

[1]Recall that beliefs are probability distributions, and we put the usual weak topology on probability distributions, the topology corresponding to "the convergence in distribution".

[2]This also requires that I do not consider whether two situations are strategically close. When $\varepsilon' > 0$ and $\bar{\theta} \neq 1/2$, $\left(\varepsilon = 0, \bar{\theta}, \bar{\theta}\right)$ is strategically distinct from any type $\left(\varepsilon', \bar{\theta}, \bar{\theta}\right)$. As $\varepsilon \to 0$, the types with $x_i = \bar{\theta}$ would not converge to the common knowledge case in strategic topologies, such as that of Monderer and Samet (1989) and Dekel, Fudenberg, and Morris (2004). (In this paper, I call two situations strategically distinct when the set of rationalizable actions are distinct, even if the $\delta$-rationalizable actions are similar, as the former concept turns out to be the relevant notion for understanding multiplicity.)

on the belief hierarchies. That is, taking limit on belief hierarchies with respect to the product topology is *equivalent* to taking limit in the original space on the corresponding states. Since I would like to be able to enlarge models in a manner in which beliefs remain continuous functions of the states, I will then use the product topology in the universal type space. Because of this isomorphism, my results will be true for any such model with respect to its own topology, provided that it is sufficiently rich.

As I will explain formally in Section 3, the open sets in product topology precisely correspond to the sets of types consistent with our observation when we make noisy observation about finite-order beliefs. (In the rest of the paper, I will consider only such observations.) In this topology, openness of a set $U$ means that if the actual type is in $U$ and our observation is sufficiently precise, then we would know that actual type is in $U$. Similarly, denseness of a set $U$ means that we could never rule out the possibility that the actual case is represented by a type in $U$. Therefore, an open and dense set contains all types, except for a nowhere-dense set of idealized situations, such as $\varepsilon = 0$ in the above example. The nowhere-dense set is just the boundary of the open and dense set, containing no open set. Relaxation $O$ of a type profile $t$, defined above, has now a special meaning: Even if the actual situation is as described by $t$, we can never rule out the possibility that, by having a more precise observation, we will learn that the actual situation is in $O$.

I consider a finite set of players and a finite set $A$ of actions. Following Carlsson and van Damme, I assume that each action becomes strictly dominant for some parameter value. I endow the game with the universal type space $T^*$ with the product topology. Then, I prove the following.

MAIN RESULT. *Generically, there exists a unique rationalizable action profile, and it is generically continuous. That is, there exist an open, dense set $U \subset T^*$ and a continuous (i.e. locally constant) function $s^* : U \to A$, such that $s^*(t)$ is the unique rationalizable action profile at $t$ for each $t \in U$. In particular, every rationalizable strategy is continuous on the open, dense set $U$.*[3]

First, suppose that there is a unique rationalizable outcome for the actual situation we want to model. Since $U$ is open and $s^*$ is locally constant, if we have a

[3]Here, $U$, the set of all type profiles with unique rationalizable action profile, is open simply because the rationalizability correspondence is upper semicontinuous (Dekel, Fudenberg, and Morris (2003)) and the action space is finite. I show that $U$ is dense, using a result of Mertens and Zamir (1985) and a variant of a construction by Weinstein and Yildiz (2004), whose main idea can be traced back to Rubinstein (1989) and Carlson and van Damme (1993).

sufficiently precise noisy observation about finite-order beliefs, then we would know what that outcome is. Second, suppose that there are multiple rationalizable outcomes for the actual situation. Then, no matter how precise our observation is, we could not rule out the possibility that, by having a more precise information, we would come to learn that there is a unique rationalizable outcome, and we would learn what it is. (This is the precise meaning of $T^{*}\backslash U$ being nowhere-dense in this topology.) In that sense, multiplicity occurs only in idealized situations, where slight relaxation of assumptions would lead to open set of situations with unique rationalizable outcome.

Now consider a usual model and enrich it by allowing sufficiently many types, maintaining the compactness, continuity, and "non-redundancy" assumptions described above. Then, according to the main result, multiplicity will occur only on a nowhere-dense set of states in the enriched model with respect to the topology on the model. In particular, there cannot be an open set of states at which there are multiple rationalizable outcomes. Hence, in order to maintain an argument based on multiplicity, one has to maintain very rigid assumptions on the set of states. Slight relaxations of these assumptions (with respect to the topology on the model) will lead to an open set of states with unique rationalizable outcomes. In particular, given any finite-action game with arbitrary payoff and information structures, we can introduce small incomplete information in such a way that the resulting game is dominance-solvable. Moreover, the dominance-solvable model will remain so, when further small perturbations are introduced. In contrast, even if there is a perturbation that leads to multiplicity, we can introduce further perturbations to reach a unique rationalizable outcome.

In application, we often use small type spaces, such as finite models with a common prior. Suppose that only a subset of models are considered to be possible, and let $T$ be the set of type profiles generated by these models. Assume that $T$ is dense; i.e., for each possible observation, there is a type profile $t \in T$ that is consistent with our observation. For example, finite models with common prior would be one such set (Mertens and Zamir (1985), Lipman (2003)). Then, since $U$ is open and dense in $T^{*}$, $U \cap T$ is dense and open in relative topology on $T$. Once again, within this smaller set of models, multiplicity occurs only in idealized situations.

In general, slight relaxation of an assumption in a given model tends to reduce the number of rationalizable actions for the perturbed types. There is a simple mathematical reason for this: the rationalizability correspondence is upper-semicontinuous. Each type profile $t$ has an open neighborhood, such that if an action profile is rationalizable for some $t'$ in this neighborhood, it must also be rationalizable for $t$. Then, when we relax an assumption—so slightly that we remain in this neighborhood, we can only get rid of some rationalizable actions.

*Why do we have multiplicity?* If an action is rationalizable when an assumption is approximately satisfied, then that action will remain rationalizable when the assumption is exactly satisfied (by upper-semicontinuity). Now, consider two open sets $U^a$ and $U^b$, where $a$ and $b$ are unique rationalizable actions, respectively, and suppose that there is a type $t$, located on the common boundary of these sets. Then, both $a$ and $b$ will be rationalizable for $t$, leading to multiplicity. Each of $U^a$ and $U^b$ describe a strategically distinct situation, and $t$ can be considered as an idealization of either situation. Keeping track of all such idealizations, rationalizability yields both actions $a$ and $b$ as possible solutions at $t$.

I show that this is the *only* way a finite type (i.e. a type from a finite model) can have multiple rationalizable actions. Given any type $t$ and any rationalizable action $a$ for $t$, we can suitably relax the assumptions of $t$ and find an open set $U^a$ of types for which $a$ is uniquely rationalizable. For example, if $t$ has rationalizable actions $a$, $b$, and $c$, then $t$ is located on the common boundary of some open sets $U^a$, $U^b$, and $U^c$, where $a$, $b$, and $c$ are uniquely rationalizable, respectively. That is to say $t$ is a simultaneous idealization of multiple strategically distinct situations. Each rationalizable action for $t$ represents such a situation idealized by type $t$. Combining this with the observation in the previous paragraph, one then uncovers a precise reason for multiplicity: a type has multiple rationalizable actions if and only if it represents simultaneous idealization of multiple strategically distinct situations. It is not surprising, then, that we tend to a have a large number of rationalizable strategies in complete-information models. Such a model idealizes all of the situations in which private information is small, which can happen in many different information structures. The information structure may have significant impact on the outcome even when players have small private information. (The above result remains valid when we impose the common-prior assumption. This is due to Lipman (2003), who shows that the common-prior assumption puts only few restrictions on the finite-order beliefs.)

The above result is closely related to robustness results of Weinstein and Yildiz (2004). Under a significantly weaker richness assumption, they have shown that, if a non-empty equilibrium refinement has a prediction that remains valid when we only know the finite-order beliefs, then the prediction must be true for all strategies that survive iterated elimination of actions that are never a strict best reply. A variant of their construction plays a central role in my proofs.

Carlsson and van Damme have extended Example 1 to all two-player, two-action supermodular games of complete information, and Morris and Shin (1998) and Frankel, Morris, and Pauzner (2003) have extended it further to all monotone supermodular games of complete information (Van Zandt and Vives (2004)). These results appear to be specific to supermodular games, in that their perturbations need not reduce the set of rationalizable outcomes in general games, such as the Matching-Pennies game. I show that, once we relax their assumptions on information structure, their conclusion generalizes to all games with arbitrary information structures.

In the next section I illustrate how one can make the Matching-Pennies game dominance-solvable by introducing a small amount of incomplete information. In Section 3, I introduce the model and preliminary results. The main results are presented in Section 4. The proof of a central lemma is presented in Section 5. Section 6 concludes.

## 2. Matching Pennies

The information structure of Carlsson and van Damme does not work in Matching-Pennies game. Focusing on this difficult case, I now illustrate how one can get rid of multiplicity by introducing incomplete information. In the Matching-Pennies game, multiplicity arises because some players have an incentive to change their actions when these actions are known by the players. Introducing incomplete information will ease this tension. For, under incomplete information, players need not know the other players' actions even if they know the others' strategies. In my construction, I will first consider an intuitive belief structure without a common prior and then obtain a belief structure with a common prior.

Example 2 (Matching Pennies—without a common prior). Consider the payoff matrix

| | $\alpha_2$ | $\beta_2$ |
|---|---|---|
| $\alpha_1$ | $\theta, 0$ | $\theta - 1, \theta$ |
| $\beta_1$ | $0, 0$ | $0, \theta - 1$ |

.

