GENERIC UNIQUENESS AND CONTINUITY OF RATIONALIZABLE STRATEGIES

MUHAMET YILDIZ

Abstract. For a finite set of actions and a rich set of fundamentals, consider the rationalizable actions on a 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) A rationalizable strategy is continuous at a finite type if and only if there is a unique rationalizable action for that type.

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

JEL Numbers: C72, C73.

1. Introduction

This paper shows that, if one considers all possible payoff and belief structures, then rationalizability generically leads to a unique solution. Moreover, when there is multiplicity, refining rationalizability implies ruling out some nearby dominance-solvable models as the true model. Formally, consider a finite-player, finite-action game with some unknown payoff parameters. The set A of action profiles is endowed with the discrete topology. Assume that each action can be strictly dominant for some parameter value, e.g., that the domain of possible payoff structures is not restricted a priori. Endow the game with the universal type space T of Mertens and

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.

Zamir (1985) and Brandenburger and Dekel (1993), where T is endowed with the usual product topology of weak convergence. 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 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]

That is, if we exclude a nowhere-dense set of type profiles, then for each remaining type profile, there is a unique rationalizable action profile, and the action profile is given by a continuous function s^*. Since a rationalizable strategy profile must agree with s^* on the open and dense set U, it must be continuous on U. Continuity of s^* means that each type profile in U has an open neighborhood on which s^* is constant. This leads to an interesting picture: the universal type space is comprised of a collection of open sets and their boundaries, such that in each of the open sets, a fixed action profile is the unique rationalizable action profile. Multiplicities and discontinuities occur only on the boundaries of these sets, where the unique rationalizable action profile potentially changes. Ubiquity of multiple rationalizable actions in usual game theoretical models suggests that our common knowledge assumptions put our models on these boundaries. This also shows that the nowhere-dense set here is not negligible, as it includes many of the models in economics literature.

What does this mathematical result tell us about economic modeling? For an answer, let us examine the universal type space more closely. In this space, a type is a coherent hierarchy of beliefs about the payoff parameters, where the first-order beliefs are about the parameters, the second-order beliefs are about the first-order

[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 the seminal works of Rubinstein (1989) and Carlson and van Damme (1993).

beliefs, and so on. The universal type space contains most type spaces as "belief-closed" subspaces (henceforth, models). For example, it contains a family of models in which the players observe the parameters with noise, where the level of noise and the prior beliefs vary across the models, as well as the complete information model with no noise. If we fix a prior and let the size of the noise go to zero, the players' beliefs at each finite order converges to that of complete information. In that case, it becomes difficult for the modeler to distinguish these models from each other in the interim stage, when one can only observe the posterior beliefs (possibly only partialy). The product topology captures this difficulty of identification. In this topology, a sequence of types converge to a fixed type if the beliefs at all orders converge. In particular, the above models converge to that of complete information as the noise vanishes.

In the ex ante stage, the modeler can, of course, find the above models quite different;[2] the prior may have substantial impact on strategic behavior even in the presence of strong information. Unfortunately, however, in most applications, the modeler faces the situation only in the interim stage. The ex ante stage is often constructed by the modeler in order the capture the situation in a coherent model. Indeed, the central question of this paper is how the modeler should proceed given that he has selected one model to analyze among many indistinguishable models. To interpret the result from that angle, assume that the modeler can make observations about finite but arbitrarily many orders of beliefs, so that after the observation he knows that the belief at each observed order is in some arbitrarily given open neighborhood. Now, genericity of uniqueness means the following:

(i) the modeler can never rule out the case that each player has a unique rationalizable action and

[2]The identification problem in the ex ante stage can be captured by the "uniform topology", which requires the belief at all orders to converge uniformly. In this topology, the above models do not converge as the noise vanishes. For, with noise, the limit of kth-order expectations as $k \to \infty$ is the ex ante expected value of parameters (Samet (1998)).

(ii) whenever the players do have unique rationalizable actions, the modeler could make sure that that is the case by making a sufficiently precise observation (by choosing sufficiently small open sets at sufficiently many orders).

Continuity of a strategy means that the modeler can know what the player will play according to the strategy if his observation is sufficiently precise. Then, genericity of continuity and uniqueness implies that the modeler can never rule out the possibility that he could have learned what the players will play according to rationalizability, by making a more precise observation. In that sense, the rationalizability is a strong solution concept.

Genericity of uniqueness provides a new perspective on refinements of rationalizability (and equilibrium). Towards establishing this, I further show that, given any finite type space and any rationalizable strategy in that type space, one can slightly perturb the players' perceptions about the payoffs to obtain a nearby dominance-solvable model in which the given strategy is uniquely rationalizable. (For each type in the original model, there will be a type in the dominance-solvable model whose beliefs are arbitrarily close to that of the original type for arbitrarily many orders.) We can therefore regard a finite type space as a model that summarizes many indistinguishable situations by abstracting away from the details that are used in computing the beliefs at very high orders. By specifying these details, one could make any rationalizable strategy uniquely rationalizable. In the detailed model, one must take the unique rationalizable strategy as the only prediction, no matter what refinement of rationalizability (or equilibrium) he believes in. Therefore, when one refines rationalizability by ignoring some rationalizable strategies, he simply ignores the dominance-solvable models that are indistinguishable from the model at hand but lead to the ignored strategies as unique solutions. In that sense, refinement is a selection among payoff and information structures, rather than an epistemic issue.

