PURIFICATION OF BAYES NASH EQUILIBRIUM WITH CORRELATED TYPES AND INTERDEPENDENT PAYOFFS

PAULO BARELLI AND JOHN DUGGAN

Abstract. We establish purification results for Bayes-Nash equilibrium in a large class of Bayesian games with finite sets of pure actions. We allow for correlated types and interdependent payoffs and for type-dependent feasible action sets. The latter feature allows us to prove existence and purification results for pure Bayes-Nash equilibria in undominated strategies. We give applications to auctions, global games, and voting to illustrate the usefulness of our results.

1. Introduction

Purification is a potentially powerful tool for obtaining existence of a pure strategy Bayes-Nash equilibrium (BNE) in games of incomplete information. It ensures, under non-atomicity of the underlying distribution of types and some other regularity conditions, that for every BNE in mixed strategies, there exists an equivalent BNE in pure strategies. Thus, insofar as existence of a BNE in mixed strategies has been established in great generality (cf. Balder (1988, 2002)), it suffices that one verifies the conditions needed for purification for a BNE in pure strategies to exist. Unfortunately, the extra regularity conditions for purification provided in the literature are quite restrictive: in particular, types are required to be independent conditional on a finite environmental state variable, and no correlation or interdependence of payoffs in addition to this environmental state variable is allowed. A further limitation of this approach is that known existence results do not adhere to common refinements used in applied modeling, such as the requirement that players use undominated strategies. We address these issues by proving a purification result that allows for general forms of correlation of types and of interdependence of payoffs, and by establishing existence (and purification) in the class of undominated BNE. We illustrate our results with applications to auctions, global games, and voting with incomplete information.

Date: May 28, 2015.

We thank anonymous referees and the advisory editor for comments that improved the presentation of our results.

We consider games of incomplete information among n players who choose from finite sets of pure actions, in which the types of each player i can be decomposed into two components, t_i and u_i. The first is a general component that affects the payoffs of every player; the second is a private-value component that affects only player i's payoffs and moreover is conditionally independent given t_i. We allow the general type components $(t_1, \ldots, t_n)$ to be highly correlated (subject only to the standard diffuseness condition of Milgrom and Weber (1985)), and because of this, our framework encompasses models in which players receive conditionally independent signals of an underlying general state variable in addition to some conditionally independent information (e.g., preference shocks). We provide two results. First, we establish purification, and thereby existence of pure strategy BNE in a large class of models with correlated types and interdependent payoffs. Second, we sharpen the first result by proving that an undominated BNE exists, and that for every such equilibrium, there is an equivalent pure-strategy BNE that also satisfies the refinement. Interestingly, the latter argument relies essentially on our initial purification result: after eliminating undominated pure actions, we obtain a mixed strategy BNE that assigns positive probability only to undominated pure actions, but this mixed BNE may well be dominated. We address this issue by purifying the mixed BNE so that, by construction, players choose only undominated pure actions. Our results deliver pure strategy undominated BNE in applications where players have correlated information that is payoff-relevant for other players, and in particular to applications such as auctions and voting where such BNE are especially important.

Our approach to purification is as follows. Given type (t_i, u_i), we assume player i mixes over a finite set of pure actions. A mixed strategy σ_i assigns a mixed action to each realization (t_i, u_i). Integrating out the private component of the type u_i, we obtain the average action $\gamma_i(t_i|\sigma_i) = \int \sigma_i(t_i, u_i)du_i$. Because payoffs are multilinear in mixed actions, the interim payoff function of player i depends on i's own action and type and on the profile of average actions $\gamma_{-i} = (\gamma_j)_{j\neq i}$. Thus, the set of interim best response actions can be written as $M_i(t_i, u_i; \gamma)$. Now consider the set $\mathscr{G}_i(\gamma)$ of selections from the correspondence $t_i \to \int M_i(t_i, u_i; \gamma)du_i$ from general types to interim best response actions. We show that each fixed point $\gamma^* \in \mathscr{G}(\gamma^*)$ (there is at least one in our framework) corresponds to a class of equivalent BNE; and using the assumption that the private-value components are non-atomically distributed, there is necessarily a pure BNE belonging to this class.[1] Our results for undominated BNE follow from a proposition showing that the set of interim undominated pure actions for a player varies with respect to types in a measurable way. We then

[1]The techniques used here are similar to those used by Duggan (2012) in establishing existence of stationary Markov perfect equilibria in a class of stochastic games.

incorporate the restriction to undominated pure actions into the players' feasible sets of actions and apply our existence and purification theorem, recalling the important role (described in the previous paragraph) of purification in the existence of an undominated BNE.

There are substantial literatures on purification and pure strategy existence that impose more structure. Radner and Rosenthal (1982) provide an early purification result, also assuming a decomposition of types, but they assume the t_i's live in a finite set, that the u_i's are (unconditionally) independent, and that payoff interdependence and correlation of types is at most finite dimensional.[2] Other results in the purification literature (e.g., Milgrom and Weber (1985), Khan, Rath, and Sun (2006), Fu et al. (2007), and Balder (2008)) explicitly assume a finite environmental variable t_0, and that player types are private values and independent conditional on t_0. As such, only the special form of correlation and interdependence of payoffs mediated by the finite state variable t_0 is allowed. In contrast, we allow the general type components $(t_1, \ldots, t_n)$ to be highly correlated (subject to diffuseness) and arbitrarily payoff relevant for all players, which permits a general space of environmental states. It should be noted that there are many recent purification results that rely on extremely diffuse environments (e.g., Podczeck (2009) and Wang and Zhang (2010)) and obtain purification results for games with general action spaces. As with the standard purification literature, this agenda does not allow for general forms of correlation of types or interdependence of payoffs, and its practical usefulness is limited by the assumption of super non-atomicity. Moreover, as recently shown by Greinecker and Podczeck (2013), purification in extremely diffuse environments is spurious as the resulting purified strategies still involve mixing. There is also a sizable literature on existence of pure strategy BNE based on modularity ideas (cf. Athey (2001), McAdams (2004), Van Zandt and Vives (2007), Van Zandt (2010), Reny (2011)), including recent work by de Castro (2012), who uses a decomposition of types that is different from ours.[3] As our applications show, our approach ensures existence of pure strategy BNE in models previously studied in this literature without the modularity restrictions needed to deduce equilibria with a monotone structure.

[2]The use of types with more than one component is a common feature in the purification literature; see, e.g., Khan, Rath, and Sun (2006), Fu et al. (2007), or de Castro (2012).

[3]As a sufficient condition for his richness condition, de Castro (2012) uses a type decomposition in which one component is a private-value type that determines a player's preferences, and the other component is a payoff-irrelevant belief type; moreover, he assumes a partial ordering of preference types and a separability condition on belief types. Our decomposition assumes a private-value component, but the remaining component is completely general, and we do not impose any ordering or separability conditions.

Less is known about existence of undominated BNE in pure or mixed strategies, and in fact simple examples (cf. Simon and Stinchcombe (1995)) illustrate that infinite normal form games may not admit undominated equilibria. The normal form of the Bayesian games we study are infinite games, but the product structure on type sets and non-atomicity of the private-value component allow us to circumvent existence counterexamples. Despite terminological similarity, our results are unrelated to those of Le Breton and Weber (1997) and Balder (2003), who consider large non-atomic games and show existence of an equilibrium that is "undominated" in the sense that it is not Pareto dominated by any other equilibrium.

The paper is organized as follows. Section 2 presents a number of applications to illustrate the practical usefulness of our results. In Section 3, we present the product Bayesian game framework, and we state our results in Section 4. Proofs of the main results are in Section 5.

2. Applications

We present several examples of models from economics and political science to illustrate the applicability of our results and to demonstrate that the decomposition of types we assume is quite natural in many special cases. We assume for simplicity that the players' general types t_i are conditionally independent signals of an underlying state s,[4] which also enters into payoffs, and that each player's payoff from an outcome (e.g., an allocation of objects in an auction or the winner of an election) is subject to an idiosyncratic shock u_i. In each example, we establish existence and purification of undominated BNE with less structure than currently imposed in the literature. The natural modeling assumption in the examples (and the one commonly used in the modeling applications) is that the state is drawn from a continuous distribution, so that previous purification results cannot be applied to these problems.

Example 2.1: *Auctions* Consider a multi-unit discriminatory auction for bonds, drilling rights, etc. Assume that there are n bidders, indexed by $i = 1, \ldots, n$, who each submit a vector $b_i = (b_i^1, \ldots, b_i^d)$ of bids, where each b_i^k lives in a finite set $B \subset [0,1]$, representing the possible bids for each of the $d > 0$ units. The auctioneer collects bids $b = (b_1, \ldots, b_n)$, and units are allocated according to an outcome function that determines a set of objects, $g_i(b_1, \ldots, b_n)$, for each bidder and may involve ties (e.g., second highest bid on each unit), and prices are given by $p_i(b_1, \ldots, b_n)$.

[4]Information structures generated in this manner automatically satisfy Milgrom and Weber's (1985) diffuseness condition.

