Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 146 (2011) 638–655
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Correlated equilibrium existence for infinite games with type-dependent strategies ☆

Maxwell B. Stinchcombe *

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

Received 25 August 2005; final version received 6 October 2010; accepted 10 November 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. This paper gives a counter-example to the main step in Cotter's [K. Cotter, Correlated equilibrium in games with type-dependent strategies, J. Econ. Theory 54 (1991) 48–69] argument for correlated equilibrium existence for this class of games, and supplies an alternative proof.

JEL classification: C62; C72; C65

Keywords: Infinite games; Correlated equilibrium existence; Informational commonalities; Nonstandard analysis

1. Introduction

This paper studies equilibrium existence for 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 help of an associate editor and anonymous referees in making this work more comprehensible has been invaluable.
* Fax: +1 512 471 3510.
E-mail address: max.stinchcombe@gmail.com.

0022-0531/$ – see front matter
doi:10.1016/j.jet.2010.12.006

Despite the continuity in actions, the *expected* utility payoff functions in these games often fail to be jointly continuous, and Simon [32] showed that Nash equilibria do not generally exist. The discontinuities are serious enough that there are correlated equilibria that are only reachable as limits of strategies that fail to be ϵ-equilibria for any ϵ close to 0. The companion paper [35] examines the properties of these discontinuities in more detail.

1.1. Correlated equilibria

The correct generalization of Aumann's [3] correlated equilibrium to this class of games is given in Cotter [8]. That paper approximates this class of measurable-continuous games by restricting players to strategies that are measurable with respect to increasingly fine sequences of finite sub-partitions of their signal spaces. The increasingly fine aspect of the sequences means that the players can nearly play any of their strategies, but, in taking limits of such sequences of strategies, correlation between the players' actions can be lost.

Cotter's [8] argument for the existence of correlated equilibria has two steps. First, the compactness of the set of finitely additive probabilities (with respect to a topology making the expected utilities continuous) guarantees the existence of convergent sequences of strategies. Second, a limit argument is given to show that the finitely additive limits can be replaced by legitimate, countably additive correlation devices.

Example A shows that the limit argument is not correct, that one needs some other representation of the correlation lost in taking the limit. A novel class of 'diagonally concentrated' probability spaces provides the representation. Theorem A uses this class to deliver the existence of correlated equilibria.

1.2. On the use of nonstandard analysis in the proof

Diagonally concentrated probability spaces are products of two probability spaces. The first is a standard space, here taken to be a player's space of signals, and the second is a nonstandard, expanded version of the same space. The reason this works so well is that nonstandard spaces contain elements that represent the fine details of sequences while taking limits suppresses them.[1] In $\mathbb{R}$, the four sequences

$$x - 1/n, \qquad x - 1/n^2, \qquad x + 1/n^2, \quad \text{and} \quad x + 1/n \tag{1}$$

all converge to x. This single limit point represents the tail behavior of all four sequences, and this suppresses the details: the first and second sequence converge from below, the third and the fourth from above; the middle two sequences converge at a much faster rate than the outer two sequences.

The nonstandard, expanded version of $\mathbb{R}$ is denoted $^*\mathbb{R}$. As well as containing every $r \in \mathbb{R}$, $^*\mathbb{R}$ contains an *infinitesimal*, denote it by ϵ, that represents the sequence $1/n$. The new number, ϵ, is both strictly positive and smaller than any $r \in \mathbb{R}_{++}$. With ϵ, the four sequences in (1) are represented by

$$x - \epsilon, \qquad x - \epsilon^2, \qquad x + \epsilon^2, \quad \text{and} \quad x + \epsilon. \tag{2}$$

[1] Both Lindström [21] and Corbae, Stinchcombe, and Zeeman [7, Chapter 11] develop nonstandard analysis building directly on intuitions from sequences and limit constructions.

Here, the first two nonstandard numbers are strictly less than x, the second two strictly larger. Further, $\epsilon^2/\epsilon = \epsilon$ is infinitesimal, so the middle two numbers are infinitely closer to x than the outer two numbers.

Suppose that a player's signal belongs to $\mathbb{R}$ and has the distribution P. A diagonally concentrated probability μ corresponding to P is a probability on $\mathbb{R} \times {}^*\mathbb{R}$ with the following property: for any measurable E,

$$\mu\big(E \times {}^*\mathbb{R}\big) = \mu\big(\mathbb{R} \times {}^*E\big) = P(E), \tag{3}$$

where ${}^*E \subset {}^*\mathbb{R}$ is the nonstandard expansion of $E \subset \mathbb{R}$. The property that for all $n \in \mathbb{N}$, $\mu(\{(x, x') \in \mathbb{R} \times {}^*\mathbb{R}\colon |x - x'| < 1/n\}) = 1$ is why these probabilities are called diagonally concentrated.

This paper combines Cotter's approximations with nonstandard representations of the sequences of finite partitions of the spaces of signals. In $\mathbb{R}$, the limit of any sequence of σ-fields generated by finite partitions that become arbitrarily fine is the usual Borel σ-field. By contrast, in ${}^*\mathbb{R}$, the limit is a σ-field based on a partition of ${}^*\mathbb{R}$ into elements having infinitesimal diameter, and the partition depends on the chosen sequence. The nonstandard representations of the sequences of finite partitions keeps the correlation from disappearing, and the standard part keeps track of the signal itself.

The difficult part of the argument is to guarantee that the information in the nonstandard representations of the limits is *innocuous*, that it reveals none of the players' private information.[2] Intuitively, taking limits of approximations in which players know less about others' private information should not give rise to situations in which they know more. However, this intuition is decidedly partial [35, Example 5.3].

1.3. Outline

The next section contains examples of the games under study as well as the notation and assumptions used throughout. Section 3 contains the counterexample, an explanation of finitely additive correlating devices, and an explanation of why they do not lead to correlated equilibrium existence for this class of games.

Section 4 is the longest. It: (1) reformulates correlated equilibria in a fashion that meshes with game models that use exhaustive, star-finite partition approximations to the agents' information; (2) shows that these star-finite approximations are innocuous informational expansions; (3) shows that they have Nash equilibria; and (4) shows that the Nash equilibria of any innocuous informational expansion of a game is a correlated equilibrium of the original game.

The companion paper to this one [35] is a more thorough study of the information structures in these games. It provides a novel strategic interpretation of the conditions known to be sufficient for the existence of a Nash equilibrium, and shows that these conditions are strongly nongeneric. Further, it shows that the discontinuities of the expected utility functions that arise are of the kind that cannot be "fixed" by the addition of ideal points. Of particular interest is the notion of a balanced and an unbalanced approximation to an infinite game, a concept that should help illuminate the differences between finite and infinite extensive form games.

[2] In signaling games and in games of almost perfect information, limits of approximations introduce cheap talk [24,12]. In normal form games, approximations may introduce specialized utility transfers by a randomizing referee [31,15]. More generally, limits of approximations can destroy information structures, forcing unwilling revelation of private information, allowing observation of what should be unobservable, or allowing concealment of what should be observable [34].

2. Examples, notation and assumptions

Under study are finite player games in which the players have differential information, identified as their types. Types are distributed to a commonly known distribution. As a function of their types, the players simultaneously pick actions in compact sets. Payoffs are measurable in information/type, continuous in actions, integrable, and, except when otherwise specified, probabilities are countably additive. This section discusses some of the uses of this class of games, sets notation and assumptions, demonstrates the discontinuity of the expected utility functions, and discusses Cotter's [8] insights into the limits it is necessary to place on correlation devices for this class of games.

2.1. Examples

The class of games under study appear in our models of firm competition with private, stochastic, cost structures, demand functions, or other payoff relevant information. They also appear in strategic signaling games, in purification interpretations of mixed strategy equilibria, and in wars of attrition. Because of the assumption that payoffs are continuous in actions, they are not natural models of auctions.[3] The known Nash existence result that applies to this class of games for all integrable $(u_i(\cdot))_{i\in I}$ depends on an *informational diffuseness* condition due to Milgrom and Weber [27].[4] As shown in [35, Theorem 1], this assumption requires that there not be any continuously distributed, commonly observable information.