The last result leads to extensions of two seemingly opposing results. Firstly, I extend, in a weaker form, the results of Carlsson and Van Damme (1993) and Frankel, Morris, and Pauzner (2003) for supermodular games to all finite-action games. For supermodular games of complete information, they showed that any perturbation within a canonical class leads to a dominance-solvable model—except for the degenerate signal values at which the strategies jump. For arbitrary finite-action games with arbitrary payoff and information structures (with possibly infinite type spaces), I show that there exists a perturbation that leads to a nearby dominance-solvable model. The dominance-solvable model will remain so, when further small perturbations are introduced. This suggests that multiplicity will become rare as we allow higher-order uncertainty at all levels. (As we successively introduce higher-order uncertainty in the form of "small" noise, the domain of dominance-solvability will grow, while the domain of multiplicity will shrink.)

Second, extending a discontinuity result of Weinstein and Yildiz (2004) for equilibrium, I obtain a characterization: a rationalizable strategy is continuous at a type that lies in a finite type space if and only if there is a unique rationalizable action for that type. At such a type, either all rationalizable strategies are continuous, or all of them are discontinuous.

Some may find the above discussion misleading. The examples refer to models with a common prior, while the universal type space contains the models without a common prior as well. The above counterintuitive results may be due the latter models. Using a result by Lipman (2003), I show that all of the above results remain intact if we restrict ourselves to the finite models with common prior. In particular, the nearby dominance-solvable models for the finite type spaces can be taken as part of a larger finite type space with a common prior, and the above characterization of continuity with dominance-solvability remains intact even if we restrict the domain of the strategies to types with a common prior.

In the next section I provide examples of nearby dominance-solvable models for 2×2 games. 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. Examples

In this section, using 2×2 games, I will illustrate how multiplicity disappears when incomplete information is introduced. I will first consider the games with multiple equilibria—analyzed by Carlsson and van Damme (1993).

Example 1. Consider the 2×2 game

	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. Without incomplete 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 $x_i > 1/2$.

While multiplicity holds for the fragile case of $\varepsilon = 0$, uniqueness prevails in an open set of parameter $\varepsilon > 0$, and for an open set of parameters for the distributions and so on. This is a reflection of a more general fact that dominance-solvability holds for an open set in the universal type space. When the type space is finite, the degenerate signal values with multiplicity, such as $x_i = 1/2$, are also easily avoided—as in the next example.

EXAMPLE 2. In the previous example, consider the case that $\theta \in \Theta = \{\theta_0, \theta_1, \ldots, \theta_{M-1}\}$, where $\theta_0 = -\varepsilon/2$, $\theta_1 = \varepsilon/2$, $\theta_2 = 3\varepsilon/2, \ldots$, $\theta_{M-1} = \bar{\theta}$ for some fixed $\bar{\theta} > 1$. Ex ante, θ is distributed uniformly on Θ. The players observe θ with noise: if $\theta = \theta_m$, then $x_i = \theta_{m-1}$ with probability $1/2$ and $x_i = \theta_{m+1}$ with probability $1/2$.[3] The signals are independent conditional on θ. In this model it is common knowledge that the signal values are in ε neighborhood of true value. Moreover, for any signal value x_i and any integer $k > 0$, it is mutually known at the kth-order that $\theta \in [x_i - (2k+1)\varepsilon, x_i + (2k+1)\varepsilon]$. Hence, as $\varepsilon \to 0$, the players' kth-order beliefs converge to that of kth-order mutual knowledge of $\theta = x_i$. As $\varepsilon \to 0$, the game converges to a model where $\theta \in [0, \bar{\theta}]$ becomes common knowledge before the players take their action. The limit game is characterized by multiple equilibria. But when $\varepsilon > 0$, except for a nowhere-dense set of parameter values for which $\theta_m = 1/2$ for some m (i.e., $(\bar{\theta} - 1/2) / [(\bar{\theta} + \varepsilon/2)\varepsilon]$ is an integer), the game is dominance-solvable with the unique rationalizable strategy

$$s_i^*(x_i) = \begin{cases} \alpha_i & \text{if } x_i > 1/2 \\ \beta_i & \text{if } x_i < 1/2. \end{cases}$$

(One can easily check this starting from the two ends.)

Since the generic 2×2 games with unique equilibrium in pure strategies are dominance-solvable already, the above examples cover all 2×2 games, except for the games with no equilibrium in pure strategies, such as the Matching-Penny game. In such a game, if the dominance considerations had led to a unique strategy for a player i and there were no payoff uncertainty, then his opponent would foresee the i's action and play a best response, against which i would have wanted to play another strategy. Every action is rationalizable in these games. The introduction of incomplete information in the above form does not render these games dominance-solvable.

[3] I use the covention that $\theta_{-1} = \theta_0 - \varepsilon$ and $\theta_M = \theta_{M-1} + \varepsilon$. In previous formulation, η_i is the random variables that takes values 1 and -1 with equal probabilities of $1/2$.

Nevertheless, the next example shows that these games, too, can be perturbed to obtain a dominance-solvable model—using a different belief structure.

Example 3 (Matching Pennies—without a common prior). Now, 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 belief 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 $(x_i, x_j) = (\theta_{m-1}, \theta_m)$, where $\gamma \in (0, \varepsilon / [2(1 - \varepsilon)])$. As before, 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 for the open set $\{(\varepsilon, \gamma) | 0 < \gamma < \varepsilon / [2(1 - \varepsilon)]\}$ of parameters, 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. 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. The new game will be dominance-solvable from these types' point of view, 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 $\theta = x_i$. Moreover, one can check 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 $y_i(k) \leq 2K$.[4] That is, the posterior beliefs in the new model are identical to that 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.[5] In particular, the game is dominance-solvable from the point of view of the types with $(x_i, y_i(1))$, which approximate the complete-information model.

