

The Alternative to Equilibrium Existence

David Rahman*
University of Minnesota

October 3, 2008

Abstract

This paper establishes and interprets a necessary and sufficient condition for existence of (countably additive) correlated equilibrium in n -person games, assuming only that utility functions are bounded, measurable. A sequence of deviation profiles is *consistent* if there exists a correlated strategy that makes every profile in the sequence unprofitable with respect to the sum of utilities. An equilibrium exists if and only if every sequence of deviation profiles has a consistent subsequence. This condition fails to characterize Nash equilibrium. As a direct corollary, existence of (communication) equilibrium is characterized in games with incomplete information on type spaces large enough to include the universal one. Exact conditions for existence of approximate correlated equilibrium are also obtained, as well as a value for two-person zero-sum games.

JEL Classification: C61, C72.

Keywords: correlated equilibrium, consistency, duality, discontinuous games.

*Please send any comments to dmr@umn.edu. I thank Ichiro Obara, Joe Ostroy, Itai Sher and Bill Zame for comments and suggestions.

1 Introduction

This paper provides necessary and sufficient conditions for existence of both correlated equilibrium and approximate equilibrium in general games, specifically games with measurable action spaces and bounded, measurable utility functions. Therefore, no topological assumptions are required for existence of equilibrium. As corollaries, existence of communication equilibrium in games with incomplete information on large type spaces (including the universal one) as well as existence of a value for zero-sum games are characterized.

Near consistency characterizes existence of approximate equilibrium (Theorem 1), and *sequential consistency* characterizes equilibrium existence (Theorem 2). Near consistency means that for every $\varepsilon > 0$ and every profile of deviation plans by the players, there exists a correlated strategy (not necessarily the same one) such that the sum of players' unilateral deviation gains is less than ε . Conversely, near consistency fails if there exists $\varepsilon > 0$ and a deviation profile such that for any correlated strategy, some player i gains at least ε/n . Sequential consistency means that every sequence of deviation profiles has a subsequence for which there is a correlated strategy that makes the sum of deviation gains not positive along the subsequence. Neither near consistency nor sequential consistency characterize Nash equilibrium.

This study builds on and extends the work of Hart and Schmeidler (1989) by studying countably additive equilibria of discontinuous games. Hart and Schmeidler use duality to prove existence of countably additive correlated equilibrium in games with compact action spaces and continuous utility functions, as well as finitely additive correlated equilibrium when utilities are just bounded, measurable. However, they argue by example that finitely additive probabilities are undesirably unintuitive. This note bridges the gap between their two results by finding necessary and sufficient conditions for existence of *countably additive* correlated equilibrium when utilities are not necessarily continuous and action spaces are not necessarily compact.

The study of (countably additive) correlated equilibrium in discontinuous games has not been explored much in the literature.¹ However, the fact that correlated equilibrium is defined by linear inequalities suggests that its existence may be easier to analyze than that of Nash equilibrium. Moreover, existence of correlated equilibrium is strictly more general than Nash equilibrium, as Example 1 below shows.

¹Yannelis and Rustichini (1992) is an exception, but it focuses on (weakly) continuous games.

Example 1. Consider the following three-player game without a Nash equilibrium (in pure or mixed strategies) but where a correlated equilibrium exists.² Player 1 has two actions, $A_1 = \{U, D\}$. Player 2 also has two actions, $A_2 = \{L, R\}$. Player 3 has *infinitely many* actions, $A_3 = \mathbb{Z}$, with typical element z . Payoffs are depicted below.

	L	R		L	R		L	R
U	2, 2, $3 + 1/z$	0, 0, 8	U	2, 2, 2	0, 0, 0	U	2, 2, 0	0, 0, 0
D	0, 0, 0	2, 2, 0	D	0, 0, 0	2, 2, 2	D	0, 0, 8	2, 2, $3 - 1/z$
	$z < 0$			$z = 0$			$z > 0$	

Player 1 chooses a row, player 2 chooses a column, and player 3 chooses a matrix. Thus, if players 1 and 2 choose (U, L) and player 3 chooses $z = -4$ then players 1 and 2 each get a payoff of 2 and player 3 gets a payoff of 2.75.

Notice first of all that there is no Nash equilibrium in this game. Indeed, if there were one then it would have to involve players 1 and 2 playing either (U, L) , (D, R) , or each mixing independently with probability $\frac{1}{2}$. However, it is easy to see that in either case player 3 has no best response. On the other hand, there does exist a correlated equilibrium, namely where players 1 and 2 perfectly correlate their play by mixing with probability $\frac{1}{2}$, say, between (U, L) and (D, R) , and player 3 plays 0. It is easy to see that no player has the incentive to deviate from this correlated strategy.

The question of existence of Nash equilibrium as been explored in depth, with many partial answers (Dasgupta and Maskin, 1986; Simon, 1987; Simon and Zame, 1990; Reny, 1999; Athey, 2001). Most approaches tended to consider mixed Nash equilibrium as a special case of pure Nash. That is, they first established sufficient conditions for existence of pure-strategy Nash equilibrium, and then reinterpreted them on the space of mixed strategies.

The approach below relies on duality to squarely ask for existence of any equilibrium, be it pure, mixed, or correlated. Technically, we apply Clark's (2006) extension of Farkas' Lemma for infinite-dimensional spaces to the incentive constraints defining correlated equilibrium. This leads to a dual system of inequalities that is equivalent to equilibrium, which lends itself to a strategic interpretation exploited below.

²Although Hart and Schmeidler (1989) provide an example without Nash equilibrium but with correlated equilibrium, they rely on there being infinitely many players, each with finitely many strategies. Here, on the other hand, there are finitely many players, only one of which has infinitely many strategies. It is also possible to construct similar examples with only two players.

The paper is organized as follows. [Section 2](#) introduces the basic model and states the main results. [Section 3](#) discusses the model, interprets the results, and relates them to the previous literature. [Section 4](#) uses these results in games of incomplete information, zero-sum games, and environments with transfers. [Section 5](#) concludes. [Appendix A](#) collects notation and preliminaries on duality in general spaces. Finally, proofs are relegated to [Appendix B](#).

2 Model

In this section we define our objects of study and state our main results. Specifically, we first introduce the general environment of games with bounded measurable payoffs, followed by definitions of correlated and approximate correlated equilibrium. Next, we define and interpret consistency, near-consistency and sequential consistency. Finally the paper’s main results are stated and discussed. The main results characterize equilibrium and approximate equilibrium.

Let $I = \{1, \dots, n\}$ be a finite set of players. For every $i \in I$, let (A_i, \mathcal{A}_i) be a measurable space³ of actions for player i , where \mathcal{A}_i is a σ -algebra on A_i . Let

$$A = \prod_{i=1}^n A_i$$

be the product space of action profiles, endowed with the product σ -algebra \mathcal{A} . To describe a game, every player is given a utility function over action profiles.

Assumption 1. Every player i ’s utility function $u_i : A \rightarrow \mathbb{R}$ is bounded, measurable.

This assumption is necessary for the all of the analysis below. Although imposing measurability is not controversial (otherwise expected utility would be meaningless), boundedness—admittedly—is.⁴ Nevertheless, let us confine ourselves to this otherwise quite general environment.