Let s be an unobservable state variable selected by nature that lives in a complete, separable metric space S and that determines a common value aspect of the objects for all bidders. Bidders perform private investigations to determine the value of different portfolios, and conditional on s, each bidder i independently draws a private signal t_i from a complete, separable metric space that reflects the bidder's signal technology and the inherent randomness in testing. In addition, assume that bidder i has a private characteristic u_i that affects the value of the objects for bidder i only (due, e.g., to aspects of the bidder's production technology or product market), and that u_i is conditionally independent from the state s and the other bidders' types (t_{-i}, u_{-i}) given t_i. Let $t = (t_1, \ldots, t_n)$ denote the profile of signals. Bidder i's utility from winning a subset $G_i \subseteq \{1, \ldots, d\}$ of objects and paying price p_i in state s is denoted by $U_i(G_i, p_i, s, u_i)$, which is assumed measurable in (s, u_i). This determines a Bayesian game with payoffs $\pi_i(b, t, u_i)$ equal to $E[U_i(g_i(b), p_i(b), s, u_i)|t]$, the induced expected payoffs over allocations and prices. Applying Theorem 4.2, we obtain existence of an undominated pure strategy BNE. This extends Athey's (2001) Theorem 1 for finite games, McAdams' (2006) Theorem 1 for multi-unit auctions to settings with a private-value component, and Reny's (2011) Proposition 5.4: the existence of (possibly non-monotonic) pure strategy BNE in these settings does not require single-crossing conditions. Moreover, we do not require independence of the general type components.[5] □

Example 2.2: *Global games* Consider n players $i = 1, \ldots, n$ playing a finite game of incomplete information, as follows. An unobserved state s is drawn by nature from a complete, separable metric space S; conditional on s, each player i receives an independent signal t_i about s from a complete, separable metric space; and each player chooses an action x_i from a finite set, and an outcome $g(x_1, \ldots, x_n, s)$ is determined. In addition, player i's payoff is subject to an idiosyncratic shock u_i, where we assume that u_i is drawn from a non-atomic distribution over a complete, separable metric space and is conditionally independent from the state s and the other players' types (t_{-i}, u_{-i}) given t_i. Assume that payoffs $U_i(x_i, g(x_1, \ldots, x_n, s), s, u_i)$ are measurable in (s, u_i). This determines a Bayesian game with payoffs $\pi_i(x, t, u_i)$ equal to $E[U_i(x_i, g(x, s), s, u_i)|t]$, which can be viewed as a version of a finite "global game," of the kind first studied by Carlsson and Van Damme (1983) and applied and generalized by many others (e.g., Morris and Shin (1998) and (2003)) to the study of bank runs, currency attacks, and other coordination problems. For instance, the

[5] Reny (2011) does allow for interdependent types, as long as the distributions are independent; we allow for highly correlated type distributions. Relying on modularity conditions, de Castro (2012) obtains pure strategy BNE allowing for correlation. Similar comments apply to the results of Van Zandt and Vives (2007) and Van Zandt (2010).

players could be currency speculators and an action x_i may be a level of short sales of the currency of a given country, the state may represent the fundamentals of the economy in this country, the outcome may indicate whether the country successfully defends its currency, and u_i may be perturb the payoff of speculator i from attacking due to variations in the opportunity cost of short selling. The focus of the global games literature is the selection of a particular pure strategy BNE based on iterated dominance, which requires assumptions on states, signals, and payoffs; but Theorem 4.2 yields existence of an undominated pure strategy BNE without this additional structure. □

Example 2.3: *Voting with private information* Consider a voting game among n voters $i = 1, \ldots, n$ who must choose from a finite set of alternatives $\{a_1, \ldots, a_d\}$. Voters cast ballots $b_i \in B_i$ simultaneously, where B_i is a finite set of possible ballots, and the outcome is determined by a voting rule $v(b_1, \ldots, b_n)$, which may involve randomization. (In the two-alternative case, this generalizes majority rule and other quota rules.) Let s be an unobserved state selected by nature from a complete, separable metric space S, and assume that conditional on s, each voter i receives a private signal t_i drawn (independently conditional on s) from a complete, separable metric space. This signal is quite general, so we obtain the model of Feddersen and Pesendorfer (1997) in which $t_i = (\sigma_i, k_i)$, where k_i is an information service to which i is assigned, and σ_i is a signal generated by that information service. In addition, assume each voter is characterized by a private preference parameter u_i that is drawn from a non-atomic distribution over a complete, separable metric space, and assume that conditional on t_i, u_i is independent of the state s and the other voters' types (t_{-i}, u_{-i}). Given preference parameter u_i and state s, the utility of voter i from outcome $x \in \{a_1, \ldots, a_d\}$ is $U_i(x, s, u_i)$, which is assumed measurable in (s, u_i). This defines a Bayesian game with payoffs $\pi_i(b, t, u_i)$ equal to $E[U_i(v(b), s, u_i)|t]$, the induced expected payoffs over ballot profiles. Applying Theorem 4.2, we obtain existence of a pure BNE in undominated strategies in a framework that generalizes that of Feddersen and Pesendorfer (1997), who assume a one-dimensional state, finite signal space, two alternatives with voting by a quota rule, and a type of single-crossing condition. Thus, the existence of pure strategy equilibria implied by their Proposition 1 does not rely on the monotonicity conditions used to obtain equilibria in cutoff strategies. □

Example 2.4: *Costly voting* Consider costly voting in an election with a finite number of candidates $\{c_1, \ldots, c_d\}$. There are n voters $i = 1, \ldots, n$ who cast ballots $b_i \in B_i$ simultaneously, where B_i is a finite set of possible ballots, and the outcome is determined by a voting rule $v(b_1, \ldots, b_n)$, which may involve randomization. To model the possibility of abstention, we augment B_i with an artificial ballot a_i, which indicates that the voter does not participate in the election. Let s be an unobserved state selected by nature from a complete, separable metric space S that determines the popularity of the different candidates. Assume that conditional on s, each voter i independently draws a private preference parameter t_i from a complete, separable metric space. Myatt (2012) assumes that s is continuously distributed over the unit interval and that the parameter t_i is binary and indicates which candidate a voter prefers, but our construction allows for general states and preference parameters. In addition, assume each voter has a private cost of voting, u_i, that is drawn from a complete, separable metric space according to a non-atomic distribution, and assume that u_i is independent of the other voters' types (t_{-i}, u_{-i}). The utility of voter i from electing candidate x and casting ballot b_i given the voter's preference type t_i is $U_i(x, s, \chi_{a_i}(b_i)u_i)$, where $\chi_{a_i}(b_i) = 1$ if $b_i \neq a_i$ and equal to zero otherwise. U_i is assumed to be jointly measurable. This defines a Bayesian game with payoffs $\pi_i(b, t, u_i)$ equal to $E[U_i(v(b), t, \chi_{a_i}(b_i)u_i)|t]$. This setup generalizes Myatt's (2012), who assumes two candidates with majority voting, two preference types, one dimensional preference and cost parameters, and linear payoffs. Although Myatt's goal is to resolve the paradox of voting, our Theorem 4.2 yields general existence of a pure BNE in undominated strategies without the assumptions needed to generate equilibria in the form of cutoff strategies. □

3. Product Bayesian Games

The subject of our analysis is the class of product Bayesian games, defined informally as follows. There is a finite number of players, each player i having a type with two components, t_i and u_i. The first is a general component that may be payoff relevant for all players, and the second is a private-value component that, conditional on t_i, is independent of the other players' types. The action sets are type-dependent, but we assume a finite bound $d > 0$ on the number of pure actions for each player. We model a pure action a_i for player i as a unit coordinate vector e^ℓ, $\ell \in \{1, \dots, d\}$, in $\mathfrak{R}^d$. A mixed action α_i, which is a probability measure over pure actions, can be viewed as an element of the unit simplex Δ in $\mathfrak{R}^d$, and then a pure action belongs to the set extΔ of extreme points of the set of mixed actions.[6] We do not distinguish notationally between α_i as a probability measure and α_i as an element of the unit simplex, as no confusion should result.

[6]This identification is convenient because it allows us to integrate over pure actions and apply known methods for vector-valued integration in our purification arguments.

Formally, a *product Bayesian game* is a tuple $((T_i,\mathcal{T}_i),(U_i,\mathcal{U}_i),d,A_i,\pi_i,\kappa,\nu_i)_{i=1}^n$ indexed by a set $\{1,\ldots,n\}$ of players such that for each player $i=1,\ldots,n$,

- $(T_i,\mathcal{T}_i)$ is the measurable space of i's general types,
- $(U_i,\mathcal{U}_i)$ is the measurable space of i's private types,
- $d>0$ is a fixed integer,
- $A_i\colon T_i\times U_i \rightrightarrows \mathrm{ext}\Delta$ is i's feasible pure action correspondence,
- $\pi_i\colon(\mathrm{ext}\Delta)^n\times T\times U_i\to\Re$ is i's payoff function,
- κ is a probability measure on $(T,\mathcal{T})$,
- $\nu_i: T_i\times\mathcal{U}_i\to[0,1]$ is a transition probability,[7]

where $T=\times_{i=1}^n T_i$ is the set of profiles of general types, denoted $t=(t_1,\ldots,t_n)$, and $\mathcal{T}=\otimes_{i=1}^n\mathcal{T}_i$ is the product sigma-algebra. The probability measure κ represents the players' common prior beliefs about the profile of general types, and the transition probability ν_i specifies the distribution of i's private type conditional on each t_i. Let κ_i denote the marginal of κ on T_i.

We construct prior beliefs over types as follows. First, let $U=\times_{i=1}^n U_i$ with $\mathcal{U}=\otimes_{i=1}^n\mathcal{U}_i$ as its product sigma-algebra. Second, define the transition probability $\nu\colon T\times\mathcal{U}\to[0,1]$ so that for each $t\in T$ and each product set $S=\times_{i=1}^n S_i\in\mathcal{U}$, we have $\nu(S|t)=\prod_{i=1}^n\nu_i(S_i|t_i)$. Finally, let $\mu=\nu(\cdot|t)\otimes\kappa$ represent the common prior of the players and μ_i the marginal of μ on $T_i\times U_i$. Thus, we can obtain the marginal as $\mu_i=\nu_i(\cdot|t_i)\otimes\kappa_i$. Note that, conditional on t_i, the random variable u_i is independent of (t_{-i},u_{-i}) and distributed according to $\nu_i(\cdot|t_i)$.

We extend payoffs to mixed actions via expected utility. That is, representing the probability measure α_i by the element $(\alpha_{i,1},\ldots,\alpha_{i,d})$ of the unit simplex, we define the extension $\pi_i\colon\Delta^n\times T\times U_i\to\Re$ by

$$\pi_i(\alpha,t,u_i) \;=\; \sum_{\ell_1=1}^{d}\cdots\sum_{\ell_n=1}^{d}\pi_i(e^{\ell_1},\ldots,e^{\ell_n},t,u_i)\alpha_{1,\ell_1}\cdots\alpha_{n,\ell_n}$$

for all t, all u_i, and all profiles $\alpha=(\alpha_j)_{j=1}^n$ of mixed actions. Clearly, $\pi_i(\alpha,t,u_i)$ is multi-linear in the mixing probabilities α_j, $j=1,\ldots,n$.

For each $i=1,\ldots,n$, we assume:

(B1) T_i and U_i, are complete, separable metric spaces endowed with their Borel sigma-algebras,

[7] That is, $\nu_i(t_i,S_i)$ is measurable in t_i for all $S_i\in\mathcal{U}_i$, and $\nu_i(t_i,\cdot)$ is a probability measure on $\mathcal{U}_i$ for all t_i.