If $\theta$ is common knowledge and is in $(0,1)$, then there is no pure strategy equilibrium. Take $\Theta = \{\theta_0, \theta_1, \ldots, \theta_{M-1}\}$, where $\theta_0 = -\varepsilon/2$, $\theta_1 = \varepsilon/2$, $\theta_2 = 3\varepsilon/2, \ldots,$ $\theta_{M-1} = \bar{\theta} < 1$, and assume that $\theta$ is uniformly distributed on $\Theta$. Each player $i$ observes a signal $x_i$, about which the players have different priors. Conditional on $\theta = \theta_m$, each player $i$ assigns probability $1-\gamma$ to $(x_i, x_j) = (\theta_m, \theta_{m-1})$ and probability $\gamma$ to $(x_i, x_j) = (\theta_{m-1}, \theta_m)$. As in Example 1, it is common knowledge that the players' signals are within $\varepsilon$-neighborhood of $\theta$, and the game converges to the the complete-information game as $\varepsilon \to 0$. For $\varepsilon = 0$, every strategy is rationalizable. But when $0 < \gamma < \varepsilon / [2(1-\varepsilon)]$, the incomplete-information game is dominance-solvable, and the unique rationalizable strategy profile is as in the following table:

| $x_i$ | $\theta_0$ | $\theta_1$ | $\theta_2$ | $\theta_3$ | $\theta_4$ | $\theta_5$ | $\theta_6$ | $\theta_7$ | $\theta_8$ | $\cdots$ |
|---|---|---|---|---|---|---|---|---|---|---|
| $s_1^*(x_1)$ | $\beta_1$ | $\alpha_1$ | $\alpha_1$ | $\beta_1$ | $\beta_1$ | $\alpha_1$ | $\alpha_1$ | $\beta_1$ | $\beta_1$ | $\cdots$ |
| $s_2^*(x_2)$ | $\alpha_2$ | $\alpha_2$ | $\beta_2$ | $\beta_2$ | $\alpha_2$ | $\alpha_2$ | $\beta_2$ | $\beta_2$ | $\alpha_2$ | $\cdots$ |

(Clearly, when $x_i = \theta_0$, player $i$ assigns high probability $1-\gamma$ to $\theta = \theta_0$, when $\beta_1$ and $\alpha_2$ are dominant actions for players 1 and 2, respectively. When, $x_i = \theta_1$, player $i$ assigns high probability to $(\theta, x_j) = (\theta_1, \theta_0)$. Given the dominant action for $j$ at $x_j = \theta_0$, the player $i$ has a unique best response; it is $\alpha_i$. One computes $s^*$ iteratively in this way.)

In this example the players do not have a common prior. This is not crucial. The elimination process in this game stops at the $M$th round, and hence the rationalizability depends only on the first $M$ orders of beliefs (Dekel, Fudenberg, and Morris (2003)). Using Lipman's (2003) method, we can then construct an incomplete-information game with a common prior and with types whose first $M$ orders of beliefs are as in the original game. These types will have unique rationalizable actions, as in the following example.

Example 3 (Matching Pennies—with a common prior). In the previous example, assume that, in addition to $x_i$, each player $i$ observes a noisy signal $y_i$ about a random variable $k$ that is correlated with $\theta$ and takes values in $\{1, 2, \ldots, 2K\}$ for some integer $K > M$. Player 1 observes the value $y_1(k)$ of the smallest odd number $y$ with $y \geq k$; e.g., $y_1(1) = 1$, $y_1(2) = 3$, $y_1(3) = 3$, etc. Player 2 observes the value $y_2(k)$ of the smallest even number $y$ with $y \geq k$, e.g., $y_2(1) = 2$, $y_2(2) = 2$, etc. Now, the players have a common prior $\bar{\mu}$ about $(\theta, x_1, x_2, k)$ as follows. Let $\mu_i(\theta, x_1, x_2)$ be the prior probability of $(\theta, x_1, x_2)$ according to player $i$ in the previous example, e.g., $\mu_1(\theta_1,\theta_1,\theta_0) = (1-\gamma)^2/M$ and $\mu_1(\theta_1,\theta_0,\theta_1) = \gamma^2/M$. Define $\bar{\mu}$ iteratively by

$$\begin{aligned} \bar{\mu}(\theta,x_1,x_2,1) &= \alpha\mu_1(\theta,x_1,x_2) \\ \bar{\mu}(\theta,x_1,x_2,k) &= L^{k-1}\alpha\mu_{i_k}(\theta,x_1,x_2) - \sum_{l<k}\bar{\mu}(\theta,x_1,x_2,l) \end{aligned}$$

for each $(\theta,x_1,x_2)$ and $k \in \{2,3,\ldots,2K\}$ where $L > (1-\gamma)/\gamma$, $\alpha = 1/L^{2K-1}$, and $i_k$ is 1 if $k$ is odd and 2 if $k$ is even. Once again, it is common knowledge that, in addition to $y_i$, each player observes a signal $x_i$ that is within $\varepsilon$-neighborhood of $\theta$. As $\varepsilon \to 0$, the belief hierarchy of each type with $(x_i, y_i(k))$ converges to that of the common knowledge of $\theta = x_i$. Lipman (2003) shows that

$$\bar{\mu}((\theta,x_1,x_2)|x_i,y_i(k)) = \mu_i((\theta,x_1,x_2)|x_i) \tag{2.1}$$

for each $y_i(k) \le 2K$. That is, the posterior beliefs in the new model are identical to those of previous example, except for the case that player 1 observes that $y_1(k) = 2K+1$. It follows from (2.1) that, for each $(x_i, y_i(k))$ with $y_i(k) \le 2K - m$ and $x_i = \theta_m$, there exists a unique rationalizable action

$$\hat{s}_i(x_i,y_i(k)) = s_i^*(x_i),$$

where $s_i^*$ is the unique rationalizable strategy of $i$ in the previous example.[4] In particular, the types with $(x_i, y_i(1))$, which approximate the complete-information model, will have unique rationalizable actions.

Notice that, in this example, the types whose belief hierarchies are far away from those of original model may have multiple rationalizable actions; for an example consider the types with $y_i(k) > 2K - m$ and $x_i = \theta_m$ for some $m$.

## 3. Model

Consider a game with finite set of players $N = \{1,2,\ldots,n\}$, finite set $A = A_1 \times \cdots \times A_n$ of action profiles $a = (a_1,a_2,\ldots,a_n)$, and utility functions $u_i : \Theta^* \times A \to \mathbb{R}$, $i \in N$, where $\Theta^*$ is a compact metric space of payoff-relevant parameters $\theta$, and $u_i$ is continuous in $\theta$. The finite set $A$ is endowed with the discrete topology. The

[4]Use induction on $m$ to check this. For $m = 0$, by (2.1), $s_i^*(\theta_m)$ is dominant action for each $(\theta_m, y_i(k))$ with $y_i(k) \le 2K$. Assuming the statement is true for $m-1$, consider any $(\theta_m, y_i(k))$ with $y_i(k) \le 2K - m$. Player $i$ knows that $y_j(k) \le 2K - m + 1$, and assigns very high probability on $\{\theta = \theta_m, x_j = \theta_{m-1}\}$. By assumption, he must assign high probability on $j$ playing $s_j^*(\theta_{m-1})$, against which the only best response is $s_i^*(\theta_m)$.

game is endowed with the universal type space. A type of a player $i$ is an infinite hierarchy of beliefs

$$t_i = \left(t_i^1, t_i^2, \ldots\right)$$

where $t_i^1 \in \Delta(\Theta^*)$ is a probability distribution on $\Theta^*$, representing the beliefs of $i$ about $\theta$, $t_i^2 \in \Delta(\Theta^* \times \Delta(\Theta^*)^n)$ is a probability distribution for $(\theta, t_1^1, t_2^1, \ldots, , t_n^1)$, representing the beliefs of $i$ about $\theta$ and the other players' first-order beliefs, and so on. Here, $\Delta(X)$ is the space of all probability distributions on $X$, endowed with the weak* topology. I assume that it is common knowledge that the beliefs are coherent (i.e., each player knows his beliefs and his beliefs at different orders are consistent with each other). The set of all such types are denoted by $T_i^*$; $T^* = T_1^* \times \cdots \times T_n^*$ denotes the set of all type profiles $t = (t_1, \ldots, t_n)$, and $T_{-i}^* = \prod_{j \neq i} T_j^*$ is the set of profiles of types $t_{-i}$ for players other than $i$. Each $T_i^*$ is endowed with the product topology, so that a sequence of types $t_{i,m}$ converges to a type $t_i$, denoted by $t_{i,m} \to t_i$, if and only if $t_{i,m}^k \to t_i^k$ for each $k$. A sequence of type profiles $t(m) = (t_{1,m}, \ldots, t_{n,m})$ converges to $t$ iff $t_{i,m} \to t_i$ for each $i$. For each type $t_i$, let $\kappa_{t_i} \in \Delta\left(\Theta^* \times T_{-i}^*\right)$ be the unique probability distribution that represents the beliefs of $t_i$ about $(\theta, t_{-i})$. Mertens and Zamir (1985) have shown that the mapping $t_i \mapsto \kappa_{t_i}$ is an isomorphism. That is, it is one-to-one, and $\kappa_{t_{i,m}} \to \kappa_{t_i}$ if and only if $t_{i,m} \to t_i$.

The product topology captures the idea that we cannot observe infinite hierarchy of beliefs. Suppose that we consider only a set of models as possible, and let $T$ be the set of belief hierarchies generated by these models. Suppose also that we have made some noisy observation about some first $k$ orders of beliefs, and for each $k' \leq k$, we find an open set of beliefs possible (w.r.t the weak topology on probability distributions). Then, the set of types that we find possible are those types in $T$ whose first $k$ orders of beliefs are in these open sets. The product topology relative to $T$ is the smallest topology under which all of such sets of types are open.

Remark 1. In my formulation, it is common knowledge that the payoffs are given by a fixed continuous function of parameters. This assumption is without loss of generality because we can take a parameter to be simply the function that maps action profiles to the payoff profiles. For example, we can take $\Theta^* = \Theta_1^* \times \cdots \times \Theta_n^*$ where $\Theta_i^* = [0,1]^A$ for each $i$, and let $u_i(\theta, a) = \theta_i(a)$ for each $(i, a, \theta)$. This model allows all possible payoff functions, and here $\theta$ is simply an index for the profile of the payoff functions. This model clearly satisfies the following richness assumption, which is also made by Carlsson and van Damme (1993).

Assumption 1 (Richness Assumption). *For each $i$ and each $a_i$, there exists $\theta^{a_i} \in \Theta^*$ such that*

$$u_i\left(\theta^{a_i}, a_i, a_{-i}\right) > u_i\left(\theta^{a_i}, a_i', a_{-i}\right) \qquad \left(\forall a_i' \neq a_i, \forall a_{-i}\right).$$

That is, the space of possible payoff structures is rich enough so that each action can be strictly dominant for some parameter value. When there are no a priori restrictions on the domain of payoff structures and the actions represent the strategies in a static game, Assumption 1 is automatically satisfied. When actions represent the strategies in a dynamic game, one needs to introduce trembles and use a reduced form to satisfy this assumption.

