



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Balance and discontinuities in infinite games with type-dependent strategies [☆]

Maxwell B. Stinchcombe ^{*}

Department of Economics, University of Texas, Austin, TX 78712-1173, USA

Received 20 March 2007; final version received 2 June 2010; accepted 8 October 2010

Available online 3 December 2010

Abstract

Under study are games in which players receive private signals and then simultaneously choose actions from compact sets. Payoffs are measurable in signals and jointly continuous in actions. Stinchcombe (2011) [19] proves the existence of correlated equilibria for this class of games. This paper is a study of the information structures for these games, the discontinuous expected utility functions they give rise to, and the notion of a balanced approximation to an infinite game with discontinuous payoffs.

© 2010 Elsevier Inc. All rights reserved.

JEL classification: C62; C72; C65

Keywords: Infinite games; Continuous informational commonalities; Balanced approximations; Prevalence; Shyness

1. Introduction

This paper studies the properties of information structures of games in which players receive private signals (their types), and then simultaneously choose actions from compact sets. By assumption, the payoffs are measurable in signals, jointly continuous in actions, and integrable. This class of games has been used to model firm competition with private information, strategic signaling, purification of mixed strategy equilibria, and wars of attrition.

[☆] Many thanks to Mark Machina, Preston McAfee, and Dale Stahl for their timely and much appreciated help, and especially to Bob Anderson, Chris Harris, Leo Simon and Bill Zame for many conversations about infinite games. The associate editor and anonymous referees deserve the credit for vast improvements in the clarity and cohesiveness of this work.

^{*} Fax: +1 512 471 3510.

E-mail address: max.stinchcombe@gmail.com.

Beyond their widespread use, there is a good theoretical reason to study this class of infinite games. They are infinite extensive form games involving serious informational considerations, yet still simple enough that normal form analyses are feasible. In this way, they provide a logical bridge between infinite normal form games and infinite extensive form games.

1.1. Diffuse information

The companion to this paper [19] showed that, despite discontinuous expected utility payoff functions, these games always have correlated equilibria. Simon [16] showed that these games need not have Nash equilibria. However, if the joint distribution of signals is absolutely continuous with respect to the product of its marginals, a diffuseness of information condition due to Milgrom and Weber [13], expected utility functions are jointly continuous and the game has an equilibrium.¹

Theorem A shows that, with at most a small set of exceptions, failing the Milgrom and Weber diffuseness condition has the following implication: the existence of a non-null set of signals, B , perhaps not in any player's information set, conditional on which two or more players can infer the value of some continuously distributed random variable. A simpler situation in which we have this failure has two or more players observing (or being able to infer), say as one component of a vector-valued signal, the value of a continuously distributed random variable, in which case we can take $P(B) = 1$. Theorem A also shows that the existence of this kind of informational commonality is strongly nongeneric in the set of all joint distributions of signals, a result in accord with some, but not all, intuitions about the richness of informational environments.

1.2. Discontinuities

Theorem B examines the discontinuities of the expected utility functions that arise if there is a commonly inferable continuously distributed random variable. The result is that, for a generic set of utility functions, the expected utility functions are discontinuous in a number of equivalent senses.

1. The strategy sets cannot be approximated by finite sets in the Fudenberg and Levine [8] “most utility difference it can make to anyone” pseudo-metric.
2. The oscillation of the utility functions across elements of any finite product partition of the strategy spaces is bounded away from 0.
3. There exists no way to embed the strategy sets into compact spaces so that the expected utility functions are jointly continuous.

In the presence of such discontinuities, many intuitions/guesses about the limit behavior of approximate games and equilibria are incorrect; Milgrom and Weber showed that limits of equilibrium strategies may not be equilibria (see Example 3.1); the limits of strategies that fail to be ϵ -equilibria for e.g. $\epsilon = 1$ turn out to be correlated equilibria (see Example 5.1). Such examples indicate that fine details of the approximations to these infinite games may matter.

¹ Milgrom and Weber [13] unified and extended a large and disparate literature. Balder [3] weakened the assumptions needed for existence in all other aspects of the description of the games except the diffuseness condition on the joint distribution of signals.

1.3. Balanced and unbalanced approximations

An approximation to an infinite game is a game in which each player is restricted to a subset of their full strategy set. Seemingly well-behaved approximations to infinite games can alter and deform information structures.² An approximation is ϵ -balanced if each player can, against any vector of the other players' strategies in the approximation, guarantee payoffs that are within ϵ of what they can achieve using their full strategy set. A net of approximations is balanced if, for all $\epsilon > 0$, it is eventually ϵ -balanced.

The essential device in the existence proof for correlated equilibria in [19] is the use of limits of nets of equilibria when the players are restricted to play the game as if their information arose from finite sub-partitions. A side-effect of restricting the players to use less than their full information is that the approximations may be unbalanced.

An equilibrium of an ϵ -balanced approximation is necessarily an ϵ -equilibrium. An equilibrium of an unbalanced approximation may fail to be an approximate equilibrium. As detailed in Section 5, the correlated equilibria studied here may be the limits of strategies that fail to be approximate equilibria in this fashion.³

1.4. Outline

The next section sets the assumptions and notation. Section 3 studies the interpretation of the informational diffuseness condition known to be sufficient for Nash equilibrium existence. Section 4 shows that the existence of continuously distributed informational commonalities matters lead, generically, to strongly discontinuous expected utility functions. Section 5 examines what this class of infinite games teach us about balanced and unbalanced approximations. Section 6 contains conclusions and complements.

2. Notation and assumptions

The notation and assumptions are exactly as in the companion paper [19], and are repeated here for ease of reference.

For each $i \in I$, I a finite set of players, the “type” ω_i belongs to a measure space $(\Omega_i, \mathcal{F}_i)$. The joint distribution of $\omega = (\omega_i)_{i \in I} \in \Omega = \prod_i \Omega_i$ is given by a countably additive probability P defined on a σ -field \mathcal{F} , $\bigotimes_i \mathcal{F}_i \subset \mathcal{F}$. Summarizing, an *information structure* is a triple, $(\prod_i (\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$.

Each $i \in I$ has a compact, metric action space A_i , and $A := \prod_i A_i$. The utility functions, u_i , are assumed to belong to $L^1(P; C(A))$, the set of integrable functions from Ω to the separable Banach space $C(A)$ (with the supnorm, $\|\cdot\|_\infty$, the associated topology and Borel σ -field). Specifically, the assumption is that for all $i \in I$, $\int_\Omega \|u_i(\omega)\|_\infty P(d\omega) < \infty$. Player i receives utility $u_i(\omega)(a)$ if ω occurs and a is chosen by the players.

² In signaling games and in games of almost perfect information, the limits do no more than introduce cheap talk, [9,12]. In normal form games, they may introduce specialized utility transfers by a randomizing referee [15]. More generally, approximations can destroy information structures, forcing unwilling revelation of private information, allowing the observation of what should be unobservable, or allowing players to hide information that they should not be able to conceal [18].

³ As a referee pointed out, if one takes correlated equilibrium to be the appropriate solution concept, then including correlation that arises as the limit of non-equilibrium phenomena is not problematic.