Definition 1. A *game with bounded measurable payoffs*, or simply a *game*, is any $\Gamma = (I, A, u)$ that satisfies all of the above.

³We assume measurable action spaces to define (mixed) correlated strategies and deviation plans.

⁴Boundedness is assumed to avoid indeterminacies like $\infty - \infty$ when calculating deviation gains. Although we may apply a monotone transformation to utilities in order to make them bounded, in general this would lead to a different von Neumann-Morgenstern utility over lotteries.

2.1 Equilibrium and Approximate Equilibrium

Next, let us define equilibrium and approximate equilibrium. A *deviation plan* for player i is any measurable map $\alpha_i : A_i \rightarrow \Delta(A_i)$.^{5,6} Thus, *obedience* for player i is the plan θ_i defined pointwise by $\theta_i(a_i) = [a_i]$.⁷ A *profile of deviation plans* or simply *deviation profile* is any n -tuple $\alpha = (\alpha_1, \dots, \alpha_n)$ of deviation plans for each player. Thus, the obedient profile is $\theta = (\theta_1, \dots, \theta_n)$. A *correlated strategy* is any $\mu \in \Delta(A)$.

Definition 2. A correlated strategy μ is called a *correlated equilibrium* or simply an *equilibrium* if for every player i and every deviation plan α_i ,

$$\int_A \int_{A_i} u_i(b_i, a_{-i}) - u_i(a) \alpha_i(db_i|a_i) \mu(da) \leq 0.$$

This definition may be interpreted as follows. A mediator selects an action profile a according to the correlated strategy μ , and privately recommends each player to play a_i . Every player i contemplates a deviation plan α_i that amounts to recommendation-contingent mixed strategies with the only restriction being that α_i be measurable.⁸ A correlated equilibrium is any μ such that no deviation plan is ex ante profitable.⁹

Definition 3. A *weakly approximate correlated equilibrium* or simply an *approximate equilibrium* is any sequence of correlated strategies $\{\mu_m\} \subset \Delta(A)$ such that for every player i and every deviation plan α_i ,

$$\limsup_{m \rightarrow \infty} \int_A \int_{A_i} u_i(b_i, a_{-i}) - u_i(a) \alpha_i(db_i|a_i) \mu_m(da) \leq 0.$$

Intuitively, an approximate equilibrium is a sequence of correlated strategies such that every player's deviation gains become arbitrarily small along the sequence, regardless of the player's deviation plan. This definition is particularly weak because it is not asking for uniform convergence across deviation plans or players.

⁵For a measurable space (X, \mathcal{X}) , $\Delta(X, \mathcal{X})$ is the set of *countably additive* probability measures on X with domain contained in \mathcal{X} . If the σ -algebra \mathcal{X} is understood, we abbreviate to $\Delta(X)$.

⁶To define measurability, endow $\Delta(A_i)$ with the σ -algebra generated by the topology associated with the total variation norm.

⁷The notation $[a_i]$ stands for Dirac measure.

⁸Otherwise, a player would not be able to evaluate deviation plans ex ante.

⁹Definition 6 is equivalent to that provided by Hart and Schmeidler (1989) for continuous utilities on compact action spaces. The main difference is that they allow μ to be *finitely* additive when utilities are not continuous, whereas here we restrict attention to the countably additive case. They argue by example that finitely additive equilibrium may have undesirably unintuitive properties.

2.2 Consistency and Approximate Equilibrium

Next, we will define consistency and near consistency of a deviation plan. The latter will be equated to existence of approximate equilibrium. To this end, for any deviation profile α and any correlated strategy $\mu \in \Delta(A)$, let

$$\Pi(\mu)(\alpha) := \int_A \sum_{i=1}^n \int_{A_i} u_i(b_i, a_{-i}) - u_i(a) \alpha_i(db_i|a_i) \mu(da).$$

denote the sum of players' ex ante unilateral deviation gains from α when recommendations are drawn according to μ . With a slight abuse of notation, write $\Pi(a)(\alpha)$ instead of $\Pi([a])(\alpha)$ for correlated strategies that happen to be pure action profiles.¹⁰

Definition 4. A deviation profile α is *consistent with equilibrium* or simply *consistent* if $\Pi(\mu)(\alpha) \leq 0$ for some correlated strategy μ . It is *nearly consistent* if for every $\varepsilon > 0$ there is a correlated strategy μ_ε such that $\Pi(\mu_\varepsilon)(\alpha) < \varepsilon$.

Intuitively, a deviation profile is consistent if it does not strictly dominate obedience with respect to the sum of utilities. It is nearly consistent if it does not *uniformly* strictly dominate obedience, i.e., there is no $\varepsilon > 0$ such that the sum of deviation gains across players is at least ε at every action profile. Figure 1 below illustrates.

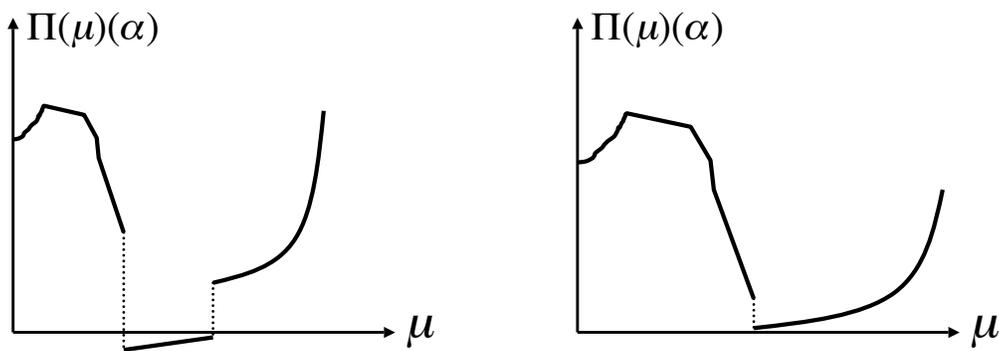


Figure 1: Consistency (left) versus near consistency (right) of α .

By convexity, we may have equivalently defined consistency and near-consistency in terms of pure action profiles rather than correlated strategies. Thus, it follows immediately that a deviation profile is consistent if and only if $\Pi(a)(\alpha) \leq 0$ for some action profile a , and it is nearly consistent if and only if for every $\varepsilon > 0$ there exists an action profile a_ε such that $\Pi(a_\varepsilon)(\alpha) < \varepsilon$.

¹⁰The choice of notation $\Pi(\mu)(\alpha)$ is explained later in [Appendix A](#).

Theorem 1. *A game has an approximate correlated equilibrium if and only if every deviation profile is nearly consistent.*

The thrust of [Theorem 1](#) is that different correlated strategies may be used to make different deviation profiles nearly consistent. Intuitively, if every deviation profile can *independently* be made nearly unprofitable then it is possible to *jointly* make every deviation profile nearly unprofitable. Conversely, to rule out approximate equilibria it is enough to find a deviation profile that fails to be nearly consistent. To illustrate, consider the following classic example due to [Sion and Wolfe \(1957\)](#).

