GENERIC UNIQUENESS OF RATIONALIZABLE ACTIONS

MUHAMET YILDIZ

ABSTRACT. For a finite set of actions and a rich set of fundamentals, consider the 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, continuity

JEL Numbers: C72, C73.

1. INTRODUCTION

Rationalizability is considered a weak solution concept, for the set of rationalizable strategies tends to be large in complete information games. In their seminal work, Carlsson and van Damme (1993) have challenged this view. They have assumed that the domain of possible payoff parameters is rich enough, so that each action can be strictly dominant for some parameter value. Then, for two-player, two-action supermodular games of complete information, they have shown that if each player observes a noisy signal about parameters, with small, additive, independent noise instead of parameters being common knowledge, then the resulting game is dominance-solvable—except for the degenerate signal values at which the strategies jump. Morris and Shin (1998) and Frankel, Morris, and Pauzner (2003) have extended this result to all monotone supermodular games of complete information (Van Zandt and Vives (2004)). These results appear to be specific to supermodular games, in that such perturbations need not reduce the set of rationalizable

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; the main ideas of 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.

outcomes in general games, such as the Matching-Pennies game. Also, multiplicity reappears when there is a sufficiently precise "public" signal (Helwig (2002)). In this paper, I show that the intuition of Carlsson and van Damme is quite general. Under their richness assumption, we can always introduce small incomplete information in such a way that the perturbed types have unique rationalizable actions. Indeed, for generic types in the universal type space, there is a unique rationalizable action. One cannot use this to select a particular equilibrium, however. By considering suitable information structures, we can select any rationalizable strategy profile in the original game.

Example 1 (Carlsson and van Damme (1993)). To be concrete, consider

	α_2	β_2
α_1	θ, θ	$\theta - 1, 0$
β_1	$0, \theta - 1$	$0, 0$

where θ is a real number. Assume that θ is unknown but each player $i \in \{1, 2\}$ observes a noisy signal $x_i = \theta + \varepsilon\eta_i$, where (η_1, η_2) is independently distributed from θ, and the support of θ contains an interval $[a, b]$ where $a < 0 < 1 < b$. When $\varepsilon = 0$, θ 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 ε 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 β_i whenever $x_i < 1/2$, and it is α_i whenever $x_i > 1/2$.

This difference between the cases of $\varepsilon = 0$ and small but positive ε becomes important when we analyze a strategic situation in the interim stage, when players already have their private information. In the interim stage, players have beliefs about fundamentals, which are called the first-order beliefs, beliefs about the first-order beliefs, which are called the second-order beliefs, and so on. Assume that we cannot observe players' beliefs perfectly, but we observe their beliefs at finitely many orders with some noise. In particular, our observation suggests that a player's kth-order belief is within an open set for finitely many k—with respect to the weak topology as in "convergence in distribution". Now, suppose that it is consistent with our observation that it is common knowledge that $\theta = \bar{\theta}$ for some $\bar{\theta} \in (0, 1)$ (and $\varepsilon = 0$). That is, if we compute the kth-order beliefs using this complete-information model, they will be in the open set for each k above. In that case, we tend to model the situation by a simple game in which it is common knowledge that $\theta = \bar{\theta}$. This simplicity tends to complicate the analysis. There are now multiple equilibria and hence multiple rationalizable outcomes. We then introduce epistemic arguments in order to reduce the set of outcomes, so that we can make predictions. On the other hand, when $\varepsilon = 0$ is consistent with our observation, it is also consistent with our observation that ε is small but positive. If we model the situation by small but positive ε, then our model will be more complicated, but our analysis will be simple. There will be a unique rationalizable strategy profile. This will also allow generating insights that were not possible in the complete-information case. For example, in the currency-attack problem, if θ is the vulnerability of the economy, then Example 1 predicts that attack becomes likelier when the economy is more vulnerable (Morris and Shin (1998)). Moreover, our epistemic arguments for selection in the case of $\varepsilon = 0$ are all muted when ε is positive, as there is a unique outcome in the latter case. The unique outcome may differ from what we select. In that case, our predictions will be false when ε is positive.

My main result shows that this intuition is quite general. Formally, consider a finite-player, finite-action game with some unknown payoff parameters. Following Carlsson and van Damme, assume that each action becomes strictly dominant for some parameter value. Endow the game with the universal type space T^* of Mertens and Zamir (1985) and Brandenburger and Dekel (1993), where T^* is endowed with the usual product topology of weak convergence. The product topology precisely captures the above observational problem in the interim stage, as the open sets here correspond to the set of type profiles we cannot rule out for some observation. Writing A for the set of actions, 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.*[1]

[1]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 construction by Weinstein and Yildiz (2004), whose main idea can be traced back to Rubinstein (1989) and Carlson and van Damme (1993).