(B2) the correspondence $A_i: T_i \times U_i \rightrightarrows \text{ext}\Delta$ is lower measurable with nonempty values,[8]

(B3) $\pi_i(\alpha, t, u_i)$ is Borel measurable in (α, t, u_i),

(B4) the mapping $t \mapsto \sup_{(\alpha,u_i)\in\Delta^n \times U_i} |\pi_i(\alpha, t, u_i)|$ is κ-integrable, i.e., $\pi_i(\alpha, t, u_i)$ is bounded in (α, u_i) for κ-almost all t, and

$$\int_T \sup_{(\alpha,u_i)\in\Delta^n \times U_i} |\pi_i(\alpha, t, u_i)| \kappa(dt) < \infty,$$

(B5) κ is absolutely continuous with respect to $\bigotimes_{i=1}^n \kappa_i$.

Observe that (B5) is the standard condition of diffuse information, introduced by Milgrom and Weber (1985), applied to general types. In what follows, we let $f: T \to \mathfrak{R}$ be a density for κ with respect to the product of the marginals over general type components.

Remark 3.1. *We obtain the model with only general types by trivializing the sets U_i. This yields a finite-action version of the standard framework of Milgrom and Weber (1985) in which players' actions sets are type-dependent.*

Remark 3.2. *The players' general types may be highly correlated, subject to (B5), and payoff functions $\pi_i(a, t, u_i)$ may depend on the profile of general types. Thus, we allow for correlated types and interdependent payoffs.*

A *strategy* for player i is a transition probability from types to feasible mixed actions, i.e., a $\mathcal{T}_i \otimes \mathcal{U}_i$-measurable function $\sigma_i: T_i \times U_i \to \Delta$ such that $\sigma_i(t_i, u_i) \in \text{co}A_i(t_i, u_i)$ for μ_i-almost all (t_i, u_i), and a *strategy profile* $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a mapping $\sigma: T \times U \to \Delta^n$ with values $\sigma(t, u) = (\sigma_i(t_i, u_i))_{i=1}^n$. Let Σ_i denote the set of strategies for i. The *ex ante expected payoff* for player i from a strategy profile σ is

$$\Pi_i(\sigma) = \int_{T\times U} \pi_i(\sigma(t, u), t, u_i) \mu(d(t, u)),$$

which, by (B4), is finite. A *Bayes-Nash equilibrium (BNE)* is a strategy profile σ^* such that for each $i = 1, \ldots, n$, $\Pi_i(\sigma^*) = \sup_{\sigma_i \in \Sigma_i} \Pi_i(\sigma_i, \sigma^*_{-i})$, and σ^* is a *pure-strategy Bayes-Nash equilibrium* if it is a Bayes-Nash equilibrium such that for each $i = 1, \ldots, n$ and for μ_i-almost all (t_i, u_i), we have $\sigma^*_i(t_i, u_i) \in A_i(t_i, u_i)$.

The next section presents existence and purification results for Bayes-Nash equilibria. To define the notion of equivalence used in our concept of purification, we

[8]Given a measurable space $(X, \mathcal{X})$ and topological space Y, a correspondence $\psi: X \rightrightarrows Y$ is *lower measurable* if for all open $G \subseteq Y$, the lower inverse $\psi^\ell(G) = \{x \in X : \psi(x) \cap G \neq \varnothing\}$ is measurable. Since A_i has finite range, lower measurability of A_i is equivalent to the notion of measurable graph.

define the *average action* of player i determined by strategy σ_i, conditional on general component t_i, as

$$\gamma_i(t_i|\sigma_i) \equiv \int_{U_i} \sigma_i(t_i,u_i)\nu_i(du_i|t_i).$$

By a generalization of Fubini's theorem (see Proposition 2.3.2 (p.47) of Rao (1993)), the function $\gamma_i(\cdot|\sigma_i)$ is κ_i-integrable. We let $\Gamma_i = \{\gamma_i(\cdot|\sigma_i) : \sigma_i \in \Sigma_i\}$ denote the space of average actions for player i as a function of the general component t_i, and we define the product set $\Gamma_{-i} = \times_{j\neq i}\Gamma_j$. The next lemma states that the set Γ_i of average actions for player i consists of all transition probabilities from T_i to Δ that for κ_i-almost all t_i, place probability one on feasible actions.[9] We then endow Γ_i with the narrow topology (see Balder (1988,2002)), and by Theorem 4.1.1 of Balder (2002), it follows that Γ_i, and therefore Γ_{-i} endowed with the product topology, is compact.

Lemma 3.1. *For $i = 1,\dots,n$ and each Borel measurable $\gamma_i: T_i \to \Re^d$, we have $\gamma_i \in \Gamma_i$ if and only if for κ_i-almost all t_i, $\gamma_i(t_i) \in \int_{U_i} coA_i(t_i,u_i)\nu_i(du_i|t_i)$.*

By Fubini's theorem (see Theorem 11.27 of Aliprantis and Border (2006)) and multilinearity of $\pi_i(\cdot,t,u_i)$ in Δ^n, we have

$$\begin{aligned}\Pi_i(\sigma) &= \int_T \int_U \pi_i(\sigma(t,u),t,u_i)\bigotimes_{k=1}^n \nu_k(du|t)\kappa(dt)\\ &= \int_T\int_{U_i} \pi_i\left(\sigma_i(t_i,u_i),\left(\int_{U_j}\sigma_j(t_j,u_j)\nu_j(du_j|t_j)\right)_{j\neq i},t,u_i\right)\nu_i(du_i|t_i)\kappa(dt)\\ &= \int_T\int_{U_i}\pi_i(\sigma_i(t_i,u_i),\gamma_{-i}(t_{-i}|\sigma_{-i}),t,u_i)\nu_i(du_i|t_i)\kappa(dt),\end{aligned}$$

where $\gamma_{-i}(t_{-i}|\sigma_{-i}) = (\gamma_j(t_j|\sigma_j))_{j\neq i}$. Because the expected payoff depends on σ_{-i} only through average actions of the other players, we can redefine the ex ante payoff function for player i as the mapping for $\Pi_i: \Sigma_i \times \Gamma_{-i} \to \Re$ given by

$$\Pi_i(\sigma_i,\gamma_{-i}) = \int_T\int_{U_i}\pi_i(\sigma_i(t_i,u_i),\gamma_{-i}(t_{-i}),t,u_i)\nu_i(du_i|t_i)\kappa(dt),$$

an abuse of notation that should cause no confusion. Accordingly, we replace the optimization of $\Pi_i(\sigma_i,\sigma^*_{-i})$ with $\Pi_i(\sigma_i,\gamma_{-i}(\cdot|\sigma^*_{-i}))$ in the definition of Bayes-Nash equilibrium.

We then say two strategy profiles σ and σ' are *equivalent* if they determine the same average actions, i.e., for each $i = 1,\dots,n$ and for κ_i-almost all t_i, we have

[9] Given a measure space $(X,\mathcal{X},\xi)$ and a correspondence $\psi: X \rightrightarrows \Re^d$, we write $\int_X \psi(x)\xi(dx)$ for the *Aumann integral* of ψ, which is the set of integrals of all integrable, ξ-almost everywhere selections of ψ.

$\gamma_i(t_i|\sigma) = \gamma_i(t_i|\sigma')$. This notion of equivalence is, in light of the foregoing remarks, perhaps stronger than it would at first seem, for it implies payoff equivalence: if σ and σ' are equivalent, then for each $i = 1, \ldots, n$, we have $\Pi_i(\cdot, \sigma_{-i}) = \Pi_i(\cdot, \sigma'_{-i})$.

It is convenient for the treatment of dominance to delve into the interim level, although our formulation of equilibrium is at the ex ante level. To this end, we use the density of κ with respect to $\bigotimes_{i=1}^n \kappa_i$ to construct a Borel probability measure $\kappa(\cdot|t_i)$ on T_{-i} for all t_i in the obvious way: for each Borel R_{-i}, we define

$$\kappa(R_{-i}|t_i) = \frac{\int_{R_{-i}} f(t_i, t_{-i})(\bigotimes_{j\neq i} \kappa_j)(dt_{-i})}{\int_{T_{-i}} f(t_i, t_{-i})(\bigotimes_{j\neq i} \kappa_j)(dt_{-i})}$$

for t_i outside the κ_i-measure zero set T_i' for which the denominator equals zero, and we define $\kappa(\cdot|t_i)$ arbitrarily (subject to measurability) for $t_i \in T_i'$. This gives us a regular conditional probability for κ with conditioning variable t_i. Then the *interim expected payoff* for player i given average actions γ_{-i} as a function of type (t_i, u_i), denoted $\Pi_i(\cdot;\gamma_{-i})\colon \Delta \times T_i \times U_i \to \mathfrak{R}$, is defined by

$$\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i}) = \int_{T_{-i}} \pi_i(\alpha_i, \gamma_{-i}(t_{-i}), t, u_i)\kappa(dt_{-i}|t_i),$$

which is measurable in (t_i, u_i). By (B4) and Proposition 2.3.2 (p.47) of Rao (1993), we have

$$\begin{aligned}\int_{T_i}\int_{T_{-i}} \sup_{(\alpha,u_i)\in\Delta^n\times U_i} |\pi_i(\alpha, t, u_i)|\kappa(dt_{-i}|t_i)\kappa_i(dt_i) &= \int_T \sup_{(\alpha,u_i)\in\Delta^n\times U_i} |\pi_i(\alpha, t, u_i)|\kappa(dt) \\ &< \infty.\end{aligned}$$

It follows that for all t_i outside a κ_i-measure zero Borel set T_i'' and for all $u_i \in U_i$, we have $\int_{T_{-i}} \sup_{\alpha\in\Delta^n} |\pi_i(\alpha, t, u_i)|\kappa(dt_{-i}|t_i) < \infty$, and therefore $\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i}) < \infty$. In fact, $\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i})$ is μ_i-integrable. Viewing γ_{-i} as a product of transition probabilities from general types to pure actions, Theorem 2.5 of Balder (1988), with (B1)–(B5), implies that for all i and all (t_i, u_i), the interim expected payoff $\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i})$ is jointly continuous in (α_i, γ_{-i}). For future reference, set $\tilde{T}_i = T_i' \cup T_i''$.