Example 2. Two players, $A_i = [0, 1]$ with the Borel σ -algebra and payoffs

$$u_1(a_1, a_2) = -u_2(a_1, a_2) = \begin{cases} -1 & \text{if } a_1 < a_2 < z_1 + 1/2 \\ 0 & \text{if } x = y \text{ or } y = x + 1/2 \\ 1 & \text{otherwise.} \end{cases}$$

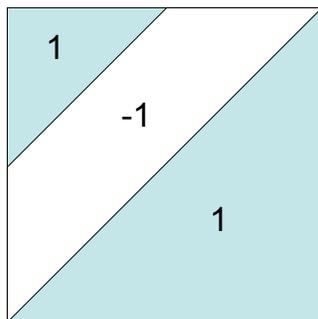


Figure 2: A zero-sum game with no value.

Let $\alpha_1(\{1\}|a_1) = 1$ if $a_1 < 1$, $\alpha_1(\{0\}|1) = 1/3$, $\alpha_1(\{1\}|1) = 2/3$ and $\alpha_2(\{0\}|a_2) = \alpha_2(\{1/2\}|a_2) = 1/2$. It is easy to verify that this deviation profile is not nearly consistent. Therefore, an approximate correlated equilibrium fails to exist.

By [Theorem 1](#), equilibrium existence in games with finitely many actions as well as those with compact metric action spaces and continuous payoffs follows easily. Indeed, with finitely many actions near consistency clearly implies consistency, which is dually equivalent to existence of correlated equilibrium (see, e.g., [Myerson, 1997](#)). For continuous games with compact metric action spaces, an approximate equilibrium has a (weak*) convergent subsequence with a limiting correlated strategy that is a correlated equilibrium by continuity of payoffs. Appealing to an argument by [Hart and Schmeidler \(1989, page 24\)](#), without loss we may restrict attention to (weak*) continuous deviation profiles. After such restriction it is easy to see that near consistency again implies consistency by continuity and compactness.

2.3 Sequential Consistency and Equilibrium

The previous paragraph begs the question of whether consistency of every deviation profile is *sufficient* for equilibrium existence.¹¹ Although it does lead to equilibrium in both finite and compact-continuous games, we will show that consistency is generally not enough for equilibrium to exist. To find what is enough for equilibrium existence, we first define consistency of a *sequence* of deviation profiles.

Definition 5. A sequence of deviation profiles $\{\alpha_m\}$ is *consistent* if

$$\bigcap_{m \in \mathbb{N}} \{\mu \in \Delta(A) : \Pi(\mu)(\alpha_m) \leq 0\} \neq \emptyset.^{12}$$

A sequence of deviation profiles is consistent if all the deviation profiles in the sequence have a correlated strategy in common that make them simultaneously unprofitable with respect to the sum of utilities. Figure 3 (left) below provides an illustration of such a consistent sequence. To help understand this definition, firstly notice that restricting attention to constant sequences recovers consistency from Definition 4. Indeed, clearly a constant sequence of deviation profiles $\{\alpha_m\}$ with $\alpha_m = \alpha$ for all m is consistent if and only if α is consistent. Sequential consistency is strictly stronger than consistency, as Figure 3 (right) suggests with an inconsistent sequence $\{\alpha_m\}$ of deviation profiles such that every deviation profile α_m is itself consistent. Intuitively, sequential consistency requires consistency as well as a minimal amount of commonality in correlated strategies that make deviation profiles unprofitable.

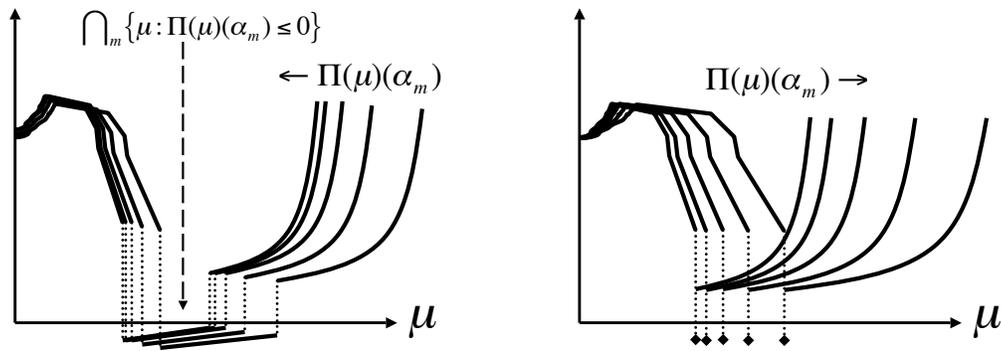


Figure 3: Satisfaction (left) versus failure (right) of consistency for $\{\alpha_m\}$.

¹¹Clearly, consistency is necessary. Otherwise there exists a deviation profile such that the sum of deviation gains is positive for every action profile. By linearity, it is still positive for every correlated strategy, i.e., somebody always gains. Hence, correlated equilibrium fails to exist.

¹²For any deviation profile α , $\{\mu \in \Delta(A) : \Pi(\mu)(\alpha) \leq 0\}$ is the set of correlated strategies whose sum of deviation gains from α is not positive.

Theorem 2. *A game has a correlated equilibrium if and only if every sequence of deviation profiles has a consistent subsequence.*

Call this condition, which characterizes equilibrium existence, *sequential consistency*. Sequential consistency bridges the gap between consistency and equilibrium. On the one hand, consistency ensures that no deviation plan can strictly dominate obedience, even though different correlated strategies may be used to make different deviation plans unprofitable. On the other hand, equilibrium requires that the same correlated strategy make unprofitable every deviation plan simultaneously.

Rather than check if all deviation plans have a common correlated strategy that makes them unprofitable, it may be easier to check for sequential consistency, both to confirm and deny equilibrium existence. Thus, to conclude that no equilibrium exists, it is enough to find a sequence of deviation profiles such that every subsequence is inconsistent. Although equilibrium discourages uncountably many deviation plans simultaneously, [Theorem 2](#) only requires focusing on them one sequence at a time.

We end this section with a corollary that offers a test of equilibrium existence that may be easier to check than sequential consistency. Call a sequence $\{\alpha_m\}$ of deviation profiles *purely consistent* if $\bigcap_m \{a \in A : \Pi(a)(\alpha_m) \leq 0\} \neq \emptyset$, i.e., there is a pure action profile that makes it consistent. [Figure 4](#) below illustrates the difference between pure consistency and consistency by sketching a consistent sequence of deviation profiles that clearly fails to be purely consistent.

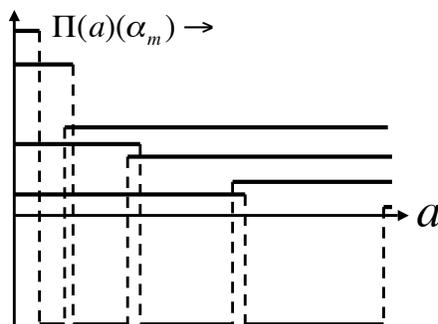


Figure 4: A consistent but not purely consistent sequence $\{\alpha_m\}$.

Since pure consistency implies consistency, it follows immediately from [Theorem 2](#) that pure sequential consistency delivers equilibrium existence.

Corollary 1. *A game has a correlated equilibrium if every sequence of deviation profiles has a purely consistent subsequence.*

3 Discussion

In this section we shall comment on the significance of [Theorems 1 and 2](#), further interpret consistency and sequential consistency, and relate them to the literature.