2.1.1. Cournot competition with private information

Firms $i \in I$ have increasing, convex cost functions $c_i(\cdot)$. As a function of the privately known, stochastic cost structures, $c_i(\cdot)$, and signals about demand conditions, ω_i, they pick their quantities $q_i \in [0, \overline{q}_i]$. Expected profits conditional on signal ω_i are $E(\pi_i(c,q)|\omega_i) = E([q_i p_i((q_j)_{j\in I}) - c_i(q_i)]|\omega_i)$ where $p_i(\cdot)$ is i's, possibly random, demand function. One might expect the cost functions and the signals about demand conditions to contain, in general, many continuously distributed common components, including input prices, technological knowledge, and other market information. If so, the question of Nash equilibrium existence is open.

2.1.2. Signaling games with finite signal spaces

Two players observe their private information, ω_1 and ω_2. Then the sender chooses a_1 in a *finite* space A_1. The receiver observes a_1 and picks a point in $K(a_1)$, a compact metric space. Thus, the receiver's action space is $A_2 = \times_{a_1\in A_1} K(a_1)$.

If the u_i are measurable in ω, continuous on each $K(a_1)$, and integrable, then the signaling game fits into the class considered here. When ω_2 is a degenerate random variable and A_1 is not finite, Manelli [24] shows that, to accomodate limits of approximations and to guarantee equilibrium existence, the appropriate strategy space for the sender includes an expansion to allow for cheap talk. When ω_1 and ω_2 have a continuously distributed, commonly observable information, the question of Nash equilibrium existence is open.

[3] Introducing discontinuities, while no longer conceptually difficult, requires a great deal of supplementary technique (Jackson et al. [15], Jackson and Swinkels [16], Stinchcombe [34]), technique that would obscure the lessons that these games offer about the treatment of infinite models of differential information in game theory.

[4] Balder [5] proves equilibrium existence after generalizing essentially every aspect of the games studied by Milgrom and Weber except the information diffuseness condition.

2.1.3. Purifications

Randomization is crucial to the existence of saddle points in 0-sum games, and, more generally, to Nash's (1950) equilibrium existence theorem for finite games. Despite its crucial role in game theory, randomization is, to many, not an attractive behavioral assumption. Bellman and Blackwell [6] and Dvoretzky et al. [11] were early studies of the extent to which randomization might not be needed in games. Harsanyi [14] shows that, generically at least, mixed strategy equilibria are observationally equivalent to pure strategy equilibria of infinitesimal perturbations of the game. A stronger version of this result is in Govindan, Reny, and Robson [12].

Fix a game $\mathrm{G}(v) = (A_i, v_i)_{i\in I}$ where each A_i is finite, $v = (v_i)_{i\in I}$, and the utilities $v_i \in \mathbb{R}^A$, $A = \bigtimes_i A_i$. For all v outside a closed set having Lebesgue measure 0, every equilibrium of $\mathrm{G}(v) = (A_i, v_i)_{i\in I}$ is regular.

Let $(\omega_i)_{i\in I}$ be an independent collection of random vectors in $\mathbb{R}^A$ assigning, for all $i \in I$ and for any fixed strategy σ_{-i} of the players, mass 0 to the event that $(\omega_i(a_i, \cdot) - \omega_i(a_i', \cdot))$ lies in the hyperplane orthogonal to σ_{-i} (e.g. if the distribution of the ω_i has a density with respect to Lebesgue measure). A *perturbation* of $\mathrm{G}(v)$ is an incomplete information game in which each $i \in I$ observes the vector ω_i, picks an $a_i \in A_i$, and payoffs are $v_i(a) + \omega_i(a)$. A perturbation is a *δ-perturbation* if the distribution of the ω_i is within δ of point mass on 0 in the weak* topology.

For regular equilibria $\bar{\sigma}$, [12] shows that for all $\epsilon > 0$, there is a $\delta > 0$ such that any δ-perturbation of $\mathrm{G}(v)$ has an essentially strict, hence pure strategy, equilibrium inducing a distribution on A that is within ϵ of $\bar{\sigma}$. The interpretation is that mixed strategies played in equilibrium are, observationally, impossible to distinguish from strict pure strategy equilibria in nearby games. These nearby games have independent idiosyncratic shocks to utilities, and the pure strategy equilibria of these nearby games "purify" the (regular) mixed equilibria of $\mathrm{G}(v)$.

The independence of the $\omega_i(a)$ rules out the existence of any continuously distributed common information, and is crucial to the existence of Harsanyi's purifications. Radner and Rosenthal [28] give a generic game $\mathrm{G}(v)$ and expand it using $\omega_i(a)$ that are uniformly distributed on a bounded triangle in $\mathbb{R}^A$. Posterior distributions are, with probability 1, atomless, which is part of what is needed for purification.[5] Because of the way that independence of the signals fails in Radner and Rosenthal's example, the game has a unique equilibrium in which the players randomize after seeing a probability 1 set of signals.

The finiteness of A is also crucial to exact purification. Khan et al. [18] present a game in which $A_i = [0, 1]$ and exact purification is not possible, even when types are smoothly and independently distributed. They further show that exact purification is possible when the A_i are countable.

2.1.4. Wars of attrition

Two players have types t_i smoothly distributed in $(0, 1)$. A bounded, increasing value function $v : (0, 1) \to [0, \bar{v}]$ gives the value of an object, in terms of the cost of time spent fighting, to a player of type t_i. A pure strategy for i is a mapping b_i from $(0, 1)$ to $A_i = [0, 2 \cdot \bar{v}]$, with $b_i(t_i)$ being the time at which the player stops fighting for the object. If player j plays a strategy giving an atomless cdf F_j on A_j, the payoff to i of fighting until a is $u_i(t_i, (a, b_j)) = \int_0^a (v(t_i) - s)\, dF_j(s) - a(1 - F_j(a))$.

[5] In related work, Aumann et al. [4] showed that each player having, with probability 1, a non-atomic posterior is sufficient for *approximate* purification.

In games with independent types, Milgrom and Weber [27] show that the first order conditions $\partial u_i/\partial a = 0$ contain a great deal of information about pure strategy equilibria. The independence of types rules out any continuously distributed common information, and is the leading special case of Milgrom and Weber's informational diffuseness requirement. They show that if the joint distribution of the types is diffuse in the sense of having a density with respect to the product of its marginals, then the game has jointly continuous expected utility functions and compact strategy sets, leading to Nash equilibrium existence.

2.1.5. *Special classes of games*

An alternative approach to equilibrium existence is to look for useful classes of games by putting conditions on the utility functions and the joint distribution of types. Mamer and Shilling [23] study 0-sum games and show that information diffuseness can be dispensed with. Athey [2] gives joint conditions on the signals and the utility functions that give rise to single crossing condition that leads to the existence of pure strategy Nash equilibria. McAdams [26] contains a large generalization of Athey's analysis. Reny [29] goes further, substituting a weaker best response monotonicity condition and showing that one does not need Milgrom and Weber's informational diffuseness in this setting.

2.2. *Notation and assumptions*

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 $\omega = (\omega_i)_{i\in I} \in \Omega = \bigtimes_i \Omega_i$ 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, $(\bigtimes_i (\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$.

Each $i \in I$ has a compact, metric action space A_i, and $A := \bigtimes_i A_i$. 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.

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) $b_i^n \to b_i$ iff $\int_\Omega \langle v_i(\omega), b_i^n(\omega)\rangle \, P(d\omega) \to \int_\Omega \langle v_i(\omega), b_i(\omega)\rangle \, P(d\omega)$ 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 [8] showed that $\mathbb{B}_i(\mathcal{F}_i)$ is compact, and metrizable if $\mathcal{F}_i$ is countably generated.

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

$$u_i^P(b) = \int_\Omega \Big\langle u_i(\omega), \bigtimes\nolimits_i b_i(\omega)\Big\rangle \, P(d\omega) \tag{4}$$

where $\langle f, \nu\rangle := \int_A f(a)\,\nu(da)$ for continuous $f : A \to \mathbb{R}$ and Borel probabilities ν, and $\bigtimes_i b_i$ is the product probability on A having b_i as the i'th marginal. For $b \in \mathbb{B}$ and $b_i' \in \mathbb{B}_i$, $(b\backslash b_i')$ 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 $b_i' \in \mathbb{B}_i$, $u^P(b) \geqslant u^P(b\backslash b_i')$.