As usual, we can decompose ex ante expected payoffs into interim expected payoffs:

$$\begin{aligned}\Pi_i(\sigma_i, \gamma_{-i}) &= \int_T \int_{U_i} \pi_i(\sigma_i(t_i, u_i), \gamma_{-i}(t_{-i}), t, u_i)\nu_i(du|t_i)\kappa(dt) \\ &= \int_{T_i}\int_{T_{-i}}\int_{U_i} \pi_i(\sigma_i(t_i, u_i), \gamma_{-i}(t_{-i}), t, u_i)\nu_i(du|t_i)\kappa(dt_{-i}|t_i)\kappa_i(dt_i) \\ &= \int_{T_i\times U_i}\int_{T_{-i}} \pi_i(\sigma_i(t_i, u_i), \gamma_{-i}(t_{-i}), t, u_i)\kappa(dt_{-i}|t_i)\mu_i(d(t_i, u_i)) \\ &= \int_{T_i\times U_i} \Pi_i(\sigma_i(t_i, u_i), t_i, u_i; \gamma_{-i})\mu_i(d(t_i, u_i)),\end{aligned}$$

where the second and third equalities follow from Proposition 2.3.2 (p.47) of Rao (1993).

Given $t_i \in T_i \setminus \tilde{T}_i$ and any u_i, a mixed action $\alpha_i \in \Delta$ is *dominated* for i at (t_i, u_i) if there exists $\alpha_i' \in \mathrm{co}A_i(t_i, u_i)$ such that (i) $\Pi_i(\alpha_i', t_i, u_i; \gamma_{-i}) \geq \Pi_i(\alpha_i, t_i u_i; \gamma_{-i})$ for all $\gamma_{-i} \in \Gamma_{-i}$, and (ii) $\Pi_i(\alpha_i', t_i, u_i; \gamma_{-i}) > \Pi_i(\alpha_i, t_i u_i; \gamma_{-i})$, for some $\gamma_{-i} \in \Gamma_{-i}$. Accordingly, a strategy σ_i is *undominated* for i if for μ_i-almost all (t_i, u_i) with $t_i \in T_i \setminus \tilde{T}_i$, $\sigma_i(t_i, u_i)$ is not dominated, and a profile σ is *undominated* if σ_i is undominated for each player i. This corresponds to the usual notion of undominated strategy, translated to the product Bayesian game framework.

4. Purification and Existence Results

Because it has finite range, lower measurability of the feasible correspondence implies that it has a measurable graph. Moreover, (B5) implies that the prior μ satisfies diffuse information.[10] Therefore, the conditions of Theorem 3.3 in Balder (1988) are satisfied, delivering existence of equilibrium in mixed strategies.

Theorem 4.1. *Under (B1)–(B5), a Bayes-Nash equilibrium exists.*

Our contribution comes next. When the private type components are atomless, we show that (a) every BNE can be purified, (b) an undominated, pure-strategy BNE exists, and (c) every undominated BNE can be purified into an undominated, pure-strategy BNE. All three results allow for highly correlated general type components and interdependent payoffs. We emphasize that in many applications the focus of analysis is equilibria in which only undominated strategies are played, and that existence and purification of BNE in undominated strategies for general Bayesian games has not been addressed in the literature.

Theorem 4.2. *Assume (B1)–(B5), and assume the probability measures* $\{\nu_i(\cdot|t_i) : t_i \in T_i\}$ *are nonatomic for each* $i = 1, \ldots, n$. *(a) For every Bayes-Nash equilibrium, there exists an equivalent pure-strategy Bayes-Nash equilibrium; (b) a pure-strategy, undominated Bayes-Nash equilibrium exists; (c) for every undominated Bayes-Nash equilibrium, there exists an equivalent pure-strategy, undominated Bayes-Nash equilibrium.*

[10] To see this, let $Q \subseteq T \times U$ be a Borel set such that $(\otimes_i \mu_i)(Q) = 0$. That is, letting $\mu(\cdot|t)$ be a regular conditional probability for μ and writing $Q_t = \{u \in U : (t, u) \in Q\}$ for the section at t, we have $\int_U \mu(Q_t|t)(\otimes_i \kappa_i)(dt) = 0$. Then $\mu(Q_t|t) = 0$ for $(\otimes_i \kappa_i)$-almost all t. By (B5), $\mu(Q_t|t) = 0$ for κ-almost all t, which yields $\mu(Q) = \int_U \mu(Q_t|t)\kappa(dt) = 0$, as required.

The proof of part (a) uses an application of Artstein's (1989) theorem for purification purposes. We begin with a Bayes-Nash equilibrium σ^*, which exists by Theorem 4.1. We then use Lyapunov's theorem to show that equilibrium average actions γ_i^* can be obtained using only pure actions, under nonatomicity: for each t_i, we replace $\sigma_i^*(t_i, \cdot)$ with a selection $\tilde{\sigma}_i(t_i, \cdot)$ from best response pure actions for player i as the private component u_i is varied, and we do so in a way that preserves the average action $\gamma_i^*(t_i)$. By multilinearity, this purification step maintains the payoffs of all players. This procedure is implemented separately for each t_i, and it remains to sew the selections of best response pure actions together in a measurable way; this is where Artstein's theorem is applied. After verifying the conditions needed to apply Artstein's theorem, we have a jointly measurable strategy $\hat{\sigma}_i$ for each player i such that for κ_i-almost all t_i, the average action $\gamma_i^*(t_i)$ is preserved, and for $\nu_i(\cdot|t_i)$-almost all u_i, $\hat{\sigma}_i(t_i, u_i)$ is a best response pure action. This is a Bayes-Nash equilibrium, and it is equivalent to the original σ^*, as required.

To gain more insight into the nature of our contribution, consider the following procedure: from Theorem 4.1, there is a BNE; for each t, one applies standard purification arguments to produce a purified strategy profile that preserves interim payoffs for the game conditional on t; then, one would apply Artstein's theorem to obtain a jointly measurable, purified replacement of the original BNE. If this procedure were correct, then part (a) of Theorem 4.2 could be viewed as a somewhat predictable, measurably parametrized form of standard purification. But the procedure just described fails because for each player i, it produces a purified strategy that depends on the entire profile t, rather than only on t_i. We avoid this problem by applying Lyapunov's theorem for each t_i to obtain average actions using only pure actions, rather than applying it for each profile t. Thus, just as the results themselves, the architecture of the argument may be of interest in itself.

The following example illustrates the usefulness of the product Bayesian game structure in facilitating purification.

Example 4.1. Consider Example 1 from Radner and Rosenthal (1982): there are two players, with payoff matrix below,

	L	M
T	$2,0$	$0,1$
M	$0,1$	$2,0$

and players have payoff-irrelevant types t_1 and t_2 uniformly distributed on the triangle of the unit square defined by $0 \le t_1 \le t_2 \le 1$.[11] The key feature of the example is that the types are not independent, a fact that enabled Radner and Rosenthal to establish that the mixed strategy equilibrium (σ_1, σ_2), where $\sigma_1(t_1)$ assigns probability $\frac{1}{2}$ to T and M for all t_1 and $\sigma_2(t_2)$ assigns probability $\frac{1}{2}$ to L and M for all t_2, *cannot* be purified. Now let us add payoff-relevant private types u_1 and u_2 yielding the payoff matrix below,

	L	M
T	$2, u_2/2$	$u_1/2, 1$
M	$u_1/2, 1$	$2, u_2/2$

where u_i is uniformly distributed on $[t_i - \epsilon, t_i + \epsilon]$ given t_i, independently of t_j and u_j, for some given $\epsilon \in (0,1)$ and $i \ne j = 1, 2$.[12] Then we have a product Bayesian game, and the analogous profile, say $(\tilde{\sigma}_1, \tilde{\sigma}_2)$, remains a mixed strategy equilibrium: now, $\tilde{\sigma}_1(t_1, u_1)$ assigns probability $\frac{1}{2}$ to T and M for all (t_1, u_1) and $\tilde{\sigma}_2(t_2, u_2)$ assigns probability $\frac{1}{2}$ to L and M for all (t_2, u_2). But this profile *can* be purified into (σ_1', σ_2') such that $\sigma_1'(t_1, u_1) = \text{T}$ (resp., $\sigma_1'(t_1, u_1) = \text{M}$) when $u_1 \ge t_1$ (resp., $u_1 < t_1$) and $\sigma_2'(t_2, u_2) = \text{L}$ (resp., $\sigma_2'(t_2, u_2) = \text{M}$) when $u_2 \ge t_2$ (resp., $u_2 < t_2$). □

The proof of parts (b) and (c) of Theorem 4.2 consists of transforming the product Bayesian game by restricting each player to mixtures over undominated pure actions contingent on each type. Once we show that this restricted set of actions varies in a lower measurable way with respect to the player's type, we obtain (B2) in the transformed game, and then existence of BNE in the transformed game is obtained by Theorem 4.1. An issue that arises, however, is that *a convex combination of undominated pure actions can be dominated,* so a BNE of the transformed game (which may in general involve mixing) need not correspond to an undominated BNE of the original game, *but under the assumption of non-atomicity, we can purify such a BNE so that players select only undominated pure actions.* We then argue that this purified profile gives us the desired undominated BNE. Thus, existence (from Theorem 4.1) and purification (from part (a) of Theorem 4.2) combine to produce existence and purification in terms of undominated strategies. To illustrate the difficulties of obtaining undominated BNE outside the class of product Bayesian games

[11]Hence, we are only focusing on the case of correlated general types to illustrate our result. It is simple to modify the example and allow also for the general types to be payoff relevant, by for instance adding $t_1 + t_2$ to some of the entries of the matrix.