A *strategy* of a player $i$ is any function $s_i : T_i^* \to A_i$.[5] For each $i \in N$ and for each belief $\pi \in \Delta\left(\Theta^* \times A_{-i}\right)$, $BR_i(\pi)$ denotes the set of actions $a_i \in A_i$ that maximize the expected value of $u_i\left(\theta, a_i, a_{-i}\right)$ under the probability distribution $\pi$.

**Interim Correlated Rationalizability**. For each $i$ and $t_i$, set $S_i^0[t_i] = A_i$, and define sets $S_i^k[t_i]$ for $k > 0$ iteratively, by letting $a_i \in S_i^k[t_i]$ if and only if $a_i \in BR_i\left(\mathrm{marg}_{\Theta^* \times A_{-i}} \pi\right)$ for some $\pi \in \Delta\left(\Theta^* \times T_{-i}^* \times A_{-i}\right)$ such that $\mathrm{marg}_{\Theta^* \times T_{-i}^*} \pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in S_{-i}^{k-1}[t_{-i}]\right) = 1$. That is, $a_i$ is a best response to a belief of $t_i$ that puts positive probability only to the actions that survive the elimination in round $k-1$. I write $S_{-i}^{k-1}[t_{-i}] = \prod_{j \neq i} S_j^{k-1}[t_j]$ and $S^k[t] = S_1^k[t_1] \times \cdots \times S_n^k[t_n]$. The set of all rationalizable actions for player $i$ (with type $t_i$) is

$$S_i^\infty[t_i] = \bigcap_{k=0}^{\infty} S_i^k[t_i].$$

A strategy $s_i : T_i^* \to A_i$ is said to be rationalizable iff $s_i(t_i) \in S_i^\infty[t_i]$ for each $t_i$.

Remark 2. The interim correlated rationalizability (Battigalli (2003), Battigalli and Siniscalchi (2003) and Dekel, Fudenberg, and Morris (2003)) is the weakest among the known notions of rationalizability. Dekel, Fudenberg, and Morris (2003) show that, for arbitrary type space and independent of whether correlations are allowed, if an action $a_i$ is rationalizable for a type with belief hierarchy $t_i$, then $a_i$ is interim correlated rationalizable for $t_i$. Using such a weak notion of rationalizability strengthens my results; they will remain valid under any stronger notion of rationalizability.

[5]I do not restrict the strategies to be measurable. Measurability restriction could lead to a non-existence problem, which can be avoided in the present interim framework (Simon, 2003).

**Mathematical Definitions and Preliminary Results**.

Definition 1 (Genericity). The *closure* of a set $T \subseteq T^*$, denoted by $\overline{T}$, is the smallest closed set that contains $T$. A set $T$ is *dense* (in $T^*$) iff $\overline{T} = T^*$, i.e., for each $t \in T^*$, there exists a sequence of type profiles $t(m) \in T$ such that $t(m) \to t$. A set $T$ is said to be *nowhere-dense* iff the interior of $\overline{T}$ is empty, i.e., $\overline{T}$ does not contain any open set. A statement is said to be *generically true* if it is true on an open, dense set of type profiles.

An open and dense set $T \subseteq T^*$ is large in the sense that its complement, $T^* \backslash T$, is nowhere-dense. That is, we can approximate each $\tilde{t} \in T \backslash T^*$ by type profiles $t \in T^*$, and we cannot approximate any $t \in T^*$ by type profiles $\tilde{t} \in T \backslash T^*$. In that case, $T^* \backslash T$ is simply the boundary of $T$, denoted by $\partial T$. Clearly, topological notions of genericity may widely differ from measure theoretical notions of genericity, which are about how commonly an event occurs, but these notions are related (Oxtoby (1980)). This paper uses a strong topological notion of genericity with respect to a canonical topology. However, the results may not be true under other topologies or under measure theoretical notions of genericity. Hence, one should not interpret the results of this paper as saying that there are few types with multiple rationalizable actions.

Definition 2 (Finite Types, Models). A subset $T \subseteq T^*$ is said to be *belief-closed* iff for each $t_i \in T_i$, $\kappa_{t_i}(\Theta^* \times T_{-i}) = 1$. A *model* is a subset $\Theta \times T \subseteq \Theta^* \times T^*$ such that $\kappa_{t_i}(\Theta \times T_{-i}) = 1$ for each $t_i \in T_i$. When it does not lead to a confusion, I will use the terms *model* and *belief-closed subset of* $T^*$ interchangeably. A model $\Theta \times T$ (or $T$) is said to be *finite* iff $|\Theta \times T| < \infty$. Let $\hat{T}$ be the union of all finite, belief-closed subspaces $T \subset T^*$. Members of $\hat{T}$ are referred to as *finite types*.

In general, the image of a model in the universal type space need not be a product set $\Theta \times T$, but one can enlarge it to a product set by adding new states. By a well-known theorem of Mertens and Zamir, the beliefs $\kappa_{t_i}$ of players will be continuous functions of states $(\theta, t)$.

Lemma 1 (Mertens and Zamir (1985)). *$\hat{T}$ is dense, i.e., $\overline{\hat{T}} = T^*$.*

Definition 3 (Dominance-Solvability). A model $T \subseteq T^*$ is said to be *dominance-solvable* if and only if $|S^\infty [t]| = 1$ for each $t \in T$.

Definition 4 (Common Prior). A model $\Theta \times T \subseteq T^*$ is said to *admit a common prior (with full support)* if and only if there exists a probability distribution $p \in \Delta (\Theta \times T)$ such that $\kappa_{t_i} = p(\cdot|\Theta^* \times \{t_i\} \times T_{-i})$ for each $t_i \in T_i$.

The set of all type profiles that comes from a model with a common prior is denoted by $T^{CPA}$; formally,

$$T_i^{CPA} = \{t_i | t_i \in T_i \text{ for some } \Theta \times T \text{ with a common prior}\}.$$

Lipman (2003) shows that, in finite models, the common-prior assumption does not put any restriction on finite-order beliefs other than full support (see also Feinberg (2000)), proving following useful result.

Lemma 2 (Lipman (2003)). *$\hat{T} \cap T^{CPA}$ is dense in the universal type space.*

Definition 5 (Continuity). A strategy $s_i$ is said to be *continuous (or locally-constant) at* $t_i$ iff $s_i$ is constant on an open neighborhood of $t_i$. A (bounded) correspondence $F : T^* \to 2^A$ is said to be *upper-semicontinuous* if its graph is closed in the product topology of $T^* \times A$. Since $A$ is finite, $F$ is upper-semicontinuous iff each $t$ has a neighborhood $\eta$ with $F[t'] \subseteq F[t]$ for each $t' \in \eta$.

Lemma 3 (Dekel, Fudenberg, and Morris (2004)). *$S^\infty$ is non-empty and upper-semicontinuous.*

Dekel, Fudenberg, and Morris (2004) prove upper-semicontinuity of interim correlated rationalizability in their framework. Since my framework is slightly different (e.g. $\Theta^*$ may be infinite), for the sake of completeness, I provide a proof in the appendix. Together with the observations in the following lemma, this lemma will provide a main step in the proof of the main result.

Lemma 4. *Given any non-empty, upper-semicontinuous $F$, let $U_F = \{t| \, |F[t]| = 1\}$. Then, $U_F$ is open, and there exists a continuous function $f^* : U_F \to A$ such that $F[t] = \{f^*(t)\}$ for each $t \in U_F$.*

*Proof.* Define $f^* : U_F \to A$ by $F[t] = \{f^*(t)\}$, $t \in U_F$. By upper-semicontinuity of $F$, each $t \in U_F$ has a neighborhood $\eta$ with $F[t'] \subseteq F[t] = \{f^*(t)\}$ for each $t' \in \eta$. Since $F[t'] \neq \varnothing$, this implies that $F[t'] = \{f^*(t)\}$ for each $t' \in \eta$, so that $\eta \subset U_F$. Therefore, $U_F$ is open. By definition, $f^*(t') = f^*(t)$ for each $t' \in \eta$, and hence $f^*$ is continuous. □

Lemmas 3 and 4 imply the following useful fact.

LEMMA 5. *For any $t$ with $S^\infty[t] = \{a\}$, there exists a neighborhood $\eta$ of $t$ such that $S^\infty[t'] = \{a\}$ for each $t' \in \eta$.*

## 4. RESULTS

In this section, I show that, generically, there exists a unique rationalizable action, and for any model, there is a perturbation that leads to a dominance-solvable model. Moreover, for each rationalizable strategy of any finite model, I show that there exists a perturbation that leads to a finite model with common prior and such that the given strategy of the original model is uniquely rationalizable for the perturbed types. The next result will be the main tool for this analysis.

LEMMA 6. *Under Assumption 1, for any $\hat{t} \in \hat{T}$, and any $a \in S^\infty[\hat{t}]$, there exists a sequence of finite, dominance-solvable models $T^m$ with type profiles $\tilde{t}(m) \in T^m$, such that $\tilde{t}(m) \to \hat{t}$ as $m \to \infty$ and $S^\infty[\tilde{t}(m)] = \{a\}$ for each $m$.*

That is, given any type and any rationalizable action $a_i$ for that type, one can find a nearby finite type for which $a_i$ is uniquely rationalizable. Since the proof of this result is somewhat involved, I will present the proof in Section 5, after exploring the important implications of the lemma for this paper.

### 4.1. **Genericity of Uniqueness.** Let

$$U = \{t \in T^* | \, |S^\infty[t]| = 1\}$$

be the set of type profiles with unique rationalizable actions. Together with Lemma 1, Lemma 6 implies that $U$ is dense in universal type space. Since $S^\infty$ is upper-semicontinuous, $U$ is also open. This yields the first main result of the paper: if one excludes a nowhere-dense set of types, there is a unique rationalizable action for each remaining type, which must be continuous in player's belief hierarchy.