2.3. Discontinuous expected utility functions

For fixed b_{-i}, the expected utility function $u_i^P(\cdot, b_{-i})$ is individually continuous, hence achieves its maximum on the compact set $\mathbb{B}_i(\mathcal{F}_i)$. If the $\mathcal{F}_i$ are not only countably generated, but generated by a countable partition, then the expected utilities are *jointly* continuous. Jointly continuous utilities and compact strategy spaces imply that $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ has a Nash equilibrium.

If 2 or more players can both infer the value of a continuously distributed random variable, then for generic u_i in $L^1(P; C(A_i))$, the expected utilities, fail joint continuity [35, Theorem 2]. Whether or not Nash equilibria exist in this case is not known. The failure of joint continuity can be seen in

Example 2.1 *(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

$$\eta\big(E \times \big\{(L_1, L_2)\big\}\big) = \frac{1}{2}P(E), \qquad \eta\big(E \times \big\{(R_1, R_2)\big\}\big) = \frac{1}{2}P(E), \tag{5}$$

$E \in \mathcal{F}$. 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 $(3,3)$.

In this example, the payoffs do not depend on ω. It is not the measurable dependence on ω that gives rise to the failure of joint continuity. Rather, the discontinuity arises because the players both observe the value of a continuously distributed signal. This allows them to play strategies that become arbitrarily tightly coordinated, but the coordination is lost in the limit.

2.4. Nonstandard signaling spaces

An outline of the use of nonstandard versions of signal spaces can be seen in Example 2.1. If $\epsilon \in {}^*\mathbb{R}$ is a non-zero infinitesimal, then $1/\epsilon \in {}^*\mathbb{R}$ is an infinitely large number. Among the infinitely large numbers in ${}^*\mathbb{R}$ are the infinitely large integers, ${}^*\mathbb{N} \subset {}^*\mathbb{R}$.

The nonstandard representation of the sequence of finite partitions, $\mathcal{P}_n = \{(k/2^n, (k+1)/2^n]\colon k = 0, \ldots, 2^n - 1\}$, is a *finite partition $\mathcal{P}_N = \{(k/2^N, (k+1)/2^N]\colon k = 0, \ldots, 2^N - 1\}$ where $N \in {}^*\mathbb{R} \setminus \mathbb{R}$ is an infinite integer. The sequence of strategies has a nonstandard representation, $b_i^N(\omega_i) = \delta_{L_i}$ if $\omega_i \in (k/2^N, (k+1)/2^N]$ with k even, $b_i^N(\omega_i) = \delta_{R_i}$ otherwise. Here, an "even" infinite integer is one belonging to the set ${}^*\{2 \cdot n\colon n \in \mathbb{N}\}$, and the ω_i are points in ${}^*\Omega_i$, the expanded, nonstandard version of Ω_i.

In general, the diagonally concentrated probabilities on $\Omega_i \times {}^*\Omega_i$ will be based on the distribution of i's signals. In the expanded games, each i will observe the vector signal, (ω_i, ω_i^f) before choosing their action. The first component, ω_i, is the original signal, the second component contains the extra correlating information. In this case, the correlating information is which element of the *finite partition contains the signal. The difficulty is to guarantee that nothing but correlating information is added.

2.5. Replacing only the lost coordination

Correlating devices can mimic the lost coordination, which is useful, but extreme care is required lest they do more than replace the lost coordination. If the correlating device fails to be independent of ω, it can force unwilling revelation of private information or allow information leakage in the form of players conditioning on unobservable events and/or variables. This explains the form of the following definition. Let $\mathbb{B}_i^\circ$ denote the set of all $\mathcal{F}$-measurable functions from Ω to Δ_i, not merely the $\mathcal{F}_i$-measurable ones.

Definition 2.2 *(Cotter).* A *strategy correlated equilibrium* (*SCE*) is a probability measure ν on $\bigtimes_{i\in I}\mathbb{B}_i^\circ$ such that $\nu(\bigtimes_{i\in I}\mathbb{B}_i(\mathcal{F}_i))=1$, and for each $i\in I$ and each measurable $\gamma_i:\mathbb{B}_i\to\mathbb{B}_i$,

$$\int\limits_{\times_{i\in I}\mathbb{B}_i^\circ} u_i(b)\,\nu(db)\geqslant\int\limits_{\times_{i\in I}\mathbb{B}_i^\circ} u_i\big(\gamma_i(b_i),b_{-i}\big)\,\nu(db).$$

If ν is a point mass on a strategy b, then b must be a Nash equilibrium.

To see how independence between ν and P is maintained, note that whatever b is recommended by ν, P is used to evaluate payoffs. The restriction that $\nu(\bigtimes_{i\in I}\mathbb{B}_i(\mathcal{F}_i))=1$ prevents other forms of information leakage. As noted in [8], many ν give rise to the limit distribution η in Example 2.1.

Example 2.2. For any $[0,\alpha]\subset[0,1]$, define the strategies that play "L up to α with R after" and "R up to α with L after" by

$$b_i^{L,\alpha}(\omega_i)=\delta_{L_i}1_{[0,\alpha]}(\omega_i)+\delta_{R_i}1_{(\alpha,1]}(\omega_i)$$

and

$$b_i^{R,\alpha}(\omega_i)=\delta_{R_i}1_{[0,\alpha]}(\omega_i)+\delta_{L_i}1_{(\alpha,1]}(\omega_i).$$

For any α, the distribution $\nu_\alpha:=\frac{1}{2}\delta_{(b_1^{L,\alpha},b_2^{L,\alpha})}+\frac{1}{2}\delta_{(b_1^{R,\alpha},b_2^{R,\alpha})}$ mixes $\frac{1}{2}:\frac{1}{2}$ over the two strategies "L up to α" and "R up to α." Each ν_α, $\alpha\in[0,1]$, induces the distribution η, though ν_0 and ν_1 seem the most intuitive.

3. Finitely additive correlation and a counter-example

If the σ-fields $\mathcal{F}_i$ are generated by a countable partition, the expected utility functions are jointly continuous and equilibrium existence is guaranteed. This suggests approximating general $\mathcal{F}_i$ by nets/sequences of larger and larger finite sub-σ-fields, taking limits, along a subsequence if necessary, of the equilibria, and representing any coordination lost in the passage to the limit by a correlating device.

A version of this strategy appears in Cotter [8]. His first step is to represent the coordination lost in the limit as a finitely additive correlating device. His second step is to come back to the countably additive 'nearby' part of the finitely additive probability.

Cotter's first step uses the main advantage of the set of finitely additive probabilities. This is its compactness in the weakest topology making integration against bounded *measurable* functions continuous, a topology that makes the expected utility functions continuous. The second advantage is the ease with which finitely additive probabilities represent the coordination in a sequence of strategies.

The disadvantage is that drawing a point according to a purely finitely additive distribution corresponds to drawing a point in a much larger compactification of the original space.[6] For this reason, one needs the second step, of turning the finitely additive correlating device back into a countably additive one having the same coordinating properties.

The counterexample given here shows that this second step does not work. The discussion of finitely additive mixtures directly below shows that the first step does work, and leads directly to the counterexample. This paper arrives at correlated equilibria using correlating devices directly linked with the signals via diagonally concentrated joint distributions. This is a different approach, neither a distribution over strategies nor a distributional strategy. The correlating devices used here indirectly, but immediately, yield distributions over strategies and distributional strategies. Finding an alternative to this indirectness seems extraordinarily difficult.