[12]The payoff relevance of u_i is for illustrative purposes only; the point we are making can already be seen with u_i not appearing in the payoff function.

and to highlight the instrumental role of purification in our arguments, consider the following modification of Example 4.1.[13]

Example 4.2. There are two players, with payoff matrix below,

	L	M	R
T	2,0	0,1	0,0
M	0,1	2,0	0,0
B	1,0	1,0	1,0

and players have payoff-irrelevant types u_1 and u_2 distributed uniformly on the triangle of the unit square defined by $0 \le u_1 \le u_2 \le 1$ (there is no general type component).[14] No pure action of player 1 is dominated, while the pure action R of player 2 is. Again, from Radner and Rosenthal (1982), we know that the BNE (σ_1, σ_2) such that $\sigma_1(u_1)$ assigns probability $\frac{1}{2}$ to T and M for all u_1 and $\sigma_2(u_2)$ assigns probability $\frac{1}{2}$ to L and M for all u_2 cannot be purified. Furthermore, player 1's strategy σ_1 is dominated by the pure strategy that chooses B with probability 1 for all u_1. Assuming instead that types are distributed independently and uniformly on the unit interval, then we have a product Bayesian game, and the equilibrium can be purified into an equilibrium where each type of player 1 plays either T or M with probability one, and likewise each type of player 2 plays either L or M with probability one. Specifically, (σ_1, σ_2) can be purified into the BNE (σ_1^*, σ_2^*) with $\sigma_1^*(u_1) = \mathrm{T}$ for $u_1 \in [0, 1/2]$, $\sigma_1^*(u_1) = \mathrm{M}$ for $u_1 \in (1/2, 1]$, $\sigma_2^*(u_2) = \mathrm{L}$ for $u_2 \in [0, 1/2]$, and $\sigma_2^*(u_2) = \mathrm{M}$ for $u_2 \in (1/2, 1]$. Now we have a BNE that is undominated, as every type plays an undominated pure action, despite the fact that player 1's strategy σ_1 in the initial BNE is dominated. □

The approach above rests on lower measurability of the set of undominated feasible actions as a function of a player's type, which is established next. Specifically, given type (t_i, u_i) of player i, let $D_i(t_i, u_i) \subseteq \mathrm{co}A_i(t_i, u_i)$ denote the undominated feasible mixed actions for i, and let $E_i(t_i, u_i) \subseteq A_i(t_i, u_i)$ be the undominated pure actions; then we show that the correspondence $(t_i, u_i) \mapsto E_i(t_i, u_i)$ is lower measurable and has nonempty values. The result may be of independent interest: by Remark 3.1, the proposition applies to general Bayesian games with finite (possibly type-dependent) actions sets. The proof is contained in the next section.

[13] Observe that the normal form of a product Bayesian game is an infinite game where existence of an undominated equilibrium cannot be guaranteed in general. See, e.g., Simon and Stinchcombe (1995).

[14] Again, it is simple to make the example more complex adding payoff relevant and interdependent general types.

Proposition 4.3. *Assume (B1)–(B5). For each $i = 1, \ldots, n$, the correspondence $E_i : T_i \times U_i \rightrightarrows ext\Delta$ is lower measurable and has nonempty values.*

Before proceeding to proofs in the next section, several remarks are in order.

Remark 4.1. *We can add a complete, separable metric space T_0 of environmental states at no cost.*

For this, we simply add an artificial player 0, whose general type component t_0 corresponds to the environmental state, with trivial action set $A(t_0, u_0) \equiv \{e^1\}$ for all (t_0, u_0). A general common value component cannot be added via an artificial player in extant papers on purification (e.g., Milgrom and Weber (1985), Khan, Rath, and Sun (2006), Fu et al. (2007), and Balder (2008)), because they assume private values. Furthermore, the purification literature assumes an environmental state variable that takes finitely many values, with private types being conditionally independent given the state. In contrast, our general types t_i, $i > 0$, are allowed to be highly correlated with the environmental state t_0 (subject to (B5)), and are not required to be conditionally independent given t_0. On the other hand, conditional on t_i, our private type u_i is independent of t_0, so it can depend on the environmental states t_0 only through the general component t_i.[15]

Remark 4.2. *We extend the literature on purification by allowing a non-private values type component that is correlated across players.*

The correlation in information and interdependence of payoffs in the purification literature are all captured by the finite environmental state variable as described above. We allow for general correlation subject only to diffuseness, capturing a general underlying metric space of states S and private signals t_i about the state (as in Examples 2.1–2.4), but more generally, because payoffs depend on the profile t of general types, we allow player j to have information that is payoff-relevant to player $i \neq j$ as well.

Remark 4.3. *To the best of our knowledge, ours is the first result establishing purification of undominated BNE.*

As a consequence, we leverage purification arguments in a way that is potentially useful in the analysis of auctions, voting games, and other applications (see Examples 2.1–2.4), where the focus is often on BNE in undominated strategies.

[15]Our arguments could also accommodate an additional finite state variable $\hat{t}_0$ and allow the private components u_i to depend on $\hat{t}_0$ at the cost of some complication of the analysis.

Our results have implications for games with infinitely many actions as well. Instead of taking the unit coordinate vectors as the underlying space of pure actions, we let X be a fixed compact metric space; we formulate feasible action correspondences $A_i\colon T_i \times U_i \rightrightarrows X$ with values in this space; and we specify payoff functions $\pi_i\colon X^n \times T \times U_i \to \mathfrak{R}$ over profiles of actions. Denote by $\mathcal{P}(Z)$ the space of Borel probability measures over any Borel measurable $Z \subseteq X$ endowed with the weak* topology. Then we replace (B2) with the assumption that for each $i = 1, \ldots, n$,

(B2′) the correspondence $A_i\colon T_i \times U_i \rightrightarrows X$ is lower measurable with nonempty, compact values,

and we extend $\pi_i(\cdot, t, u_i)$ to $\mathcal{P}(X)^n$ by taking expected payoffs. Thus, with $\mathcal{P}(X)$ in the place of Δ in Section 3, we have a well-defined product Bayesian game with infinitely many actions. In addition, we assume that for each $i = 1, \ldots, n$,

(B6) for all $(t, u_i) \in T \times U_i$, $\pi_i(a, t, u_i)$ is continuous in $a = (a_j)_{j=1}^n \in X^n$.

A strategy for player i is a $\mathcal{T}_i \otimes \mathcal{U}_i$-measurable function $\sigma_i\colon T_i \times U_i \to \mathcal{P}(X)$ such that $\sigma_i(t_i, u_i) \in \mathcal{P}(A_i(t_i, u_i))$ for κ_i-almost all (t_i, u_i). A Bayes-Nash equilibrium is defined as above, and we define *ϵ-Bayes-Nash equilibrium* in the usual way, which requires only that each player's strategy be an ϵ-best response.

Remark 4.4. *Assume (B1), (B2′), and (B3)–(B6) hold in the infinite-action product Bayesian game. (a) A Bayes-Nash equilibrium exists; (b) if the probability measures $\{\nu_i(\cdot|t_i) : t_i \in T_i\}$ are nonatomic for each $i = 1, \ldots, n$, then for every $\varepsilon > 0$, there exists a pure strategy ϵ-Bayes-Nash equilibrium.*

The second part of the remark follows by taking limits of finite approximations of X. Because the argument is a standard one, we provide only a brief sketch. For each $\delta > 0$, let $X^\delta \subseteq X$ be a δ-dense, finite subset of X. Define a product Bayesian game with finitely many actions as the restriction of the original game to X^δ. More precisely, suppose there are d_i elements of X^δ within a distance of δ from $A_i(t_i, u_i)$, and set $d \equiv \max_i d_i$. We let $A_i^\delta(t_i, u_i)$ consist of the unit coordinate vectors in $\mathfrak{R}^d$, where a feasible action a_i corresponds to an element of X^δ within distance δ of $A_i(t_i, u_i)$. The main technical issue is to ensure lower measurability of A_i^δ, and for this we rely on results for measurability of the distance function of a correspondence. Letting ρ denote a metric for X, set $Y_i^\delta(t_i, u_i) \equiv \{x \in X^\delta : \rho(x, Y_i(t_i, u_i)) < \delta\}$, so that each element of $Y_i^\delta(t_i, u_i)$ can be identified with a feasible action in $A_i^\delta(t_i, u_i)$. By Theorem 18.5 of Aliprantis and Border (2006), the mapping $(x, t_i, u_i) \mapsto \rho(x, Y_i(t_i, u_i))$ is continuous in x and Borel measurable in (t_i, u_i), i.e., Carathéodory. By Lemma 18.7 of Aliprantis and Border (2006), Y_i^δ is lower measurable. Then Theorem 3 of Himmelberg and Van Vleck (1975) implies that Y_i^δ is measurable, and by Lemma 18.2 of Aliprantis and Border (2006), it is lower measurable.

Again, part (a) follows from Balder (1988). For part (b), use part (a) of Theorem 4.2 and the approximation scheme defined above to obtain a net $\{\sigma^\delta\}$ of pure strategy BNE with limit $\sigma = \lim_{\delta\to 0}\sigma^\delta$. Then σ is a BNE, and σ^δ for δ sufficiently small is an ϵ-BNE in pure strategies.[16] Observe that we cannot ensure that the limit σ is a pure strategy, even if X is a convex space.

5. Proofs

We begin with some ancillary results. Throughout this section, we maintain assumptions (B1)–(B5), and we rely on several results from Aliprantis and Border (2006), henceforth abbreviated AB. We proceed to part (a) of Theorem 4.2, then to Proposition 4.3, and finally to parts (b) and (c) of Theorem 4.2. We start with the proof of Lemma 3.1.

Lemma 3.1. *For $i = 1,\dots,n$ and each Borel measurable $\gamma_i: T_i \to \Re^d$, we have $\gamma_i \in \Gamma_i$ if and only if for κ_i-almost all t_i, $\gamma_i(t_i) \in \int_{U_i} coA_i(t_i,u_i)\nu_i(du_i|t_i)$.*

Proof. The "only if" direction is immediate from the definition of average action for i. Indeed, letting $\gamma_i \in \Gamma_i$ be determined as $\gamma_i = \gamma_i(\cdot|\sigma_i)$ for the strategy σ_i for i, it follows that for κ_i-almost all $t_i \in T_i$, $\sigma_i(t_i,\cdot): U_i \to \Delta$ is a Borel measurable, $\nu_i(\cdot|t_i)$-almost everywhere selection from $\mathrm{co}A_i(t_i,\cdot): U_i \rightrightarrows \Delta$, and thus $\gamma_i(t_i) = \int_{U_i}\sigma_i(t_j,u_i)\nu_i(du_i|t_i) \in \int_{U_i}\mathrm{co}A_i(t_i,u_i)\nu_i(du_i|t_i)$, as claimed. For the "if" direction, let γ_i be a κ_i-almost everywhere selection from $t_i \mapsto \int_{U_i}\mathrm{co}A_i(t_i,u_i)\nu_i(du_i|t_i)$. Then the theorem of Artstein (1989) yields a Borel measurable mapping $\sigma_i: T_i\times U_i \to \Delta$ such that for κ_i-almost all t_i, we have: $\gamma_i(t_i) = \int_{U_i}\sigma_i(t_i,u_i)\nu_i(du_i|t_i)$, and for $\nu_i(\cdot|t_i)$-almost all u_i, $\sigma_i(t_i,u_i) \in \mathrm{co}A_i(t_i,u_i)$. In particular, his assumptions (i)–(vi) are fulfilled by, respectively, (B1) applied to T_i, (B1) applied to U_i, the assumption that $\nu_i(\cdot|\cdot): \mathcal{U}_i \times T_i \to [0,1]$ is a transition probability (AB, Theorem 19.7), (B2) and Lemma 18.2 of AB, and because the co-domain of $\mathrm{co}A_i$ is Δ. Thus, γ_i is determined as $\gamma_i = \gamma_i(\cdot|\sigma_i)$ for the strategy σ_i for i. □