Since U is dense, in each open set there is a type profile with unique rationalizable outcome. Hence, in the modelling problem discussed above, no matter how precise our observation is, there will be a type profile (from some model) that is consistent with our observation. If one chooses to represent the situation by that type profile, then rationalizability (and the chosen model) will predict a unique strategy profile as the outcome for the situation at hand. Moreover, U is open, and s^* is locally constant. Now, suppose that in the actual situation that we try to model, there is indeed a unique rationalizable outcome, i.e., the players' beliefs are accurately represented by some $t \in U$. Then, this implies that there is a neighborhood of t at which we have a unique rationalizable outcome (openness), and all type profiles in this neighborhood have the same action (continuity). In that case, if we had a sufficiently precise (but not necessarily perfect) observation about sufficiently many orders of beliefs, we could know precisely what players will play according to rationalizability—independent of the type profile we choose to model the situation. This establishes that rationalizability is a strong solution concept in the following sense. Without precise knowledge of entire infinite hierarchy of beliefs, we can never rule out the possibility that, by having a more precise (but not perfect) observation about the actual case, we could have known what players will do using only rationalizability.

Using this genericity result, we can generalize the result of Carlsson and van Damme (1993) to arbitrary games. For arbitrary finite-action games with arbitrary payoff and information structures (with possibly infinite type spaces), I show that we can introduce small incomplete information in such a way that the resulting game is dominance-solvable, where the incomplete information need not satisfy the assumptions of Carlsson and van Damme on the noise structure. Moreover, the dominance-solvable model will remain so, when further small perturbations are introduced. In particular, even if there were perturbations that lead to multiplicity, we can introduce further perturbations to regain uniqueness, and the reverse is not true for perturbations with unique rationalizable strategies.

In particular, when presence of a "public" signal leads to multiplicity, we can introduce further small incomplete information to obtain a dominance-solvable model, which will be robust to further small perturbations. The intuition for this comes from Carlsson and van Damme. By designating a signal public, one assumes that its value is common knowledge. That makes it possible to coordinate on different equilibria. When we slightly relax this common-knowledge assumption, coordination becomes difficult. Common-knowledge assumptions lead to multiplicity in different games for different reasons. For example, in the Matching-Pennies game, some players have an incentive to change their actions when these actions are known by the players. In that case, there must be multiple rationalizable outcomes in the complete-information case. 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.

More broadly, slight relaxation of an assumption in a given model (if anything) reduces 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. Indeed, if we relax this assumption in a suitable way, we can get rid of all but one rationalizable action—as this paper shows.

This leads to a natural question: should we then use a dominance-solvable model that is consistent with our observation to make predictions about what players will do? Indeed, since all the types in their perturbed model play according to risk dominance, Carlsson and van Damme (1993) have proposed that we select the risk-dominant equilibrium in the original game. My second result, which builds upon an earlier result of Weinstein and Yildiz (2004), answers this question. I show that, given any finite type space and any rationalizable strategy in that type space, one can slightly perturb the players' interim beliefs to obtain a nearby dominance-solvable model in which the given strategy is uniquely rationalizable for the nearby types. That is, we can select any rationalizable strategy profile that we want by focusing on a suitable dominance-solvable model that is consistent with observation. Therefore, we cannot use this method for selecting a particular equilibrium. The following example illustrates this.

Example 2 (Izmalkov and Yildiz (2005)). In Example 1, drop the common-prior assumption, and assume that

$$\Pr_i \left(\eta_j > \eta_i | \eta_i\right) = q \qquad (\forall i \neq j) \tag{1.1}$$

for some $q \in (0,1)$, where $\Pr_i$ is the probability according to player i. That is, player i assigns probability q to the event that the other player is more optimistic ex-post.

Under the common-prior assumption $q = 1/2$. Here, $q - 1/2$ measures the level of optimism of j according to i. For each $x_i \neq 1 - q$, there is a unique rationalizable action $s_i^*(x_i)$, given now by

$$s_i^*(x_i) = \begin{cases} \alpha_i & \text{if } x_i > 1 - q \\ \beta_i & \text{if } x_i < 1 - q. \end{cases}$$

Hence, given any x_i, we can make action α_i uniquely rationalizable by choosing the level q of optimism sufficiently high, or make action β_i uniquely rationalizable by choosing the level q of optimism sufficiently low. But when ε is very small, the value of q has very small impact on players' beliefs,[2] and the players' beliefs converge to that of common knowledge at all orders as $\varepsilon \to 0$.

Now consider the currency-attack problem of Morris and Shin. As they discuss, in their model there is no role for investor sentiments, which were given a prominent role in earlier informal arguments based on multiple equilibria. Example 2 shows that by dropping the common-prior assumption (about small aspects of the problem), we can develop a richer theory. In the new theory, measured as the likelihood of the fellow players' optimism, investor sentiments play an important role along with the fundamentals. We can generate a richer set of monotone comparative statics: the attack becomes likelier if investors are more optimistic, or the economy is more vulnerable. We cannot, however, do this by simply focusing on risk-dominant equilibrium.

This result also establishes that rationalizability is a strong solution concept in another sense: without having precise information about the entire hierarchy of beliefs, we cannot refine rationalizability any further to obtain sharper predictions. For each rationalizable action profile, there will always be a type profile that is consistent with our observation and for which the strategy profile is uniquely rationalizable. Since our refinement must select that outcome for that type profile, the set of outcomes that cannot be ruled out by our refinement and observation will contain the set of all rationalizable outcomes. This issue has been extensively studied for equilibrium refinements by Weinstein and Yildiz (2004).