3.1. The role of finitely additive mixtures

If one allows finitely additive mixtures over $\bigtimes_{i\in I}\mathbb{B}_i$, there are intriguing alternative representations for the limits of sequences of strategies such as the b_i^n in Example 2.1. The easiest class of representations corresponds to purely finitely additive 'point masses.' Let μ be a purely finitely additive $\{0,1\}$-valued measure on the set of all subsets of $\mathbb{N}$ such that $\mu(\mathbb{N})=1$, and μ is equal to 0 on all finite sets.[7]

To fix the ideas, consider a sequence, $x_n \downarrow \frac{1}{2}$ in $[0,1]$. The sequence is a mapping from $\mathbb{N}$ to $[0,1]$, and the image distribution of μ under the sequence is a purely finitely additive probability on the subsets of $[0,1]$. Specifically, define $\mu x^{-1}(E)=\mu(\{n\colon x_n\in E\})$. This is a purely finitely additive probability that puts mass 1 on every open neighborhood of $\frac{1}{2}$ but puts mass 0 on $\frac{1}{2}$ itself. Let $A_m=(\frac{1}{2}-\frac{1}{m},\frac{1}{2}+\frac{1}{m})$, $m\in\mathbb{N}$, so that $\mu x^{-1}(A_m)=1$ while $\mu x^{-1}(\bigcap_m A_m)=\mu x^{-1}(\{\frac{1}{2}\})=0$.

For any continuous $f:[0,1]\to\mathbb{R}$, $\int f(r)\,\mu x^{-1}(dr)=f(\frac{1}{2})$. For continuous functions, there is no difference between μx^{-1} and the point mass $\delta_{\frac{1}{2}}$. In this sense, $\delta_{\frac{1}{2}}$ is the unique countably additive probability that is 'nearby' μx^{-1}.

In Example 2.1, a sequence of pure strategy, 'saw-tooth' b_i^n weak* converges to $b_i^\infty(\omega)\equiv\frac{1}{2}\delta_L+\frac{1}{2}\delta_R$. For $E\subset\mathbb{B}_1(\mathcal{F}_1)\times\mathbb{B}_2(\mathcal{F}_2)$, define $\nu_{fa}(E)=\mu(\{n\colon b_n\in E\})$. Let $B_\epsilon(b_i^\infty)$ be the ϵ-ball around b_i^∞ in $\mathbb{B}_i(\mathcal{F}_i)$. ν_{fa} puts mass 1 on each set $B_\epsilon(b_1^\infty)\times B_\epsilon(b_2^\infty)$ even though $\nu_{fa}(\{(b_1^\infty,b_2^\infty)\})=0$.

[6] In stochastic process theory [20] gives an early use of these larger spaces, and Kingman [19] uses them to explain *inter alia* how finitely additive probabilities on the set of polynomials represent pure jump processes. [33] uses the extra points in these larger spaces to resolve the money pump paradoxes of finitely additive decision theory, and discusses a number of the other uses of these points in economic theory. Applications of these points in game theory include [25,13,34].

[7] Zorn's Lemma implies the existence of a free ultrafilter, $\mathcal{U}$, on $\mathbb{N}$. Setting $\mu(A)=1$ if $A\in\mathcal{U}$, and $\mu(A)=0$ otherwise gives a measure with these properties. Conversely, any such measure determines a free ultrafilter. Since ultrafilters can be identified with points in the Stone–Čech compactification of $\mathbb{N}$, each $\{0,1\}$-valued μ corresponds to point mass on this much larger space.

To see that ν_{fa} replaces the coordination lost in the passage to the limit, fix $E \in \mathcal{F}$ and $\epsilon > 0$. Let $H(E, \epsilon)$ denote the set of strategy pairs $b = (b_1, b_2)$ for which the induced outcome, η_b, is within ϵ of satisfying (5), that is, for which

$$\left|\eta_b\bigl(E \times \bigl\{(L, L)\bigr\}\bigr) - \frac{1}{2}P(E)\right| < \epsilon,$$

and

$$\left|\eta_b\bigl(E \times \bigl\{(R, R)\bigr\}\bigr) - \frac{1}{2}P(E)\right| < \epsilon. \tag{6}$$

By construction, $\nu_{fa}(H(E, \epsilon)) = 1$ and $\int_{\mathbb{B}} u^P(b)\, \nu_{fa}(db) = (6, 6)$.

Finitely additive correlating devices are difficult to interpret. The failure of countable additivity means that there exists a sequence of events $E_n \downarrow \emptyset$ with the property that $\nu_{fa}(E_n) \downarrow \delta$ for some $\delta > 0$. This means that there is no point or set of points to support a probability mass of size δ. In the discussion above, take $E_n = \{b^n, b^{n+1}, \ldots\}$ and $\delta = 1$.

3.2. The counterexample

The basic limit result in Cotter's work is [8, Theorem 4.3]. The following is a special case.

Limit Claim. *For each $i \in I$, let $\mathcal{F}_i^n$ be an increasing sequence of finite sub-σ-fields of $\mathcal{F}_i$ such that $\mathcal{F}_i = \sigma(\{\mathcal{F}_i^n\colon n \in \mathbb{N}\})$. For each $n \in \mathbb{N}$, let ν^n be an SCE for the game $(\mathbb{B}_i(\mathcal{F}_i^n), u_i)_{i\in I}$. Then there exist a subsequence $\{\nu^{n_k}\}$ of $\{\nu^n\}$ and a countably additive probability, ν, on $\times_{i\in I} \mathbb{B}_i^\circ$, such that*

(a) *ν is an SCE for the game $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$, and*
(b) *for every $v \in L^1(P; C(A))$ and continuous $\gamma_i : \mathbb{B}_i \to \Delta_i$,*

$$\lim_{k\to\infty} \int\limits_{\times_{i\in I} \mathbb{B}_i^\circ} \left[\int\limits_{\Omega}\left[\int\limits_{A_{-i}}\left[\int\limits_{A_i} v(\omega)(a_i, a_{-i})\, \gamma_i(b_i)(da_i)\right] b_{-i}(da_{-i})\right] P(d\omega)\right] \nu^{n_k}(db)$$
$$= \int\limits_{\times_{i\in I} \mathbb{B}_i^\circ} \left[\int\limits_{\Omega}\left[\int\limits_{A_{-i}}\left[\int\limits_{A_i} v(\omega)(a_i, a_{-i})\gamma_i(b_i)(da_i)\right] b_{-i}(da_{-i})\right] P(d\omega)\right] \nu(db). \tag{7}$$

The subsequence arises from compactness of the set of finitely additive probabilities. The limit, ν, should be the countably additive probability that is 'nearby' the finitely additive limit. The counterexample works with the sequence of Nash equilibria in Example 2.1.

Example A. For $n \in \mathbb{N}$, let ν^n be point mass on the vector of strategies b^n from Example 2.1, that is, $b_i^n(\omega_i) = \delta_{L_i}$ (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. Each (b_1^n, b_2^n) is a Nash equilibrium, hence each ν^n is an SCE. For $S \subset A$, define $v_S \in L^1(P; C(A))$ by $v_S(\omega)(a) = 1_\Omega(\omega) \cdot 1_S(a)$. For each $\epsilon > 0$, $i \in I$ and $a_i \in A_i$, pick a continuous γ_{i,a_i}^ϵ satisfying $\gamma_{i,a_i}^\epsilon(b_i) = \delta_{a_i}$ if b_i is within ϵ of b_i^∞ and $\gamma_{i,a_i}^\epsilon(b_i) = x_i$ if b_i is more than 2ϵ from b_i^∞, for some $x_i \in \Delta_i$.

The γ_{i,a_i}^ϵ are possible deviations from an SCE. They specify playing a_i for sure if the correlating device tells player i to play a strategy b_i that is within ϵ of b_i^∞ and to play an arbitrary

x_i outside of the 2ϵ-neighborhood of b_i^∞. Calculations below will show that the limits of the integrals in (7) are the same along any subsequence, and independent of x_i.

Having the same limit along each subsequence arises because, for any $\epsilon > 0$, for large n, each ν^n puts mass 1 on the ϵ-neighborhood of (b_1^∞, b_2^∞). If (7) is to hold, then integration against ν must give the limit, and this is independent of x_i. Since ν is countably additive, this implies that $\nu(\{(b_1^\infty, b_2^\infty)\}) = 1$. A correlated equilibrium which recommends only one strategy must recommend a Nash equilibrium, but we have seen that (b_1^∞, b_2^∞) is not a Nash equilibrium. This shows that (a) in the Limit Claim cannot be true.