Δ_i is the set of (countably additive) Borel probabilities on A_i with the weak* topology and the corresponding σ -field. $\mathbb{B}_i(\mathcal{F}_i)$ is i 's set of behavioral strategies, the \mathcal{F}_i -measurable functions from Ω_i to Δ_i . $\mathbb{B}_i(\mathcal{F}_i)$ is given the weak* topology, so that a sequence (or net if need be) $b_i^n \rightarrow b_i$ iff $\int_{\Omega} \langle v_i(\omega), b_i^n(\omega) \rangle P(d\omega) \rightarrow \int_{\Omega} \langle v_i(\omega), b_i(\omega) \rangle P(d\omega)$ for all $v_i \in L^1(P; C(A_i))$ where $\langle f, \mu_i \rangle := \int_{A_i} f(a_i) \mu_i(da_i)$ for $f \in C(A_i)$ and $\mu_i \in \Delta_i$. Cotter [6] showed that $\mathbb{B}_i(\mathcal{F}_i)$ is compact, and metrizable if \mathcal{F}_i is countably generated.

Given a vector $b = (b_i)_{i \in I} \in \mathbb{B} := \prod_i \mathbb{B}_i(\mathcal{F}_i)$, player i 's expected utility if b is played is defined by

$$u_i^P(b) = \int_{\Omega} \left\langle u_i(\omega), \prod_i b_i(\omega) \right\rangle P(d\omega) \tag{1}$$

where $\langle f, \nu \rangle := \int_A f(a) \nu(da)$ for continuous $f : A \rightarrow \mathbb{R}$ and Borel probabilities ν , and $\prod_i b_i$ is the product probability on A having b_i as the i 'th marginal. For $b \in \mathbb{B}$ and $b'_i \in \mathbb{B}_i$, $(b \setminus b'_i)$ denotes the strategy vector b with b'_i substituted into the i 'th component. $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i \in I}$ denotes the normal form game.

Definition 2.1. A (Nash) equilibrium for $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i \in I}$ is a vector $b \in \mathbb{B}$ such that for all $i \in I$ and all $b'_i \in \mathbb{B}_i$, $u^P(b) \geq u^P(b \setminus b'_i)$.

3. Continuously distributed informational commonalities

A continuously distributed informational commonality (CIC) arises if there is a continuously distributed random variable, defined on a non-null set B of joint signals, and two or more players can, after being told that B has occurred and observing their own signal, infer its value. The known sufficient condition for Nash existence is that the joint distribution of the players' signals is absolutely continuous with respect to the product of the marginal distributions. The information diffuseness condition rules out any CIC.

Theorem A shows that the set of joint distributions that have a CIC is generic in a very strong sense, and is a subset of the joint distributions that fail the diffuseness condition. This provides a strategic interpretation of the diffuseness condition, and shows that the known results for Nash existence only apply to a nongeneric class of games. The last part of this section examines several intuitive and non-intuitive aspects of the smallness of this nongeneric set.

3.1. Continuously distributed informational commonalities

For $B \in \mathcal{F}$ and \mathcal{G} a sub- σ -field of \mathcal{F} , $\mathcal{G}|_B := \{E \cap B : E \in \mathcal{G}\}$ is the trace of \mathcal{G} on B . For Q a probability and f a measurable function, $f(Q)$ denotes the image law of Q under the function f . For $B \in \mathcal{F}$ with $P(B) > 0$, $P_B(\cdot) := P(\cdot|B)$.

Definition 3.1. A joint distribution of signals, P , has a continuously distributed informational commonality (CIC) if there exists a $B \in \mathcal{F}$, $P(B) > 0$, and $\varphi : B \rightarrow (0, 1]$ that is $(\mathcal{F}_i \cap \mathcal{F}_j)|_B$ -measurable, and $\varphi(P_B)$ is non-atomic.

The continuous distribution of the commonly observed signal allows for arbitrarily fine coordination between players. This gives rise to discontinuous expected utility functions.

Example 3.1 (Milgrom and Weber, Cotter). For the two players, $\Omega_1 = \Omega_2 = [0, 1]$, P is the uniform distribution on the diagonal so that, with probability 1, the common value of the ω_i is known to each player. The action spaces are $A_i = \{L_i, R_i\}$. Payoffs are (10, 10) if the players coordinate on (L_1, L_2) , (2, 2) if they coordinate on (R_1, R_2) , and (0, 0) otherwise. Expected payoffs are (6, 6) if both play the strategy $b_i^n(\omega_i) = \delta_{L_i}$ (pointmass on L_i) if $\omega_i \in (k/2^n, (k+1)/2^n]$ with k even, $b_i^n(\omega_i) = \delta_{R_i}$ otherwise.

Let η^n denote the distribution on $\Omega \times A$ induced by b^n . The η^n have a unique weak limit, η , determined by the equalities

$$\eta(E \times \{(L_1, L_2)\}) = \frac{1}{2}P(E), \quad \eta(E \times \{(R_1, R_2)\}) = \frac{1}{2}P(E), \tag{2}$$

$E \in \mathcal{F}$. η represents a public signal correlated equilibrium, and has expected payoffs (6, 6). By contrast, the unique weak* limit of the b_i^n is uncoordinated strategy (b_1^∞, b_2^∞) where $b_i^\infty(\omega) \equiv \frac{1}{2}\delta_{L_i} + \frac{1}{2}\delta_{R_i}$. This uncoordinated, non-equilibrium strategy vector delivers payoffs of (3, 3).

The informational commonality in Example 3.1 has B being the diagonal in $\Omega = [0, 1] \times [0, 1]$, and has φ equal to projection onto either axis. The measurability of φ with respect to both \mathcal{F}_i and \mathcal{F}_j arises from P putting mass 1 on the diagonal. Conditioning on an unobservable set B can be seen in the following example.

Example 3.2. $\Omega = [0, 1] \times [0, 1]$ with the usual Borel σ -field. With probability $\frac{1}{2}$, (ω_1, ω_2) is uniformly distributed on the line, B , joining $(0, 0)$ to $(\frac{1}{2}, \frac{1}{2})$. With probability $\frac{1}{2}$, (ω_1, ω_2) is uniformly distributed on the complement of the rectangle $(0, \frac{1}{2}] \times (0, \frac{1}{2}]$. When $\omega_i < \frac{1}{2}$, player i does not know if $\omega_i = \omega_j$ or if ω_j is uniformly distributed over the interval $(\frac{1}{2}, 1]$. However, conditional on B , a set in neither player's information set, projection onto either axis provides the continuously distributed φ the value of which both players can infer.

There are informational commonalities that do not obviously involve non-atomic distributions on the graphs of invertible functions from a subset of Ω_i to a subset of Ω_j . The following informational commonality is based on [4, Ex. 31.1, pp. 407–408], which introduces, for $r \in (0, 1)$, $F_r(\cdot)$, the strictly increasing, continuous cdf of the random variable $X^r = \sum_{n=1}^\infty X_n 2^{-n}$ where X_1, X_2, \dots is an i.i.d. sequence with $P(X_n = 1) = r = 1 - P(X_n = 0)$. X^r is, with probability 1, concentrated on the set of $x \in (0, 1)$ having, in the limit, r of the terms in their non-terminating binary expansion equal to 1. The mapping $(r, s) \mapsto F_r^{-1}(s)$ is jointly continuous. If s is uniformly distributed on $(0, 1)$, then $F_r^{-1}(s)$ is distributed as X^r .