Example 2 may suggest that the above results are due to the type spaces without a common prior. That is not the case. Using a result by Lipman (2003), I show

[2]For example, starting from the uniform distribution on $[-1, 1]$ for both η_i and η_j, shift the distribution of η_j up by the amount of y where $(1 - y/2)^2/2 = 1 - q$, so that (1.1) is satisfied. Then, the conditional distribution of x_j on x_i is shifted only by εy.

that all of the above results remain intact if we restrict ourselves to the finite models with common prior. This will hold with the exception that, when we impose the common-prior assumption, the perturbed model may contain some far-away types with multiple rationalizable actions. All of the nearby types will have unique rationalizable actions, which will be in accordance with the fixed rationalizable strategy of the original model in my second result.

As I discussed already, the common-knowledge assumptions in our models tend to produce extra rationalizable actions. The above multiplicity at far away types may also be a factor in multiplicity of rationalizable strategies in our models. Dominance-solvability of the entire model is a much more stringent condition than having unique rationalizable actions for types relevant to the analysis of a particular case. For example, consider a class of games of incomplete information, such that in each of them most of the types have unique rationalizable actions but some types have multiple rationalizable actions. Clearly, each of these games will have multiple rationalizable strategies, but rationalizability will lead to unique outcomes for most of the types considered in this class.

In the next section I illustrate how one can make the Matching-Pennies game dominance-solvable by introducing small 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. In order to illustrate how one can introduce incomplete information in a general game and obtain a dominance-solvable model, I now consider the difficult case of Matching-Pennies game. I will first consider a belief structure without a common prior and then reinstate this assumption.

Example 3 (Matching Pennies—without a common prior). Consider the payoff matrix

	α_2	β_2
α_1	$\theta, 0$	$\theta - 1, \theta$
β_1	$0, 0$	$0, \theta - 1$

.

If θ 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 θ is uniformly distributed on Θ. Players have different priors on the signals (x_1, x_2). Conditional on $\theta = \theta_m$, each player i assigns probability $1 - \gamma$ to $(x_i, x_j) = (\theta_m, \theta_{m-1})$ and probability γ to $(x_i, x_j) = (\theta_{m-1}, \theta_m)$. As in Example 1, it is common knowledge that the players' signals are within ε-neighborhood of θ, 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	θ_0	θ_1	θ_2	θ_3	θ_4	θ_5	θ_6	θ_7	θ_8	$\cdots$
$s_1^*(x_1)$	β_1	α_1	α_1	β_1	β_1	α_1	α_1	β_1	β_1	$\cdots$
$s_2^*(x_2)$	α_2	α_2	β_2	β_2	α_2	α_2	β_2	β_2	α_2	$\cdots$

(Clearly, when $x_i = \theta_0$, player i assigns high probability $1 - \gamma$ to $\theta = \theta_0$, when β_1 and α_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 α_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 Mth 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 4 (Matching Pennies—with a common prior). In the previous example, assume that, in addition to x_i, each player i partially observes a random variable k that is correlated with θ 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 (θ, x_1, x_2, k) as follows. Let $\mu_i(\theta, x_1, x_2)$ be the prior probability of (θ, 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 (θ,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 ε-neighborhood of θ. 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 $y_i(k) \leq 2K$. 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 $y_i(k) \leq 2K - m$ where $x_i = \theta_m$, 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.[3] 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 way 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 Θ is a compact metric space of payoff-relevant parameters θ, and u_i is continuous in θ. The finite set A is endowed with the discrete topology. The game is

[3] 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 $y_i(k) \leq 2K$. Assuming the statement is true for $m-1$, consider any $(\theta_m, y_i(k))$ with $y_i(k) \leq 2K - m$. Player i knows that $y_j(k) \leq 2K - m + 1$, 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 $s_i^*(\theta_m)$.

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 Θ, representing the beliefs of i about θ, $t_i^2 \in \Delta(\Theta \times \Delta(\Theta)^n)$ is a probability distribution for $(\theta, t_1^1, t_2^1, \ldots, , t_n^1)$, representing the beliefs of i about θ 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 (θ, 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 $t_{i,m} \to t_i$.

A *strategy* of a player i is any function $s_i : T_i^* \to A_i$.[4] For each $i \in N$ and for each belief $\pi \in \Delta(\Theta \times A_{-i})$, $BR_i(\pi)$ denotes the set of actions $a_i \in A_i$ that maximize the expected value of $u_i(\theta, a_i, a_{-i})$ under the probability distribution π.

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, θ). This model allows all possible payoff functions, and here θ 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*

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

[4]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).

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.

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 profile $s : T^* \to A$ (resp. a strategy $s_i : T_i^* \to A_i$) is said to be rationalizable iff $s(t) \in S^\infty[t]$ for each t (resp., $s_i(t_i) \in S_i^\infty[t_i]$ 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 a weak notion of rationalizability strengthens my results; they will remain valid under any stronger notion of rationalizability.