Calculations. Since the ν^n are point masses on the strategies b^n, and the $b^n(\omega)$'s are also point masses, calculating the integrals in (7) is not too difficult. For example, take $S = \{(L_1, L_2)\}$. For n large enough that b_1^n is ϵ or closer to b_1^∞, $\gamma_{1,L_1}^\epsilon(b_1^n) = \delta_{L_1}$, which yields

$$\int_{\times_{i\in I}\mathbb{B}_i^\circ}\left[\int_\Omega\left[\int_{A_2}\left[\int_{A_1} v_S(\omega)(a_1,a_2)\gamma_{1,L_1}^\epsilon(b_1)(da_1)\right]b_2(da_2)\right]P(d\omega)\right]\nu^n(db)$$

$$\begin{aligned}
&= \int_\Omega\left[\int_{A_2}\left[\int_{A_1} 1_{\{(L_1,L_2)\}}(a_1,a_2)\gamma_{1,L_1}^\epsilon\big(b_1^n\big)(da_1)\right]b_2^n(\omega)(da_2)\right]P(d\omega)\\
&= \int_\Omega\left[\int_{A_2}\left[\int_{A_1} 1_{\{(L_1,L_2)\}}(a_1,a_2)\,\delta_{L_1}(da_1)\right]b_2^n(\omega)(da_2)\right]P(d\omega)\\
&= \int_\Omega\left[\int_{A_2} 1_{\{L_2\}}(a_2)\,b_2^n(\omega)(da_2)\right]P(d\omega) = \frac{1}{2}.
\end{aligned} \tag{8}$$

The first equality follows from ν^n being a point mass on (b_1^n, b_2^n). The second equality follows from n being large enough that $\gamma_{1,L_1}^\epsilon(b_1^n) = \delta_{L_1}$. The third equality comes from $1_{\{(L_1,L_2)\}}(L_1, a_2) = 1_{\{L_2\}}(a_2)$. The fourth, and final, equality comes from $b_2^n(\omega)$ playing L_2 and R_2 with probability $\frac{1}{2}$ each. Similar calculations for other $S \subset A$ and γ_{i,a_i}^ϵ show that the limit of the integrals are the same along any subsequence and independent of x_i.

4. Correlated equilibrium existence

The section is devoted to proving Theorem A, which shows that all games $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ have correlated equilibria. The proof has three conceptual steps. The first is to introduce a related class of games with extra signals. If the extra signals do not convey too much information, they are called innocuous informational expansions. The second step involves defining a special class of nonstandard informational expansions. These can be shown to be innocuous, and to have Nash equilibria. The final step is to show that the Nash equilibria of any innocuous informational expansion is a correlated equilibrium of the original game.

4.1. Notation

The following conventions serve to lessen the notational burden for σ-fields and sub-σ-fields on product spaces: let $X = \times_k X_k$ and for each k, let $\mathcal{X}_k$ be a σ-field of subsets of X_k; $\mathcal{X} = \bigotimes_k \mathcal{X}_k$ denotes the product σ-field on X; $\mathcal{X} = \times_k \mathcal{X}_k$ denotes the product field (not σ-field) of finite unions of measurable rectangles X; in a piece of context-dependent notational ambiguity, the symbol $\mathcal{X}_k$ also denotes the sub-σ-field of $\mathcal{X}$ given by $\{\pi_k^{-1}(E_k)\colon E_k \in \mathcal{X}_k\}$ where $\pi_k : X \to X_k$ is the canonical projection map; $\mathcal{X}_j \otimes \mathcal{X}_k$ denotes the smallest sub-σ-field of $\mathcal{X}$ containing both $\mathcal{X}_j$ and $\mathcal{X}_k$. In particular, a set $E_k \in \mathcal{X}_k$ can be regarded either as a subset of X_k, or as $\pi_k^{-1}(E_k)$, a "cylindrical" subset of X.

$([0,1], \mathcal{B}, \lambda)$ is the unit interval with the usual Borel σ-field and Lebesgue measure. The following two results can be found in e.g. [10, Chapter 13]: for M an uncountable, complete separable metric space, E is an uncountable Borel subset of M if and only if it is measurably isomorphic to $[0,1]$; and any Borel probability, η, on M is of the form $\eta(E) = \lambda(f^{-1}(E))$ for some measurable $f : [0,1] \to M$.

As above, Δ_i is the set of Borel probabilities on the compact metric space A_i. As is well-known, Δ_i is compact in the weak* topology. $\mathcal{D}_i$ denotes the Borel σ-field on Δ_i generated by the weak* topology.

4.2. Correlated equilibria

The following replaces a SCE's recommendation of a vector $(b_i)_{i\in I}$ of complete contingent plans with each player i observing their own ω_i and an $r \in [0,1]$.

Definition 4.1. A *correlated strategy for i* is a measurable function φ_i on the probability space $(\Omega_i \times [0,1], \mathcal{F}_i \otimes \mathcal{B}, P \times \lambda)$ to Δ_i. A *correlated strategy* is a mapping $(\omega, r) \mapsto (\varphi_i(\omega_i, r))_{i\in I}$ where each φ_i is a correlated strategy for i.

The expected utility associated with a correlated strategy φ is

$$u^P(\varphi) := \int_{[0,1]} u^P\big(\varphi(\cdot, r)\big)\,\lambda(dr) = \int_{\Omega\times[0,1]} \Big\langle u(\omega, \cdot), \textstyle\bigtimes_i \varphi_i(\omega_i, r)\Big\rangle\, P \times \lambda(d\omega, dr). \tag{9}$$

If ψ_i is an $\mathcal{F}_i \otimes \mathcal{D}_i/\mathcal{D}_i$-measurable function from $\Omega_i \times \Delta_i$ to Δ_i, and φ is a correlated strategy, $\langle\varphi \| \psi_i\rangle$ denotes the correlated strategy with j'th component $(\omega, r) \mapsto \varphi_j(\omega_j, r)$, $j \neq i$, and i'th component $(\omega, r) \mapsto \psi_i(\omega_i, \varphi_i(\omega_i, r))$.

Definition 4.2. A correlated strategy is a *correlated equilibrium* if for all $i \in I$ and all measurable ψ_i from $\Omega_i \times \Delta_i$ to Δ_i, $u_i^P(\varphi) \geqslant u_i^P(\langle\varphi\|\psi_i\rangle)$.

It is straightforward to show that $f(r) := (\varphi_i(\cdot, r))_{i\in I}$ is measurable, and that $\nu(E) := \lambda(f^{-1}(E))$ is a SCE if and only if the correlated strategy $(\varphi_i(\cdot, r))_{i\in I}$ is a correlated equilibrium.

4.3. Innocuous informational expansions

Adding an independent observation of a point in $[0,1]$ to the domain of the players' strategies in a payoff-irrelevant way is an example of an informational expansion.

Definition 4.3. An *informational expansion of $\Gamma = (\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$* is a game Γ' with information structure $(\bigtimes_i(\Omega_i \times \Omega_i', \mathcal{F}_i \otimes \mathcal{F}_i'), \mathcal{F} \otimes \mathcal{F}', P')$, action sets A_i, and utility function u' where

(1) each $(\Omega'_i, \mathcal{F}'_i)$ is a measure space,
(2) P' is a probability on $\mathcal{F} \otimes \mathcal{F}'$ with $P'(E \times \Omega') = P(E)$ for all $E \in \mathcal{F}$, and
(3) $u'(\omega, \omega') = u(\omega)$,

where $\Omega' = \times_{i\in I} \Omega'_i$.

Informational expansions can substantively change a game.

Example 4.1. Set $\Omega'_i = \times_{j\neq i} \Omega_j$ with P' being the image of P under the mapping $\omega \mapsto (\omega_i, \omega_{-i})_{i\in I}$. This expansion makes public everyone's private information.

To prevent such changes we have

Definition 4.4. An informational expansion is *innocuous* if for all $E \in \mathcal{F}$ and for all players $i \in I$, $P'(E \times \Omega' | \mathcal{F}_i \otimes \mathcal{F}'_i) = P'(E \times \Omega' | \mathcal{F}_i)$ P'-a.e.

The product measure structure of $P \times \lambda$ renders innocuous the information in Definitions 4.1 and 4.2.

4.4. Outcome equivalence

The proof of the existence of correlated equilibria introduces a class of innocuous expansions for which Nash equilibrium existence is easily shown. The remaining question is whether the Nash equilibria of all innocuous informational expansions 'are' correlated equilibria. They are in the sense that their outcome distributions are equal.