Example 3.3. Define $(\omega_1, \omega_2) = (F_r^{-1}(s_1), F_r^{-1}(s_2))$ where r, s_1, s_2 are i.i.d. and uniform, and let P be the joint distribution of ω_1 and ω_2 . To see that the value of r provides a continuous informational continuity, let $b_m(x)$ denote the m 'th digit of x 's non-terminating binary expansion, and let $\varphi_m(x) = \#\{n \leq m: b_n(x) = 0\}/m$. We have $r = \limsup_m \varphi_m(\omega_1) = \limsup_m \varphi_m(\omega_2)$ P -a.e. Take $\varphi: \Omega \rightarrow (0, 1]$ to be the uncountable-to-one common value of these limits.

If instead, $(\omega_1, \omega_2) = (r, F_r^{-1}(s_1))$, r is again commonly known, but the function φ has very different formulations for the two players, one as projection, the other as a limit of ratios.

3.2. Diffuseness and jointly continuous expected utility

The discontinuity of expected utilities in Example 3.1 requires correlation between players' actions at arbitrarily fine scales. The following condition rules out this kind of coordination.

Definition 3.2 (Milgrom and Weber). A joint distribution P is *continuous* if it is absolutely continuous with respect to the product of its marginals.

This condition delivers jointly continuous expected utility functions because integration against conditional densities is, essentially, the same operation that defines the weak* topology that makes the strategy sets compact.

If ω_i and ω_j are smoothly distributed, absolute continuity with respect to the product of marginals requires that i 's posterior distribution about ω_j have a density. This kind of diffuse information rules out the coordination at arbitrarily fine scales. Neither “continuity” nor “diffuseness” seem perfectly apt as names.

Example 3.4. If all but at most one of the ω_i have countable support, then P is continuous.

3.3. CIC's are generic

The following is based on Anderson and Zame's [1] finite shyness/prevalence.⁴

Definition 3.3. A set $S \subset \Delta(\mathcal{F})$ is *one-dimensionally tiny* if there exists a non-degenerate line $L = \{\alpha P + (1 - \alpha)P' : \alpha \in [0, 1]\} \subset \Delta(\mathcal{F})$ such that for all finite signed measures x on \mathcal{F} , $(L + x) \cap S$ contains at most one point. T is *one-dimensionally full* if $\Delta(\mathcal{F}) \setminus T$ is one-dimensionally tiny.

Suppose that S is a subset of C , C an n -dimensional convex subset of \mathbb{R}^ℓ , $n \geq 2$, e.g. $C = \Delta(\mathcal{F})$ if \mathcal{F} is generated by a finite partition of $n + 1$ non-null sets. S being one-dimensionally tiny implies that S is a Lebesgue null set relative to the affine hull of C in \mathbb{R}^ℓ , but the reverse is not true. If the signal spaces are topologically complete metric spaces, that is, the σ -field \mathcal{F}_i is the Borel σ -field for some metric making the space Ω_i a complete metric space, then the one-dimensionally tiny sets are a special subclass of Anderson and Zame's shy subsets of $\Delta(\mathcal{F})$.⁵

Let $CIC \subset \Delta(\mathcal{F})$ denote the set of joint distributions of signals for which there is a CIC and let $DIS \subset \Delta(\mathcal{F})$ denote the set of joint distributions that are not absolutely continuous with respect to the product of their marginals.

Theorem A. If \mathcal{F}_i and \mathcal{F}_j support countably additive non-atomic probabilities for at least two players, $i \neq j$, then:

⁴ Instead of requiring that $(L + x) \cap S$ contain at most one point as in the following definition, Anderson and Zame's one-dimensional shyness requires only that it be a Lebesgue null set of the α .

⁵ “Exceptional” or small sets in a form related to 1-shyness were first introduced by Aronszajn [2] to study the points of non-differentiability of Lipschitz mappings between Banach spaces. A larger class of exceptional sets was studied more extensively in Christensen [5] under the name of “Haar null sets.” They were independently described in Hunt, Sauer and Yorke [11], from where we have the names “shy” and “prevalent.” Anderson and Zame [1] solved the rather delicate problem of extending the definition of shy sets to subsets of convex sets that are themselves shy in the larger ambient space. Stinchcombe [17] discusses some interpretational issues for this class of small sets.

- (1) $CIC \subset DIS$, and
- (2) CIC is one-dimensionally full.

The following lemma will lead directly to Theorem A(1).

Lemma 3.1. *Suppose that, for each $i \in I$, $(\Omega'_i, \mathcal{F}'_i)$ is a measurable space, $\mathcal{F}' = \otimes_{i \in I} \mathcal{F}'_i$, $g_i : \Omega_i \rightarrow \Omega'_i$ is measurable, $Q_i = g_i(P_i)$, $g(\omega) = (g_i(\omega_i)_{i \in I})$, and $Q = g(P)$. Then $[\times_{i \in I} P_i \gg P] \Rightarrow [\times_{i \in I} Q_i \gg Q]$.*

Proof. The set $\mathcal{E}' = \{E' \in \mathcal{F}' : \times_{i \in I} Q_i(E') = \times_{i \in I} P_i(g^{-1}(E'))\}$ is a σ -field containing the measurable rectangles, so that $\mathcal{E}' = \mathcal{F}'$. Suppose that $\times_{i \in I} Q_i(E') = 0$. Showing that $Q(E') = 0$ will complete the proof. $\times_{i \in I} Q_i(E') = 0$ iff $\times_{i \in I} P_i(g^{-1}(E')) = 0$. Since $\times_{i \in I} P_i \gg P$, $P(g^{-1}(E')) = 0$, which in turn implies that $Q(E') = 0$. \square

Proof of Theorem A. (1) Suppose P has a CIC but that $\times_{i \in I} P_i \gg P$. Let $B \in \mathcal{F}$, $P(B) > 0$, with a measurable $\varphi : B \rightarrow (0, 1]$ such that $g := E^{P_B}(\varphi | \mathcal{F}_i) = E^{P_B}(\varphi | \mathcal{F}_j)$ P_B -almost everywhere for some pair of players $i \neq j$, and $g(P_B)$ is non-atomic. Define $g_i(\omega_i) = E^{P_B}(\varphi | \mathcal{F}_i)$, $g_j(\omega_j) = E^{P_B}(\varphi | \mathcal{F}_j)$, and $g(\omega_i, \omega_j) = (g_i(\omega_i), g_j(\omega_j))$. By Lemma 3.1, $g(P_B)$ must be absolutely continuous with respect to the product of its marginals. This contradicts $g(P_B)$ being a non-atomic distribution on the diagonal in $(0, 1] \times (0, 1]$, completing the proof of the first statement.

(2) Let $S = \Delta(\mathcal{F}) \setminus CIC$. There are two steps to showing that S is one-dimensionally shy: (A) there exist $P \neq P'$ having a CIC, and (B) when arbitrarily translated, the line joining P and P' intersects S at most once.

(A) There exist $P \neq P'$ having a CIC. There are three parts to this argument: (i), identifying a set B ; (ii), constructing a P ; and (iii), varying the construction of P to get a P' .