In finite games, the strategic content of equilibrium existence is arguably consistency: equilibrium exists if and only if every deviation plan is consistent, although different correlated strategies may be used to make different deviation plans consistent. Since every deviation plan *is* consistent in a finite game, equilibrium exists. A similar logic applies to compact-continuous games. In general games without topological structure, the strategic content of equilibrium existence becomes sequential consistency. With regard to approximate equilibrium, however, near consistency suffices for existence: approximate equilibrium exists as long as every deviation profile is nearly consistent, where different correlated strategies may be used to make nearly consistent different deviation profiles. Strategically, this means that as long as every deviation plan can be approximately discouraged individually then all can be discouraged simultaneously.

[Theorems 1 and 2](#) interpret the dual inequalities to those describing equilibrium existence. Since equilibrium is described linearly and the games under consideration are relatively general, the conditions of these theorems cannot be relaxed. At the same time, given the generality of the problem it is not surprising that the conditions are general, difficult to check from a practical point of view. However, these conditions, although generally impractical, do suggest a way to interpret equilibrium existence.¹³ Moreover, adding specificity to the model does add specificity, hence practical value, to the characterizing conditions, sometimes even without loss of generality.

One way to add specificity to the model is to consider only *simple* deviation plans, i.e., those with finite range, which is clearly without loss. Another specification that does incur some loss is to endow action spaces with topologies. In this case, by Lusin's Theorem (see, e.g., [Folland, 1999](#)) attention may be restricted to *continuous* deviation plans, since they are dense in the space of all deviation plans. Such specifications facilitate comparison with the literature. For instance, neither compactness nor metrizable-ness of action spaces is required for sequential consistency, in contrast with the literature. As a result, this condition incorporates the relevant aspects—and *only* the relevant aspects—of assumptions such as compactness for equilibrium.

¹³To illustrate, in one-player games sequential consistency is equivalent to the player's utility function possessing a maximum, without requiring reference to continuity or compactness.

Specifically, compare sequential consistency with other conditions in the literature. Firstly, consider [Nikaido and Isoda \(1955\)](#). They introduced the auxiliary function $\Pi(a, b) = \sum_i u_i(b_i, a_{-i}) - u_i(a)$ and emphasized its use for studying Nash equilibrium, establishing existence under certain continuity assumptions that were subsequently relaxed by [Baye et al. \(1993\)](#). [Dasgupta and Maskin \(1986\)](#) required that (i) $\sum_i u_i(a)$ be upper semicontinuous, (ii) each $u_i(b_i, a_{-i})$ be “weakly” lower semicontinuous in b_i , and (iii) discontinuities lie along “diagonal” sets for mixed Nash equilibrium to exist.

Use of these “aggregators” has been criticized (e.g., [Reny, 1999](#), footnotes 11 and 12) for being cardinal in nature, and more “ordinal” generalizations of these conditions have emerged ([Simon, 1987](#); [Reny, 1999](#)).¹⁴ However, [Theorems 1 and 2](#) show that Π is as ordinal as possible.¹⁵ In this paper, Π follows from the Lagrangian associated with finding correlated equilibrium, where deviation plans are (proportional to) multipliers on incentive constraints. Hence, there is nothing ad hoc about restricting attention to Π when studying equilibrium existence.

Furthermore, Π delivers and clarifies [Reny’s](#) sufficient condition, called better-reply security, for existence of mixed Nash equilibrium (which in turn generalizes [Simon’s](#)) with a proof similar to that by [Baye et al. \(1993\)](#). Indeed, let $\Pi(\sigma, \tau)$ be the extension of Π to the (weak* compact) space of mixed strategy profiles. For every such profile τ , let $F_\tau = \{\sigma : \Pi(\sigma, \tau) \leq 0\}$ be the set of profiles σ that make τ unprofitable with respect to Π . Clearly, equilibrium exists if and only if $\bigcap_\tau F_\tau \neq \emptyset$. If F_τ is closed for every τ (which follows if $\Pi(\sigma, \tau)$ is lower semicontinuous in σ) then we are done by compactness and Nash’s Theorem. To see this, for any finite subfamily $\{\tau_1, \dots, \tau_m\}$, we must have $\bigcap_k F_{\tau_k} \neq \emptyset$, since every finite game has an equilibrium point. Therefore, $\{F_\tau\}$ has the finite intersection property and by compactness equilibrium exists. In case F_τ is not necessarily closed, [Reny’s](#) better-reply security implies that $\bigcap_\tau F_\tau = \bigcap_\tau \overline{F_\tau}$, which leads to equilibrium since clearly $\{\overline{F_\tau}\}$ has the finite intersection property. To see this, assume the contrary, i.e., there exists $\sigma \in \bigcap_\tau \overline{F_\tau} \setminus \bigcap_\tau F_\tau$. Hence, $\sigma \notin F_\tau$ for some τ , i.e., $\Pi(\sigma, \tau) > 0$, so σ is not an equilibrium. By better-reply security, there is a player i and a strategy $\hat{\tau}_i$ such that $u_i(\hat{\tau}_i, \hat{\sigma}_{-i}) \geq \alpha > \bar{u}_i$ for all $\hat{\sigma}_{-i}$ in some open neighborhood of σ_{-i} , where $\bar{u}_i = \lim u_i(\sigma_m)$ for some sequence $\{\sigma_m\}$ converging to σ .

¹⁴E.g., reciprocal upper semicontinuity ([Simon, 1987](#)) generalizes [Dasgupta and Maskin’s \(1986\)](#) requirement that the sum of utilities be upper semicontinuous ([Reny, 1999](#), page 1034).

¹⁵When looking exclusively for pure strategy equilibria, seeking conditions that are ordinal in the sense that they are invariant to monotone transformations of preferences can be justified. However, when mixed or correlated strategies are considered, the meaningful family of preference-representing utility functions is significantly reduced to just affine transformations.

But this implies that $\sigma \notin \overline{F_{\hat{\tau}}}$, where $\hat{\tau} = (\hat{\tau}_i, \sigma_{-i})$, a contradiction. Otherwise, there would exist some sequence $\{\sigma_m\}$ converging to σ with $\Pi(\sigma_m, \hat{\tau}) \leq 0$ for all m , i.e., $u_i(\hat{\tau}_i, \sigma_{-im}) \leq u_i(\sigma_m)$. But this contradicts better-reply security, since it implies that $u_i(\hat{\tau}_i, \sigma_{-im}) \geq \alpha > u_i(\sigma_m)$ for all m large enough that $|u_i(\sigma_m) - \bar{u}_i| < \alpha - \bar{u}_i$.