Definition 4.5. The *outcome* associated a strategy s' for the expansion Γ' is the probability $Q_{s'}$ on $\Omega \times \Delta$ defined by $Q_{s'}(E) = P'\{(\omega, \omega')\colon (\omega, s'(\omega, \omega')) \in E\}$, $E \in \mathcal{F} \otimes \mathcal{D}$. Two strategy profiles, s' and s'', for two information expansions, Γ' and Γ'' are *outcome equivalent* if $Q_{s'} = Q_{s''}$.

Lemma 4.1. *If s' is a Nash equilibrium of an innocuous information expansion Γ', then s' is outcome equivalent to a correlated equilibrium of Γ.*

Proof. Let s' be an equilibrium of an innocuous information expansion Γ' of Γ. The proof has two steps: first, show that any outcome equivalent correlated strategy is a correlated equilibrium; and second, show that there is a correlated strategy that is outcome equivalent to s'.

First step: Suppose that φ is a correlated strategy outcome equivalent to s'. Pick $i \in I$ and let ψ_i be an $\mathcal{F}_i \otimes \mathcal{D}_i$-measurable function from $\Omega_i \times \Delta_i$ to Δ_i. In the game Γ', consider the strategy $s' \backslash t'_i$ with $t'_i(\omega, \omega') = \psi_i(\omega, s'_i(\omega, \omega'))$. Because s'_i is $\mathcal{F}_i \otimes \mathcal{H}'_i$-measurable, so is the strategy t'_i. Checking measurable rectangles shows that the strategy $s' \backslash t'_i$ is outcome equivalent to the strategy $\langle \varphi \| \psi_i \rangle$. Because s' is an equilibrium of Γ', $u_i^P(\varphi) \geqslant u_i^P(\langle \varphi \| \psi_i \rangle)$.

Second step: For $i \in I$, let $\mathcal{G}_i^\circ$ be the field generated by measurable rectangles, $E_i \times H_i$, $E_i \in \mathcal{F}_i$ and $H_i \in \mathcal{H}'_i$. Since $\mathcal{G}_i^\circ$ is a field generating $\mathcal{G}_i$, each s'_i is the P'-a.e. limit of a sequence, $s_i'^n$, of simple $\mathcal{G}_i^\circ$-measurable functions. There is no loss in assuming that the $(s'_i)^n$ are measurable with respect to fields $\mathcal{G}_i^n$ generated by an increasingly fine sequence of product partitions, $\mathcal{P}_i^n \times \mathcal{Q}_i^n$, of $\Omega \times \Omega'$ into measurable rectangles. Let $\mathcal{P}^n \times \mathcal{Q}'^n$ be the coarsest common refinement of the partitions $(\mathcal{P}_i^n \times \mathcal{Q}_i^n)_{i \in I}$.

The correlated strategy φ the limit of an inductively constructed sequence φ^n that replicates the sequence s'^n. To begin the inductive construction, define $s'^0 = \mu^0$ as a constant function on the rectangle $\Omega \times \Omega'$ with $\mathcal{P}^0 \times \mathcal{Q}'^0 := \{\Omega \times \Omega'\}$. The rectangle $\Omega \times \Omega'$ is associated with the rectangle $\Omega \times [0, 1)$ with $\mathcal{P}^0 \times \mathcal{B}'^0 := \{\Omega \times \Omega'\}$. The function φ^0 on $\Omega \times [0, 1)$ is defined to match s'^0 on associated rectangles. For $n = 0$, this requires $\varphi^0 \equiv \mu^0$.

For $n \geqslant 1$, note that each rectangle, $E \times H$ in $\mathcal{P}^{n-1} \times \mathcal{Q}'^{n-1}$ is a associated with a rectangle $E \times B$, $B = [a, b)$, in $\mathcal{P}^{n-1} \times \mathcal{B}^{n-1}$. Further $E \times H$ is a disjoint union of rectangles in $\mathcal{P}^n \times \mathcal{Q}'^n$, enumerable as $E \times H = \bigcup_{k=1}^K \bigcup_{\ell=1}^L E_k \times H_\ell$. For each $k \in \{1, \ldots, K\}$, partition $B = [a, b)$ into ℓ disjoint half open intervals, $(B_\ell)_{\ell=1}^L$, such that $P'(E_k \times H_\ell) = P \times \lambda(E_k \times B_\ell)$. This can be done because λ is non-atomic, the marginal of P' on Ω is equal to P, and $P'(E \times H) = P \times \lambda(E \times B)$. Associate each $E_k \times H_\ell$ with $E_k \times B_\ell$ and define φ^n on $E_k \times B_\ell$ to equal s'^n on $E_k \times H_\ell$.

On a set with P'-probability 1, each $s_i'^n$ is convergent. This implies that on a set with $P \times \lambda$-probability 1, each φ_i^n is convergent. The measurability of the φ_i and the outcome equivalence of s' and φ are immediate. □

4.5. *Finitistic versions of information structures*

The star-finite sets that replace the information structure $(\bigtimes_i(\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$ are constructions from Robinson's [30] nonstandard analysis. Star-finite (or *-finite) measure spaces were first examined in [22], their use entails little restriction in generality [1].

By assumption, all the nonstandard objects used here belong to an $\aleph$-saturated extension, $^*V(X)$, of a superstructure $V(X)$ where the base set, X, contains $\mathbb{R}$, each Ω_i and A_i, $i \in I$, and $\aleph$ is a cardinal greater than the cardinality of $V(X)$. If $Y \in V(X)$, the elements of $^*Y$ in $^*V(X)$ are called internal. The class of finite subsets of any $Z \in V(X)$ is denoted $\mathcal{P}^f(Z)$, and is itself an element of $V(X)$.

Definition 4.6. The **-finite or star-finite* subsets of Y are $^*\mathcal{P}^f(Y)$.

Definition 4.7. A partial order $\succcurlyeq$ on $Y \in V(X)$ is *finitely satisfiable* if for all finite $\{y_1, \ldots, y_N\} \subset Y$, there exists a $y \in Y$ such that $y \succcurlyeq y_n$, $n = 1, \ldots, N$.

A consequence of saturation is that if $\succcurlyeq$ is finitely satisfiable, then there exists an $y' \in {}^*Y$ such that for all $y \in Y$, $y' \, {}^*\!\!\succcurlyeq y$. One of the most fruitful applications of this is the existence of exhaustive *-finite sets.

Definition 4.8. For $Z \in V(X)$, an element Z^f of $^*\mathcal{P}^f(Z)$ is *exhaustive for* Z or is *a finitistic version of* Z if for any $z \in Z$, $^*z \in Z^f$.

Take Y to be $\mathcal{P}^f(Z)$ and $y_1 \succcurlyeq y_2$ to be $y_1 \supset y_2$. This gives a finitely satisfiable partial order on Y. Hence, there exists a $y' \in {}^*Y$ such that for all $y' \in Y$, $y' \supset {}^*y$. Let $Z^f = y'$, and we see that Z^f is exhaustive for Z.

Fix an information structure $(\bigtimes_i(\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$. The construction of a *-finite version of $(\bigtimes_i(\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$ which provides an innocuous information expansion is a four-step process. Notationally, "^{f}" replaces the "$'$" of information expansions in Definition 4.3.

Definition 4.9. A *star finite expansion of* $(\bigtimes_i(\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$ is an information structure of the form $(\Omega \times \Omega^f, (\mathcal{F}_i \otimes \mathcal{F}_i^f)_{i \in I}, \mathcal{F} \otimes \mathcal{F}^f, P^f)$ where

(1) $\Omega_i^f = {}^*\Omega_i$,
(2) $\mathcal{F}_i^f = \sigma(\mathcal{N}_i)$ where $\mathcal{N}_i \in {}^*\mathbb{F}_i$, for each $i \in I$, $\mathbb{F}_i$ is the set of finite sub-fields of $\mathcal{F}_i$, and $\mathcal{N}_i$ is exhaustive for $\mathcal{F}_i$,
(3) $\mathcal{F} = \sigma(\mathcal{N})$ where $\mathcal{N} \in {}^*\mathbb{F}$, $\mathbb{F}$ is the set of finite sub-fields of $\mathcal{F}$, $\mathcal{N}$ contains $\bigtimes_{i \in I} \mathcal{N}_i$ and is exhaustive for $\mathcal{F}$, and
(4) for $E \in \mathcal{F}$ and $D \in \mathcal{F}^f$, $P^f(E \times D) = L({}^*P)({}^*E \cap D)$ on measurable rectangles where $L({}^*P)$ is the Loeb measure on $\sigma(\mathcal{F}^f)$ generated by ${}^*P$.