Proposition 1. *Generically, there exists a unique rationalizable action, and it is generically continuous. That is, there exist an open, dense set $U$ and a continuous function $s^* : U \to A$, such that $S^\infty[t] = \{s^*(t)\}$ for each $t \in U$. In particular, every rationalizable strategy is continuous on the open and dense set $U$.*

*Proof.* Since $S^\infty[t]$ is upper-semicontinuous, by Lemma 4, $U$ is open, and there exists a continuous function $s^* : U \to A$ with $S^\infty[t] = \{s^*(t)\}$ for each $t \in U$. To show that $U$ is dense, first observe that, by Lemma 6, for any $\hat{t} \in \hat{T}$, there exists a sequence $\tilde{t}(m) \to \hat{t}$ with $S^\infty\left[\tilde{t}(m)\right] = \{a\}$ for some $a \in S^\infty\left[\hat{t}\right]$. By definition, $\tilde{t}(m) \in U$ for each $m$. Hence, $\bar{U} \supseteq \hat{T}$. But $\overline{\hat{T}} = T^*$ by Lemma 1. Therefore, $\bar{U} \supseteq \overline{\hat{T}} = T^*$, showing that $U$ is dense. □

By Proposition 1, we can partition the universal type space to an open and dense set $U$ and its nowhere-dense boundary $T^*\backslash U$. On $U$, each type has a unique rationalizable action, and every rationalizable strategy is continuous. On the boundary, each type profile has multiple rationalizable action profiles. Assumption 1 is not superfluous. For example, a complete-information game can be modeled with $|\Theta^*| = 1$, when $T^*$ consists of a single common-knowledge type profile. When the original game is not dominance-solvable, $U = \varnothing$.

One may wonder if the genericity result above applies to smaller type spaces of interest, such as the space of finite types and space of types consistent with common prior assumption. The next result shows that the same genericity result is true for any dense type space, including the mentioned spaces.

Corollary 1. *For any dense model $T \subseteq T^*$, the set $U \cap T$ is dense and open with respect to the relative topology on $T$. In particular, $U \cap (\hat{T} \cap T^{CPA})$ is dense and open with respect to the relative topology on $\hat{T} \cap T^{CPA}$.*

*Proof.* Since $U$ is open and dense and $T$ is dense, $U \cap T$ is dense. Since $U$ is open, $U \cap T$ is open with respect to the relative topology on $T$—by its definition. □

Remark 3 (Redundant Types). When there are distinct types with identical belief hierarchies in a model, there are different ways to define rationalizability (see e.g. Ely and Peski (2004)), and $S^\infty$ is the largest among them (Dekel, Fudenberg, and Morris (2003)). Proposition 1 establishes that, generically, $|S_i^\infty[t_i]| = 1$, and hence for types with generic belief hierarchies, all these rationalizability concepts are equivalent, and all give the same unique solution.

4.2. **Nearby dominance-solvable models.** Since $U$ is dense, for any usual game with a large set of rationalizable strategy profiles, there is a model such that if a player's interim beliefs and payoffs are similar to that of a player in the original game, then he has a unique rationalizable action. I will now show a stronger fact. Given any model, one can find a nearby dominance-solvable model, where *every* type has a unique rationalizable action.

PROPOSITION 2. *Under Assumption 1, for any model $T \subseteq T^*$, and any integer $m$, there exist a dominance-solvable model $T^m$ and a mapping $\tau(\cdot, m) : T \to T^m$ such that $\tau(t, m) \to t$ as $m \to \infty$.*

*Proof.* First, take any $t \in T^*$. By Lemma 1, there exists a sequence of type profiles $\hat{t}(m) \in \hat{T}$ with $\hat{t}(m) \to t$. By Lemma 6, for all integers $m$ and $k$, there exists a dominance-solvable model $T^{m,k}$ with member $\tilde{t}(m, k)$ such that $\tilde{t}(m, k) \to \hat{t}(m)$ as $k \to \infty$. Define $T^{t,m} \equiv T^{m,m}$ and $\tau(t, m) \equiv \tilde{t}(m, m)$. Clearly, $\tau(t, m) \to t$. Now, define $T^m$ by

$$T_i^m = \bigcup_{t \in T} T_i^{t,m}.$$

Since each $T^{t,m}$ is dominance-solvable, so is $T^m$. For each $t \in T$, $\tau(t, m) \in T^m$. □

Proposition 2 extends the result of Carlsson and van Damme to arbitrary games. It states that, given any model, we can perturb the model by introducing a small noise in players' perceptions of the payoffs in such a way that the new model is dominance-solvable. Moreover, since $U$ is open, the perturbed model will remain dominance-solvable when we introduce new small perturbations.

I will now turn to question of whether one can enlarge a model continuously so that the multiplicity occurs only on a nowhere-dense set of states. By Proposition 1, one can always embed any model $\Theta \times T$ in a larger model, namely $\Theta^* \times T^*$, in which multiplicity occur only on a nowhere-dense set. Since the latter model is large, the nowhere-dense set may be large, too. Using Proposition 2, I next show that we can pick the enlarged model so that the new types all have unique rationalizable actions. The enlarged model will be "small" whenever $\Theta \times T$ is "small".

COROLLARY 2. *For any model $\Theta \times T$, there exists a model $\Theta' \times T'$ such that multiplicities can occur only in a nowhere-dense subset $\Theta' \times T''$ with respect to the relative topology on $\Theta' \times T'$ where $T'' \subseteq T$. Moreover, whenever $\Theta \times T$ is countable, so is $\Theta' \times T'$.*

*Proof.* For each $t \in T$ and integer $m$, there is a finite, dominance-solvable model $\Theta^{t,m} \times T^{t,m}$ as in the proof of Proposition 2. Define $\Theta' \times T'$ by $\Theta' = (\cup_{t,m}\Theta^{t,m}) \cup \Theta$ and $T'_i = \left(\cup_{t,m}T_i^{t,m}\right) \cup T_i$. Let $T''$ be the set of all type profiles $t \in T'$ with multiple rationalizable actions. By construction, $T'' \subseteq T$. As in Proposition 2, for each $(\theta, t) \in \Theta' \times T''$, there exists $\tau(t) \in T' \backslash T''$, converging to $t$, so that $(\theta, \tau(t)) \to (\theta, t)$. Hence, $\Theta' \times T''$ has empty interior with respect to the topology on $\Theta' \times T'$. Moreover, $T' \backslash T'' = U \cap T'$ is open (and hence $T''$ is closed) w.r.t. the relative topology. Hence, $\Theta' \times T''$ is closed in this topology. Therefore, $\Theta' \times T''$ is nowhere-dense. Since each $\Theta^{t,m} \times T^{t,m}$ is finite, whenever $\Theta \times T$ is countable, $\Theta' \times T'$ is also countable, for it is then a countable union of finite sets. □

I will now turn to the main question of why we have multiple rationalizable actions at the first place. Focusing on finite models, the upcoming proposition, which is the second main result of this paper, provides an answer to this question. It states that, for each finite model and for each rationalizable strategy profile $s_T$ in this model, we can perturb the beliefs and find a new model such that $s_T$ is the unique rationalizable strategy profile for the perturbed types. We can pick the new model with a common prior. In this proposition, there are two perturbations. The first perturbation leads to a finite, dominance-solvable model $T^{s_T,m}$ where the unique rationalizable actions of perturbed types $\tau(\cdot, s_T, m)$ agree with $s_T$. The second perturbation leads to a finite model $\tilde{T}^{s_T,m}$ that admits a common prior, which may not be dominance-solvable, but the perturbed types $\tilde{\tau}(t, s_T, m)$ have all unique rationalizable actions, and these actions agree with $s_T$.

Proposition 3. *Let $T \subseteq \hat{T}$ be any finite model and $s_T : T \to A$ be any rationalizable strategy profile, with $s_T(t) \in S^\infty[t]$ for each $t \in T$. Then, under Assumption 1, there exist sequences of finite models $T^{s_T,m}$ and $\tilde{T}^{s_T,m}$ and one-to-one mappings $\tau(\cdot, s_T, m) : T \to T^{s_T,m}$ and $\tilde{\tau}(\cdot, s_T, m) : T \to \tilde{T}^{s_T,m}$ such that*

(1) *$T^{s_T,m}$ is dominance-solvable, and $\tilde{T}^{s_T,m}$ admits a common prior,*
(2) *$S^\infty[\tau(t, s_T, m)] = S^\infty[\tilde{\tau}(t, s_T, m)] = \{s_T(t)\}$, and*
(3) *$\tau(t, s_T, m) \to t$ and $\tilde{\tau}(t, s_T, m) \to t$ as $m \to \infty$ for each $t \in T$.*

*Proof.* By Lemma 6, for each $t \in T$ and $m$, there exists a finite, dominance-solvable model $T^{t,s_T,m}$ with $\tau(t, s_T, m) \in T^{t,s_T,m}$ as in the proposition. As in the proof of Proposition 2, define the finite model $T^{s_T,m}$ by

$$T_i^{s_T,m} = \bigcup_{t \in T} T_i^{t,s_T,m}.$$

Since $\tau(t, s_T, m) \to t$ for each $t \in T$ and $T$ is finite, there exists $\bar{m}$ such that, for any distinct $t, t'$ and any $m > \bar{m}$, we have $\tau(t, s_T, m) \neq \tau(t', s_T, m)$. Hence, $\tau(\cdot, s_T, m)$ is one-to-one for $m > \bar{m}$. (Consider only $m > \bar{m}$.)