- (i) Identifying a set B : Suppose that $\mathcal{F}_i, \mathcal{F}_j$ both support non-atomic distributions, P_i, P_j . Construct a function φ_i on Ω_i and a function φ_j on Ω_j , both taking values in $(0, 1]$ such that $\varphi_i(P_i) = \varphi_j(P_j) = U$, where U is the uniform distribution on $(0, 1]$. Let $I_{k,n} = (\frac{k}{2^n}, \frac{k+1}{2^n}]$ be the half-open k 'th dyadic interval of order n . Define $B = \bigcap_{n \in \mathbb{N}} \bigcup_{k=0}^{2^n-1} [\varphi_i^{-1}(I_{k,n}) \times \varphi_j^{-1}(I_{k,n})]$.
- (ii) Constructing a P : For non-disjoint $(a, b], (c, d] \subset (0, 1]$, define $p_U(\varphi_i^{-1}((a, b]) \times \varphi_j^{-1}((c, d])) = r - s$ where $r = \max(a, c)$ and $s = \min(b, d)$. For disjoint $(a, b], (c, d] \subset (0, 1]$, define $p_U(\varphi_i^{-1}((a, b]) \times \varphi_j^{-1}((c, d])) = 0$. Let \mathcal{G}_B be the small σ -field of subsets of B containing the sets $\varphi_i^{-1}((a, b]) \times \varphi_j^{-1}((c, d])$. By the Carathéodory extension Theorem, there is a unique countably additive extension of p_U to a probability \widehat{p}_U on \mathcal{G}_B . By the Hahn–Banach Theorem, \widehat{p}_U has at least one extension, P , to \mathcal{F} . $P(\varphi_i = \varphi_j) = \widehat{p}(\varphi_i = \varphi_j) = 1$. Letting φ denote the common value of the functions, $\varphi(P)$ is the non-atomic uniform distribution on $(0, 1]$.
- (iii) Constructing a P' : Replace the uniform distribution, U , with a probability $V \neq U$ that is mutually absolutely continuous with U . This gives rise to a \widehat{p}_V that is mutually absolutely continuous with \widehat{p}_U . Let P' be an extension of \widehat{p}_V to \mathcal{F} .

(B) Arbitrary translations of the line joining P and P' intersect S at most once. Pick an arbitrary signed measure x on \mathcal{F} and set $L = \{\alpha P + (1 - \alpha)P' : \alpha \in [0, 1]\}$. Let $A_x = \{\alpha \in [0, 1] : x + (\alpha P + (1 - \alpha)P') \in \Delta(\mathcal{F})\}$. Because membership in $\Delta(\mathcal{F})$ is determined by weak linear inequalities, A_x is a closed, possibly empty, convex interval. To finish the proof, it is sufficient to show that at most one $\alpha \in A_x$ can belong to S .

Let $x|_B$ be the restriction of x to $\mathcal{G}|_B$, $x|_B^{ac}$ the part of $x|_B$ that is absolutely continuous with respect to \widehat{p}_U , hence with respect to \widehat{p}_V , and $x|_B^\perp$ the part of $x|_B$ that is singular with respect to \widehat{p}_U , hence with respect to \widehat{p}_V . There is at most one $\alpha^\circ \in A_x$ such that $x|_B^{ac} + (\alpha^\circ P + (1 - \alpha^\circ P'))$ is the 0 measure. Let $B' \subset B$ have the property that $x|_B^\perp(B') = 0$ while $P(B') = P'(B') = 1$. For any $\alpha \in A_x \setminus \{\alpha^\circ\}$, B' and $\varphi : B' \rightarrow (0, 1]$ are a CIC. \square

I conjecture that having a continuous informational continuity and failing absolutely continuity with respect to the product of marginals are not only generically equivalent, but completely equivalent.

3.4. Context and genericity

Whether or not genericity conclusions are compelling depends on whether or not the context is correct. For example, the set $S = \{(x_1, x_2, 0) : x_1 \neq x_2\}$ is a small subset of \mathbb{R}^3 , but is a very large subset of the projection of \mathbb{R}^3 onto its first two coordinates. The question is whether Theorem A's conclusion about the smallness of the set of information structures not having a CIC is driven by asking the question in too large a context, namely the set of all joint distributions of signals. A pair of examples may help illuminate the issues, but neither seems entirely conclusive because, in each case, the nongeneric set is quite rich.

Example 3.5. Suppose that each $i \in I$ observes S_i signals in \mathbb{R} , S_i finite, that is, suppose that Ω_i is a subset of \mathbb{R}^{S_i} . If the joint distribution of the signals has a density with respect to Lebesgue measure, a very rich class of models, then the information structure has no CIC.

However, with the same distribution, if any component of the S_i is common between 2 or more players, then the information structure has a CIC. Such a common component arises if there is any continuous random variable in the environment, e.g. the continuous time arrival of some event, that both players observe before making their choice of action.

Events unobservable to the players may also lead to a CIC.

Example 3.6. Suppose that each $i \in I$ observes $v + \epsilon_i$, where v is the unknown value of some object and ϵ_i is a measurement error. If the ϵ_i have a joint density with respect to Lebesgue measure, again a very rich class of models, the information structure has no CIC.

However, if there is a non-null set B , perhaps unobservable by any player, conditional on which $\epsilon_i = \epsilon_j$ for some pair of players, then there is a CIC. Such a set B might arise if the players hire experts to evaluate v and the experts sometimes economize on their expenses by giving the same report to more than one player.

4. The one-dimensional fullness of deep discontinuities

A subset S of a separable Banach space \mathfrak{X} is one-dimensionally tiny if there is a non-zero $v \in \mathfrak{X}$ such that for all $u \in \mathfrak{X}$, there is at most a single $r \in \mathbb{R}$ such that $u + rv \in S$. A subset T

is one-dimensionally full if its complement is one-dimensionally tiny. Theorem B shows that for any P having a CIC, there is a one-dimensionally full set of utility functions in $L^1(P; C(A))$ giving rise to badly discontinuous expected utility functions.

For strategy sets T_i and bounded utilities $v_i, i \in I$, the Fudenberg and Levine [8] “most utility difference it can make to anyone” pseudo-metric is defined by

$$d_i(s_i, t_i) = \max_{k \in I} \sup_{t \in T} |v_k(t \setminus s_i) - v_k(t \setminus t_i)|. \tag{3}$$

The d_i -distance between s_i and t_i is equal to 0 if and only if s_i and t_i are strategically equivalent so that d_i is a metric on equivalence classes of strategies. A game is *finitely approximable* if for every $\epsilon > 0$ and every $i \in I$, there is a finite $T_i^f \subset T_i$ such that $d_i(T_i^f, T_i) < \epsilon$.

Theorem B. *If P has a CIC, then $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i \in I}$ is not finitely approximable except for a one-dimensionally tiny set of u_i in $L^1(P; C(A))$.*

Finite approximability fails only for games where the strategic content of the infinite strategy sets is different than the strategic content of finite subsets. Harris, Stinchcombe, and Zame [10] show that finite approximability has equivalent topological and measure theoretic formulations.⁶ From [18, Theorem 13], these are equivalent to satisfying the oscillation condition given below. For finite partitions \mathcal{P}_i of the spaces T_i , the product partition, \mathcal{P} , is the class of sets of the form $\times_i E_i, E_i \in \mathcal{P}_i$.