How does sequential consistency compare with better-reply security? By [Example 1](#) and [Theorem 2](#), better-reply security implies sequential consistency but not otherwise. For a version of better-reply security corresponding to correlated equilibrium, let $F_\alpha = \{\mu : \Pi(\mu)(\alpha) \leq 0\}$. Given a game with compact action spaces, correlated equilibrium exists if $\bigcap_\alpha F_\alpha = \bigcap_\alpha \overline{F_\alpha}$, where $\overline{F_\alpha}$ is the weak* closure of F_α . This yields the following generalization of better-reply security: if $\mu \in \Delta(A)$ is not a correlated equilibrium then there is a player i and a deviation plan α_i such that $u_i(\alpha_i, \hat{\mu}) \geq \alpha > \bar{u}_i$ for all $\hat{\mu}$ in some open neighborhood of μ , where $\bar{u}_i = \lim u_i(\mu_m)$ for some sequence $\{\mu_m\}$ converging to μ . To see that (not even this version of) better-reply security is unnecessary for sequential consistency consider the following example of a one-person game. Let the player's action space be the interval $A = [0, 1]$ with the usual topology, with his utility function equal to 1 if $a \in (1/3, 2/3)$, and 0 otherwise. Clearly, better-reply security fails (since the set of equilibria is not closed) yet sequential consistency holds, since equilibrium exists.

This discussion helps to clarify the content of [Theorem 2](#). Since $\Delta(A)$ (with the topology of [Appendix A](#)) is Lindeloff, it inherits the countable intersection property, so equilibrium existence is implied by every sequence $\{\alpha_m\}$ satisfying $\bigcap_m F_{\alpha_m} \neq \emptyset$. In order to establish equilibrium existence, i.e., $\bigcap_\alpha F_\alpha \neq \emptyset$, it is enough that every sequence $\{\alpha_m\}$ to have a subsequence $\{\alpha_{m_k}\}$ with $\bigcap_k F_{\alpha_{m_k}} \neq \emptyset$. (Necessity is obvious: if some sequence has no consistent subsequence then clearly $\bigcap_\alpha F_\alpha = \emptyset$.)

As a final comment on sequential consistency, note that it is *not* equivalent to *finite consistency*, i.e., that every finite family of deviation plans is consistent with the same correlated strategy. By [Theorem 2](#), sequential consistency implies finite consistency. For an example to show that the converse fails, consider a one-person game with $A = \mathbb{R}$ and $u(\alpha_m, a) > 0$ if and only if $a \leq m$. (Examples also exist with A compact.)

We end this section with a technical corollary of [Theorem 2](#) and the Riesz Representation Theorem, characterizing regular equilibrium (see, e.g., [Folland, 1999](#)).

Corollary 2. *If every A_i is locally compact Hausdorff with the Borel σ -algebra then sequential consistency also characterizes existence of regular correlated equilibrium.*

4 Applications

In this section the main results of the paper, [Theorems 1 and 2](#), are applied and extended to three important settings. Firstly, we characterize existence of a value for zero-sum games. Secondly, we characterize implementability allocations with linear transfers. Finally, we characterize existence of equilibrium in games with incomplete information on large type spaces, including the universal one.

4.1 Zero-Sum Games

In this subsection we explore equilibrium existence in two-person zero sum games. We will adopt the following notation. Let X be the measurable space of actions for player 1 and Y the actions of player 2, with typical elements x and y . Let $u : X \times Y \rightarrow \mathbb{R}$ be player 1's contingent payment function from player 2, assumed bounded, measurable. A deviation plan for player 1 is denoted by α and one for player 2 is denoted by β .

An immediate corollary of [Theorem 2](#) is the following necessary and sufficient condition for existence of correlated equilibrium.

Corollary 3. *A two-person zero-sum game has a correlated equilibrium if and only if every sequence $\{(\alpha_m, \beta_m)\}$ of deviation profiles has a subsequence $\{(\alpha_{m_k}, \beta_{m_k})\}$ with*

$$\bigcap_{k \in \mathbb{N}} \left\{ \mu \in \Delta(X \times Y) : \int \int_Y u(x, z) \beta_{m_k}(dz|y) - \int_X u(w, y) \alpha_{m_k}(dw|x) d\mu \leq 0 \right\} \neq \emptyset.$$

Given the particular structure of two-person zero-sum games, it is possible to characterize Nash equilibrium with similar techniques, as the next result shows.

Theorem 3. *A two-person zero-sum game has a Nash equilibrium if and only if every sequence of mixed-strategy profiles has a product-consistent subsequence, i.e., for each sequence $\{(\sigma_m, \tau_m)\} \subset \Delta(X) \times \Delta(Y)$ there is a subsequence $\{(\sigma_{m_k}, \tau_{m_k})\}$ with*

$$\bigcap_{k \in \mathbb{N}} \left\{ (\sigma, \tau) \in \Delta(X) \times \Delta(Y) : \int_X \int_Y u(x, y) [\tau_{m_k}(dy) \sigma(dx) - \sigma_{m_k}(dx) \tau(dy)] \leq 0 \right\} \neq \emptyset.$$

Call this property sequential product-consistency.

Existence of approximate correlated as well as Nash equilibria can also be characterized in the same spirit as [Theorem 1](#) above—the details are available on request.

4.2 Mechanism Design

Let us first define a *type space*. For every player i , let T_i be a measurable space of *types* with typical element t_i , and

$$T := \prod_{i=1}^n T_i$$

the product space of *type profiles*, endowed with the product σ -algebra, and with typical element $t = (t_1, \dots, t_n)$. Every player i 's *posterior beliefs* conditional on his type are collected in the measurable map $p_i : T_i \rightarrow \Delta(T_{-i})$.¹⁶

Let X denote a measurable space of social *choices*, and endow every player i with a bounded, measurable *type-contingent* utility function $u_i : X \times T \rightarrow \mathbb{R}$. An *allocation* is any measurable map $\mu : T \rightarrow \Delta(X)$.¹⁷ An allocation μ is *incentive compatible* if for every player i and every ordered pair of types (t_i, s_i) ,

$$\pi_i(t_i, s_i) := \int_{T_{-i}} \int_X u_i(x|t) [\mu(dx|s_i, t_{-i}) - \mu(dx|t)] p_i(t_{-i}|t_i) \leq 0.$$

Clearly, a *trivial* allocation (with $\mu(t) = \mu(s)$ for all t and s) is always incentive compatible. Given an allocation μ we will ask whether or not it is *implementable with linear transfers* or simply *implementable*, i.e., there exists a bounded, measurable map $\xi : I \times T \rightarrow \mathbb{R}$ such that for every player i and every ordered pair of types (t_i, s_i) ,

$$\int_{T_{-i}} \int_X u_i(x|t) [\mu(dx|s_i, t_{-i}) - \mu(dx|t)] p_i(t_{-i}|t_i) \leq \int_{T_{-i}} \xi_i(s_i, t_{-i}) - \xi_i(t) p_i(t_{-i}|t_i).$$

We will characterize implementation as follows. A *type-report outcome* for player i is any $\rho_i \in \Delta(T_i \times T_i)$, i.e., a probability measure over type-report pairs for player i . A *type-report outcome profile* or simply *outcome profile* is any n -tuple $\rho = (\rho_1, \dots, \rho_n)$. An outcome profile is *undetectable* if for every player i and every measurable $S \subset T$,

$$q_i(S) := \int \int_{T_{-i}} \mathbf{1}_S(t) - \mathbf{1}_S(s_i, t_{-i}) p_i(dt_{-i}|t_i) d\rho_i = 0.$$

A *sequence of outcome profiles* or simply an *outcome sequence* is denoted by $\{\rho_m\}$. An outcome sequence is *asymptotically undetectable* if $\lim q_{im} = 0$ for every player i , and *asymptotically unprofitable* if $\lim \int \pi_i d\rho_{im} \leq 0$.