I will now construct a finite model $\tilde{T}^{s_T,m}$ that admits a common prior with full support and has the desired properties. For transparency, I will use elementary techniques. Since $\hat{T} \cap T^{CPA}$ is dense, for each $\tau(t, s_T, m)$, there exists a sequence of finite models $T^{t,m,k}$ with common priors $p^{t,m,k}$ (with full support) and members $\bar{\tau}(t, m, k)$ such that $\bar{\tau}(t, m, k) \to \tau(t, s_T, m)$ as $k \to \infty$. By Lemma 5, there exists a $\bar{k}$ such that for each $k > \bar{k}$, $S^\infty[\bar{\tau}(t, m, k)] = S^\infty[\tau(t, m, k)] = \{s_T(t)\}$. Hence, without loss of generality, pick each $\bar{\tau}(t, m, k)$ with

$$S^\infty[\bar{\tau}(t, m, k)] = \{s_T(t)\}. \tag{4.1}$$

For each $\varepsilon \in [0, 1]$, I will now construct a usual finite type space $T^{m,k,\varepsilon}$ in which the types are denoted by integers. For each $i$, let $\hat{\tau}_i$ be any one-to-one mapping that maps types $\tilde{t}_i$, $\tilde{t}_i \in T_i^{t,m,k}$, $t \in T$, to integers. (Recall that there are only finitely many such types.) Define

$$T_i^{m,k,\varepsilon} = \left\{\hat{\tau}_i\left(\tilde{t}_i\right) | \tilde{t}_i \in T_i^{t,m,k}, t \in T\right\} \qquad (i \in N).$$

Let $\Theta$ be the set of all $\theta \in \Theta^*$ on which some type $\tilde{t}_i \in T_i^{t,m,k}$ puts positive probability. I will now define a common prior $p^{m,k,\varepsilon}$ on $\Theta \times T^{m,k,\varepsilon}$ with full support. Since each $p^{t,m,k}$ has full support, given any $t$ and $t'$, either $T^{t,m,k} = T^{t',m,k}$ or $T^{t,m,k} \cap T^{t',m,k} = \varnothing$. Let $K$ be the number of disjoint sets $T^{t,m,k}$ and $L = \left|\Theta \times T^{m,k,\varepsilon}\right|$. Define $p^{m,k,\varepsilon}$ by setting

$$p^{m,k,\varepsilon}\left(\theta, \bar{t}\right) = \begin{cases} \varepsilon/L + (1-\varepsilon)\, p^{t,m,k}\left(\theta, \tilde{t}\right)/K & \text{if } \bar{t} = \left(\hat{\tau}_1\left(\tilde{t}_1\right), \ldots, \hat{\tau}_n\left(\tilde{t}_n\right)\right) \text{ for some } \tilde{t} \in T^{t,m,k} \text{ and } t, \\ \varepsilon/L & \text{otherwise} \end{cases}$$

at each $(\theta, \bar{t}) \in \Theta \times T^{m,k,\varepsilon}$. According to $p^{m,k,\varepsilon}$, with probability $\varepsilon$, we have a uniform distribution on $\Theta \times T^{m,k,\varepsilon}$, and with probability $(1-\varepsilon)$ one of the type spaces $T^{t,m,k}$ is selected, each with equal probability. Let $h_i\left(\hat{\tau}_i\left(\tilde{t}_i\right); m, k, \varepsilon\right)$ be the belief hierarchy of $\hat{\tau}_i\left(\tilde{t}_i\right)$ under $p^{m,k,\varepsilon}$. By a well-known result of Mertens and Zamir, $h_i(\cdot; m, k, \varepsilon)$ is continuous, and as $\varepsilon \to 0$, $h_i\left(\hat{\tau}_i\left(\tilde{t}_i\right); m, k, \varepsilon\right) \to h_i\left(\hat{\tau}_i\left(\tilde{t}_i\right); m, k, 0\right)$ in the product topology for each $\hat{\tau}_i\left(\tilde{t}_i\right)$. Moreover, $h_i\left(\hat{\tau}_i\left(\tilde{t}_i\right); m, k, 0\right) = \tilde{t}_i$ by construction. (A type's belief hierarchy cannot change when we add an exante stage to choose between type spaces.) Therefore, $h_i\left(\hat{\tau}_i\left(\tilde{t}_i\right); m, k, \varepsilon\right) \to \tilde{t}_i$ for each $\tilde{t}_i$. Since types $\tilde{t}_i$ are all distinct hierarchies (by the disjointness above and no redundancy in $T^{t,m,k}$), this implies that, for some $\bar{\varepsilon} > 0$, there will be no redundant types under $p^{m,k,\varepsilon}$ whenever $\varepsilon < \bar{\varepsilon}$. It also implies by Lemma 5 and (4.1) that, for each $(t, m, k)$, there exists

$\varepsilon^{t,m,k} < \bar{\varepsilon}$ such that $S^\infty [h(\hat{\tau}(\bar{\tau}(t,m,k)), m, k, \varepsilon)] = \{s_T(t)\}$ whenever $\varepsilon < \varepsilon^{t,m,k}$. Fix $\varepsilon < \min_{t \in T} \varepsilon^{t,m,k}$ and set $\tilde{T}^{s_T,m} = h\left(T^{m,m,\varepsilon/m}; m, m, \varepsilon/m\right)$. For each $t$ and $m$, set also $\tilde{\tau}(t, s_T, m) = h(\hat{\tau}(\bar{\tau}(t,m,m)), m, m, \varepsilon/m)$. By construction, $T^{m,m,\varepsilon/m}$ under $p^{m,m,\varepsilon/m}$ does not have any redundant types; $S^\infty[\tilde{\tau}(t, s_T, m)] = \{s_T(t)\}$ for each $(t,m)$, and $\tilde{\tau}(t, s_T, m) \to t$ as $m \to \infty$. □

Building on a result of Weinstein and Yildiz (2004), Proposition 3 uncovers a striking structure of rationalizability on the set $\hat{T}$ of finite types. This structure remains intact when one imposes the common-prior assumption (i.e. on $\hat{T} \cap T^{CPA}$). One can divide $\hat{T}$ into finitely many open sets

$$U^a = \left\{\hat{t} \in \hat{T} | S^\infty\left[\hat{t}\right] = \{a\}\right\} \qquad (a \in A),$$

and their boundaries $\partial U^a \equiv \overline{U^a} \backslash U^a$, where $\overline{U^a}$ is the closure of $U^a$, all with respect to the relative topology on $\hat{T}$. The open sets form a partition of an open, dense set $U \cap \hat{T}$, while their boundaries cover the boundary of $U \cap \hat{T}$, i.e., $\hat{T} \backslash U = \bigcup_{a \in A} \partial U^a$, which is a nowhere-dense set with respect to the relative topology. On each open set $U^a$, $a$ is the unique rationalizable action profile. Since $S^\infty$ is upper-semicontinuous, $a \in S^\infty\left[\hat{t}\right]$ for each $\hat{t} \in \partial U^a$. Hence, at any $\hat{t} \in \partial U^a \cap \partial U^{a'}$ with distinct $a$ and $a'$, both $a$ and $a'$ are rationalizable. Here, there are multiple rationalizable actions $a$ and $a'$ because $\hat{t}$ can be thought of idealization of two strategically distinct relaxed assumptions, under which $a$ and $a'$ are unique solutions respectively, and the set of rationalizable actions reflects this fact. Proposition 3 shows that the converse is also true:

$$\hat{t} \in \bigcap_{a \in S^\infty\left[\hat{t}\right]} \overline{U^a} \qquad \left(\forall \hat{t} \in \hat{T}\right).$$

That is, the set $S^\infty\left[\hat{t}\right]$ tells us which actions could be uniquely rationalizable when we slightly relax the assumptions of $\hat{t}$ using various information structures. When $\hat{t}$ has multiple rationalizable actions, it cannot be in the interior of any of these sets, and hence

$$\hat{t} \in \bigcap_{a \in S^\infty\left[\hat{t}\right]} \partial U^a.$$

That is, whenever there are multiple rationalizable actions at $\hat{t}$, $\hat{t}$ embodies an idealization of some relaxed assumptions with distinct strategic implications. Each rationalizable action $a$ represent such a relaxed assumption, which leads to $a$ as the unique solution. Therefore, there are multiple rationalizable actions at $\hat{t}$ *if and only if* $\hat{t}$ embodies an idealization of two strategically distinct situations. As a solution concept, rationalizability leads to a generically unique and locally constant theory that yields multiple solutions when, and only when, the theory changes its prescribed behavior for players.

Proposition 3 also provides a new perspective on refining rationalizability. It implies that a finite model summarizes various dominance-solvable situations by abstracting away from the details that would have mattered mostly for computing the beliefs at very high orders. By specifying these details appropriately, any rationalizable strategy could have been made uniquely rationalizable. But then, refining rationalizability tantamount to ruling out some of these nearby models as the true model. Hence, selection of a refinement is tied to which information structures one finds more reasonable—more so than which epistemic arguments make more sense on what beliefs players should form on other players' strategies. Weinstein and Yildiz (2004) have proved a similar result by considering only equilibrium refinements and "strictly rationalizable" actions.

Upper-semicontinuity is a desirable property of solution concepts, as it allows one to keep track of idealizations one makes. It is customary to check whether a solution concept is upper-semicontinuous. Proposition 3 immediately yields a very useful fact about upper-semicontinuity of general solution concepts with respect to the types. It shows that any solution concept that is stronger than the correlated rationalizability on the set of finite models necessarily fails upper-semicontinuity. Therefore, almost none of the known solution concepts is upper-semicontinuous with respect to the types. In other words, correlated rationalizability is a very strong solution concept, in that it does not have any proper refinement that is upper-semicontinuous.

Corollary 3. *Let $R : \hat{T} \cap T^{CPA} \to 2^A$ be any non-empty, upper-semicontinuous correspondence with $R(t) \subseteq S^\infty[t]$ for each $t \in \hat{T} \cap T^{CPA}$. Then, $R = S^\infty$.*

*Proof.* Suppose that $R \neq S^\infty$. Then, there exists $\hat{t}$ with $a \in S^\infty[\hat{t}]$ such that $a \notin R(\hat{t})$. By Proposition 3, there exists a sequence $t(m) \in \hat{T} \cap T^{CPA}$ such that $t(m) \to \hat{t}$ and for each $m$, $S^\infty[t(m)] = \{a\}$, i.e., $a \in R(t(m))$. Then, by upper-semicontinuity of $R$, $a \in R(\hat{t})$—a contradiction. □

## 5. Proof of Lemma 6

Now, I will prove Lemma 6. A substantial part of the proof utilizes the following stronger notion of rationalizability, analyzed by Weinstein and Yildiz (2004).

**Strict Interim Rationalizability**. Let $W_i^0[t_i] = A_i$ and, for each $k > 0$, let $a_i \in W_i^k[t_i]$ if and only if $BR_i\left(\mathrm{marg}_{\Theta^* \times A_{-i}}\pi\right) = \{a_i\}$ for some $\pi \in \Delta\left(\Theta^* \times T_{-i}^* \times A_{-i}\right)$ such that $\mathrm{marg}_{\Theta^* \times T_{-i}^*}\pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in S_{-i}^{k-1}[t_{-i}]\right) = 1$. Finally, let

$$W_i^\infty[t_i] = \bigcap_{k=0}^{\infty} W_i^k[t_i]$$

be the set of all *strictly rationalizable* actions for $t_i$. Notice that an action is eliminated if it is not a strict best-response to any belief on the remaining strategies of the other players. Clearly, $W_i^k \subseteq S_i^k$, and $W_i^k[t_i]$ may be empty.