Definition 4.1. For a product partition \mathcal{P} , \mathcal{P} -oscillation of a bounded function $v = (v_i)_{i \in I}$ from $\times_i T_i$ to \mathbb{R}^I is

$$osc(f, \mathcal{P}) = \max_{E \in \mathcal{P}} \sup_{s, t \in E} \{ \|v(s) - v(t)\| \}. \tag{4}$$

A game $(T_i, v_i)_{i \in I}$ satisfies the *oscillation condition* if for all $\delta > 0$, there is a finite product partition \mathcal{P} such that $osc(f, \mathcal{P}) < \delta$.

For general I person games, the patterns of failure of the oscillation condition can involve all subsets of I of size 2 or larger. To avoid this notational burden, the proof only discusses two person subsets, $\{1, 2\} \subset I$, but the general case is an immediate consequence.

Proof of Theorem B. Fix an arbitrary P with a CIC. Step 1 constructs a specific $v \in L^1(P; C(A))$. Step 2 shows that $\Gamma(P, v)$ fails the oscillation condition. Step 3 completes the proof by showing that for any $u \in L^1(P; C(A))$, there is at most one r such that $\Gamma(P, u + r \cdot v)$ satisfies the oscillation condition.

Because P has a CIC, there exists a $B \in \mathcal{F}$ with $P(B) > 0$ and a $\varphi : B \rightarrow (0, 1]$ that is $(\mathcal{F}_1 \cap \mathcal{F}_2)|_B$ -measurable, and $\varphi(P_B)$ is non-atomic. Extend φ to Ω by setting $\varphi(B^c) \equiv 0$.

Step 1: Pick arbitrary distinct points $a_i, b_i \in A_i$ and $\epsilon > 0$ such that the ϵ -balls around the four points $(a_1, a_2), (a_1, b_2), (b_1, a_2),$ and (b_1, b_2) are disjoint. Define the continuous function $f : A_1 \times A_2 \rightarrow \mathbb{R}^2$ by

⁶ Topologically, a game is finitely approximable iff it is possible to embed each strategy space as a dense subset of a compact metric space in such a fashion that the utility functions have *jointly* continuous extensions. Measure theoretically, a game is finitely approximable iff the utility function is integrable with respect to all products of finitely additive probabilities on the strategy sets.

$$f(x, y) = (2, 2) \cdot \max\left(1 - \frac{1}{\epsilon}d((x, y), (a_1, a_2)), 0\right) + (4, 4) \cdot \max\left(1 - \frac{1}{\epsilon}d((x, y), (b_1, b_2)), 0\right). \tag{5}$$

Define the utility function $v : \Omega \rightarrow C(A : \mathbb{R}^2)$ by $v(\omega) = f \cdot 1_B(\omega)$.

Step 2: Each $g_i := E(\varphi|\mathcal{F}_i)$ is a function defined for ω_i in some set $B_i \subset \Omega_i$. By [7, Theorem 4.2.5, p. 97], g_i can be extended to a measurable $(0, 1]$ -valued function on all of Ω_i . Let G be the continuous cdf of the distribution $g_i(P_B)$. For $n \in \mathbb{N}$ and $0 \leq k \leq 2^n$, pick $r_{k,n} \in G^{-1}(k/2^n)$. For each $n \in \mathbb{N}$, define the strategies

$$b_i^n(\omega_i) = \begin{cases} \delta_{a_i} & \text{if } g_i(\omega_i) \in (r_{k,n}, r_{k+1,n}], k \text{ even,} \\ \delta_{b_i} & \text{if } g_i(\omega_i) \in (r_{k,n}, r_{k+1,n}], k \text{ odd.} \end{cases} \tag{6}$$

For all n , $v^P(b_1^n, b_2^n) = (1 - P(B)) \cdot (0, 0) + P(B) \cdot (3, 3)$, and for all $m \neq n$, $v^P(b_1^n, b_2^m) = P(B) \cdot (3/2, 3/2)$. Define $\delta = |v^P(b_1^n, b_2^n) - v^P(b_1^n, b_2^m)| = 3P(B)/2$.

Let \mathcal{P}_i be an arbitrary finite partition of \mathbb{B}_i , $i \in I$. Some $E_1 \times E_2$, $E_i \in \mathcal{P}_i$ must contain two distinct strategy vectors, (b_1^n, b_2^n) , (b_1^m, b_2^m) , $n \neq m$. This means that both (b_1^n, b_2^n) and (b_1^m, b_2^m) belong to $E_1 \times E_2$. Therefore, the oscillation of $v^P(\cdot)$ on the product partition $\mathcal{P}_1 \times \mathcal{P}_2$ is δ .

Step 3: Pick arbitrary $u \in L^1(P; C(A))$. Either (a) for all $r \in \mathbb{R}$, $\Gamma(P, u + r \cdot v^P)$ fails the oscillation condition, or (b), there exists at least one $r^\circ \in \mathbb{R}$ such that $\Gamma(P, u + r^\circ \cdot v^P)$ satisfies it. If (a) holds, there is nothing to prove. Assume that (b) holds, and pick arbitrary $s \neq 0$. Showing that $\Gamma(u + (r^\circ + s) \cdot v^P)$ violates the oscillation condition will complete the proof.

Pick an arbitrary product partition $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2$ of $\mathbb{B}_1 \times \mathbb{B}_2$. Refining the partitions \mathcal{P}_1 and \mathcal{P}_2 if necessary, we can assume that the integrable function $u^P(\cdot) + r^\circ \cdot v^P(\cdot)$ oscillates at most $\frac{1}{4}\delta|s|$ over \mathcal{P} . This implies that oscillation of $u^P(\cdot) + (r^\circ + s) \cdot v^P(\cdot)$ across \mathcal{P} must be at least $\frac{1}{2}\delta|s|$. Since this number is strictly positive and independent of \mathcal{P} , $\Gamma(P, u + (r^\circ + s) \cdot v^P)$ fails the oscillation condition. \square

5. Balanced approximations

The failure of a game to be finitely approximable in the Fudenberg and Levine pseudo-metric of Eq. (3) does not imply that finite approximations are useless for equilibrium analysis. At issue is whether or not an individual player can, in the finite approximations, guarantee themselves, to within any $\epsilon > 0$, the same payoffs as they can in the infinite game. This is an implication of finite approximability, but is a strictly weaker condition. The difference is between own payoff approximation and approximation of payoff differences for all players simultaneously.

The correlated equilibria of the companion paper [19] represent the limits of (generalized) sequences of approximate equilibria of games in which the players are *restricted* to play strategies that are measurable with respect to finite sub-fields of their information. The problem is that this restriction is artificial, and may matter. In particular, the correlated equilibria that arise may represent the limits of strategies that are not ϵ -equilibria because the approximations are not balanced.

5.1. Balance

Fix a game $(T_i, v_i)_{i \in I}$ with pure strategy sets T_i and payoffs $v_i : T \rightarrow \mathbb{R}$, $T := \times_{i \in I} T_i$. Let $d_H(\cdot, \cdot)$ be the Hausdorff pseudo-metric for bounded subsets of \mathbb{R} .