Given average action profile $\gamma \in \Gamma$, we write $M_i(t_i,u_i;\gamma)$ for the set of mixed actions maximizing player i's interim expected payoff conditional on (t_i,u_i). Formally,

[16]Here, we use Theorem 4.1.1 of Balder (2002), which establishes that the space Σ of strategy profiles is compact in the narrow topology, and (B5) and Theorem 2.5 of Balder (1988), which together imply joint continuity of ex ante expected payoffs.

we define the correspondence $M_i(\cdot;\gamma): T_i \times U_i \rightrightarrows \Delta$ as follows. Recall that there is a κ_i-measure zero set $\tilde{T}_i$ such that for all $t_i \in T_i \setminus \tilde{T}_i$, $\kappa(\cdot|t_i)$ is a conditional probability on T_{-i} and for all $u_i \in U_i$, we have $\int_{T_{-i}} \sup_{\alpha \in \Delta^n} |\pi_i(\alpha, t, u_i)| \kappa(dt_{-i}|t_i) < \infty$. Our arguments will focus on $t_i \in T_i \setminus \tilde{T}_i$, specifying actions in an arbitrary measurable way for $t_i \in \tilde{T}_i$. Thus, for $t_i \in T_i \setminus \tilde{T}_i$, set

$$M_i(t_i, u_i; \gamma) \quad = \quad \arg\max\{\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i}) : \alpha_i \in \mathrm{co}(A_i(t_i, u_i))\}.$$

For the remaining case, note that by (B2), coA_i is a lower measurable correspondence with nonempty, closed values, and it therefore admits a measurable selection (AB, Theorem 18.13). Denoting such a selection by $\tilde{\sigma}_i$, we define $M_i(t_i, u_i; \gamma) = \{\tilde{\sigma}_i(t_i, u_i)\}$ when $t_i \in \tilde{T}_i$. An implication of the next lemma is that $M_i(\cdot;\gamma)$ is compact- and convex-valued.

Lemma 5.1. *For $i = 1, \ldots, n$ and each $\gamma \in \Gamma$, (a) for all (t_i, u_i) with $t_i \in T_i \setminus \tilde{T}_i$, $M_i(t_i, u_i; \gamma)$ is a nonempty face of Δ; (b) $M_i(\cdot;\gamma) : T_i \times U_i \rightrightarrows \Delta$ is lower measurable.*

Proof. We first prove (a). Given (t_i, u_i) with $t_i \in T_i \setminus \tilde{T}_i$, co$A_i(t_i, u_i)$ is clearly a nonempty face of Δ. Furthermore, by multi-linearity of expected payoffs, the interim payoff function $\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i})$ is linear in α_i. Therefore, $M_i(t_i, u_i; \gamma)$ is also a nonempty face of Δ. To prove (b), recall that $\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i})$ is Borel measurable in (t_i, u_i). Restricting $\Pi_i(\alpha_i, \cdot; \gamma_{-i})$ to $(T_i \setminus \tilde{T}_i) \times U_i$, the function is real-valued. Moreover, as α_i is a convex combination of pure actions $a_i \in \mathrm{ext}\Delta$, (B3) implies that for $t_i \in T_i \setminus \tilde{T}_i$,

$$\Pi_i(\alpha_i, t_i, u_i; \gamma) \quad = \quad \sum_{\ell=1}^{d} \alpha_{i,\ell} \Pi_i(e^{\ell}, t_i, u; \gamma),$$

where $\alpha_i = (\alpha_{i,\ell})_\ell$. Thus, the interim expected payoff of player i is continuous in α_i, so $\Pi_i(\cdot;\gamma)$ is a Carathéodory function on $(T_i \setminus \tilde{T}_i) \times U_i$. By (B2), the restriction of the correspondence coA_i to $(T_i \setminus \tilde{T}_i) \times U_i$ is lower measurable with nonempty, compact values, so we can apply a measurable version of the maximum theorem (AB, Theorem 18.19) to conclude that $M_i(\cdot;\gamma)$ is measurable, and therefore (AB, Lemma 18.2) lower measurable, on $(T_i \setminus \tilde{T}_i) \times U_i$. It then follows from the construction that $M_i(\cdot;\gamma)$ is lower measurable on $T_i \times U_i$, as required. □

The next lemma provides the usual interim formulation of Bayes-Nash equilibrium in terms of the players' action correspondences and average actions.

Lemma 5.2. *Let σ^* be a strategy profile, and let $\gamma^* = \gamma(\cdot|\sigma^*)$. Then σ^* is a BNE if and only if for μ_i-almost all (t_i, u_i), we have $\sigma_i^*(t_i, u_i) \in M_i(t_i, u_i; \gamma^*)$.*

Proof. First, assume σ^* is a Bayes-Nash equilibrium, and suppose that for some player i, σ_i^* is not a μ_i-almost everywhere selection from $M_i(\cdot;\gamma_i^*)$. Note that for each $\ell = 1,\ldots,d$, the mapping $(t_i,u_i) \mapsto \Pi_i(e^\ell, t_i, u_i; \gamma_{-i}^*)$ is Borel measurable, as is the composite mapping $(t_i,u_i) \mapsto \Pi_i(\sigma_i^*(t_i,u_i), t_i, u_i; \gamma_{-i}^*)$. Thus, the set

$$Q^\ell \quad = \quad \left\{(t_i,u_i) \in T_i \times U_i : \Pi_i(\sigma_i^*(t_i,u_i), t_i, u_i; \gamma_{-i}^*) < \Pi_i(e^\ell, t_i, u_i; \gamma_{-i}^*)\right\}$$

is Borel measurable, as is $\bigcup_{\ell=1}^d Q^\ell$. This is union is just the set $\{(t_i,u_i) \in T_i \times U_i : \sigma_i^*(t_i,u_i) \notin M_i(t_i,u_i;\gamma^*)\}$, and thus by supposition we have $\mu_i(Q^\ell) > 0$ for some $\ell = 1,\ldots,d$. But then we can define the strategy σ_i by splicing σ_i^* with the pure action e^ℓ on the set Q^ℓ, i.e., for all $(t_i,u_i) \in T_i \times U_i$,

$$\sigma_i(t_i,u_i) \quad = \quad \begin{cases} e^\ell & \text{if } (t_i,u_i) \in Q^\ell, \\ \sigma_i^*(t_i,u_i) & \text{else.} \end{cases}$$

Then $\Pi_i(\sigma_i, \sigma_i^*) > \Pi_i(\sigma^*)$, a contradiction. Second, assume that σ^* is not a BNE. Then there is some player i with a strategy σ_i such that $\Pi_i(\sigma_i, \sigma_{-i}^*) > \Pi_i(\sigma^*)$, so we have

$$\begin{aligned} \Pi_i(\sigma_i, \gamma_{-i}^*) \quad &= \quad \int_T \int_{U_i} \pi_i(\sigma_i(t_i,u_i), \gamma_{-i}^*(t_{-i}), t, u_i) \nu_i(du|t_i) \kappa(dt) \\ &> \quad \int_T \int_{U_i} \pi_i(\sigma_i^*(t_i,u_i), \gamma_{-i}^*(t_{-i}), t, u_i) \nu_i(du|t_i) \kappa(dt) \\ &= \quad \Pi_i(\sigma_i^*, \gamma_{-i}^*). \end{aligned}$$

But then there is a Borel set $Q \subseteq T_i \times U_i$ with positive μ_i-measure such that for all $(t_i,u_i) \in Q$, we have $t_i \notin \tilde{T}_i$ and

$$\Pi_i(\sigma_i(t_i,u_i), t_i, u_i; \gamma^*) \quad > \quad \Pi_i(\sigma_i^*(t_i,u_i), t_i, u_i; \gamma^*).$$

Thus, for all $(t_i,u_i) \in Q$, we have $\sigma_i^*(t_i,u_i) \notin M_i(t_i,u_i;\gamma^*)$, contradicting the choice of σ_i^*. □

We now define a useful correspondence for each player i. For each $\gamma \in \Gamma$, define $M_i^*(\cdot;\gamma) \colon T_i \rightrightarrows \Re^d$ by

$$M_i^*(t_i;\gamma) \quad = \quad \int_{U_i} M_i(t_i,u_i;\gamma) \nu_i(du_i|t_i),$$

with values equal to the Aumann integral of the choice correspondence $M_j(t_i,u_i;\gamma)$ with respect to u_i. Our purification argument rests on the next lemma, which as a special case provides access to the inner structure of the correspondence $M_i^*(\cdot;\gamma)$ of average optimal actions when private types are non-atomically distributed. Let $\overline{U}_i \in \mathcal{U}_i$ contain the atoms of $\{\nu_i(\cdot|t_i) : t_i \in T_i\}$.