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. In that case, $T^* \backslash T$ is simply the boundary of T, denoted by ∂T. Clearly, topological notions of genericity may widely differ from measure theoretical notions of genericity, but they 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. This caveat also applies to the discussions. Clearly, all what matters is what these results mean in terms of economic modeling, as discussed in the introduction.

Definition 2 (Finite Types, Models). A subset $T \subseteq T^*$ is said to be *belief-closed* iff for each $t_i \in T_i$, $\text{supp}(\kappa_{t_i}) \subseteq \Theta \times T_{-i}$. A belief-closed $T \subseteq T^*$ is said to be *finite* iff T contains finitely many members and t_i^1 has finite support for each $t_i = (t_i^1, t_i^2, \ldots) \in T_i$. 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*. I will use the terms *model* and *belief-closed subset of* T^* interchangeably.

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 $T \subseteq T^*$ 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 $\text{supp}(p) = \Theta' \times T$ for some $\Theta' \subseteq \Theta$ and $\kappa_{t_i} = p(\cdot|\Theta \times \{t_i\} \times T_{-i})$ 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, $T_i^{CPA} = \{t_i \in T_i | T \text{ is belief-closed and admits a common prior}\}$. The next result by Lipman (2003) shows that, given any finite model "with full support", one can obtain a nearby finite model that admits a common prior. This is because the common-prior assumption does not put any restriction on finite-order beliefs other than full support (see also Feinberg (2000)).[5]

[5]Lipman (2003) uses a partitional model. If one takes $\Omega = \Theta \times T^*$ as the state space and $\{\{(\theta, t) | t_i = \tilde{t}_i\} | \tilde{t}_i \in T_i^*\}$ as the partition of player i, then the condition in the lemma immediately implies his weak-consistency condition, which characterizes the finite-order implications of the common-prior assumption.

Lemma 2 (Lipman (2003)). *Let $T \subseteq \hat{T}$ be a finite model with* $\text{supp}(\kappa_{t_i}) = \Theta' \times T_{-i}$ *for some $\Theta' \subseteq \Theta$ and for each $t_i \in T_i$. Then, for each m, there exists a finite model $T^m \subseteq \hat{T}$ that admits a common prior with full support and a one-to-one mapping $\tau(\cdot, m) : T \to T^m$ such that $\tau(t, m) \to t$ as $m \to \infty$.*

Lemmas 1 and 2 immediately implies the following result.

Lemma 3 (Lipman (2003)). *$\hat{T} \cap T^{CPA}$ is dense.*

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, i.e.,

$$t_{i,m} \to t_i \Rightarrow s_i(t_{i,m}) \to s_i(t_i) \tag{3.1}$$

for each sequence of types $t_{i,m}$. 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 η with $F[t'] \subseteq F[t]$ for each $t' \in \eta$.

Lemma 4 (Dekel, Fudenberg, and Morris (2004)). *S^∞ is non-empty and upper-semicontinuous.*

Dekel, Fudenberg, and Morris (2004) proves upper-semicontinuity of interim correlated rationalizability in their framework. Since my framework is slightly different (e.g. Θ 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 5. *Given any non-empty, upper-semicontinuous F, let $U_F = \{t \mid |F[t]| = 1\}$. 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 η 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. □

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 $a \in S^{\infty}\left[\hat{t}\right]$, there exists a sequence of finite 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 $S^{\infty}\left[\tilde{t}(m)\right] = \{a\}$ for each m. Moreover, T^m can be chosen to be dominance-solvable or with a common prior with full support.*

That is, given any type and any rationalizable action a_i for that type, one can find a nearby type for which a_i is uniquely rationalizable. Moreover the new type can be found in a dominance-solvable model or in a (possibly not dominance-solvable) model with a common prior. 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^{∞} 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 5, 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, $U = \varnothing$.

Proposition 1 uncovers an interesting structure of the universal type space T^*. One can divide T^* into finitely many open sets

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

and their boundaries $\partial U^a \equiv \overline{U^a}\backslash U^a$, where $\overline{U^a}$ is the closure of U^a. The open sets form a partition of an open, dense set U, while their boundaries cover the boundary of U, i.e., $T^*\backslash U = \bigcup_{a \in A} \partial U^a$, which is a nowhere-dense set. On each open set U^a, a is the unique rationalizable action profile. Since S^∞ is upper-semicontinuous, $a \in S^\infty[t]$ for each $t \in \partial U^a$. At any $t \in \partial U^a \cap \partial U^{a'}$ with distinct a and a', both a and a' are rationalizable. At any such t with multiple rationalizable action profiles, every rationalizable strategy profile s must be discontinuous, as there are sequences $t(a,m) \to t$ and $t(a',m) \to t$ with $s(t(a,m)) = a$ and $s(t(a',m)) = a'$, where $t(a,m) \in U^a$ and $t(a',m) \in U^{a'}$. Here, all rationalizable strategies are rendered discontinuous at t by the fact that the generically unique rationalizable theory changes its prescribed behavior at t.[6]