LEMMA 7. *Given any belief-closed $T$, consider any family $V_i[t_i] \subseteq A_i$, $t_i \in T_i$, $i \in N$, such that each $a_i \in V_i[t_i]$ is a strict best reply to a belief $\pi \in \Delta(\Theta^* \times T_{-i} \times A_{-i})$ of $t_i$ with $\pi(a_{-i} \in V_{-i}[t_{-i}]) = 1$. Then, $V_i[t_i] \subseteq W_i^\infty[t_i]$ for each $t_i$.*

*Proof.* It follows from the fact that no $a_i \in V_i[t_i]$ is ever eliminated for $t_i$. □

The proof of Lemma 6 has two main steps, which are presented as the following two lemmas. The first step (namely, Lemma 8) shows that, when we focus on strictly rationalizable strategies, Lemma 6 is true for each $t_i \in \hat{T}_i$. The second step (namely, Lemma 9) will state that for any finite type and any rationalizable action, there is a nearby finite type for which the action is strictly rationalizable. Combining these two steps immediately yields Lemma 6.

The following lemma is similar to Proposition 1 of Weinstein and Yildiz (2004). They show that if $a_i \in W_i^k[t_i]$, one can change the beliefs at order $k+1$ and higher so that $a_i$ is played by the new type in equilibrium. The lemma states that one can select the new type $\tilde{t}_i$ so that $a_i$ is the only member of $S_i^{k+1}\left[\tilde{t}_i\right]$. To prove this lemma, I use their construction but make sure that the new type $\tilde{t}_i$ assigns positive probability only on types $t_{-i}$ that come from finite models that are solved by $k$ rounds of iterated dominance (i.e., $S^k$ is singleton-valued on these models). In that case, I show that $\tilde{t}_i$ also comes from a finite model that is solved by $k+1$ rounds of iterated dominance.

LEMMA 8. *Under Assumption 1, for each $i, k$, for each $\hat{t}_i \in \hat{T}_i$, and for each $a_i \in W_i^k[t_i]$, there exists $\tilde{t}_i$ such that (i) $\tilde{t}_i^l = \hat{t}_i^l$ for each $l \leq k$, (ii)*

$$S_i^{k+1}\left[\tilde{t}_i\right] = \{a_i\},$$

*and* $\tilde{t}_i \in T_i^{\tilde{t}_i}$ *for some finite model* $T^{\tilde{t}_i} = T_1^{\tilde{t}_i} \times \cdots \times T_1^{\tilde{t}_i}$ *such that* $\left|S^{k+1}[t]\right| = 1$ *for each* $t \in T^{\tilde{t}_i}$. *For any* $a_i \in W_i^\infty\left[\hat{t}_i\right]$ *and integer* $m$*, there exists a finite, dominance-solvable model* $T^m$ *with type* $t_{i,m} \in T_i^m$*, such that* $S_i^\infty[t_{i,m}] = \{a_i\}$ *and* $t_{i,m} \to \hat{t}_i$ *as* $m \to \infty$*.*

*Proof.* For $k = 0$, let $\tilde{t}$ be the type profile according to which it is common knowledge that each $j$ assigns probability 1 to $\{\theta = \theta^{a_j}\}$, where $\theta^{a_j}$ is as defined in Assumption 1. By Assumption 1, $S_i^1\left[\tilde{t}_i\right] = \{a_i\}$, and it is vacuously true that $\tilde{t}_i^l = \hat{t}_i^l$ for each $l \le k$. Clearly, the type space $\left\{\tilde{t}\right\}$ is belief-closed.

Now fix any $k > 0$ and any $i$. Write each $t_{-i}$ as $t_{-i} = (l, h)$ where $l = \left(t_{-i}^1, t_{-i}^2, \ldots, t_{-i}^{k-1}\right)$ and $h = \left(t_{-i}^k, t_{-i}^{k+1}, \ldots\right)$ are the lower and higher-order beliefs, respectively. Let $L = \left\{l \mid \exists h : (l, h) \in T_{-i}^*\right\}$. The inductive hypothesis is that for each finite $t_{-i} = (l, h)$ and each $a_{-i} \in W_{-i}^{k-1}[t_{-i}]$, there exists finite $\tilde{t}_{-i}[a_{-i}] = \left(l, \tilde{h}[l, a_{-i}]\right) \in T_{-i}^{\tilde{t}_{-i}[a_{-i}]}$ such that

$$S_{-i}^k\left[\tilde{t}_{-i}[a_{-i}]\right] = \{a_{-i}\}, \tag{IH}$$

and $T^{\tilde{t}_{-i}[a_{-i}]} = T_1^{\tilde{t}_{-i}[a_{-i}]} \times \cdots \times T_n^{\tilde{t}_{-i}[a_{-i}]}$ is a finite model with $\left|S^k[t]\right| = 1$ for each $t \in T^{\tilde{t}_{-i}[a_{-i}]}$. Take any $a_i \in W_i^k\left[\hat{t}_i\right]$. I will construct a type $\tilde{t}_i$ as in the lemma. By definition, $BR_i\left(\text{marg}_{\Theta^* \times A_{-i}}\pi\right) = \{a_i\}$ for some $\pi \in \Delta\left(\Theta^* \times T_{-i}^* \times A_{-i}\right)$ such that $\text{marg}_{\Theta^* \times T_{-i}^*}\pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in W_{-i}^{k-1}[t_{-i}]\right) = 1$. Using the inductive hypothesis, define mapping $\mu : \text{supp}\left(\text{marg}_{\Theta^* \times L \times A_{-i}}\pi\right) \to \Theta^* \times T_{-i}^*$, by

$$\mu : (\theta, l, a_{-i}) \mapsto \left(\theta, l, \tilde{h}[l, a_{-i}]\right), \tag{5.1}$$

where type $\tilde{t}_{-i}[a_{-i}] = \left(l, \tilde{h}[l, a_{-i}]\right)$ is as in (IH). Define $\tilde{t}_i$ by

$$\kappa_{\tilde{t}_i} \equiv \left(\text{marg}_{\Theta^* \times L \times A_{-i}}\pi\right) \circ \mu^{-1} = \pi \circ \text{proj}_{\Theta^* \times L \times A_{-i}}^{-1} \circ \mu^{-1},$$

where $\text{proj}_X$ denotes the projection mapping to $X$. Notice that $\text{proj}_{\Theta^* \times L} \circ \mu \circ \text{proj}_{\Theta^* \times L \times A_{-i}} = \text{proj}_{\Theta^* \times L}$. Then, $\text{marg}_{\Theta^* \times L}\kappa_{\tilde{t}_i} = \text{marg}_{\Theta^* \times L}\kappa_{t_i}$, and hence the first $k$ orders beliefs will be identical under $t_i$ and $\tilde{t}_i$ (see Weinstein and Yildiz (2004) for a detailed derivation). Moreover, by (IH), each $(\theta, t_{-i}) \in \text{supp}\left(\kappa_{\tilde{t}_i}\right)$, which is of the form $\left(\theta, l, \tilde{h}[l, a_{-i}]\right)$, has a unique action $a_{-i} \in S_{-i}^{k-1}\left[\tilde{t}_{-i}[a_{-i}]\right]$. Thus, there exists a unique $\tilde{\pi} \in \Delta\left(\Theta^* \times T_{-i}^* \times A_{-i}\right)$ such that $\text{marg}_{\Theta^* \times T_{-i}^*}\pi = \kappa_{\tilde{t}_i}$ and $\pi\left(a_{-i} \in S_{-i}^{k-1}[t_{-i}]\right) = 1$; it is $\tilde{\pi} = \kappa_{\tilde{t}_i} \circ \gamma^{-1} = \pi \circ \text{proj}_{\Theta^* \times L \times A_{-i}}^{-1} \circ \mu^{-1} \circ \gamma^{-1}$ where $\gamma : \left(\theta, l, \tilde{h}[l, a_{-i}]\right) \mapsto \left(\theta, l, \tilde{h}[l, a_{-i}], a_{-i}\right)$. Clearly, $\text{proj}_{\Theta^* \times A_{-i}} \circ \gamma \circ \mu \circ \text{proj}_{\Theta^* \times L \times A_{-i}} = \text{proj}_{\Theta^* \times A_{-i}}$. Hence, $\text{marg}_{\Theta^* \times A_{-i}}\tilde{\pi} = \text{marg}_{\Theta^* \times A_{-i}}\pi$. But $a_i$ is the only best reply to this belief. Therefore, $S_i^{k+1}\left[\tilde{t}_i\right] = \{a_i\}$.

Now, I will define $T^{\tilde{t}_i}$ as in the lemma. Define

$$\begin{aligned} T_i^{\tilde{t}_i} &= \left\{\tilde{t}_i\right\} \cup \left( \bigcup_{(\theta, t_{-i}[a_{-i}]) \in \text{ supp}\left(\kappa_{\tilde{t}_i}\right)} T_i^{t_{-i}[a_{-i}]} \right), \\ T_j^{\tilde{t}_i} &= \bigcup_{(\theta, t_{-i}[a_{-i}]) \in \text{ supp}\left(\kappa_{\tilde{t}_i}\right)} T_j^{t_{-i}[a_{-i}]} \qquad (j \neq i). \end{aligned}$$

Since $\text{supp}\left(\text{marg}_{\Theta^* \times L \times A_{-i}} \pi\right) \subseteq \text{supp}\left(\kappa_{\hat{t}_i}\right) \times A_{-i}$ is finite, the range of $\mu$ is finite, rendering $\text{supp}\left(\kappa_{\tilde{t}_i}\right)$ finite. Hence, $T^{\tilde{t}_i}$ is finite. For any $t_j \in T_j^{\tilde{t}_i} \setminus \left\{\tilde{t}_i\right\}$, $t_j \in T_j^{t_{-i}[a_{-i}]}$ for some $t_{-i}[a_{-i}]$, and since $T^{t_{-i}[a_{-i}]}$ is belief-closed, $\text{supp}\left(\kappa_{t_j}\right) \subseteq \Theta^* \times T_{-j}^{t_{-i}[a_{-i}]} \subseteq \Theta^* \times T_{-j}^{\tilde{t}_i}$. On the other hand, $\text{supp}\left(\kappa_{\tilde{t}_i}\right) \subseteq \Theta^* \times T_{-i}^{\tilde{t}_i}$, as $t_{-i}[a_{-i}] \in T_{-i}^{t_{-i}[a_{-i}]}$ for each $(\theta, t_{-i}[a_{-i}]) \in \text{supp}\left(\kappa_{\tilde{t}_i}\right)$. Hence, $T^{\tilde{t}_i}$ is belief-closed. Finally, since $S_i^{k+1}\left[\tilde{t}_i\right] = \{a_i\}$, $\left|S_i^{k+1}\left[\tilde{t}_i\right]\right| = 1$, and by construction, for each $t_j \in T_j^{\tilde{t}_i} \setminus \left\{\tilde{t}_i\right\}$, $\left|S^{k+1}[t_j]\right| = \left|S^k[t_j]\right| = 1$.