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, complete and separable 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 endowed with the universal type space. A type of a player i

[4] To see this, notice that, for $y_1(k) = 1$, player 1 knows that $k = 1$, and $\bar{\mu}$ is proportional to μ_1. For $y_1(k) = 2$, player 2 knows that $k \in \{1, 2\}$, and $\bar{\mu}(\cdot, 1) + \bar{\mu}(\cdot, 2) = L^2 \alpha \mu_2(\cdot)$ is proportional to $\mu_2(\cdot)$. For any $y_i(k) \in \{3, \ldots, 2K\}$, player i knows that $k \in \{y_i(k) - 1, y_i(k)\}$, and $\bar{\mu}(\cdot, y_i(k) - 1) + \bar{\mu}(\cdot, y_i(k)) = \left(L^2 - 1\right) \alpha L^{y_i(k)-2} \mu_i(\cdot)$, which is proportional to $\mu_i(\cdot)$. (See Lipman (2003) for a complete proof.)

[5] 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 $s_i^*(\theta_m)$.

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.[6] 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 $\kappa_{t_i} \in \Delta(\Theta \times T_{-i})$ 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 $s_i : T_i \to A_i$.[7] 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

[6]**Notation:** In general, I write $x = (x_1, \ldots, x_n) = (x_i, x_{-i}) \in X = X_1 \times \cdots \times X_n$ and $X_{-i} = \prod_{j \neq i} X_j$. Similarly, for functions, I write $f(x) = (f_1(x_1), \ldots, f_n(x_n))$ and $f_{-i}(x_{-i}) = (f_1(x_1), \ldots, f_{i-1}(x_{i-1}), f_{i+1}(x_{i+1}), \ldots, f_n(x_n))$, and for set-valued functions, I write $F[x] = F_1[x_1] \times \cdots \times F_n[x_n]$ and $F_{-i}[x_{-i}] = \prod_{j \neq i} F_j[x_j]$.

[7]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, 200x).

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*

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

and $\theta^{a_i} \neq \theta^{a_i'}$ whenever $a_i \neq a_i'$.

That is, the space of possible payoff structures is rich enough so that each action can be strictly dominant for some parameter value. In developing a unified theory of games, one would want to avoid a priori restrictions on the domain of payoff structures. When there are no such restrictions and the actions represent the strategies in a one-shot, simultaneous-move game, Assumption 1 is automatically satisfied. When actions represent the strategies in a dynamic game, a player will be indifferent between any two strategies that differ only on information sets that are ruled out by the strategies themselves, assuming that the player believes that he does not make any mistake (or does not "tremble") in playing these strategies. Hence, Assumption 1 may appear to rule out all these games. But it is possible that the player thinks that each player may make a mistake at each information set with positive probability, as game theorists typically assume in their analyses of such games. The latter case is modeled by another game. In that case, he will not necessarily be indifferent between those strategies. Indeed, Assumption 1 will be satisfied for a reduced-form representation, if one does not rule out the possibility of such mistakes a priori and allows all payoff vectors at terminal nodes.

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 $a_i \in BR_i\left(\text{marg}_{\Theta\times A_{-i}}\pi\right)$ for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\text{marg}_{\Theta\times T_{-i}}\pi = \kappa_{t_i}$ 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 $k-1$. (As described in Footnote 6, 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). The set of rationalizable strategies is denoted by R_i, and $R = R_1 \times \cdots \times R_n$.

Remark 2. When there is incomplete information, rationalizability can be defined in many different ways, leading to different sets of rationalizable strategies. I use a version of interim correlated rationalizability (Battigalli (2003), Battigalli and Siniscalchi (2003) and Dekel, Fudenberg, and Morris (2003)). The interim correlated rationalizability is the weakest among the known interim 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 both positive generic uniqueness and the negative discontinuity results; these results will remain valid under any stronger notion of rationalizability. To simplify the exposition, I formulate rationalizability slightly different from Dekel, Fudenberg, and Morris (2003). They define rationalizability through the beliefs on functions $f : \Theta \times T_{-i} \to A_{-i}$, rather than $\Theta \times T_{-i} \times A_{-i}$. By definition, my formulation, if anything, allows more actions to be rationalizable, which only strengthens my results, and the two formulations are equivalent when κ_{t_i} has a finite support.

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 $\partial T'$. Clearly, topological notions of genericity may widely differ from measure theoretical notions of genericity. Since this paper is about the topological properties of rationalizable strategies, the topological notions seem to be appropriate. (To see how these notions are related, see Oxtoby (1980).)

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 $\kappa_{t_i} = p\left(\cdot|\Theta \times \{t_i\} \times T'_{-i}\right)$ 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, $T_i^{CPA} = \{t_i \in T'_i | T'$ 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)).[8]

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 $m \to \infty$.*

DEFINITION 5 (Continuity). A strategy s_i is said to be *continuous at t_i* iff

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

for each sequence of types $t_{i,m}$. Since A_i is endowed with the discrete topology, if s_i is continuous at t_i, then s_i is constant on a 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 endowed with the discrete topology, F is upper semicontinuous iff each t has a neighborhood η with $F[t'] \subseteq F[t]$ for each $t' \in \eta$.

LEMMA 3 (Dekel, Fudenberg, and Morris (2004)). *S^∞ is non-empty and upper-semicontinuous.*

[8] Lipman (2003) uses a partitional model. If one takes $\Omega = \Theta \times T$ as the state space and $\left\{\{(\theta, t)\,|t_i = \tilde{t}_i\}\,|\tilde{t}_i \in T_i\right\}$ 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.