In summary, Proposition 1 establishes that, if one excludes a nowhere-dense set of types, then there will be a unique rationalizable strategy profile for the remaining types, and it will be continuous with respect to players' beliefs. Discontinuities or multiplicities arise only on the nowhere-dense boundary of the open and dense set

[6] It is also a general possibility that $t \in \partial U^a \backslash \cup_{a' \neq a} \partial U^{a'}$ for some a. But Lemma 6 implies that there cannot be such a finite type; it implies that $\hat{t} \in \cap_{a \in S^\infty[\hat{t}]} \overline{U^a}$ for each $\hat{t} \in \hat{T}$. At any $t \in \partial U^a \backslash \cup_{a' \neq a} \partial U^{a'}$, there are multiple rationalizable action profiles (as $t \in T^*\backslash U$), but a is the only action profile that remains rationalizable on an open neighborhood of t, and some rationalizable strategies may be continuous at t.

U, where the unique rationalizable strategy above potentially changes its prescribed behavior for players. Hence, from a theoretical point of view, for generic situations, rationalizability leads to quite robust predictions: we can know the players' actions if we know their beliefs sufficiently well. We do not need to know their beliefs about the strategies for this prediction; common knowledge of rationality suffices.

This is a theoretical robustness, however. The usual practical problems with dominance-solvability and other robustness results do apply here. One may have to specify the players' beliefs with such a high precision that it may be impractical to make any prediction with any reasonable level of precision. For example, a finitely-repeated prisoners' dilemma game with many repetitions will become dominance-solvable if we introduce small trembles and use a suitable reduced-form representation, but it is well known that the equilibrium predictions will dramatically change when the probability of an "irrational" type exceeds a very low threshold, such as 0.001, as shown by Kreps, Milgrom, Roberts, and Wilson (1982). Moreover, in application, we typically have a large set of rationalizable actions, suggesting that our common knowledge assumptions lead us to the boundary of U, and the present economic theories are about these nowhere-dense set of types.

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). In some type spaces, there may be distinct types with identical belief hierarchies. In such type spaces with "redundant types", there may be equilibrium strategies that are not rationalizable for the corresponding belief hierarchy in the universal type space if one insists on independence of strategies from θ. One needs a larger type space to capture the strategically relevant information encoded in the redundant types (Ely and Peski (2004)). On the other hand, even when there are "redundant types", if the belief hierarchy of a type is t_i, then all the rationalizable actions of that type are contained in $S_i^{\infty}[t_i]$ (Dekel, Fudenberg, and Morris (2003)). Proposition 1 establishes that, generically, $|S_i^{\infty}[t_i]| = 1$, and hence a unique action is rationalizable for all types that come from arbitrary spaces but have the same generic belief hierarchy. Then, the universal type space suffices to capture the strategic behavior of types with generic belief hierarchies. (The set U may not be open and dense in the larger space of Ely and Peski (2004).)

Remark 4 (Unified Theories). A strategy profile in this paper simultaneously describes an outcome for every model embedded in the universal type space. It can then be regarded as a *unified theory*. Proposition 1 implies that, if we assume common knowledge of rationality, then we can have only one unified theory for generic cases, and each of these unified theories will be continuous (prescribing the same behavior for indistinguishable models) at generic type profiles. Kohlberg and Mertens (1986) and Govindan and Wilson (2004) seek equilibrium refinements that depend only on the reduced-form representation and are independent to certain "irrelevant transformations," including the introduction of mixed strategies as pure strategies, a transformation that is ruled out here by the richness assumption. I take a complementary approach to the same conceptual problem they have addressed. Towards a unified theory of games, they focus on developing a uniform equilibrium refinement, while I show that generically there is only one such theory.

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 economic 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. The next result, which is the second main result of this paper, states that, when the original type space is finite, the dominance-solvable model can be taken to be part of a model that admits a common prior with full support.[7] Moreover, we can do this for each rationalizable strategy profile s_T in the finite model, so that s_T is the unique rationalizable strategy profile in the perturbed model. 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

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

[7] As in the matching-penny game, this result does not rule out the possibility that some far away types in the common-prior model have multiple rationalizable actions. (This is rather due to the method of proof.)

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 $m > \bar{m}$.)

Since $T^{s_T,m}$ is finite, by Lemma 11 in Section 5, for each integer k, there exist a finite model $\tilde{T}^{m,k}$ that admits a common prior and a one-to-one mapping $\tau'(\cdot, k) : T^{s_T,m} \to \tilde{T}^{m,k}$ such that $S^\infty[\tau'(\bar{t}, k)] = S^\infty[\bar{t}]$ and $\tau'(\bar{t}, k) \to \bar{t}$ as $k \to \infty$ for each $\bar{t} \in T^{s_T,m}$. Pick $\tilde{T}^{s_T,m} = \tilde{T}^{m,m}$ and $\tilde{\tau}(\cdot, s_T, m) = \tau'(\cdot, m) \circ \tau(\cdot, s_T, m)$. □

Building on a result of Weinstein and Yildiz (2004), Proposition 3 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, which will be defined in the next section.