To prove the last statement in the lemma, take any $a_i \in W_i^\infty\left[\hat{t}_i\right]$. For each $m$, since $a_i \in W_i^\infty\left[\hat{t}_i\right] \subseteq W_i^m\left[\hat{t}_i\right]$, by the first part of the lemma, there exists $t_{i,m}$ such that $t_{i,m}^l = \hat{t}_i^l$ for each $l \leq m$ and $S_i^{m+1}[t_{i,m}] = S_i^\infty[t_{i,m}] = \{a_i\}$. Clearly, for any fixed $k$, $t_{i,m}^k = \hat{t}_i^k$ for each $m > k$, showing that $t_{i,m}^k \to \hat{t}_i^k$ as $m \to \infty$. By the first part, $t_{i,m} \in T_i^{t_{i,m}}$ for some finite model $T^{t_{i,m}}$ with $|S^\infty[t]| = |S^{m+1}[t]| = 1$ for each $t \in T^{t_{i,m}}$. Pick $T^m = T^{t_{i,m}}$ as the dominance-solvable model in the lemma. □

The next lemma states that any rationalizable strategy of a finite model is strictly rationalizable for nearby types in a nearby finite model.

Lemma 9. *Under Assumption 1, for any finite model $T \subseteq \hat{T}$ and any integer $m$, there exist a finite model $T^m$ and a one-to-one and onto mapping $\tau(\cdot, m)$ that maps each $(t, a)$ with $a \in S^\infty[t]$ and $t \in T$ to $\tau(t, a, m) = (\tau_1(t_1, a_1, m), \ldots, \tau_n(t_n, a_n, m)) \in T^m$ such that (i) $a \in W^\infty[\tau(t, a, m)]$ for each $(t, a, m)$, and (ii) $\tau(t, a, m) \to t$ as $m \to \infty$ for each $(t, a)$.*

*Proof.* The new type space $T^m$ will consist of types $\tau_i(t_i, a_i, m)$, for $i \in N$, $t_i \in T_i$, and $a_i \in S_i^\infty[t_i]$. Let $\delta_x$ denote the probability distribution that puts probability 1 on $\{x\}$ and $\Theta'$ be the finite set of all parameter values that some type $t_j \in T_j$ assigns positive probability. I will define $\tau(\cdot, m)$ by simultaneously defining the beliefs of each $\tau_i(t_i, a_i, m)$ about $\theta$ and the others' types $\tau_{-i}(t_{-i}, a_{-i}, m)$.[6] Now, since $a_i \in$

[6]Notice that I am simply defining a finite type space. Hence, it suffices to define the belief of each type about $\theta$ and the other players' types. At the end of the proof, I will show that there are

$S_i^\infty[t_i]$, there exists a belief $\pi^{t_i,a_i} \in \Delta(\Theta' \times T_{-i} \times A_{-i})$ with finite support and such that $a_i \in BR_i\left(\text{marg}_{\Theta' \times A_{-i}} \pi^{t_i,a_i}\right)$, $\pi^{t_i,a_i}\left(a_{-i} \in S_{-i}^\infty[t_{-i}]\right) = 1$, and $\text{marg}_{\Theta^* \times T^*_{-i}} \pi^{t_i,a_i} = \kappa_{t_i}$, where we also view $\pi^{t_i,a_i}$ as a probability distribution on $\Theta^* \times T^*_{-i} \times A_{-i}$. Define $\tau_i(t_i, a_i, m)$ by

$$\kappa_{\tau_i(t_i,a_i,m)} = \frac{1}{m}\delta_{\left(\theta^{a_i},\tau_{-i}\left(\tilde{t}_{-i},\tilde{a}_{-i},m\right)\right)} + \left(1 - \frac{1}{m}\right)\pi^{t_i,a_i} \circ \hat{\tau}_{-i,m}^{-1}$$

where $\tau_{-i}\left(\tilde{t}_{-i}, \tilde{a}_{-i}, m\right)$ is some fixed type profile in the new type space, and $\hat{\tau}_{-i,m} : (\theta, t_{-i}, a_{-i}) \mapsto (\theta, \tau_{-i}(t_{-i}, a_{-i}, m))$. The beliefs of $\tau_i(t_i, a_i, m)$ correspond to a mixture: with probability $1 - 1/m$, each $(\theta, \tau_{-i}(t_{-i}, a_{-i}, m))$ occurs with the probability of $(\theta, t_{-i}, a_{-i})$ according to $\pi^{t_i,a_i}$, and with probability $1/m$ there is a point mass at $\left(\theta^{a_i}, \tau_{-i}\left(\tilde{t}_{-i}, \tilde{a}_{-i}, m\right)\right)$. For each new type $\tau_i(t_i, a_i, m)$, define the belief

$$\tilde{\pi} = \kappa_{\tau_i(t_i,a_i,m)} \circ \gamma^{-1} \in \Delta\left(\Theta^* \times T^*_{-i} \times A_{-i}\right)$$

where $\gamma : (\theta, \tau_{-i}(t_{-i}, a_{-i}, m)) \mapsto (\theta, \tau_{-i}(t_{-i}, a_{-i}, m), a_{-i})$. This belief is generated by $\kappa_{\tau_i(t_i,a_i,m)}$ and the pure strategy profile $s_{-i}$ with $s_{-i}(\tau_{-i}(t_{-i}, a_{-i}, m)) = a_{-i}$ at each $(\theta, \tau_{-i}(t_{-i}, a_{-i}, m))$. Clearly, $\text{proj}_{\Theta^* \times A_{-i}} \circ \gamma \circ \hat{\tau}_{-i,m} = \text{proj}_{\Theta^* \times A_{-i}}$. Hence,

$$\text{marg}_{\Theta^* \times A_{-i}} \tilde{\pi} = \frac{1}{m}\delta_{(\theta^{a_i},\tilde{a}_{-i})} + \left(1 - \frac{1}{m}\right)\text{marg}_{\Theta^* \times A_{-i}} \pi^{t_i,a_i}.$$

That is, the belief of $\tau_i(t_i, a_i, m)$ about $\Theta^* \times A_{-i}$ is also a mixture. With probability $(1 - 1/m)$, $\tau_i(t_i, a_i, m)$ faces the same uncertainty as $t_i$ does when $t_i$ holds the belief $\pi^{t_i,a_i}$, in which case $a_i$ is a best reply. With probability $1/m$, the equality $\theta = \theta^{a_i}$ holds, in which case $a_i$ is the unique best reply. Then, by the Sure-thing Principle, $a_i$ is a strict best reply, i.e., $BR_i\left(\text{marg}_{\Theta^* \times A_{-i}} \tilde{\pi}\right) = \{a_i\}$. Hence, by Lemma 7, $a_i \in W_i^\infty[\tau_i(t_i, a_i, m)]$ for each $\tau_i(t_i, a_i, m)$.

I will use induction to show that $\tau_i(t_i, a_i, m) \to t_i$, i.e., each $k$th order belief $\tau_i^k(t_i, a_i, m)$ converges to $t_i^k$, as $m \to \infty$. Firstly, the first-order belief is

$$\tau_i^1(t_i, a_i, m) = \text{marg}_{\Theta^*} \kappa_{\tau_i(t_i,a_i,m)} = \frac{1}{m}\delta_{\theta^{a_i}} + \left(1 - \frac{1}{m}\right)\text{marg}_{\Theta^*} \pi^{t_i,a_i},$$

which converges to

$$\text{marg}_{\Theta^*} \pi^{t_i,a_i} = \text{marg}_{\Theta^*} \kappa_{t_i} = t_i^1$$

no redundant types in the constructed type space, so that it can be represented as a subspace of the universal type space.

as $m \to \infty$. Now, fix some $k > 0$. Let $L$ be the set of all beliefs $t_{-i}^{k-1}$ at order $k-1$, and assume that $\tau_j^{k-1}(t_j, a_j, m) \to t_j^{k-1}$ for each $(t_j, a_j) \in T_j \times A_j$. Then,

$$\tau_i^k(t_i, a_i, m) = \frac{1}{m}\delta_{\left(\theta^{a_i}, \tau_i^{k-1}(t_i,a_i,m), \tau_{-i}^{k-1}\left(\tilde{t}_{-i}, \tilde{a}_{-i}, m\right)\right)} + \left(1 - \frac{1}{m}\right)\delta_{\tau_i^{k-1}(t_i,a_i,m)} \times \text{marg}_{\Theta^* \times L}\pi^{t_i,a_i} \circ \hat{\tau}_{-i,m}^{-1}.$$

As $m \to \infty$, the right-hand side converges to

$$\begin{aligned} \lim_{m\to\infty} \delta_{\tau_i^{k-1}(t_i,a_i,m)} \times \text{marg}_{\Theta^* \times L}\pi^{t_i,a_i} \circ \hat{\tau}_{-i,m}^{-1} &= \lim_{m\to\infty} \delta_{\tau_i^{k-1}(t_i,a_i,m)} \times \pi^{t_i,a_i} \circ \hat{\tau}_{-i,m}^{-1} \circ \text{proj}_{\Theta^* \times L}^{-1} \\ &= \delta_{t_i^{k-1}} \times \text{marg}_{\Theta^* \times L}\pi^{t_i,a_i} = t_i^k. \end{aligned}$$

[To obtain the penultimate equality, observe that $\text{proj}_{\Theta^* \times L}(\hat{\tau}_{-i,m}(\theta, t_{-i}, a_{-i})) = \text{proj}_{\Theta^* \times L}(\theta, \tau_{-i}(t_{-i}, a_{-i}, m)) = \left(\theta, \tau_{-i}^{k-1}(t_{-i}, a_{-i}, m)\right)$, which converges to $\left(\theta, t_{-i}^{k-1}\right)$. That is, $\text{proj}_{\Theta^* \times L} \circ \hat{\tau}_{-i,m}$ point-wise converges to $\text{proj}_{\Theta^* \times L}$. Then, $\pi^{t_i,a_i} \circ \hat{\tau}_{-i,m}^{-1} \circ \text{proj}_{\Theta^* \times L}^{-1}$ converges to $\pi^{t_i,a_i} \circ \text{proj}_{\Theta^* \times L}^{-1} = \text{marg}_{\Theta^* \times L}\pi^{t_i,a_i}$ in weak topology.]