Dekel, Fudenberg, and Morris (2004) proves upper-semicontinuity of interim correlated rationalizability in their framework. 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,*

(1) *U_F is open;*
(2) *there exists a continuous function $f^* : U_F \to A$ such that $F[t] = \{f^*(t)\}$ for each $t \in U_F$, and*
(3) *for any function $f : T \to A$, if $f(t) \in F[t]$ for each t, then f is continuous on 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. Finally, any f as in part 3 coincides with f^* on the open neighborhood η and hence is continuous at t. □

4. RESULTS

In this section, I analyze the continuity and uniqueness properties of rationalizable strategies. I show that, generically, there exists a unique rationalizable action. Whenever there is a unique rationalizable action for a type, every rationalizable strategy is continuous at that type. For finite types, I show that the converse is also true: a rationalizable strategy is continuous at a finite type if and only if there is a unique rationalizable action for that type. I further show that for any model, there is a perturbation that leads to a dominance-solvable model. Using these results, I then present characterizations for the topologies generated by rationalizable strategies. The next result will be the main tool for this analysis.

LEMMA 5. *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 dominance-insolvable) 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. **Equivalence of continuity and uniqueness.**

PROPOSITION 1. *Under Assumption 1, a rationalizable strategy $s_i \in R_i$ is continuous at a finite type $\hat{t}_i \in \hat{T}_i$ if and only if $\hat{t}_i$ has a unique rationalizable action, i.e., $\left|S_i^{\infty}\left[\hat{t}_i\right]\right| = 1$. This characterization remains intact if the domain of strategy profiles is restricted to T^{CPA}, by imposing the common prior assumption.*

Proof. The "if" part immediately follows from Lemma 3 and part 3 of Lemma 4. To prove the "only if" part, take any rationalizable strategy s_i and any $a_i \in S_i^{\infty}\left[\hat{t}_i\right]$ with $s_i\left(\hat{t}_i\right) \neq a_i$. By Lemma 5, there exists a sequence of types $t_{i,m} \to \hat{t}_i$ with $s_i\left(t_{i,m}\right) \in S_i^{\infty}\left[t_{i,m}\right] = \{a_i\}$. Since $s_i\left(t_{i,m}\right) = a_i$ for each $t_{i,m}$, $s_i\left(t_{i,m}\right)$ does not converge to $s_i\left(\hat{t}_i\right)$. To prove the last statement of the proposition, one picks $t_{i,m} \in T_i^{CPA}$ (by Lemma 5). □

Proposition 1 establishes that, at a finite type, either all rationalizable strategies are continuous, or all of them are discontinuous. The set of finite types can be put into two groups. For the types in one group, the game is "dominance-solvable", and all rationalizable strategies are continuous at these types. For the types in the other group, there are multiple rationalizable actions, and each rationalizable strategy is discontinuous at each type in this group. Since there are typically multiple rationalizable actions, the finite types in applications typically fall into the second group. Assumption 1 is not superfluous; without Assumption 1, some rationalizable strategy may be continuous at a type with multiple rationalizable actions.

Under weaker assumptions, Weinstein and Yildiz (2004) have shown that every equilibrium is discontinuous at a type for which multiple actions survive iterated elimination of strategies that are never a strict best reply. Proposition 1 drops the equilibrium and strictness requirements in their conclusion. This extension is important because equilibrium need not exist in general, and in some important games, such as perfect-information games, there are multiple rationalizable actions, but no action survives the elimination process above. The strictness requirement is not binding in generic complete-information games.

4.2. **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 5 implies that U is dense in universal type space. Since S^{∞} is upper-semicontinuous, U is also open. This yields the 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 2. *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 part 1 of Lemma 4, U is open. To show that U is dense, first observe that, by Lemma 5, 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 part 2 of Lemma 4, there exists a continuous function $s^*: U \to A$ with $S^\infty[t] = \{s^*(t)\}$ for each $t \in U$. The last part of the proposition is by part 3 of Lemma 4. □

By Proposition 2, we can partition the universal type space to an open and dense set U and its boundary $T \backslash U$. On U, each type has a unique rationalizable action, and every rationalizable strategy is continuous. On the boundary, each type has multiple rationalizable actions. By Proposition 1, every rationalizable strategy is discontinuous at each finite type on the boundary. 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 2 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 $t \in \partial U^a \cap \partial U^{a'}$, 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.[9]

In summary, Proposition 2 establishes that, if one excludes a nowhere-dense set of types, then there will be a unique unified theory of rational behavior 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 U, where the unique unified theory above 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, 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.

[9]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 5 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.

Also, the result is true for a (strong) topological notion of genericity with respect to a (canonical) topology. As discussed earlier, it need not be true for other notions of genericity. This caveat applies the following remarks as well.

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 2 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.

REMARK 4 (Epistemic Types). In a strategic situation, a player's beliefs can be put into two groups: the beliefs regarding the payoffs, called Harsanyi type, and the beliefs regarding the payers' actions, called epistemic type. In the traditional methodology, pioneered by Harsanyi, one specifies the former beliefs as parts of the problem and infers the latter beliefs, as parts of the solution, from the former using rationality postulates. In traditional type spaces, there are often a multitude of epistemic types consistent with a given Harsanyi type and rationality. In epistemic literature, the distinction between these two types has been blurred. Proposition 2 establishes that there is indeed a unique epistemic type for a given generic Harsanyi type if we assume that players are rational throughout the model. Hence, under common knowledge of rationality, generically, there is no distinction between Harsanyi types and epistemic types, and a player's Harsanyi type uniquely determines both the decision problem and its solution.

Remark 5 (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 2 implies that, if we assume common knowledge of rationality, then we can have only one unified theory for generic cases, and each of his 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.

4.3. **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. The game is dominance-solvable from this player's point of view. In that sense, one can find "dominance-solvable" games nearby any economic model, although it may be difficult to describe the belief structures in these games. I will now show that one can indeed find a nearby dominance-solvable model in the usual sense.

Proposition 3. *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 5, 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 $t \in T'$, $\tau(t,m) \in T^m$. $\square$

Proposition 3 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 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.[10] 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.