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 directly 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 three main steps, which are presented as the following three lemmas. The first step (namely, Lemma 8) shows that, when we focus on strictly rationalizable strategies and do not require a common prior, 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 without a common prior. Finally, using the result of Lipman (2003), namely Lemma 2, and the second step one more time, one can show that the common-prior requirement can also be met (as stated in Lemma 11).

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 $S_i^{k+1}[\tilde{t}_i]$. 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)*

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

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 $a_i \in W_i^\infty[\hat{t}_i]$ 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 θ^{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 $l=\left(t_{-i}^1,t_{-i}^2,\dots,t_{-i}^{k-1}\right)$ and $h=\left(t_{-i}^k,t_{-i}^{k+1},\dots\right)$ are the lower and higher-order beliefs, respectively. Let $L=\left\{l|\exists h:(l,h)\in T_{-i}^*\right\}$. The inductive hypothesis is that for each finite $t_{-i}=(l,h)$ and each $a_{-i}\in W_{-i}^{k-1}\left[t_{-i}\right]$, there exists finite $\tilde{t}_{-i}\left[a_{-i}\right]=\left(l,\tilde{h}\left[l,a_{-i}\right]\right)\in T_{-i}^{\tilde{t}_{-i}[a_{-i}]}$ such that

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

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 $\left|S^k\left[t\right]\right|=1$ 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, $BR_i\left(\mathrm{marg}_{\Theta\times A_{-i}}\pi\right)=\{a_i\}$ 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 $\pi\left(a_{-i}\in W_{-i}^{k-1}\left[t_{-i}\right]\right)=1$. Using the inductive hypothesis, define mapping $\mu:\mathrm{supp}\left(\mathrm{marg}_{\Theta\times L\times A_{-i}}\pi\right)\to\Theta\times T_{-i}^*$, by

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

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

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

the probability distribution induced on $\Theta\times T_{-i}^*$ by the mapping μ and the probability distribution π, where proj_X denotes the projection mapping to X. Notice that $\mathrm{proj}_{\Theta\times L}\circ\mu\circ\mathrm{proj}_{\Theta\times L\times A_{-i}}=\mathrm{proj}_{\Theta\times L}$. Then, $\mathrm{marg}_{\Theta\times L}\kappa_{\tilde{t}_i}=\mathrm{marg}_{\Theta\times L}\kappa_{t_i}$, 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 $(\theta,t_{-i})\in\mathrm{supp}\left(\kappa_{\tilde{t}_i}\right)$, which is of the form $\left(\theta,l,\tilde{h}\left[l,a_{-i}\right]\right)$, has a unique action $a_{-i}\in S_{-i}^{k-1}\left[\tilde{t}_{-i}\left[a_{-i}\right]\right]$. Thus, there exists a unique $\tilde{\pi}\in\Delta\left(\Theta\times T_{-i}^*\times A_{-i}\right)$ such that $\mathrm{marg}_{\Theta\times T_{-i}^*}\pi=\kappa_{\tilde{t}_i}$ and $\pi\left(a_{-i}\in S_{-i}^{k-1}\left[t_{-i}\right]\right)=1$; it is $\tilde{\pi}=\kappa_{\tilde{t}_i}\circ\gamma^{-1}=\pi\circ\mathrm{proj}_{\Theta\times L\times A_{-i}}^{-1}\circ\mu^{-1}\circ\gamma^{-1}$ where $\gamma:\left(\theta,l,\tilde{h}\left[l,a_{-i}\right]\right)\mapsto\left(\theta,l,\tilde{h}\left[l,a_{-i}\right],a_{-i}\right)$. Clearly, $\mathrm{proj}_{\Theta\times A_{-i}}\circ\gamma\circ\mu\circ\mathrm{proj}_{\Theta\times L\times A_{-i}}=\mathrm{proj}_{\Theta\times A_{-i}}$. Hence, $\mathrm{marg}_{\Theta\times A_{-i}}\tilde{\pi}=\mathrm{marg}_{\Theta\times A_{-i}}\pi$. 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

$$T_i^{\tilde{t}_i} = \left\{\tilde{t}_i\right\} \cup \left(\bigcup_{(\theta, t_{-i}[a_{-i}]) \in \operatorname{supp}\left(\kappa_{\tilde{t}_i}\right)} T_i^{t_{-i}[a_{-i}]} \right),$$

$$T_j^{\tilde{t}_i} = \bigcup_{(\theta, t_{-i}[a_{-i}]) \in \operatorname{supp}\left(\kappa_{\tilde{t}_i}\right)} T_j^{t_{-i}[a_{-i}]} \qquad (j \neq i).$$