Partially ordering $\mathbb{F}$ and the $\mathbb{F}_i$ by inclusion gives a finitely satisfiable relation and shows that $\mathcal{F}^f$ and the $\mathcal{F}_i^f$ exist. Condition (4) is crucial because it gives rise to a property called diagonal concentration.

4.6. Diagonal concentration

Let X be a non-empty set, $\mathcal{X}$ a σ-field of subsets of X, $X^f = {}^*X$, and let $\mathcal{X}^f$ be an *-finite sub-field of ${}^*\mathcal{X}$ that is exhaustive for $\mathcal{X}$.

Definition 4.10. A probability μ^f on $\mathcal{X} \otimes \sigma(\mathcal{X}^f)$ is *diagonally concentrated* if for all finite, measurable partitions $E_1, \ldots, E_N$ of X, $\mu^f(\bigcup_n (E_n \times {}^*E_n)) = 1$.

The next result shows that condition (4) of Definition 4.9 specifies a unique probability. This implies that P^f is diagonally concentrated because $P^f(E_n \times {}^*E_n) := L({}^*P)({}^*E_n \cap {}^*E_n)$.

Lemma 4.2. *If ν is a probability on $\mathcal{X}$ and μ^f is defined by $\mu^f(E \times D) = L({}^*\nu)({}^*E \cap D)$ on measurable rectangles in $X \times X^f$, then μ^f has a unique extension to the product σ-field $\mathcal{X} \otimes \sigma(\mathcal{X}^f)$.*

Proof. By Caratheodory's extension theorem, it is sufficient to show that μ^f is countably additive. Let $\mathcal{H} = \{H_1, H_2, \ldots\}$ be a countable subset of $\mathcal{X}$. Define the Blackwell pseudo-metric on X by

$$d_{\mathcal{H}}(x, y) = \frac{1}{\min\{n\colon 1_{H_n}(x) \neq 1_{H_n}(y)\}}.$$

Let $X_{\mathcal{H}}$ be the quotient space under the equivalence relation $x \sim_{\mathcal{H}} y$ iff $d_H(x, y) = 0$, so that $(X_{\mathcal{H}}, d_{\mathcal{H}})$ is a separable metric space. Let $\mathrm{st}_{\mathcal{H}} : X^f \to X_{\mathcal{H}}$ be standard part mapping in the $d_{\mathcal{H}}$ metric topology on $X_{\mathcal{H}}$. The Borel σ-field on $X_{\mathcal{H}}$ is generated by $\mathcal{H}$, and each H_n is both closed and open. Since $\mathrm{st}_{\mathcal{H}}^{-1}(H_n) = {}^*H_n \in \mathcal{X}^f$, $\mathrm{st}_{\mathcal{H}}^{-1}$ is measurable, implying that the mapping $x^f \mapsto (\mathrm{st}_H(x_f), x_f)$ is measurable. Let $\mu_{\mathcal{H}}^f$ be the image of $L(\nu^f)$ under this mapping. μ_H^f is countably additive on the σ-field $\sigma(\mathcal{H}) \otimes \sigma(\mathcal{X}^f)$.

If $\mathcal{X}$ is separable, taking $\mathcal{H}$ to be a generating set completes the proof. If $\mathcal{X}$ is not separable, partially order the countable subsets of $\mathcal{X}$ by inclusion. If $\mathcal{H}' \succ \mathcal{H}$, then $\mu_{\mathcal{H}'}^f$ restricts to $\mu_{\mathcal{H}}^f$, and both μ_H^f and $\mu_{\mathcal{H}'}^f$ agree with μ^f on the field $\sigma(\mathcal{H}) \times \sigma(\mathcal{X}^f)$. Thus, for all $\mathcal{H}$, $\mu_{\mathcal{H}}^f$ is the unique countably additive extension of μ^f from the field $\sigma(\mathcal{H}) \times \sigma(\mathcal{X}^f)$ to the σ-field $\sigma(\mathcal{H}) \otimes \sigma(\mathcal{X}^f)$. Take $A^n \downarrow \emptyset$ in $\mathcal{X} \times \sigma(\mathcal{X}^f)$. For any $\mathcal{H} \supset \{A_1, A_2, \ldots\}$, $\mu^f_{\mathcal{H}}$ is countably additive, implying that $\mu^f(A_n) \downarrow 0$. □

The main implication of diagonal concentration is that an $\mathcal{X} \otimes \mathcal{X}^f$-measurable function depends only on its second component μ^f-a.e.

Lemma 4.3. *If S is a complete metric space, $f : X \times X^f \to S$ is $\mathcal{X} \otimes \mathcal{X}^f$-measurable, and μ^f is diagonally concentrated, then $f(x, x^f) = g(x^f)$, μ^f-a.e. for some $\sigma(\mathcal{X}^f)$-measurable $g : X^f \to S$.*

Proof. The set of functions f for which the statement is true includes the simple $\mathcal{X} \times \sigma(\mathcal{X}^f)$-measurable functions and is closed under μ^f-a.e. convergence. □

4.7. Finitistic expansions of games

Fix a game $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ with information structure $(\bigtimes_i(\Omega_i, \mathcal{F}_i), \mathcal{F}, P)$.

Definition 4.11. A *star-finite expansion of* $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ is denoted Γ^f and defined as a game with a star-finite expansion, $(\Omega \times \Omega^f, (\mathcal{F}_i \otimes \sigma(\mathcal{N}_i))_{i\in I}, \mathcal{F} \otimes \sigma(\mathcal{N}), P^f)$, of the information structure, action sets A_i for $i \in I$, and utilities defined by $u^f(\omega, \omega^f) = u(\omega)$.

Lemma 4.4. *Any star-finite expansion, Γ^f, of $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ is innocuous.*

Intuitively, this holds because the $\mathcal{N}_i$ represent the "limits" of sub-fields of the information contained in the σ-field $\mathcal{F}_i$.

Proof. Pick $E \in \mathcal{F}$. Definition 4.4 requires

$$\gamma := P^f\big(E \times \Omega^f \big| \mathcal{F}_i \otimes \mathcal{F}_i^f\big) = \eta := P^f\big(E \times \Omega^f \big| \mathcal{F}_i\big) \quad P^f\text{-a.e.}$$

The set of $s \in [0, 1]$ such that either $P(\gamma = s) > 0$ or $P(\eta = s) > 0$ is at most a countable set so that its complement is dense. Pick arbitrary r in the complement. It is sufficient to show that $P(\{\gamma > r\}\Delta\{\eta > r\}) = 0$.

Since γ is $\mathcal{F}_i \otimes \mathcal{F}_i^f$-measurable, Lemma 4.3 implies that it is equal, P^f-almost everywhere, to a function that depends only on ω_i^f. Since η is $\mathcal{F}_i$-measurable, Doob's Theorem (e.g. [9, Chapter I.18, pp. 12–13]) implies that it is equal, P^f-almost everywhere, to a function that depends only on ω_i. It is diagonal concentration that allows two such functions to be equal P^f-almost everywhere.

Since $A_i := \{\eta > r\} \in \mathcal{F}_i$ and $\mathcal{N}_i$ is exhaustive for $\mathcal{F}_i$, ${}^*A_i \in \mathcal{N}_i$. From Definition 4.9(4), $P^f({}^*A_i \Delta A_i) = 0$. Therefore, it is sufficient to show that $\{\gamma > r\} = {}^*A_i$ P^f-almost everywhere.