Proposition 4. *Under Assumption 1, for any finite model* $T' \subseteq \hat{T}$, *any rationalizable strategy profile* $s_{T'} : T' \to A$ *with* $s_{T'}(t) \in S^\infty[t]$ *for each* $t \in T'$, *and any integer* m, *there exist 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 with full support,*
(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 5, 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

[10] 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.)

Proposition 3, 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 $t \in T'$, for any distinct t, t', $\tau(t, s_{T'}, m) \neq \tau(t', s_{T'}, m)$ whenever $m > \bar{m}$ for some $\bar{m}$. Since T' is finite, $\bar{m}$ can be chosen uniformly for all types, so that $\tau(\cdot, s_{T'}, m)$ is one-to-one for $m > \bar{m}$. (Since we can change the index m, we can assume that $\tau(\cdot, s_{T'}, m)$ is one-to-one for all m without loss of generality.)

Since $T^{s_{T'},m}$ is finite, by Lemma 9 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 $\tilde{\tau}(\cdot, s_{T'}, m) = \tau'(\cdot, m) \circ \tau(\cdot, s_{T'}, m)$. □

4.4. **Strategic Equivalence and Strategic Topologies.** Now, I will use the previous results to characterize the strategies under which the rationalizability correspondence and the rationalizable strategies are continuous.[11] I will show that these topologies are closely related to the product topology.

Fix a player i. Define

$$U^{a_i} = \{t_i | S_i^\infty[t_i] = \{a_i\}\}, \qquad a_i \in A_i \tag{4.1}$$

$$M_i = T_i \backslash \textstyle\bigcup_{a_i \in A_i'} U^{a_i}.$$

Under Assumption 1, by Lemma 3, each U^{a_i} is open, and by Proposition 2, M_i is nowhere-dense, consisting of the boundaries ∂U^{a_i} of open sets U^{a_i}. Each strategic topology above will be generated by the sets U^{a_i} and some partition of M_i. In a way, the latter partition will be formed by partitions of the boundaries ∂U^{a_i}.

Let $\mathcal{T}_i^S$ be the topology on T_i generated by the rationalizability correspondence; $\mathcal{T}_i^S$ is the smallest topology on T_i with respect to which S_i^∞ is continuous. $\mathcal{T}_i^S$ is

[11] A correspondence $F : T \to 2^A$ is *lower semicontinuous* iff for each $a \in F[t]$ and each sequence $t(m) \to t$, there exists $\bar{m}$ such that $a \in F[t(m)]$ for each $m > \bar{m}$. A correspondence is *continuous* if it is both upper and lower semicontinuous.

the smallest topology that contains all the sets of the form

$$T_i^{B_i} = \{t_i | S_i^\infty [t_i] = B_i\}, \qquad B_i \subseteq A_i.$$

Clearly, the set

$$\mathcal{P}_i^S = \left\{T_i^{B_i} | B_i \subseteq A_i\right\} \setminus \{\varnothing\}$$

is a partition of T_i. The topology $\mathcal{T}_i^S$ is the set of all sets that can be written as the union of some sets in $\mathcal{P}_i^S$ and the empty set. If $T_i' \in \mathcal{T}_i^S$, then $T_i \backslash T_i' \in \mathcal{T}_i^S$. Since $U^{a_i} = T_i^{\{a_i\}} \in \mathcal{P}_i^S$ for each a_i, $\mathcal{P}_i^S$ consists of the open sets U^{a_i} and a partition of M_i. Types t_i and t_i' are said to be *strategically equivalent* iff $S_i^\infty [t_i] = S_i^\infty [t_i']$. Equivalently, t_i and t_i' are strategically equivalent iff $t_i, t_i' \in T_i^{B_i}$ for some $T_i^{B_i} \in \mathcal{P}_i^S$. Finally, let $\hat{\mathcal{T}}_i^S = \left\{T_i' \cap \hat{T}_i | T_i' \in \mathcal{T}_i^S\right\}$ be the relative topology on $\hat{\mathcal{T}}_i^S$, which defines strategic equivalence for finite types and continuity of the correspondence S_i^∞ on $\hat{T}_i$.

PROPOSITION 5. *There exists a family of closed sets C^{a_i}, $a_i \in A_i$, such that $\mathcal{T}_i^S$ is the smallest topology that contains both C^{a_i} and its complement for each a_i. Moreover, under Assumption 1, $\hat{\mathcal{T}}_i^S$ is the smallest topology that contains both $\overline{U^{a_i}} \cap \hat{T}_i$ and $\hat{T}_i \backslash \overline{U^{a_i}}$ for each a_i, where U^{a_i} is as defined in (4.1).*

Proof. Define

$$C^{a_i} = \{t_i | a_i \in S_i^\infty [t_i]\}, \qquad a_i \in A_i.$$

Since S_i^∞ is upper-semicontinuous, C^{a_i} is closed. Let $\mathcal{T}_i$ be the smallest topology that contains both C^{a_i} and $T_i \backslash C^{a_i}$ for each a_i. By definition, for each a_i, $C^{a_i} = \bigcup_{B_i \subseteq A_i} T^{B_i \cup \{a_i\}} \in \mathcal{T}_i^S$. Since C^{a_i} is in $\mathcal{T}_i^S$, so is $T_i \backslash C^{a_i}$. Hence, $\mathcal{T}_i \subseteq \mathcal{T}_i^S$. On the other hand, for each $B_i \subseteq A_i$, $T^{B_i} = \left(\bigcup_{a_i \in B_i} C^{a_i}\right) \setminus \left(\bigcup_{a_i \notin B_i} C^{a_i}\right) \in \mathcal{T}_i$, showing that $\mathcal{T}_i^S \subseteq \mathcal{T}_i$. Therefore, $\mathcal{T}_i = \mathcal{T}_i^S$. This also implies the second statement in the proposition because, under Assumption 1, by Lemma 5, $C^{a_i} \cap \hat{T}_i = \overline{U^{a_i}} \cap \hat{T}_i$ for each a_i. □