Lemma 5.3. *For each $i = 1, \ldots, n$, each $\gamma \in \Gamma$, and each $t_i \in T_i$, we have*

$$\int_{U_i \setminus \overline{U}_i} extM_i(t_i, u_i; \gamma)\nu_i(du_i|t_i) \quad = \quad \int_{U_i \setminus \overline{U}_i} M_i(t_i, u_i; \gamma)\nu_i(du_i|t_i).$$

Proof. Fix $\gamma \in \Gamma$ and $t_i \in T_i$. It is clear that the integral of extreme points of $M_i(t_i, u_i; \gamma)$ is contained in the integral of $M_i(t_i, u_i; \gamma)$. For the opposite inclusion, note that $u_i \mapsto M_i(t_i, u_i; \gamma)$ is integrably bounded. Then

$$\begin{aligned} \int_{U_i \setminus \overline{U}_i} \text{ext}M_i(t_i, u_i; \gamma)\nu_i(du_i|t_i) &= \int_{U_i \setminus \overline{U}_i} \text{coext}M_i(t_i, u_i; \gamma)\nu_i(du_i|t_i) \\ &= \int_{U_i \setminus \overline{U}_i} M_i(t_i, u_i; \gamma)\nu_i(du_i|t_i), \end{aligned}$$

where the first equality follows from nonatomicity of $\nu_i(\cdot|t_i)$ on $U_i \setminus \overline{U}_i$ and a version of Lyapunov's theorem (see Theorem 8.6.3 of Aubin and Frankowska (1990)), and the second follows from Lemma 5.1, which characterizes $M_i(t_i, u_i; \gamma)$ as a face of Δ. This completes the proof of the lemma. □

The final lemma of this section verifies that the correspondence of extreme optimal actions, $u_i \mapsto \text{ext}M_i(t_i, u_i; \gamma)$, is lower measurable given any general type $t_i \in T_i \setminus \tilde{T}_i$ and any average action profile γ.[17]

Lemma 5.4. *For each $i = 1, \ldots, n$, each $\gamma \in \Gamma$, and each $t_i \in T_i \setminus \tilde{T}_i$, the correspondence $u_i \mapsto extM_i(t_i, u_i; \gamma)$ is lower measurable.*

Proof. Given γ and $t_i \in T_i \setminus \tilde{T}_i$, Lemma 5.1 implies that $M_i(t_i, u_i; \gamma)$ is a face of Δ, and it follows from strict convexity of the Euclidean norm that

$$\text{ext}M_i(t_i, u_i; \gamma) \quad = \quad \arg\max\{||\alpha_i|| : \alpha_i \in M_i(t_i, u_i; \gamma)\}.$$

Fixing t_i and γ, and using Lemma 5.1, a measurable version of the theorem of the maximum (AB, Theorem 18.19) implies that $\text{ext}M_i(t_i, \cdot; \gamma)$ is measurable and therefore (AB, Lemma 18.2) lower measurable. □

We can now address part (a) of Theorem 4.2 on purification, assuming that $\overline{U}_i = \varnothing$ for all players. Let σ^* be a BNE with corresponding average actions γ^*. By Lemma 5.2, it follows that γ^* is a κ_i-almost everywhere selection from $M_i^*(\cdot; \gamma^*)$. And Lemma 5.3 and $\overline{U}_i = \varnothing$ imply that for all $t_i \in T_i \setminus \tilde{T}_i$,

$$M_i^*(t_i; \gamma^*) \quad = \quad \int_{U_i} \text{ext}M_i(t_i, u_i; \gamma^*)\nu_i(du|t_i).$$

[17]This result can be derived from Theorem 3 of Himmelberg and Van Vleck (1975), but it is easily proved directly using the finite dimensionality of our problem.

Thus, for κ_i-almost all t_i, we have $\gamma_i^*(t_i) \in \int_{U_i} \text{ext} M_i(t_i, u_i)\nu_i(du|t_i)$. Since the correspondence $u_i \mapsto \text{ext} M_i(t_i, u_i; \gamma)$ is lower measurable, by Lemma 5.4, we can then apply the theorem of Artstein (1989) with this correspondence in place of his F (satisfying Artstein's other conditions as in the proof of Lemma 3.1). Therefore, there exists a Borel measurable mapping $\hat{\sigma}_i: T_i \times U_i \to \Delta$ such that for κ_i-almost all t_i, we have: $\gamma_i^*(t_i) = \int_{U_i} \hat{\sigma}_i(t_i, u_i)\nu_i(du|t_i) = \gamma_i(\cdot|\hat{\sigma}_i)$, and for $\nu_i(\cdot|t_i)$-almost all u_i, $\hat{\sigma}_i(t_i, u_i) \in \text{ext} M_i(t_i, u_i)$. By Lemma 5.2, it follows that $\hat{\sigma}$ is a BNE. This completes the proof of part (a).

We now complete the proof of parts (b) and (c) of Theorem 4.2 by first proving Proposition 4.3 as an intermediate step. First, we show that for each player i and all (t_i, u_i), there is an undominated pure action. For technical purposes, define the relation Wdom on Δ so that α_i' Wdom α_i if and only if for all γ_{-i}, we have $\Pi_i(\alpha_i', t_i, u_i; \gamma_{-i}) \geq \Pi_i(\alpha_i, t_i, u_i; \gamma_{-i})$. Obviously, this relation is transitive. Since $\Pi_i(\alpha_i, t_i, u_i; \gamma_{-i})$ is continuous in (α_i, γ_{-i}), it follows that the relation is continuous, i.e., $\{(\alpha_i', \alpha_i) : \alpha_i' \text{ Wdom } \alpha_i\}$ is closed as a subset of $\Delta \times \Delta$. Thus, for all α_i, the set $\{\alpha_i' \in \Delta : \alpha_i' \text{ Wdom } \alpha_i\}$ is nonempty and compact, and we conclude that there is a mixed action α_i^* that is maximal, in the sense that for all α_i, α_i Wdom α_i^* implies α_i^* Wdom α_i (see, e.g., Proposition A.1 of Banks, Duggan, and Le Breton (2006)). In particular, α_i^* is undominated for i at (t_i, u_i). Finally, note that α_i^* cannot put positive weight on a dominated pure action, i.e., if $\alpha_{i,\ell}^* > 0$, then e^ℓ is undominated. Therefore, $E_i(t_i, u_i) \neq \emptyset$, as required.

To show lower measurability of E_i for each player i, we first define the mapping $f: \Delta \times \Delta \times T_i \times U_i \times \Gamma_{-i} \to \mathfrak{R}$ by

$$f_i(\alpha_i, \alpha_i', t_i, u_i; \gamma_{-i}) = \Pi_i(\alpha_i, t_i, u_i; \gamma_{-i}) - \Pi_i(\alpha_i', t_i, u_i; \gamma_{-i}),$$

which intuitively gives a measure of the advantage of α_i over α_i' under certain conditions. Given (t_i, u_i) and $\alpha_i \in \text{co}A_i(t_i, u_i)$, we define three programs:

$$U_i(\alpha_i, \alpha_i'; t_i, u_i) = \max_{\gamma_{-i} \in \Gamma_{-i}} f_i(\alpha_i, \alpha_i', t_i, u_i; \gamma_{-i})$$

$$V_i(\alpha_i, \alpha_i'; t_i, u_i) = \min_{\gamma_{-i} \in \Gamma_{-i}} f_i(\alpha_i, \alpha_i', t_i, u_i; \gamma_{-i})$$

and

$$W_i(\alpha_i; t_i, u_i) = \min_{\alpha_i' \in \text{co}A_i(t_i, u_i)} V_i(\alpha_i, \alpha_i'; t_i, u_i) \quad \text{s.t. } U_i(\alpha_i, \alpha_i'; t_i, u_i) \leq 0.$$

To see that these programs are well-defined, recall that Γ_{-i} compact and the interim expected payoff function $\Pi_i(\alpha_i, \gamma_{-i}; t_i, u_i)$ jointly continuous in (α_i, γ_{-i}). Thus, $f_i(\alpha_i, \alpha_i', t_i, u_i; \gamma_{-i})$ is jointly continuous in $(\alpha_i, \alpha_i', \gamma_{-i})$, and by the theorem of the maximum (AB, Theorem 17.31), $U_i(\alpha_i, \alpha_i'; t_i, u_i)$ is jointly continuous in (α_i, α_i'). Similarly, $V_i(\alpha_i; t_i, u_i)$ is jointly continuous in (α_i, α_i'). Note that $V_i(\alpha_i; t_i, u_i) \le 0$, for setting $\alpha_i' = \alpha_i$ yields a value of zero for the objective function. Therefore, the objective function of the third program is continuous, the constraint set is a nonempty (by setting $\alpha_i' = \alpha_i$) and compact subset of $\mathrm{co}A_i(t_i, u_i)$. Importantly, note that $\alpha_i \in \mathrm{co}A_i(t_i, u_i)$ is undominated at (t_i, u_i) if and only if $W_i(\alpha_i; t_i, u_i) = 0$.

For each pure action $e^\ell \in \mathrm{ext}\Delta$, let $Q_i^\ell = \{(t_i, u_i) \in T_i \times U_i : e^\ell \in A_i(t_i, u_i)\}$ be the set of types for player i such that e^ℓ is feasible. Since A_i is compact-valued and lower measurable, it is measurable (AB, Lemma 18.2), so Q_i^ℓ is Borel measurable. For each $\ell = 1, \ldots, d$ with $Q_i^\ell \neq \emptyset$, we claim that the correspondence $B_i^\ell : Q_i^\ell \rightrightarrows \Delta$ defined by

$$B_i^\ell(t_i, u_i) \quad = \quad \{\alpha_i' \in \Delta : U_i(e^\ell, \alpha_i'; t_i, u_i) \le 0\}$$

is lower measurable. Observe that the first objective function $f_i(e^\ell, \alpha_i', t_i, u_i; \gamma_{-i})$ is continuous in γ_{-i} and measurable in (α_i, t_i, u_i), so it is a Carathéodory function. By the measurable maximum theorem (AB, Theorem 18.19), it follows that $U_i(e^\ell, \alpha_i'; t_i, u_i)$ is measurable in (t_i, u_i). Because it is continuous in α_i', we conclude that $U_i(e^\ell, \cdot)$ is itself a Carathéodory function; similar remarks hold for $V_i(e^\ell, \cdot)$. Then the correspondence $(t_i, u_i) \mapsto \{\alpha_i' \in \Delta : U_i(e^\ell, \alpha_i'; t_i, u_i) < 0\}$ is lower measurable (AB, Lemma 18.7), and since Δ is compact, the correspondence $(t_i, u_i) \mapsto \{\alpha_i' \in \Delta : U_i(e^\ell, \alpha_i'; t_i, u_i) = 0\}$ is lower measurable (AB, Corollary 18.8 and Lemma 18.2). As the union of lower measurable correspondences, we conclude that B_i is lower measurable (AB, Lemma 18.4), as claimed. Furthermore, the values of B_i are nonempty, by setting $\alpha_i' = e^\ell$, and compact, by continuity of U_i in a_i'.

Let $W_i^\ell : Q_i^\ell \to \mathfrak{R}$ be the restriction of $W_i(e^\ell, \cdot)$ to Q_i^ℓ. By (B2), the restriction of the correspondence A_i to Q_i^ℓ is lower measurable with nonempty, compact values. Fixing $\alpha_i = e^\ell$, it follows that the constraint correspondence $C_i^\ell : Q_i^\ell \rightrightarrows \Delta$ for the third program, given by

$$C_i^\ell(t_i, u_i) \quad = \quad A_i(t_i, u_i) \cap B_i^\ell(t_i, u_i),$$

is lower measurable (AB, Lemma 18.4). Moreover, the values of C_i are compact and, since $e^\ell \in A_i(t_i, u_i)$ and $U_i(e^\ell, e^\ell; t_i, u_i) = 0$ for all $(t_i, u_i) \in Q_i^\ell$, nonempty as well. Then, because the objective function $V_i(e^\ell, \cdot)$ is a Carathéodory function, the measurable maximum theorem (AB, Theorem 18.19) implies that W_i^ℓ is Borel measurable. In particular, the set

$$\{(t_i, u_i) \in Q_i^\ell : W_i^\ell(t_i, u_i) = 0\} \quad = \quad \{(t_i, u_i) \in T_i \times U_i : e^\ell \in E_i(t_i, u_i)\}$$

is Borel measurable. Then the graph of E_i,

$$\text{graph}E_i = \bigcup_{\ell=1}^{d}\{(t_i,u_i)\in T_i\times U_i : e^\ell \in E_i(t_i,u_i)\}\times\{e^\ell\},$$

is Borel measurable. Because the range of E_i is a finite set, we conclude that E_i is in fact lower measurable. This completes the proof of Proposition 4.3.