Definition 5.1. For $\times_{i \in I} S_i \subset \times_{i \in I} T_i$ and $\epsilon \geq 0$, $(S_i, v_i)_{i \in I}$, is an ϵ -balanced approximation to $(T_i, v_i)_{i \in I}$ if for all $j \in I$ and all $s \in \times_{i \in I} S_i$,

$$d_H(v_j(s \setminus \Delta(S_j)), v_j(s \setminus \Delta(T_j))) \leq \epsilon. \tag{7}$$

$(T_i, v_i)_{i \in I}$ has *balanced finite approximations*, or more simply, is *finitely balanced*, if for all $\epsilon > 0$ and all finite $S'_i \subset T_i$, there are finite $S_i \subset T_i$, $S'_i \subset S_i$, such that $(S_i, v_i)_{i \in I}$ is an ϵ -balanced approximation to $(T_i, v_i)_{i \in I}$.

In studying the question of equilibrium existence for discontinuous games, Reny [14, §8] suggests that only the upper end of the utility range needs to be approximated rather than the whole range. If balanced approximations exist for a game, then they also exist for any version of the game with uniformly continuous transformations of the utility functions.

5.2. Balance for infinite games with type-dependent strategies

The proof of Theorem B can be adapted to show that in most games with a CIC, there are unbalanced approximations. This can be seen in the following, two person game, Γ .

Γ : Suppose that $I = \{1, 2\}$, $\Omega = \Omega_1 \times \Omega_2$, $\Omega_i = [0, 1]$ (with the usual Borel σ -field), P is the uniform distribution on the diagonal in Ω , each A_i is a two-point set $\{a_i, b_i\}$, and $u_i(\omega, a)$ is independent of ω and given in the matrix

	a_2	b_2
a_1	(6, 6)	(3, 0)
b_1	(0, 3)	(9, 9)

At each $\omega = (\omega_1, \omega_2)$, the three Nash equilibria, (a_1, a_2) , (b_1, b_2) , and $((\frac{1}{2}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{2}))$, give the utilities (6, 6), (9, 9), and $(4\frac{1}{2}, 4\frac{1}{2})$ respectively. Players' utilities are equal in all of these equilibria.

Example 5.1 (*Unbalanced approximations to Γ*). Let \mathcal{F}_1^n be the field generated by the sets $\{(\frac{3k}{3^n}, \frac{3k+2}{3^n}), (\frac{3k+2}{3^n}, \frac{3k+3}{3^n})\}$ while \mathcal{F}_2^n is generated by the partially overlapping sets $\{(\frac{3k}{3^n}, \frac{3k+1}{3^n}), (\frac{3k+1}{3^n}, \frac{3k+3}{3^n})\}$, $1 \leq k \leq 3^{n-1} - 1$. For each $n \in \mathbb{N}$, the following strategies are equilibria if the players are *constrained* to play \mathcal{F}_i^n -measurable strategies:

$$b_1^n(\omega_1) = \begin{cases} \delta_{a_1} & \text{if } \omega_1 \in (\frac{3k}{3^n}, \frac{3k+2}{3^n}], \\ \delta_{b_1} & \text{if } \omega_1 \in (\frac{3k+2}{3^n}, \frac{3k+3}{3^n}], \end{cases} \quad b_2^n(\omega_2) = \begin{cases} \delta_{a_2} & \text{if } \omega_2 \in (\frac{3k}{3^n}, \frac{3k+1}{3^n}], \\ \delta_{b_2} & \text{if } \omega_2 \in (\frac{3k+1}{3^n}, \frac{3k+3}{3^n}]. \end{cases} \tag{8}$$

The associated utilities are (6, 5), which does not correspond to anything in the convex hull of the Nash equilibria.

The approximations $\mathbb{B}_i(\mathcal{F}_i^n)$ are *not* ϵ -balanced for $\epsilon < 1$ — each player would strictly prefer to change their actions in the middle third of each interval $(\frac{3k}{3^n}, \frac{3k+3}{3^n}]$, thereby gaining a utility of either 3 or 6 with probability $\frac{1}{3}$.

Example 5.2 (*Balanced approximations to Γ*). Fix any finite set, S'_i , of pure strategies for $i = 1, 2$. There is a smallest finite field, $\mathcal{F}'_i \subset \mathcal{F}_i$, making each strategy in S'_i measurable. Let \mathcal{G} be the smallest field of subsets of $[0, 1]$ containing $\mathcal{F}'_i \cup \mathcal{F}'_j$. Let $S_i \supset S'_i$ be the set of \mathcal{G} -measurable

pure strategies. If i is restricted to strategies in S'_i , then mixtures over S_j allow j to achieve exactly the same set of utilities as if they used their pure behavioral strategies, \mathbb{B}_j^{pure} , that is, these approximations are 0-balanced. To see this, note that for all $s_i \in S_i$, $u_j(\Delta(S_j), s_i) = u_j(\mathbb{B}_j^{pure}, s_i)$.

In any equilibrium for a balanced approximation, the utilities of the two players are equal. This corresponds to the correlated equilibria that are in the convex hull of the Nash equilibria of the matrix game.

For the class of infinite games with type-dependent strategies studied in this paper, the existence of finitely balanced approximations is an open question. For the games in which finitely balanced approximations do exist, an immediate corollary of the existence result in [19] is the existence of a non-empty closed subset of correlated equilibria. In Examples 3.1 and the game Γ , this is the set of measurable functions from Ω to the convex hull of the Nash equilibria of the matrix game, and this yields a strict subset of the set of correlated equilibria.

5.3. Some general balance considerations

Further examples give more information about balanced approximations to infinite games.

Example 5.3. If each T_i is compact and each v_i is jointly continuous, then $(T_i, v_i)_{i \in I}$ is balanced. Further, if S_i^α is any (generalized if need be) sequence of finite approximations converging to T_i for each $i \in I$, then for all large α , $(S_i^\alpha, v_i)_{i \in I}$ is ϵ -balanced.

More generally, we have

Lemma 5.1. *If $(T_i, v_i)_{i \in I}$ is finitely approximable, then it is balanced.*

Proof. From [10, Theorem 1], a game is finitely approximable iff it is nearly compact and continuous, that is, iff each T_i can be embedded as a dense subset of a compact metric space, C_i , so that the utility functions have a jointly continuous extension. Take S_i^n to be any sequence of finite subsets of the dense image of the T_i in the C_i that becomes dense in C_i . \square

As seen above, some discontinuities in the payoffs allow for balance, but only with careful, joint choice of the S_i . The following is a more direct example of a game which is balanced but not finitely approximable.

Example 5.4. Let $T_1 = T_2 = [0, 1]$, let $D \subset T_1 \times T_2$ be the diagonal, and let $v_i(t_i, t_j) = t_i \cdot 1_D(t_i, t_j)$. The game $(T_i, v_i)_{i \in I}$ is not finitely approximable because the metric given in (3) reduces to $d_i(s_i, t_i) = \max\{s_i, t_i\}$. This implies, for example, that there are uncountably many points in T_i at distance greater than $\frac{1}{2}$ from each other, so that no finite set can ϵ -approximate for $\epsilon < \frac{1}{2}$.

If S^n is a sequence of finite sets converging to $[0, 1]$, then setting $S_1^n = S_2^n = S^n$ gives a sequence that is ϵ -balanced for large n so that $(T_i, v_i)_{i \in I}$ is balanced.

Some games are neither finitely balanced nor finitely approximable.