Proposition 5 links the topology $\mathcal{T}_i^S$, generated by the rationalizability correspondence, to the product topology, by showing that $\mathcal{T}_i^S$ is generated by finitely many sets that are closed in the product topology. This is because S_i^∞ is upper-semicontinuous. Under Assumption 1, for the finite types, the link is stronger. The closed sets that generate $\hat{\mathcal{T}}_i^S$ intersect each other only on their boundaries, so that their interiors constitute an open and dense set. $\hat{\mathcal{T}}_i^S$ is formed by partitioning the boundary of each set. This is due to Lemma 5.

Now, let $\mathcal{T}_i^{SS}$ be the topology generated by rationalizable strategies $s_i \in R_i$, so that $\mathcal{T}_i^{SS}$ is the smallest topology on T_i with respect to which all rationalizable strategies are continuous. It is the smallest topology that contains all the sets

$$s_i^{-1}(a_i) = \{t_i | s_i(a_i) = t_i\}, \qquad s_i \in R_i, a_i \in A_i.$$

When there are two distinct rationalizable actions $a_i, a_i' \in S_i^\infty[t_i]$, one can find two rationalizable strategies $\hat{s}_i, \tilde{s}_i \in R_i$ with $\hat{s}_i^{-1}(\hat{s}_i(t_i)) = U^{a_i} \cup \{t_i\}$ and $\tilde{s}_i^{-1}(\tilde{s}_i(t_i)) = U^{a_i'} \cup \{t_i\}$. In that case, $\{t_i\} = \hat{s}_i^{-1}(a_i) \cap \tilde{s}_i^{-1}(a_i') \in \mathcal{T}_i$. Hence, each singleton $\{t_i\}$ is open in topology $\mathcal{T}_i^{SS}$ for each $t_i \in M_i$. That is, $\mathcal{T}_i^{SS}$ is generated by the open sets U^{a_i}, $a_i \in A_i$, and the discrete topology on their boundaries. This topology is closely related to the following strong notion of strategic equivalence: t_i and t_i' are said to be *strongly, strategically equivalent* iff

$$s_i(t_i) = s_i(t_i') \qquad \forall s_i \in R_i.$$

When two types are equivalent in this sense, the players treat these types equivalently under any rationalizable theory. If they are not equivalent, then they will be treated differently by some rationalizable strategy. Let $\mathcal{P}_i^{SS}$ be the partition of T_i associated with this equivalence relation. One can easily show that $\mathcal{T}_i^{SS}$ is the smallest topology that contains $\mathcal{P}_i^{SS}$, and hence $\mathcal{P}_i^{SS}$ is formed of the open sets U^{a_i} and the singletons $\{t_i\}$ with $t_i \in M_i$. But by Proposition 2, under Assumption 1, M_i is a nowhere-dense set, consisting of the boundaries of open sets U^{a_i}. This establishes the following link between $\mathcal{T}_i^{SS}$ and the product topology.

PROPOSITION 6. *Under Assumption 1, there exists a partition $\{U^{a_i} | a_i \in A_i\} \cup \{M_i\}$ of T_i where U^{a_i} is open for each a_i and M_i is a nowhere-dense set (with respect to* *the product topology) and such that*

$$\mathcal{P}_i^{SS} = \{U^{a_i} | a_i \in A_i\} \cup \{\{t_i\} | t_i \in M_i\}, \tag{4.2}$$

$$\mathcal{T}_i^{SS} = \left\{\bigcup_{a_i \in B_i} U^{a_i} \cup T_i' | B_i \subseteq A_i, T_i' \subseteq M_i\right\}. \tag{4.3}$$

That is, we can partition the universal type space into finitely many open sets U^{a_i}, $a_i \in A_i$, and their boundaries ∂U^{a_i} in the following way. Two distinct types are strongly strategically equivalent if and only if they are both in an open set U^{a_i} for some a_i, and a type $t_i \in \partial U^{a_i}$ is strategically equivalent only to itself. $\mathcal{T}_i^{SS}$ is generated by open sets U^{a_i} and the discrete topology on their boundaries.

Strong strategic equivalence is a stringent condition. As the opposite benchmark, consider the weakest form of strategic equivalence: types t_i and t_i' are said to be *weakly strategically equivalent* iff there exists a rationalizable strategy s_i with $s_i(t_i) = s_i(t_i')$. This condition is equivalent to $S_i^\infty[t_i] \cap S_i^\infty[t_i'] \neq \varnothing$. When two types are not weakly strategically equivalent, they are treated differently under every rationalizable theory. The next result shows that this notion of strategic equivalence is also closely related to the closed sets $\overline{U^{a_i}}$ that only intersect each other on their boundaries with respect to the product topology:

PROPOSITION 7. *Under Assumption 1, two finite types $\hat{t}_i$ and $\hat{t}_i' \in \hat{T}_i$ are weakly strategically equivalent if and only if $\hat{t}_i, \hat{t}_i' \in \overline{U^{a_i}}$ for some a_i.*

Proof. If $\hat{t}_i$ and $\hat{t}_i'$ are weakly strategically equivalent, then there exists $a_i \in S_i^\infty\left[\hat{t}_i\right] = S_i^\infty\left[\hat{t}_i'\right]$. Then, by Lemma 5, $\hat{t}_i, \hat{t}_i' \in \overline{U^{a_i}}$. Conversely, if if $\hat{t}_i, \hat{t}_i' \in \overline{U^{a_i}}$ for some a_i, then by Lemma 3, $a_i \in S_i^\infty\left[\hat{t}_i\right]$ and $a_i \in S_i^\infty\left[\hat{t}_i'\right]$, showing that $\hat{t}_i$ and $\hat{t}_i'$ are weakly strategically equivalent. □

Dekel, Fudenberg, and Morris (2004) analyze the topologies under which ε-rationalizability exhibits the basic properties of ε-optimization in usual Euclidean spaces. They call such topologies as strategic topologies.