Theorem 4. *An allocation is implementable if and only if every asymptotically undetectable outcome sequence is asymptotically unprofitable.*

¹⁶Measurability is defined by endowing $\Delta(T_{-i})$ with the Borel σ -algebra generated by the total variation norm on $\Delta(T_{-i})$ as a subspace of $M(T_{-i})$.

¹⁷Footnote 16 applies with X instead of T_{-i} .

4.3 Communication Equilibrium

As usual, A_i denotes the set of actions available to player i , and A the product space of action profiles. With reference to the type space defined previously, endow every player i with a bounded, measurable type-contingent utility function $u_i : A \times T \rightarrow \mathbb{R}$. Assume that a player's type does not affect his available actions.¹⁸ A *game with incomplete information* is a tuple (I, T, p, A, u) satisfying all of the above.

Next, we provide a notion of equilibrium for games with incomplete information.¹⁹ To this end, call a *communication mechanism* any measurable map $\mu : T \rightarrow \Delta(A)$.²⁰ Intuitively, a mediator asks players to report their types t and as a result makes recommendations according to the probability measure $\mu(t)$. We will draw a distinction between *disobedience*, i.e., playing a different action from the mediator's recommendation, and *dishonesty*, i.e., misreporting one's type. A *disobedience* is any measurable map $\alpha_i : A_i \rightarrow \Delta(A_i)$. Intuitively, $\alpha_i(a_i)$ is played by i when a_i is recommended. Let D_i be the set of all such disobediences.²¹ Let $\Theta_i = T_i \times T_i \times D_i$ be the space of types, reports, and disobediences, with typical element $\theta_i = (t_i, s_i, \alpha_i)$ called an *outcome* for player i . An *overall outcome* or simply an *outcome* is any n -tuple $\theta = (\theta_1, \dots, \theta_n)$.

Definition 6. A communication mechanism μ is called a *communication equilibrium* or simply an *equilibrium* if for every player i , type t_i , report s_i and disobedience α_i

$$\int_{T_{-i}} \int_A \int_{A_i} u_i(b_i, a_{-i}|t) \alpha_i(db_i|a_i) \mu(da|s_i, t_{-i}) - \int_A u_i(a|t) \mu(da|t) p_i(dt_{-i}|t_i) \leq 0. \quad (*)$$

Denote the left-hand side of $(*)$ by $\Phi_i(\mu)(\theta_i)$, and define $\Phi(\mu)(\theta) := \sum_i \Phi_i(\mu)(\theta_i)$.

With the same method of proof as for [Theorem 2](#), we obtain the following exact condition for existence of communication equilibrium on arbitrary type spaces.

Theorem 5. *A game with incomplete information has a communication equilibrium if and only if for every sequence $\{\alpha_m\}$ and every measure $q \in M(T)$,*

$$\int_{\Theta} \Phi(\mu)(\theta) \alpha_m(d\theta) \geq q(S)$$

for every measurable subset $S \subset T$, every measurable map $\mu : S \rightarrow \Delta(A)$ and all except at most finitely many m implies that $q(T) \leq 0$.

¹⁸This restriction is without loss by making unavailable actions extremely undesirable.

¹⁹See Forges (1993) for a discussion of equilibrium in games with incomplete information.

²⁰Footnote 16 applies with A instead of T_{-i} .

²¹Using the norm $\|\alpha_i\| = \sup_{a_i} |\alpha_i|(A_i|a_i)$, D_i becomes a topological space. Endow D_i with the associated Borel σ -algebra to make it a measurable space.

5 Conclusion

This paper derived the alternative system of inequalities associated with existence of correlated equilibrium as well as approximate equilibrium, provided an interpretation of the alternative, briefly discussed related literature, and suggested applications.

On the one hand, near consistency and sequential consistency provide a way of understanding equilibrium, more than necessarily providing a way to check for its existence. On the other, this paper delineates the limits of what characterizations of equilibrium can be obtained with use of duality. It would be interesting to find ways of applying sequential consistency as a useful tool for establishing equilibrium.

A Preliminaries

In this appendix we present Clark's (2006) extension of Farkas' Lemma as well as notation that will be used to establish our main results.

Let X and Y be ordered, locally convex real vector spaces, with positive cones X_+ and Y_+ and topological dual spaces X^* and Y^* such that $X^{**} = X$ and $Y^{**} = Y$. Let $A : X \rightarrow Y$ be a continuous linear operator with adjoint operator $A^* : Y^* \rightarrow X^*$ and fix any $b \in Y$. Finally, for any set S let \bar{S} denote its closure.

Lemma A.1 (Clark, 2006, page 479). *For any $b \in Y$, there exists $x \in X_+$ such that $A(x) = b$ if and only if $A^*(y_0^*) \in \overline{X_+^* - \{A^*(y^*) : y^*(b) = 0\}}$ implies that $y_0^*(b) \geq 0$.*

Now for some notation. Let $M(A)$ or simply M be the space of all finite (signed) measures over A with respect to the σ -algebra \mathcal{A} , and $M(A_i)$ or simply M_i the same space of finite measures over A_i . Write $B(A)$ or just B for the space of all bounded, measurable, real-valued functions on A , and $B(A_i)$ or simply B_i for the corresponding space on A_i . We will view M and B as a dual pair with respect to Lebesgue integration. That is, we will view every $f \in B$ as a linear functional on M according to the operation $f(\mu) = \int f d\mu$, and similarly every $\mu \in M$ as a linear functional on B according to $\mu(f) = \int f d\mu$. Endow M with the weakest topology²² that makes the elements of B continuous linear functionals as stated and similarly endow B with the weakest topology that makes every $\mu \in M$ continuous. Both topologies are locally convex, since they clearly separate points, and the spaces are mutually dual (Schaefer, 1970, page 52).

²²This is contained in the topology generated by the total variation norm, hence is first countable, so without loss it is described by sequences rather than nets.

Let $B(A_i, M_i)$ denote the space of bounded, measurable M_i -valued maps with the sup norm, where M_i is endowed with the Borel σ -algebra generated by the topology induced by the total variation norm. For instance, given $\lambda_i \in B(A_i, M_i)$, its norm is

$$\|\lambda_i\| = \sup_{a_i \in A_i} |\lambda_i|(A_i|a_i).$$

This is the space in which player i 's deviation plans will live. It is easy to see that $B(A_i, M_i)$ is a Banach space under this norm.²³ Write $\mathcal{D} = \prod_i B(A_i, M_i)$ for the vector space (with the product topology) where feasible deviation profiles will live. Let \mathcal{D}^* be the topological dual of \mathcal{D} and \mathcal{D}_+ the positive cone of \mathcal{D} .