Example 5.5. Let $T_1 = T_2 = (0, 1]$ and set $v_i(t_i, t_j)$ equal to the sign of $(t_j - t_i)$ (the “pick the smallest positive number” game). Again, the game $(T_i, v_i)_{i \in I}$ is not finitely approximable, the

d_i -metric of (3) yields $d_i(s_i, t_i) = 1$ if $s_i \neq t_i$. No finite sets can give an ϵ -balanced game for $\epsilon < 1$ since each player, i , has, in T_i , a number smaller than any element of a finite subset of T_j .

6. Conclusions and complements

The present paper is a piece of a larger project — the development of a general theory of infinite extensive form games. For the larger project, this paper provides two main lessons. The first is that approximating information about moves of Nature by finite sub-fields is a workable and fruitful strategy. Whether or not this approach to information will be as useful for information about moves of other players remains to be seen.

The second lesson comes from the distinction between finite balance and finite approximability. Failure of finite approximability indicates that we have an infinite game in which the infinite strategy sets do not resemble any finite approximations in a strong and well-known sense. However, this does not rule out the existence of finite approximations useful for equilibrium analysis. Balanced approximations focus on the ability of an individual player to guarantee their own set of payoffs, and this is closer to what is needed for equilibrium analysis.

For the smaller project of analyzing the present class of games, there are two complementary points to be made: Section 6.1 shows that approximating Ω by finite subsets rather than by finite partitions may lead to unacceptable information leakage; Section 6.2 shows that the assumption of a product structure for the set of signals is without loss.

6.1. Finite subset approximations to signals

Approximating Ω by finite subsets requires approximating the joint distribution of signals, P , by finitely supported probabilities. We start by defining τ_{sf} , the strong finite topology on measures, which has the property that the finitely supported probabilities are dense, but no tighter (i.e. larger or finer) topology has this property. Example 6.1 shows that sequences of τ_{sf} -approximations may not be close enough to P to approximate information structures.

Definition 6.1. τ_{sf} , the strong finite topology on $\Delta(\mathcal{F})$, is the topology generated by classes of the form

$$G_{sf}(P; (E_n)_{n=1}^N) = \bigcap_{n=1}^N \{Q \in \Delta(\mathcal{F}): |Q(E_n) - P(E_n)| = 0\}, \tag{9}$$

where $P \in \Delta(\mathcal{F})$ and the E_n belong to \mathcal{F} .

A net of probabilities, P^α , on (Ω, \mathcal{F}) converges in τ_{sf} to P iff for every $E \in \mathcal{F}$, there exists α' such that for all $\alpha > \alpha'$, $P^\alpha(E) = P(E)$. This can also be said as P^α converges to P iff for each $E \in \mathcal{F}$, $P^\alpha(E)$ converges finitely to $P(E)$.

There are two useful comparisons, the strong topology and the weak* topology. A basis for τ_s , the strong topology, is the class of sets of the form

$$G_s(P; (E_n, \epsilon_n)_{n=1}^N) = \bigcap_{n=1}^N \{Q \in \Delta(\mathcal{F}): |Q(E_n) - P(E_n)| < \epsilon_n\}, \tag{10}$$

where $P \in \Delta(\mathcal{F})$, the E_n belong to \mathcal{F} , and the ϵ_n are strictly positive. Whether or not $\epsilon_n = 0$ is allowed distinguishes the strong and the strong finite topologies.

If Ω is a metric space and \mathcal{F} is its Borel σ -field, the weak* topology has, as a basis, the class of sets of the form

$$G_{w^*}(P; (E_n, \epsilon_n)_{n=1}^N) = \bigcap_{n=1}^N \{Q \in \Delta(\mathcal{F}): |Q(E_n^{\epsilon_n}) - P(E_n)| < \epsilon_n, \text{ and } |P(E_n^{\epsilon_n}) - Q(E_n)| < \epsilon_n\}, \quad (11)$$

where for any $E \in \mathcal{F}$ and $\epsilon > 0$, $E^\epsilon = \bigcup_{\omega \in E} B(\omega, \epsilon)$ is the ϵ -ball around E , and again, $P \in \Delta(\mathcal{F})$, the E_n belong to \mathcal{F} , and the ϵ_n are strictly positive.

A σ -field, \mathcal{F} , is said to be Hausdorff if $\{\omega\} \in \mathcal{F}$ for each $\omega \in \Omega$. As strong as it is, the finitely supported probabilities are still τ_{sf} -dense.

Lemma 6.1. *If \mathcal{F} is Hausdorff, then the finitely supported probabilities are τ_{sf} -dense in $\Delta(\mathcal{F})$.*

Proof. Let $G = G_{sf}(P; (E_n)_{n=1}^N)$ be a basis set for τ_{sf} . Let $\{F_m: m \leq M\}$ be the partition generated by $(E_n)_{n=1}^N$. For each m , pick $\omega_m \in F_m$. The finitely supported probability, Q , defined by $Q(\omega_m) = P(F_m)$ belongs to G . \square

One of the reasons to like the strong finite topology is that, for any finite collection of events E_1, \dots, E_N , all of the conditional probabilities $P(E_n|E_m)$, $n, m \leq N$, are matched by $P^\alpha(E_n|E_m)$ for all large α . Matching conditional probabilities might seem to lead to matching the informational structures of the players. The difficulty is that a P^α may agree on a given finite partition and still be far different for finer partitions. Exploiting this, the next example constructs a net P^α converging to P in the strong finite topology and a set D with the following property: $P(\{\omega_j: P(D|\omega_j) = 1\}) = 0$ yet $P^\alpha(\{\omega_j: P^\alpha(D|\omega_j) = 1\}) = \frac{1}{2}$. In words, in the continuous game, no-one assigns mass 1 to the set D , but in the approximations to P , half of the time, both players assign mass 1 to D .

Example 6.1. Let $\Omega_i = [0, 1]$, $i = 1, 2$, let λ_2 be the uniform distribution on the two-dimensional Ω , let λ_1 be the uniform distribution on the one-dimensional diagonal, $D = \{(\omega_1, \omega_2): \omega_1 = \omega_2\}$, and let $P = \frac{1}{2}\lambda_2 + \frac{1}{2}\lambda_1$ so that P has a CIC. For each ω_i , each player's posterior distribution puts mass $\frac{1}{2}$ on D . Let E_j be the event that $P(D|\omega_j) = \frac{1}{2}$ and note that $P(E_j|\omega_i) \equiv 1$. Strong finite approximations will be shown to allow half of the posterior distributions to put mass 1 on D .

Let $\mathcal{P} = \{F_m: m \in M\}$ be a finite partition of Ω . For each $F_m \in \mathcal{P}$, pick $\omega_m = (\omega_{1,m}, \omega_{2,m}) \in F_m$. Define $Q^\alpha(\{\omega_m\}) = P(F_m)$ so that Q^α agrees with P on \mathcal{P} . Let $A \subset [0, 1]$ be the set of numbers of the form $\omega_{i,m}$, $i = 1, 2$, and let $\Omega^f = A \times A$. Since $\{(\omega_{1,m}, \omega_{2,m}): m \in M\} \subset \Omega^f$, Q^α is supported on Ω^f .