5. Proof of Lemma 5

Now, I will prove Lemma 5. A substantial part of the proof utilizes the following stronger notion of rationalizability, used also by Weinstein and Yildiz (2004).

Strict 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 $BR_i\left(\text{marg}_{\Theta\times A_{-i}}\pi\right) = \{a_i\}$ for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\text{marg}_{\Theta\times T_{-i}}\pi = \kappa_{t_i}$ 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, $W_i^k \subseteq S_i^k$. For some t_i, $W_i^k[t_i]$ may be empty.[12] Finally, 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 of t_i on functions $f : \Theta \times T_{-i}' \to A_{-i}$ with $f(\theta, t_{-i}) \in V_{-i}[t_{-i}]$. Then, $V_i[t_i] \subseteq W_i^\infty[t_i]$ for each t_i.

The proof of Lemma 5 has three main steps, which are presented as the following three lemmas. The first step (namely, Lemma 6) shows that, when we focus on strictly rationalizable strategies and do not require a common prior, Lemma 5 is true for each $t_i \in \hat{T}_i$. The second step (namely, Lemma 7) 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 5 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 9).

The following lemma is similar to the main result 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

[12] One can show that, if $\Theta = \Theta_1 \times \cdots \times \Theta_n$ where $\Theta_i = [0,1]^A$ for each i, and $u_i(\theta, a) = \theta_i(a)$ for each (i, a, θ), then W^∞ is empty only on category 1 set, i.e., union of countably many nowhere-dense sets.

so that a_i is played by the new type in equilibrium. The lemma states that one can indeed make a_i the only member of $S_i^{k+1}\left[\tilde{t}_i\right]$ for the new type $\tilde{t}_i$. 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 after 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 after $k+1$ rounds of iterated dominance.

LEMMA 6. *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 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 $l \leq k$. 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 $L = \{l | \exists h : (l, h) \in T_{-i}\}$. 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

$$\text{(IH)} \qquad W_{-i}^k\left[\tilde{t}_{-i}[a_{-i}]\right] = S_{-i}^k\left[\tilde{t}_{-i}[a_{-i}]\right] = \{a_{-i}\},$$

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 $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\text{marg}_{\Theta\times T_{-i}}\pi = \kappa_{t_i}$ and $\pi\left(a_{-i} \in S^{k-1}_{-i}[t_{-i}]\right) = 1$. 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

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

the probability distribution induced on $\Theta \times T_{-i}$ by the mapping μ and the probability distribution π. 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, μ is measurable, and $\kappa_{\tilde{t}_i}$ is well-defined. Since μ leaves (θ, l) intact, the first k orders of beliefs (about (θ, l)) are identical under $\hat{t}_i$ and $\tilde{t}_i$:

$$\begin{aligned}
\text{marg}_{\Theta\times L}\kappa_{\tilde{t}_i} &= \text{marg}_{\Theta\times L}\left(\text{marg}_{\Theta\times L\times A_{-i}}\pi\right) \circ \mu^{-1} = \text{marg}_{\Theta\times L}\left(\text{marg}_{\Theta\times L\times A_{-i}}\pi\right) \\
&= \text{marg}_{\Theta\times L}\pi = \text{marg}_{\Theta\times L}\kappa_{\hat{t}_i}.
\end{aligned}$$

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 $a_{-i} \in S^{k-1}_{-i}\left[\tilde{t}_{-i}[a_{-i}]\right]$. Thus, there exists a unique $\tilde{\pi} \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\text{marg}_{\Theta\times T_{-i}}\pi = \kappa_{\tilde{t}_i}$ and $\pi\left(a_{-i} \in S^{k-1}_{-i}[t_{-i}]\right) = 1$; it is $\tilde{\pi} = \kappa_{\tilde{t}_i} \circ \gamma^{-1}$ where $\gamma : \left(\theta, l, \tilde{h}[l, a_{-i}]\right) \mapsto \left(\theta, l, \tilde{h}[l, a_{-i}], a_{-i}\right)$. But $\gamma(\mu(\theta, l, a_{-i})) = \gamma\left(\theta, l, \tilde{h}[l, a_{-i}]\right) = \left(\theta, l, \tilde{h}[l, a_{-i}], a_i\right)$, so that $\text{proj}_{\Theta\times L\times A_{-i}} \circ \gamma \circ \mu$ is the identity mapping, where proj is the projection mapping. Therefore,

$$\begin{aligned}
\text{marg}_{\Theta\times L\times A_{-i}}\tilde{\pi} &= \kappa_{\tilde{t}_i} \circ \gamma^{-1} \circ \text{proj}^{-1}_{\Theta\times L\times A_{-i}} = \left(\text{marg}_{\Theta\times L\times A_{-i}}\pi\right) \circ \mu^{-1} \circ \gamma^{-1} \circ \text{proj}^{-1}_{\Theta\times L\times A_{-i}} \\
&= \text{marg}_{\Theta\times L\times A_{-i}}\pi.
\end{aligned}$$

This, of course, yields

$$\text{marg}_{\Theta\times A_{-i}}\tilde{\pi} = \text{marg}_{\Theta\times A_{-i}}\pi.$$

But a_i is the only best reply to this belief. Therefore, $S^{k+1}_i\left[\tilde{t}_i\right] = \{a_i\}$.