Since $\operatorname{supp}\left(\operatorname{marg}_{\Theta \times L \times A_{-i}} \pi\right) \subseteq \operatorname{supp}\left(\kappa_{\hat{t}_i}\right) \times A_{-i}$ is finite, the range of μ is finite, rendering $\operatorname{supp}\left(\kappa_{\tilde{t}_i}\right)$ 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, $\operatorname{supp}\left(\kappa_{t_j}\right) \subseteq \Theta \times T_{-j}^{t_{-i}[a_{-i}]} \subseteq \Theta \times T_{-j}^{\tilde{t}_i}$. On the other hand, $\operatorname{supp}\left(\kappa_{\tilde{t}_i}\right) \subseteq \Theta \times T_{-i}^{\tilde{t}_i}$, as $t_{-i}[a_{-i}] \in T_{-i}^{t_{-i}[a_{-i}]}$ for each $(\theta, t_{-i}[a_{-i}]) \in \operatorname{supp}\left(\kappa_{\tilde{t}_i}\right)$. 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 $t_j \in T_j^{\tilde{t}_i} \setminus \left\{\tilde{t}_i\right\}$, $\left|S^k[t_j]\right| = 1$, and hence $\left|S^{k+1}[t_j]\right| = 1$.

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 $l \le m$ 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 δ_x denote the probability distribution that puts probability 1 on $\{x\}$ and Θ' 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 θ and the others' types $\tau_{-i}(t_{-i}, a_{-i}, m)$.[8] Now, since $a_i \in$

[8]Notice that I am simply defining a finite type space. Hence, it suffices to define the belief of each type about θ 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 π^{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 (θ, t_{-i}, a_{-i}) according to π^{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 π^{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 kth 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 $\text{proj}_{\Theta\times L}\left(\hat{\tau}_{-i,m}(\theta, t_{-i}, a_{-i})\right) = \text{proj}_{\Theta\times L}\left(\theta, \tau_{-i}(t_{-i}, a_{-i}, m)\right) = \left(\theta, \tau_{-i}^{k-1}(t_{-i}, a_{-i}, m)\right)$, which converges to $\left(\theta, t_{-i}^{k-1}\right)$. That is, $\text{proj}_{\Theta\times L} \circ \hat{\tau}_{-i,m}$ pointwise 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.) □

In the previous lemma, if the original model T is dominance-solvable, then the new model will also be dominance-solvable. In the new model, S^∞ and W^∞ will coincide. This is stated in the next lemma.

LEMMA 10. *Under Assumption 1, for any finite, dominance-solvable model $T \subseteq \hat{T}$ and any m, there exists a finite, dominance-solvable model $T^m \subseteq \hat{T}$ and a one-to-one and onto mapping $\tau(\cdot, m) : T \to T^m$ such that (i) $W^\infty[\tau(t,m)] = S^\infty[\tau(t,m)] = S^\infty[t]$ for each (t,m), and (ii) $\tau(t,m) \to t$ as $m \to \infty$.*

Proof. Take T^m and $\tau(\cdot, m)$ as in Lemma 9. Since T is dominance-solvable, $\tau(\cdot, m)$ is simply defined on type profiles. Since T is dominance-solvable and $\tau(t,m) \to t$, by Lemmas 4 and 5, there exists $\bar{m}$ such that for each $m > \bar{m}$, $S^\infty[\tau(t,m)] = S^\infty[t]$. Since T is finite, $\bar{m}$ can be chosen uniformly for all t. Moreover, by Lemma 9, $S^\infty[t] = W^\infty[\tau(t,m)]$. □

Together with the result of Lipman (2003) and upper-semicontinuity of S^{∞}, this implies that a dominance-solvable model can be approximately embedded in a larger model with a common prior without affecting the rationalizable strategies. This is stated in the next lemma.

LEMMA 11. *Under Assumption 1, for any finite, dominance-solvable model $T \subseteq \hat{T}$ and any m, there exist a finite model T^m that admits a common prior with full support and a one-to-one mapping $\tau(\cdot, m) : T \to T^m$ such that (i) $S^{\infty}[\tau(t,m)] = S^{\infty}[t]$ for each (t,m), and (ii) $\tau(t,m) \to t$ as $m \to \infty$.*