Every game Γ induces an operator $\Pi : M \rightarrow \mathcal{D}^*$ as follows, where \mathcal{D}^* is endowed with the weak* topology. For any measure $\mu \in M$, $\Pi(\mu) : \mathcal{D} \rightarrow \mathbb{R}$ is a continuous linear functional defined pointwise for every $\lambda = (\lambda_1, \dots, \lambda_n) \in \mathcal{D}$ by

$$\Pi(\mu)(\lambda) := \int_A \sum_{i=1}^n \int_{A_i} u_i(b_i, a_{-i}) - u_i(a) \lambda_i(db_i|a_i) \mu(da).$$

For every μ it is clear that $\Pi(\mu)$ is a linear functional of λ . As for continuity, if $\lambda^m \rightarrow \lambda$ in \mathcal{D} then $\max_i \|\lambda_i^m\| \rightarrow \max_i \|\lambda_i\|$. This convergence is in the total variation norm, which is the same as convergence with respect to every bounded, measurable function, by the Bounded Convergence Theorem. Hence, $\Pi(\mu)$ is continuous, so $\Pi(\mu) \in \mathcal{D}^*$.

Notice also that Π is linear and continuous as an operator $M \rightarrow \mathcal{D}^*$. Linearity is obvious. Continuity follows because if $\mu^m \rightarrow \mu$ then $\Pi(\mu^m)(\lambda) \rightarrow \Pi(\mu)(\lambda)$ for all λ , again by the Bounded Convergence Theorem.²⁴

The *adjoint* operator $\Pi^* : \mathcal{D} \rightarrow B$ of Π is defined as follows. For any $\lambda \in \mathcal{D}$, let $\Pi^*(\lambda) \in B$ be defined pointwise by

$$\Pi^*(\lambda)(a) = \sum_{i=1}^n \int_{A_i} u_i(b_i, a_{-i}) - u_i(a) \lambda_i(db_i|a_i).$$

Finally, since we also consider B as the dual space of M , we extend the domain of $\Pi^*(\lambda)$ from A to $\Delta(A)$ as follows. For every $\mu \in \Delta(A)$, define $\Pi^*(\lambda)(\mu)$ pointwise according to

$$\Pi^*(\lambda)(\mu) = \int_A \sum_{i=1}^n \int_{A_i} u_i(b_i, a_{-i}) - u_i(a) \lambda_i(db_i|a_i) \mu(da).$$

Of course, the equation $\Pi^*(\lambda)(\mu) = \Pi(\mu)(\lambda)$ follows by definition of the adjoint operator.

²³Hence, its topology is first-countable, so is described by convergent sequences rather than nets.

²⁴Let $\hat{\mu} = \sum_m 2^{-m} \mu^m$. Now all the μ^m 's are absolutely continuous with respect to $\hat{\mu}$ and so have a Radon-Nikodym density in $L^\infty(\hat{\mu})$, and now apply the Bounded Convergence Theorem.

B Proofs

Theorem 1. The proof proceeds in three steps, and assumes the notation of [Appendix A](#). Firstly, a system of infinite dimensional linear inequalities is posited in Step 1. Subsequently, [Clark's \(2006\)](#) extension of Farkas' Lemma to infinite dimensional spaces is applied in Step 2 to the system posited in Step 1. Finally, in Step 3 it is shown that the system of inequalities in Step 1 fails to have a solution if and only if approximate equilibrium exists.

– *Step 1.* Consider the following primal system of linear inequalities:

$$\text{Find } \kappa \in \mathbb{R} \text{ and } \lambda \in \mathcal{D}_+ \text{ such that } \Pi^*(\lambda)(\mu) \geq \kappa \text{ for every } \mu \in \Delta(A) \text{ and } \kappa \geq 1. \quad (\text{P})$$

We will transform Problem (P) into a canonical system of linear equations for which we shall seek a non-negative solution. Define $T : \mathbb{R} \times \mathcal{D} \times B \rightarrow \mathbb{R} \times B$ pointwise by

$$T(\kappa, \lambda, \eta) = (\kappa, \Pi^*(\lambda) - \eta - \kappa \mathbf{1}),$$

and write T^* for its adjoint operator. Clearly, T is a continuous linear operator since Π^* is, so its adjoint is, too. Clearly, problem (P) is equivalent to the following canonical version:

$$\text{Find } (\kappa, \lambda, \eta) \in (\mathbb{R} \times \mathcal{D} \times B)_+ \text{ such that } \kappa = 1 \text{ and } \Pi^*(\lambda) - \eta - \kappa \mathbf{1} = \mathbf{0}, \quad (\text{P}^*)$$

where $\mathbf{1} \in B$ is the function identically equal to 1 and $\mathbf{0} \in B$ the zero function.

– *Step 2.* Let $F = \{T^*(0, \mu) : \mu \in M\}$ and $J = (\mathbb{R} \times \mathcal{D}^* \times M)_+ - F$. By [Lemma A.1](#), problem (P*) has a solution if and only if for every (δ, μ_0) such that $\delta \in \mathbb{R}$ and $\mu_0 \in M$, whenever $T^*(\delta, \mu_0) = (\delta - \mu_0(A), \Pi(\mu_0), -\mu)$ belongs to the cone \bar{J} , it follows that $\delta \geq 0$, where for any set S , \bar{S} denotes its closure. Hence, (P) has a solution if and only if for every (δ, μ_0) there is a sequence $\{(\delta_m, \lambda_m^*, \mu_m)\} \subset (\mathbb{R} \times \mathcal{D}^* \times M)_+ - F$ with $\lim \delta_m = \delta - \mu_0(A)$, $\lim \lambda_m^* = \Pi(\mu_0)$, and $\lim \mu_m = -\mu_0$ only if $\delta \geq 0$.

Any element of J can be written as $(\delta^+ + \mu(A), \lambda^{*+} - \Pi(\mu), \mu^+ + \mu)$, where $\delta^+ \in \mathbb{R}_+$, $\mu \in M$, $\lambda^{*+} \in \mathcal{D}_+^*$ and $\mu^+ \in M_+$. Therefore, problem (P*) has a solution if and only if for every $\delta \in \mathbb{R}$ and $\mu_0 \in M$, if there is a sequence $\{(\delta_m^+, \mu_m, \lambda_m^{*+}, \mu_m^+)\} \subset \mathbb{R}_+ \times M \times \mathcal{D}_+^* \times M_+$ such that $\lim \delta_m^+ + \mu_m(A) = \delta - \mu_0(A)$, $\lim \lambda_m^{*+} - \Pi(\mu_m) = \Pi(\mu_0)$ and $\lim \mu_m^+ + \mu_m = -\mu_0$ then $\delta \geq 0$. Plugging the third limit into the first one (together with continuity **1**) yields

$$\begin{aligned} 0 &= \lim \delta_m^+ + \mu_m(A) - \delta + \mu_0(A) = \\ \lim \delta_m^+ + \mu_m(A) - \delta - \mu_m^+(A) - \mu_m(A) &= \lim \delta_m^+ - \delta - \mu_m^+(A). \end{aligned}$$

Clearly, this last limit equals zero if and only if $\liminf \mu_m^+(A) \geq -\delta$. Also, plugging the third limit into the second (with continuity and linearity of Π) yields

$$\begin{aligned} 0 &= \lim \lambda_m^{*+} - \Pi(\mu_m) - \Pi(\mu_0) = \\ \lim \lambda_m^{*+} - \Pi(\mu_m) + \Pi(\mu_m^+ + \mu_m) &= \lim \lambda_m^{*+} + \Pi(\mu_m^+). \end{aligned}$$