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

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

Clearly, $T^{\tilde{t}_i}$ finite. For any $t_j \in T_j^{\tilde{t}_i} \backslash \left\{\tilde{t}_i\right\}$, $t_j \in T_j^{t_{-i}[a_{-i}]}$ 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 $t_j \in T_j^{\tilde{t}_i} \backslash \left\{\tilde{t}_i\right\}$, $\left|S^k[t_j]\right| = 1$, 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 $a_i \in W_i^\infty\left[\hat{t}_i\right] \subseteq W_i^k\left[\hat{t}_i\right]$, 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 each k, $t_{i,m}^k = \hat{t}_i^k$ for each $m > k$, showing that $t_{i,m}^k \to \hat{t}_i^k$. 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 within a finite model is strictly rationalizable within a nearby finite model.

Lemma 7. *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 $\tau_{-i}(t_{-i}, a_{-i}, m)$. Now, since $a_i \in S_i^\infty[t_i]$, there exists a belief $\pi^{t_i,a_i} \in \Delta\left(\Theta' \times T'_{-i} \times A_{-i}\right)$ 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 $\text{marg}_{\Theta \times T_{-i}}\pi^{t_i,a_i} = \kappa_{t_i}$. 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 the new type $\tau_i(t_i, a_i, m)$, define the belief

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

where $\gamma : (\theta, \tau_{-i}(t_{-i}, a_{-i}, m)) \mapsto (\theta, \tau_{-i}(t_{-i}, a_{-i}, m), a_{-i})$, which is generated by 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, $\text{proj}_{\Theta \times A_{-i}}(\gamma(\hat{\tau}_{-i,m}(\theta, t_{-i}, a_{-i}))) = (\theta, a_{-i})$, and 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, $a_i \in W_i^\infty[\tau_i(t_i, a_i, m)]$ 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

$$\begin{aligned} \tau_i^1(t_i, a_i, m) &= \text{marg}_\Theta \kappa_{\tau_i(t_i,a_i,m)} = \frac{1}{m}\delta_{\theta^{a_i}} + \left(1 - \frac{1}{m}\right)\text{marg}_\Theta \pi^{t_i,a_i} \\ &\to \text{marg}_\Theta \pi^{t_i,a_i} = \text{marg}_\Theta \kappa_{t_i} = t_i^1. \end{aligned}$$

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,

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

[To obtain the penultimate equality, observe that $\text{proj}_{\Theta \times L}(\hat{\tau}_{-i,m}(\theta, t_{-i}, a_{-i})) = \text{proj}_{\Theta \times L}(\theta, \tau_{-i}(, t_{-i}, a_{-i}, m)) = \left(\theta, \tau_{-i}^{k-1}(, t_{-i}, a_{-i}, m)\right) \to \left(\theta, t_{-i}^{k-1}\right)$.]

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', $\theta^{a_i} \neq \theta^{a_i'}$, 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 $\bar{m}$ uniform. □

In the previous lemma, if the original model T' is dominance-solvable, then the new model will also be dominance-solvable, and the rationalizability and the strict rationalizability will coincide in the new model. This is stated in the next lemma.

LEMMA 8. *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 7. 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 Lemma 3 and part 2 of 4, there exists $\bar{m}$ such that for each $m > \bar{m}$, $S^\infty[\tau(t, m)] = S^\infty[t]$. Since T' is finite, $\bar{m}$ is uniform for all t. Moreover, by Lemma 7, $S^\infty[t] = W^\infty[\tau(t, m)]$. □

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.

LEMMA 9. *Under Assumption 1, for any finite, dominance-solvable model $T' \subseteq \hat{T}$ and any m, there exist a finite, dominance-solvable 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 8, 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 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) $\text{supp}\left(\kappa_{\bar{\tau}_i(\tilde{\tau}_i(t_i,m),k)}\right) = \tilde{\Theta} \times T_{-i}^{m,k}$ (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 Lemma 3 and part 2 of 4, 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 $l > \bar{l}$, $S^{\infty}[\hat{\tau}(\bar{\tau}(\tilde{\tau}(t,m), k), l)] = S^{\infty}[t]$, and $T^{m,k,l}$ is dominance-solvable. 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 5. Take any $\hat{t} \in \hat{T}$, and any $a \in S^{\infty}[\hat{t}]$. By Lemma 7, 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 6, 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 $T^m = T^{m,m}$ fit the bill. Now suppose we need a common prior. By Lemma 9, 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 $T^m = T^{m,m,m}$. □

6. Conclusion

Usual game theoretical models typically have a multitude of rationalizable actions. The predictions of these models then crucially depend on the model's assumptions about the players' beliefs—except for the few predictions that are true for all rationalizable strategies. The multiplicity may be, however, a property of the present models, rather than a property of rational behavior. Indeed, theoretically, rationalizability generically leads to quite robust predictions: there exists a unique rationalizable outcome, and it is continuous with respect to the players' beliefs.

The finite models accomplish what one would expect from a model. Each of them summarizes 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. In that case, when refining rationalizability, a researcher ought to explain why he can rule out those nearby payoff and information structures that are nearly indistinguishable from his model at the interim stage, rather than providing epistemic arguments for the refinement.

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 $B_i^k : \Delta\left(\Theta \times Gr\left(\hat{S}_{-i}^{k-1}\right)\right) \to 2^{A_i}$ by

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

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 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, $\text{marg}_{\Theta \times T_{-i}}\pi$ 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 $\phi : \Delta\left(\Theta \times T_{-i}\right) \to T_i$. 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 $S_i^k[t_i]$ 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] 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] 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`