Consider any $B_i \in \mathcal{F}_i$ with $B_i \subset A_i$ and $P(B_i) > 0$. By the definition of conditional probabilities, we know that $\int_{B_i} P(E|\mathcal{F}_i)\,dP = P(B_i \cap E)$. Since $B_i \subset A_i$, $\int_{B_i} P(E|\mathcal{F}_i)\,dP > \int_{B_i} r\,dP = r \cdot P(B_i)$. Combining, $P(B_i \cap E) > r \cdot P(B_i)$. By transfer, for any $B_i \in {}^*\mathcal{F}_i$ with $B_i \subset {}^*A_i$

and ${}^*P(B_i) > 0$,

$$ {}^*P\big(B_i \cap {}^*E\big) > r \cdot {}^*P(B_i). \tag{10} $$

Since (10) is also true for any subset of B_i, any $B_i \subset {}^*A_i$ is a subset of $\{\gamma \geqslant r\}$. Since $P^f(\gamma = r) = 0$, this establishes that ${}^*A_i \subset \{\gamma > r\}$ P^f-almost everywhere. The same arguments applied to $A_i^c = \{\eta < r\}$ establishes that $\{\gamma > r\} \subset {}^*A_i$ P^f-almost everywhere. □

4.8. Equilibrium existence

The prerequisites are now in place.

Theorem A. *All games $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ have correlated equilibria.*

Proof. Let Γ^f be any star-finite expansion of $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$. By Lemmas 4.1 and 4.4, it is sufficient to show that Γ^f has a Nash equilibrium.

Consider the internal version of $(\mathbb{B}_i(\mathcal{F}_i), u_i^P)_{i\in I}$ associated with Γ^f, that is, the game with information structure $(\bigtimes_i({}^*\Omega_i, \mathcal{N}_i), \mathcal{N}, {}^*P)$, actions sets ${}^*A_i$, and utility functions ${}^*u_i \in {}^*L^1(P; C(A))$. Transfer of Fan's or Glicksberg's existence result implies that this internal game has a Nash equilibrium, $\beta = (\beta_i)_{i\in I}$ (see also Khan and Sun [17]). Defining $b_i(\omega_i, \omega_i^f) = \operatorname{st}\beta_i(\omega_i^f)$ pushes down the strategies, and yields $b = (b_i)_{i\in I}$, a vector of $\mathcal{F}_i \otimes \sigma(N_i)$-measurable functions from $\Omega_i \times \Omega_i^f$ to Δ_i. It is sufficient to show that b is a Nash equilibrium for Γ^f.

Consider any $i \in I$ and strategy $(\omega_i, \omega_i^f) \mapsto c_i(\omega_i, \omega_i^f)$ for Γ^f. By Lemma 4.3, there is no loss in assuming that c_i is a function of ω_i^f only. Let γ_i be any lifting of c_i. Since β is a Nash equilibrium, is sufficient to show that

$$ \int\limits_{\Omega\times\Omega^f} \big\langle u(\omega), (b\backslash c_i)\big(\omega^f\big)\big\rangle P^f\big(d\big(\omega,\omega^f\big)\big) \simeq \int\limits_{\Omega^f} \big\langle {}^*u\big(\omega^f\big), (\beta\backslash\gamma_i)\big(\omega^f\big)\big\rangle {}^*P\big(d\omega^f\big), $$

and this follows from Lemmas 4.2 and 4.3. □

References

[1] R.M. Anderson, Star-finite representations of measure spaces, Trans. Amer. Math. Soc. 271 (1982) 667–687.
[2] S. Athey, Single crossing properties and the existence of pure strategy equilibria in games of incomplete information, Econometrica 69 (2001) 861–889.
[3] R. Aumann, Subjectivity and correlation in randomized strategies, J. Math. Econ. 1 (1974) 67–96.
[4] R. Aumann, Y. Katznelson, R. Radner, R. Rosenthal, B. Weiss, Approximate purification of mixed strategies, Math. Oper. Res. 8 (1983) 327–341.
[5] E. Balder, Generalized equilibrium results for games with incomplete information, Math. Oper. Res. 13 (1988) 265–276.
[6] R. Bellman, D. Blackwell, Some two-person games involving bluffing, Proc. Natl. Acad. Sci. USA 35 (1949) 600–605.
[7] D. Corbae, M.B. Stinchcombe, J. Zeeman, An Introduction to Mathematical Analysis for Economic Theory and Econometrics, Princeton Univ. Press, Princeton, NJ, 2009.
[8] K. Cotter, Correlated equilibrium in games with type-dependent strategies, J. Econ. Theory 54 (1991) 48–69.
[9] C. Dellacherie, P.A. Meyer, Probabilities and Potential, North-Holland, Amsterdam, 1978.
[10] R. Dudley, Real Analysis and Probability, Wadsworth & Brooks/Cole, Pacific Grove, CA, 1989.
[11] A. Dvoretzky, A. Wald, J. Wolfowitz, Elimination of randomization in certain statistical decision procedures and zero-sum two-person games, Ann. Math. Statist. 22 (1951) 1–21.

[12] S. Govindan, P.J. Reny, A.J. Robson, A short proof of Harsanyi's purification theorem, Games Econ. Behav. 45 (2003) 369–374.
[13] 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.
[14] J. Harsanyi, Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points, Int. J. Game Theory 2 (1973) 1–23.
[15] M. Jackson, L.K. Simon, J. Swinkels, W.R. Zame, Communication and equilibrium in discontinuous games of incomplete information, Econometrica 70 (2002) 1711–1740.
[16] M. Jackson, J. Swinkels, Existence of equilibrium in single and double private value auctions, Econometrica 73 (2005) 93–140.
[17] M. Ali Khan, Y. Sun, Non-cooperative games on hyperfinite Loeb spaces, J. Math. Econ. 31 (1999) 455–492.
[18] M. Ali Khan, K. Rath, Y. Sun, On a private information games without pure strategy equilibria, J. Math. Econ. 31 (1999) 341–359.
[19] J.F.C. Kingman, Additive set functions and the theory of probability, Proc. Cambridge Philos. Soc. 63 (1967) 767–775.
[20] L. Le Cam, Convergence in distribution of stochastic processes, Univ. Calif. Publ. Statist. 2 (1957) 207–236.
[21] T. Lindstrøm, An invitation to nonstandard analysis, in: Nigel Cutland (Ed.), Nonstandard Analysis and Its Applications, Cambridge University Press, Cambridge, 1988, pp. 1–105.
[22] P.A. Loeb, Conversion from nonstandard to standard measure spaces and applications in probability theory, Trans. Amer. Math. Soc. 211 (1975) 113–122.
[23] J. Mamer, K. Schilling, A zero-sum game with incomplete information and compact action spaces, Math. Oper. Res. 11 (1986) 627–631.
[24] A. Manelli, Cheap talk and sequential equilibria in signaling games, Econometrica 64 (1996) 917–942.
[25] M. Marinacci, Finitely additive and epsilon Nash equilibria, Int. J. Game Theory 26 (1997) 315–333.
[26] D. McAdams, Isotone equilibrium in games of incomplete information, Econometrica 71 (2003) 1191–1214.
[27] P. Milgrom, R. Weber, Distributional strategies for games with incomplete information, Math. Oper. Res. 10 (1985) 619–632.
[28] R. Radner, R. Rosenthal, Private information and pure-strategy equilibria, Math. Oper. Res. 7 (1982) 401–409.
[29] P.J. Reny, On the existence of monotone pure strategy equilibria in Bayesian games, Econometrica, forthcoming.
[30] A. Robinson, Nonstandard Analysis, North-Holland, Amsterdam, 1966, second revised edition 1974.
[31] L. Simon, W. Zame, Discontinuous games and endogenous sharing rules, Econometrica 58 (1990) 861–872.
[32] R.S. Simon, Games of incomplete information ergodic theory, and the measurability of equilibria, Israel J. Math. 138 (2003) 73–92.
[33] M.B. Stinchcombe, Countably additive subjective probabilities, Rev. Econ. Stud. 64 (1997) 125–146.
[34] M.B. Stinchcombe, Nash equilibrium and generalized integration for infinite normal form games, Games Econ. Behav. 50 (2005) 332–365.
[35] M.B. Stinchcombe, Balance and discontinuities in infinite games with type-dependent strategies, J. Econ. Theory (2010), doi:10.1016/j.jet.2010.12.009, in press.