Proof. By Lemma 10, for each m, there exist a dominance-solvable model $\tilde{T}^m \subseteq \hat{T}$ and a one-to-one and onto mapping $\tilde{\tau}(\cdot, m) : T \to \tilde{T}^m$ with $W^{\infty}[\tilde{\tau}(t,m)] = S^{\infty}[\tilde{\tau}(t,m)] = S^{\infty}[t]$, and such that $\tilde{\tau}(t,m) \to t$ as $m \to \infty$. Since each type $\tilde{\tau}_i(t_i, m)$ plays a strict best reply to his unique belief, one can perturb $\tilde{\tau}_i(t_i,m)$ by assigning positive but small probability at each $(\theta, \tilde{\tau}_{-i}(t_{-i}, m)) \in \tilde{\Theta} \times \tilde{T}^m$ on which $\tilde{\tau}_i(t_i,m)$ puts zero probability without affecting $W^{\infty}[\tilde{\tau}(t,m)]$ or $S^{\infty}[\tilde{\tau}(t,m)]$, where $\tilde{\Theta}$ is the finite set of all parameters on which some type $\tilde{t}_j \in \tilde{T}_j^m$ puts positive probability. Hence, there exist sequences of dominance-solvable models $T^{m,k} \subseteq \hat{T}$ and one-to-one mappings $\bar{\tau}(\cdot, k) : \tilde{T}^m \to T^{m,k}$, such that for each $\bar{\tau}(\tilde{\tau}(t,m), k)$, (i) $\text{supp}\big(\kappa_{\bar{\tau}_i(\tilde{\tau}_i(t_i,m),k)}\big) = \tilde{\Theta} \times T_{-i}^{m,k}$ (ii) $W^{\infty}[\bar{\tau}(\tilde{\tau}(t,m),k)] = S^{\infty}[\bar{\tau}(\tilde{\tau}(t,m),k)] = S^{\infty}[t]$, and (iii) $\bar{\tau}(\tilde{\tau}(t,m),k) \to \tilde{\tau}(t,m)$ as $k \to \infty$. But by Lemma 2, for each l, there exists a finite model $T^{m,k,l} \subseteq \hat{T}$ that admits a common prior and a one-to-one mapping $\hat{\tau}(\cdot, l) : \tilde{T}^{m,k} \to T^{m,k,l}$ such that $\hat{\tau}(\bar{\tau}(\tilde{\tau}(t,m),k),l) \to \bar{\tau}(\tilde{\tau}(t,m),k)$ as $l \to \infty$. But since $\tilde{T}^{m,k}$ is dominance-solvable, by Lemmas 4 and 5, this implies that $S^{\infty}[\hat{\tau}(\bar{\tau}(\tilde{\tau}(t,m),k),l)] = S^{\infty}[\bar{\tau}(\tilde{\tau}(t,m),k)]$ when $l > \bar{l}$ for some $\bar{l}$. Hence, when $l > \bar{l}$, $S^{\infty}[\hat{\tau}(\bar{\tau}(\tilde{\tau}(t,m),k),l)] = S^{\infty}[t]$. By setting $T^m \equiv T^{m,m,m}$ and $\tau(\cdot, m) \equiv \hat{\tau}(\cdot, m) \circ \bar{\tau}(\cdot, m) \circ \tilde{\tau}(\cdot, m)$ for $m > \bar{l}$, one completes the proof. □

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 $k \to \infty$. If we only need dominance-solvability, then $\tilde{t}(m) = t(m,m)$ and $T^m = T^{m,m}$ satisfy the desired properties. Now suppose we need a common prior. By Lemma 11, for each m, k, l, there exist a finite model $T^{m,k,l}$ that admits a common prior and a one-to-one mapping $\tau(\cdot, l) : T^{m,k} \to T^{m,k,l}$, such that $\tau(t(m,k),l) \to t(m,k)$ as $l \to \infty$, and $S^{\infty}[\tau(t(m,k),l)] = S^{\infty}\left[\hat{t}\right] = \{a\}$

for every $t(m,k)$ and l. We then obtain a model with a common prior, by setting $\tilde{t}(m) = \tau(t(m,m), m)$ and $T^m = T^{m,m,m}$. $\square$

6. Conclusion

Usual game theoretical models typically have a multitude of rationalizable strategies, suggesting that rationalizability is a weak solution concept. The multiplicity may be, however, a property of the present models, rather than a property of rationalizability. Generically, rationalizability leads to strong robust predictions: there exists a unique rationalizable outcome, and it is continuous with respect to the players' beliefs. Whenever we have only partial information about a strategic situation (as described in the Introduction), we can find a type profile that is consistent with our information and with a unique rationalizable outcome. Moreover, when there is a unique rationalizable outcome for a given situation, if our partial information is sufficiently precise, we can know which action profile is played according to rationalizability. Hence, under our partial information, we can never rule out the possibility that by obtaining more precise partial information, we could have learned what each player plays according to rationalizability. Rationalizability is a strong solution concept in this sense. It is also a strong solution concept in the sense that we could not refine it to obtain sharper predictions under our partial information.

Appendix A. Proof of Lemma 4

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

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

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 $E_\pi\left[u_i(a_i, a_{-i}, \theta)\right]$ is a continuous function of π (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, $\operatorname{marg}_{\Theta \times T_{-i}^*} \pi$ is a continuous function of π. 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 $\psi : (\pi, a_i) \mapsto \left(\phi\left(\text{marg}_{\Theta \times T_{-i}^*}\pi\right), a_i\right)$. 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 ψ 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 $\text{marg}_{\Theta \times T_{-i}^*}\pi = \kappa_{t_i}$, so that $S_i^k\left[t_i\right]$ is non-empty.

Finally, since $S_i^k\left[t_i\right]$ is non-empty for each $k < \infty$ and A_i is finite, $S_i^\infty\left[t_i\right] = \bigcap_{k<\infty} S_i^k\left[t_i\right] \neq \varnothing$. 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] Govindan, S. and R. Wilson (2004): "Axiomatic Justification of Stable Equilibria," mimeo.

[11] Kohlberg, E., and J. Mertens (1986), "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003-1038.

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

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

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

[16] Samet, D. (1998): "Iterated Expectations and Common Priors," *Games and Economic Behavior*, 24, 131-141.

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

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

[19] 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`