Let \mathcal{P}_D be the trace of the partition \mathcal{P} on the diagonal set D . There are $M' \subset M$ elements of \mathcal{P}_D with $\lambda_1(E_{m'}) > 0$. Since A is finite, for each $m' \in M'$, there is a $b_{m'} \notin A$ such that $(b_{m'}, b_{m'}) \in E_{m'}$. Let $B = \{b_{m'} \in [0, 1]: m' \in M'\}$. If $m' \in M'$, transfer the mass on $\omega_{m'}$ to $(b_{m'}, b_{m'})$, i.e. define $P^\alpha(\{(b_{m'}, b_{m'})\}) = Q^\alpha(\{(\omega_{1,m'}, \omega_{2,m'})\})$. For $m \in (M \setminus M')$, set $P^\alpha = Q^\alpha$. For each F_m , $Q^\alpha(F_m) = P^\alpha(F_m)$. Since $P(D) = \frac{1}{2}$, the switch from Q^α to P^α moves mass at least $\frac{1}{2}$ from $A \times A$ to $B \times B$.

Finally, consider the finite signal structure $\Omega^f = (A \cup B) \times (A \cup B)$ with probability distribution P^α . When $\omega_i \in B$, an event having probability $\frac{1}{2}$, both players posterior assigns mass 1 to the diagonal, D .

6.2. The generality of the product structure

The assumption that Ω is a product space is without loss of generality in the class of games under study. An alternative formulation of the information structure starts with a probability space (Ω, \mathcal{F}, P) and a collection $(\mathcal{F}_i)_{i \in I}$ of sub- σ -fields of $\mathcal{F} \supset \sigma((\mathcal{F}_i)_{i \in I})$. With such a structure, the game would have i 's strategies being the \mathcal{F}_i -measurable functions from Ω to Δ_i . It is possible to pass back and forth from this non-product structure to the product measure space formulation so that all strategic and expected utility structures are preserved.

If $\mathcal{F} \supsetneq \sigma((\mathcal{F}_i)_{i \in I})$, add a dummy player $i = 0$ to I and set $\mathcal{F}_0 = \mathcal{F}$. For each $i \in I$, define $\omega \sim_i \omega'$ if for all $E_i \in \mathcal{F}_i$, $1_{E_i}(\omega) = 1_{E_i}(\omega')$, let $\dot{\Omega}_i$ be the quotient space Ω / \sim_i , κ_i the canonical mapping of Ω onto $\dot{\Omega}_i$, and let $\dot{\mathcal{F}}_i$ be the σ -field $\kappa_i(\mathcal{F}_i)$. The product formulation of the information structure is $(\prod_i (\dot{\Omega}_i, \dot{\mathcal{F}}_i), \otimes_i \dot{\mathcal{F}}_i, \dot{P})$ where \dot{P} is the image of P under the embedding $\omega \mapsto (\kappa_i(\omega))_{i \in I}$.

Let M be a separable metric space and G the class of measurable functions from M to $[0, 1]$. To every $\dot{\mathcal{F}}_i$ -measurable $\dot{f} : \dot{\Omega}_i \rightarrow M$, associate the \mathcal{F}_i -measurable function $f(\omega) = \dot{f}(\kappa_i(\omega))$, and to every \mathcal{F}_i measurable function $f : \Omega \rightarrow M$, associate the $\dot{\mathcal{F}}_i$ -measurable function $\dot{f}(\dot{\omega}_i) = f(\kappa_i^{-1}(\dot{\omega}_i))$. Let \dot{P}_i be the restriction of \dot{P} to $\dot{\mathcal{F}}_i$, i.e. $\dot{P}_i = \kappa_i(P)$. The routine proof of the following change-of-variable lemma is omitted.

Lemma 6.2. *The associations just described are inverse images of each other and for all measurable $g : M \rightarrow [0, 1]$, $\int_{\Omega} g(f) dP = \int_{\dot{\Omega}_i} g(\dot{f}) d\dot{P}_i$.*

An immediate consequence of Lemma 6.2 is that $u^P(b) = \dot{u}^{\dot{P}}(\dot{b})$ for any strategy vector b . Thus, $(b_i)_{i \in I}$ is a (correlated) equilibrium for the game with the information structure $((\Omega, \mathcal{F}, P), (\mathcal{F}_i))$ if and only if the associated $(\dot{b}_i)_{i \in I}$ is a (correlated) equilibrium for the game with the information structure $(\prod_i (\dot{\Omega}_i, \dot{\mathcal{F}}_i), \otimes_i \dot{\mathcal{F}}_i, \dot{P})$.

References

- [1] R.M. Anderson, W.R. Zame, Genericity with infinitely many parameters, *Adv. Theor. Econ.* 1 (2001), Article 1.
- [2] N. Aronszajn, Differentiability of Lipschitz mappings between Banach spaces, *Studia Math.* LVII (1976) 147–190.
- [3] E. Balder, Generalized equilibrium results for games with incomplete information, *Math. Oper. Res.* 13 (1988) 265–276.
- [4] P. Billingsley, *Probability and Measure*, 3rd ed., John Wiley and Sons, New York, 1995.
- [5] J.P.R. Christensen, *Topology and Borel Structure*, North-Holland, Amsterdam, 1974.
- [6] K. Cotter, Correlated equilibrium in games with type-dependent strategies, *J. Econ. Theory* 54 (1991) 48–69.
- [7] R. Dudley, *Real Analysis and Probability*, Wadsworth & Brooks/Cole, Pacific Grove, CA, 1989.
- [8] D. Fudenberg, D. Levine, Limit games and limit equilibria, *J. Econ. Theory* 38 (1983) 261–279.
- [9] C. Harris, P. Reny, A. Robson, The existence of subgame-perfect equilibrium in continuous games with almost perfect information: A case for public randomization, *Econometrica* 63 (1995) 507–544.
- [10] C.J. Harris, M.B. Stinchcombe, W.R. Zame, Nearly compact and continuous normal form games: Characterizations and equilibrium existence, *Games Econ. Behav.* 50 (2005) 208–224.
- [11] B.R. Hunt, T. Sauer, J.A. Yorke, Prevalence: A translation-invariant ‘almost every’ on infinite-dimensional spaces, *Bull. Amer. Math. Soc. (N.S.)* 27 (1992) 217–238.
- [12] A. Manelli, Cheap talk and sequential equilibria in signaling games, *Econometrica* 64 (1996) 917–942.
- [13] P. Milgrom, R. Weber, Distributional strategies for games with incomplete information, *Math. Oper. Res.* 10 (1985) 619–632.
- [14] P.J. Reny, On the existence of pure and mixed strategy Nash equilibria in discontinuous games, *Econometrica* 67 (1999) 1029–1056.
- [15] L. Simon, W. Zame, Discontinuous games and endogenous sharing rules, *Econometrica* 58 (1990) 861–872.

- [16] R.S. Simon, Games of incomplete information, ergodic theory, and the measurability of equilibria, *Israel J. Math.* 138 (2003) 73–92.
- [17] M.B. Stinchcombe, The gap between probability and prevalence: Loneliness in vector spaces, *Proc. Amer. Math. Soc.* 129 (2001) 451–457.
- [18] M.B. Stinchcombe, Nash equilibrium and generalized integration for infinite normal form games, *Games Econ. Behav.* 50 (2005) 332–365.
- [19] M.B. Stinchcombe, Correlated equilibrium existence for infinite games with type-dependent strategies, *J. Econ. Theory* 146 (2) (2011) 638–655.