Finally, one can choose $m$ large enough so that $\tau(\cdot, m)$ is one-to-one, in which case $T^m$ does not have redundant types, as I will show now. For any two distinct $a_i$ and $a_i'$, by definition, $\theta^{a_i} \neq \theta^{a_i'}$, rendering $\tau_i(t_i, a_i, m) \neq \tau_i(t_i, a_i', m)$ for each $t_i$ and $m$. On the other hand, for any distinct $t_i$ and $t_i'$, since $\tau_i(t_i, a_i, m) \to t_i$ and $\tau_i(t_i', a_i', m) \to t_i'$, there exists some $\bar{m}$ such that $\tau_i(t_i, a_i, m) \neq \tau_i(t_i', a_i', m)$ for each $(a_i, a_i')$ and each $m > \bar{m}$. Since there are only finitely many types, one can choose $\bar{m}$ uniformly. (Hence, by changing the index $m$, we can take $\bar{m} = 0$ without loss of generality.) □

*Proof of Lemma 6.* Take any $\hat{t} \in \hat{T}$, and any $a \in S^\infty\left[\hat{t}\right]$. By Lemma 9, for each $m$, there exists $\bar{t}(m) \in \hat{T}$ such that $a \in W^\infty[\bar{t}(m)]$ and $\bar{t}(m) \to \hat{t}$ as $m \to \infty$. But by Lemma 8, since $a \in W^\infty[\bar{t}(m)]$, for each $m$ and $k$, there exists a finite, dominance-solvable model $T^{m,k}$ with a type profile $t(m,k)$, such that $S^\infty[t(m,k)] = \{a\}$ and $t(m,k) \to \bar{t}(m)$ as $k \to \infty$. Set $\tilde{t}(m) = t(m,m)$ and $T^m = T^{m,m}$, which satisfy the desired properties. □

## 6. Conclusion

In complete-information games, typically, there are many rationalizable strategies. This paper tries to understand why this is the case and how we should address this multiplicity. Firstly, it shows that for an open and dense set of types there is a unique rationalizable action, which is continuous with respect to players' beliefs. Multiplicity occurs only on the boundary of this set. Then, whenever we have only partial information about a strategic situation, we could not rule out the possibility that we could have found exactly what the unique rationalizable outcome is by having a more precise observation. This suggests that multiplicity is a property of these games, rather than an inherent property of rationalizability. The paper provides a characterization of rationalizable actions that uncovers a precise, intuitive reason for multiplicity. In defining a type, we often make idealized assumptions, which are meant to be approximately true. If our assumption can be thought of a simultaneous idealization of multiple strategically distinct situations, then rationalizability must yield the solutions at these non-idealized situations as possible solutions to the idealized case—due to upper-semicontinuity. This paper shows that multiplicity occurs *only* in such cases, and each rationalizable action represents such an idealization of an open set of situations in which the action is the unique solution. Therefore, we should consider the set of rationalizable actions as the summary of which idealizations we have made and what our results would have been if we have not made such idealized assumptions. Then, when we try to refine rationalizability, it is imperative to think about which information structures are more reasonable descriptions of the actual situation, rather than invoking seemingly compelling epistemic arguments. This is because any refinement is a selection among the information structures.

## Appendix A. Proof of Lemma 3

Definition 6. For any correspondence $F : X \to 2^Y$, $Gr(F) = \{(x,y) \, | y \in F[x]\}$ denotes the graph of $F$. For each $k$, define $B_i^k : \Delta\left(\Theta^* \times Gr\left(S_{-i}^{k-1}\right)\right) \to 2^{A_i}$ by

$$B_i^k(\pi) = \arg\max_{a_i'} E_\pi\left[u_i\left(a_i', a_{-i}, \theta\right)\right] = \arg\max_{a_i'} BR_i\left(\mathrm{marg}_{\Theta^* \times A_{-i}} \pi\right).$$

For $k = 0$, $S_i^k$ is upper-semicontinuous and non-empty by definition. Towards an induction, fix a $k > 0$, and assume that $S_{-i}^{k-1}$ is upper-semicontinuous and non-empty. I will show that $Gr\left(S_i^k\right)$ is closed. By the inductive hypothesis, $\Theta^* \times Gr\left(S_{-i}^{k-1}\right) \subseteq \Theta^* \times T_{-i}^* \times A_{-i}$ is closed and non-empty. Since $\Theta^* \times T_{-i}^* \times A_{-i}$ is compact, $\Theta^* \times Gr\left(S_{-i}^{k-1}\right)$ is also compact. Thus, $\Delta\left(\Theta^* \times Gr\left(S_{-i}^{k-1}\right)\right)$ is compact. Moreover, $u_i$ is continuous and bounded (by compactness of $\Theta^* \times A$), so that $E_\pi\left[u_i\left(a_i, a_{-i}, \theta\right)\right]$ is a continuous function of $\pi$ (by definition of weak convergence). Therefore, by Berge's Maximum Theorem, $Gr\left(B_i^k\right) \subseteq \Delta\left(\Theta^* \times Gr\left(S_{-i}^{k-1}\right)\right) \times A_i$ is closed. Since $\Delta\left(\Theta^* \times Gr\left(S_{-i}^{k-1}\right)\right) \times A_i$ is compact, $Gr\left(B_i^k\right)$ is also compact. Now, by definition of weak convergence, $\mathrm{marg}_{\Theta^* \times T_{-i}^*}\pi$ is a continuous function of $\pi$. Since $T_i^*$ is isomorphic to $\Delta\left(\Theta^* \times T_{-i}^*\right)$ (Mertens and Zamir (1985)), there also exists a continuous function $\phi : \Delta\left(\Theta^* \times T_{-i}^*\right) \to T_i^*$, such that $\phi\left(\kappa_{t_i}\right) = t_i$ for each $t_i$. Consider the continuous mapping $\psi : (\pi, a_i) \mapsto \left(\phi\left(\mathrm{marg}_{\Theta^* \times T_{-i}^*}\pi\right), a_i\right)$. By definition,

$Gr\left(S_i^k\right)=\psi\left(Gr\left(B_i^k\right)\right)$. But, since $Gr\left(B_i^k\right)$ is compact and $\psi$ is continuous, $\psi\left(Gr\left(B_i^k\right)\right)$ is closed. Moreover, since $\Theta^*\times Gr\left(S_{-i}^{k-1}\right)$ is closed (and $A_{-i}$ is finite), for each $t_i$, one can easily construct a $\pi\in\Delta\left(\Theta^*\times Gr\left(S_{-i}^{k-1}\right)\right)$ such that $\mathrm{marg}_{\Theta^*\times T_{-i}^*}\pi=\kappa_{t_i}$, so that $S_i^k[t_i]$ is non-empty.

Finally, since $S_i^k[t_i]$ is non-empty for each $k<\infty$ and $A_i$ is finite, $S_i^\infty[t_i]=\bigcap_{k<\infty}S_i^k[t_i]\neq\varnothing$. Moreover, since $Gr\left(S_i^k\right)$ is closed for each $k<\infty$, $Gr\left(S_i^\infty\right)=\bigcap_{k<\infty}Gr\left(S_i^k\right)$ is closed.

## References


[1] Battigalli, P. (2003): "Rationalizability in infinite, dynamic games with complete information," *Research in Economics*, 57, 1-38.

[2] Battigalli, P. and M. Siniscalchi (2003) "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, Vol. 3: No. 1, Article 3.

[3] Brandenburger, A. and E. Dekel (1993): "Hierarchies of Beliefs and Common Knowledge," *Journal of Economic Theory*, 59, 189-198.

[4] Carlsson, H. and E. van Damme (1993): "Global Games and Equilibrium Selection," *Econometrica*, 61, 989-1018.

[5] Dekel, E. D. Fudenberg, S. Morris (2003): "Interim Rationalizability," mimeo.

[6] Dekel, E. D. Fudenberg, S. Morris (2004): "Dense Types," mimeo.

[7] Ely, J. and M. Peski (2004): "Hierarchies of beliefs and interim rationalizability," mimeo.

[8] Feinberg, Y. (2000): "Characterizing Common Priors in the Form of Posteriors," *Journal of Economic Theory*, 91, 127-179.

[9] Frankel, D.M., S. Morris, A. Pauzner (2003):"Equilibrium selection in global games with strategic complementarities," *Journal of Economic Theory* 108, 1-44.

[10] Harsanyi, J. (1967): "Games with Incomplete Information played by Bayesian Players. Part I: the Basic Model," *Management Science* 14, 159-182.

[11] Morris, S. and H. Shin (1998): "Unique Equilibrium in a Model of Self-Fulfilling Attacks," *American Economic Review* 88, 587-597.

[12] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982): "Rational Cooperation in the Finitely-Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27, 245-52.

[13] Lipman, B. (2003): "Finite Order Implications of Common Priors," *Econometrica*, 71, 1255-1267.

[14] Mertens, J. and S. Zamir (1985): "Formulation of Bayesian Analysis for Games with Incomplete Information," *International Journal of Game Theory*, 10, 619-632.

[15] Monderer, D. and D. Samet (1989): "Approximating Common Knowledge with Common Beliefs," *Games and Economic Behavior*, 1, 170-190.

[16] Oxtoby, J. (1980): *Measure and Category*, Springer-Verlang, Berlin.

[17] Rubinstein, A. (1989): "The Electronic Mail Game: Strategic Behavior Under 'Almost Common Knowledge'," *The American Economic Review*, Vol. 79, No. 3, 385-391.

[18] Simon, R. (2003): "Games of incomplete information, Ergodic theory, and the measurability of equilibria," *Israeli Journal of Mathematics* 138, 73-92.

[19] Van Zandt, T. and X. Vives (2004): "Monotone Equilibria in Bayesian Games of Strategic Complementarities."

[20] Weinstein, J. and M. Yildiz (2004): "Finite-order implications of any equilibrium," *MIT Department of Economics Working Paper* 03-14.

MIT

*URL*: `http://econ-www.mit.edu/faculty/myildiz/index.htm`