Returning to Theorem 4.2, define the feasible action correspondence $A_i' : T_i \times U_i \rightrightarrows \Delta$ by $A_i'(t_i,u_i) = E_i(t_i,u_i)$, which is lower measurable. To prove part (b) of the theorem, we replace A_i with A_i' for each player i, yielding a product Bayesian game satisfying (B1)–(B5), and then Theorem 4.1 yields a BNE σ of the transformed game. Under non-atomicity, there is an equivalent pure-strategy BNE σ_i' of the transformed game, and by construction, we have $\sigma_i'(t_i,u_i) \in A_i'(t_i,u_i) = E_i(t_i,u_i)$ for each player i and all (t_i,u_i). Thus, σ' is a pure-strategy, undominated BNE of the original game. Indeed, if it were not a BNE, then there would exist a profitable deviation $\hat{\sigma}_i$ in the original game for some player i. We argue that we could then construct a profitable deviation $\tilde{\sigma}_i$ for i in the transformed game, a contradiction. Let Q be the measurable set of types (t_i,u_i) such that $\hat{\sigma}_i(t_i,u_i)$ assigns positive probability to some dominated pure action. Note that for each $(t_i,u_i) \in Q$, it follows that $\hat{\sigma}_i(t_i,u_i)$ is dominated and thus not maximal with respect to Wdom. By Proposition A.2 of Banks, Duggan, and Le Breton (2006), there is a mixed action $\alpha_i \in \text{co}A_i(t_i,u_i)$ that is maximal with respect to Wdom and such that α_i Wdom $\hat{\sigma}_i(t_i,u_i)$. Moreover, maximality of α_i implies that it puts positive probability only on undominated pure actions, i.e., $\alpha_i \in \text{co}E_i(t_i,u_i)$.

We then define the correspondences $F : Q \rightrightarrows \Delta$ and $H : \rightrightarrows \Delta$ by

$$\begin{aligned} F(t_i,u_i) &= \{\alpha_i \in \Delta : \alpha_i \text{ Wdom } \hat{\sigma}_i(t_i,u_i)\} \\ H(t_i,u_i) &= F(t_i,u_i) \cap \text{co}E_i(t_i,u_i), \end{aligned}$$

which have nonempty, closed values. From our earlier observations regarding weak domination, the correspondence $\alpha_i' \mapsto \{\alpha_i \in \Delta : \alpha_i \text{ Wdom } \alpha_i'\}$ is upper hemicontinuous, so given any closed set $M \subseteq \Delta$ of mixed actions, the set $L = \{\alpha_i' \in \Delta : \exists \alpha \in M \text{ s.t. } \alpha \text{ Wdom } \alpha_i'\}$ is closed (AB, Lemma 17.4), and the lower inverse of M under F is $\hat{\sigma}_i^{-1}(L)$, which is measurable. We conclude that F is measurable and thus lower measurable (AB, Lemma 18.2). We have shown that E_i is lower measurable, and thus $\text{co}E_i$ is lower measurable as well. Combining these observations, the correspondence H is lower measurable (AB, Lemma 18.4). By the Kuratowski-Ryll-Nardzewski selection theorem (AB, Theorem 18.13), the correspondence H admits a measurable selection, $\phi: Q \to \Delta$. Define the strategy $\tilde{\sigma}_i$ so that

$$\tilde{\sigma}_i(t_i, u_i) = \begin{cases} \phi(t_i, u_i) & \text{if } (t_i, u_i) \in Q, \\ \hat{\sigma}_i(t_i, u_i) & \text{else.} \end{cases}$$

By construction, we then have $\Pi_i(\tilde{\sigma}_i, \sigma'_{-i}) \geq \Pi_i(\hat{\sigma}_i, \sigma'_{-i}) > \Pi_i(\sigma')$, contradicting the assumption that σ' is a BNE of the transformed game.

To prove part (c), consider any undominated BNE $\hat{\sigma}$ of the original game. Then we have $\hat{\sigma}_i(t_i, u_i) \in D_i(t_i, u_i) \subseteq \text{co}A'_i(t_i, u_i)$ for each player i and all (t_i, u_i), so $\hat{\sigma}$ is a BNE of the transformed game. We apply Theorem 4.1 to obtain an equivalent pure-strategy BNE $\hat{\sigma}'$ of the transformed game, which (by an argument analogous to that for part (b)) is a pure-strategy, undominated BNE of the original game.

References

[1] C. Aliprantis and K. Border (2006) *Infinite Dimensional Analysis: A Hitchhiker's Guide, Third Edition,* New York: Springer.

[2] S. Athey (2001) "Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information," *Econometrica,* 69: 861–889.

[3] Z. Artstein (1979) "A Note on Fatou's Lemma in Several Dimensions," *Journal of Mathematical Economics,* 6: 277–282.

[4] Z. Artstein (1989) "Parametrized Integration of Multifunctions with Applications to Control and Optimization," *SIAM Journal of Control and Optimization,* 27: 1369–1380.

[5] J.-P. Aubin and H. Frankowska (1990) *Set-Valued Analysis,* Boston: Birkhauser.

[6] E. Balder (1988) "Generalized Equilibrium Results for Games with Incomplete Information," *Mathematics of Operations Research,* 13: 265–276.

[7] E. Balder (2002) "A Unifying Pair of Cournot-Nash Equilibrium Existence Results," *Journal of Economic Theory,* 102: 437–470.

[8] E. Balder (2003) "On Undominated Nash Equilibria for Games with a Measure Space of Players," *Economics Letters,* 80: 137–140.

[9] E. Balder (2008) "Comments on Purification in Continuum Games," *International Journal of Game Theory,* 37: 73–92.

[10] E. Balder and A. Rustichini (1994) "An Equilibrium Result for Games with Private Information and Infinitely Many Players," *Journal of Economic Theory,* 62: 385–393.

[11] J. Banks, J. Duggan, and M. Le Breton (2006) "Social Choice and Electoral Competition in the General Spatial Model," *Journal of Economic Theory,* 126: 194–234.

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

[13] L. de Castro (2012) "Correlation of Types in Bayesian Games," working paper.

[14] J. Duggan (2012) "Noisy Stochastic Games," *Econometrica,* 80: 2017–2045.

[15] N. Dunford and J. Schwartz (1957) *Linear Operators: Part 1,* New York: Interscience.

[16] T. Feddersen and W. Pesendorfer (1997) "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica,* 65: 1029–1058.

[17] H. Fu, Y. Sun, N. Yannelis, and Z. Zhang (2007) "Pure Strategy Nash Equilibrium in Games with Private and Public Information", *Journal of Mathematical Economics*, 43, 523-531.

[18] M. Greinecker and K. Podczeck (2013) "Spurious Purification," unpublished, University of Vienna.

[19] W. Hildenbrand (1974) *Core and Equilibria of a Large Economy,* Princeton: Princeton University Press.

[20] C. Himmelberg and F. Van Vleck (1975) "Multifunctions with Values in a Space of Probability Measures," *Journal of Mathematical Analysis and Applications,* 50: 108–112.

[21] M. Kahn, K. Rath, and Y. Sun (2006) "The Dvoretzky-Wald-Wolkowitz Theorem and Purification in Atomless Finite-Action Games," *International Journal of Game Theory,* 34: 91–104.

[22] T. Kim and N. Yannelis (1997) "Existence of Equilibrium in Bayesian Games with Infinitely Many Players," *Journal of Economic Theory,* 77: 330–353.

[23] M. Le Breton and M. Weber (1997) "On Existence of Undominated Pure Strategy Nash Equilibria in Anonymous Nonatomic Games," *Economics Letters,* 56: 171–175.

[24] D. McAdams (2006) "Monotone Equilibrium in Multi-unit Auctions," *Review of Economic Studies,* 73: 1039–1056.

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

[26] S. Morris and H. Shin (2003) "Global Games: Theory and Applications," in M. Dewatripont, L. Hansen and S. Turnoksky (Eds.), Advances in Economic Theory and Econometrics: Proceedings of the Eighth World Congress of the Econometric Society, Cambridge University Press, Cambridge.

[27] P. Milgrom and R. Weber (1985) "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research,* 619–632.

[28] D. Myatt (2012) "A Rational Choice Theory of Voter Turnout," working paper, London Business School.

[29] K. Podczeck (2009) "On Purification of Measure-valued Maps," *Economic Theory,* 38: 399–418.

[30] R. Radner and W. Rosenthal (1982) "Private Information and Pure-Strategy Equilibria," *Mathematics of Operations Research,* 7, 401-409.

[31] M. Rao (1993) *Conditional Measures and Applications,* New York: Marcel Dekker.

[32] P. Reny (2011) "On the Existence of Monotone Pure Strategy Equilibria in Bayesian Games," *Econometrica,* 79, 499-553.

[33] L. Simon and M. Stinchcombe (1995) "Equilibrium Refinements for Infinite Normal Form Games," *Econometrica,* 63, 1421-1443.

[34] T. Van Zandt and X. Vives (2007) "Monotone Equilibria in Bayesian Games of Strategic Complementarities," *Journal of Economic Theory,* 134, 339-360.

[35] T. Van Zandt (2011) "Interim Bayesian Nash Equilibrium on Universal Type Spaces for Supermodular Games," *Journal of Economic Theory,* 145, 249-263.

[36] J. Wang and Y. Zhang (2010) "Purification, Saturation, and the Exact Law of Large Numbers," *Economic Theory,* forthcoming.

DEPT. OF ECONOMICS, UNIVERSITY OF ROCHESTER.

DEPT. OF POLITICAL SCIENCE AND DEPT. OF ECONOMICS, UNIVERSITY OF ROCHESTER.