Clearly, this last limit equals zero if and only if $\limsup \Pi(\mu_m^+) \leq 0$. Finally, all this implies that problem (P*) has a solution if and only if

$$\text{Given } \delta \in \mathbb{R}, \{\mu_m^+\} \subset M_+, \text{ if } \liminf \mu_m^+(A) \geq -\delta, \limsup \Pi(\mu_m^+) \leq 0 \text{ then } \delta \geq 0. \quad (\text{D}^*)$$

– *Step 3.* Finally, we equate (P) with failure of consistency and (D*) with failure of existence of approximate equilibrium. By Step 2, (D*) holds if and only if no sequence $\{\mu_m\} \subset M_+$ exists with $\liminf \mu_m(A) > 0$ and $\limsup \Pi(\mu_m)(\alpha) \leq 0$ for every deviation profile α . Clearly, this last condition is equivalent to the nonexistence of an approximate equilibrium. Finally, we will show that failure of (P*), i.e., that for every $\lambda \in \mathcal{D}_+$ there exists $\mu_\lambda \in \Delta(A)$ with $\Pi^*(\lambda)(\mu_\lambda) < 1$, is equivalent to near consistency. Given any deviation profile α and any $\varepsilon > 0$, let $\lambda_\varepsilon = \alpha / \min\{\varepsilon, 1\}$. For sufficiency, given any deviation profile α and any $\varepsilon > 0$ it is clear that $\Pi^*(\alpha)(\mu_{\lambda_\varepsilon}) < \varepsilon$. For necessity, given $\lambda \in \mathcal{D}_+$ and $\varepsilon > 0$, for every player i let α_i be the deviation plan defined pointwise by $\alpha_i(a_i) = p\lambda_i(a_i) + (1-p)\theta_i(a_i)$, where $p = \min\{\varepsilon, \min\{\max_i\{\|\lambda_i\|\}, \max_i\{1/\|\lambda_i\|\}\}\} \in (0, 1]$. By hypothesis, the deviation profile $\alpha = (\alpha_1, \dots, \alpha_n)$ satisfies $\Pi^*(\alpha)(\mu_\varepsilon) < \varepsilon$ for some μ_ε . By linearity and the fact that θ_i is obedient, $\Pi^*(\alpha)(\mu_\varepsilon) < \varepsilon$ if and only if $\Pi^*(\lambda)(\mu_\varepsilon) < \varepsilon/p \leq 1$, as required. \square

Theorem 2. The proof is “dual” to that of [Theorem 1](#), and proceeds in three similar steps. Firstly, the problem of equilibrium existence is formulated as a family of infinite dimensional linear inequalities in Step 1. Subsequently, [Clark’s \(2006\)](#) extension of Farkas’ Lemma to infinite dimensional spaces is applied in Step 2 to the system derived in Step 1. Finally, the dual inequalities of Step 2 are related to sequential consistency in Step 3.

– *Step 1.* With reference to [Appendix A](#), consider the following primal problem:

$$\text{Find } \mu \in \Delta(A) \text{ such that } \Pi(\mu)(\lambda) \leq 0 \text{ for every } \lambda \in \mathcal{D} \text{ with } \lambda \geq 0. \quad (\text{P})$$

Problem (P) is equivalent to finding $\mu \in \Delta(A)$ such that $\Pi(\mu) \leq 0$ in \mathcal{D}^* , as well as finding $\mu \in M$ with $\mu \geq 0$ such that $\mu(A) \geq 1$ and $\Pi(\mu) \leq 0$ in \mathcal{D}^* .

Clearly, problem (P) characterizes correlated equilibrium, i.e., there exists a correlated equilibrium if and only if (P) has a solution. This follows because for every player i , we may set all other λ_j ’s equal to always obeying and obtain (*).

We will transform (P) into a canonical system of linear equations for which we shall seek a non-negative solution. Let $T : M \times \mathcal{D}^* \rightarrow \mathbb{R} \times \mathcal{D}^*$ be defined pointwise by

$$T(\mu, \nu) = (-\mu(A), \Pi(\mu) + \nu)$$

for every $\mu \in M$ and $\nu \in \mathcal{D}^*$.

Clearly, (P) is equivalent to the following version:

$$\text{Find } (\mu, \nu) \in M \times \mathcal{D}^* \text{ with } (\mu, \nu) \geq 0 \text{ and } T(\mu, \nu) = (-1, \mathbf{0}), \quad (\text{P}^*)$$

where $-1 \in \mathbb{R}$ and $\mathbf{0} \in \mathcal{D}^*$. Here, ν is a slack variable.

– *Step 2.* Let $F = \{T^*(0, \lambda) : \lambda \in \mathcal{D}\}$ and $J = (B \times \mathcal{D})_+ - F$. By [Lemma A.1](#), there exists (μ, ν) to solve (P*) if and only if given any $\kappa \in \mathbb{R}$ and $\lambda \in \mathcal{D}$, $T^*(\kappa, \lambda) \in \bar{J}$ implies that $\kappa \leq 0$, where for any set S , \bar{S} denotes its closure. Hence, there exists correlated equilibrium if and only if for every (κ, λ) , there is a sequence $\{(\eta_m, \lambda_m)\} \subset J$ with $\lim \eta_m = \Pi^*(\lambda) - \kappa \mathbf{1}$ and $\lim \lambda_m = \lambda$ (where $\mathbf{1}$ is the constant function identically equal to unity) only if $\kappa \leq 0$.

Any $(\eta, \lambda) \in J$ can be written as $(\eta^+ - \Pi^*(\lambda), \lambda^+ - \lambda)$, where $\eta^+ \in B_+$, $\lambda^+ \in \mathcal{D}_+$, and $\lambda \in \mathcal{D}$. Therefore, equilibrium exists if and only if for every $\kappa \in \mathbb{R}$ and $\lambda \in \mathcal{D}$, if there is a sequence $\{(\eta_m^+, \lambda_m, \lambda_m^+)\} \subset B_+ \times \mathcal{D} \times \mathcal{D}_+$ such that $\lim \Pi^*(\lambda) - \kappa \mathbf{1} - \eta_m^+ + \Pi^*(\lambda_m) = 0$ and $\lim \lambda_m^+ - \lambda_m - \lambda = 0$ then $\kappa \leq 0$. Plugging the second limit into the first (and using continuity and linearity of Π^* below) yields

$$\begin{aligned} \lim \Pi^*(\lambda) - \kappa \mathbf{1} - \eta_m^+ + \Pi^*(\lambda_m) &= \\ \lim \Pi^*(\lambda_m^+ - \lambda_m) - \kappa \mathbf{1} - \eta_m^+ + \Pi^*(\lambda_m) &= \\ \lim \Pi^*(\lambda_m^+) - \kappa \mathbf{1} - \eta_m^+ &= 0. \end{aligned}$$

Clearly, this last limit obtains if and only if $\liminf \Pi^*(\lambda_m^+)(\mu) \geq \kappa$ for every $\mu \in \Delta(A)$, since $\eta_m^+ \geq 0$ for every m . Therefore, equilibrium exists if and only if:

$$\text{Given } \kappa \in \mathbb{R} \text{ and } \{\lambda_m^+\} \subset \mathcal{D}_+, \liminf \Pi^*(\lambda_m^+)(\mu) \geq \kappa \text{ for all } \mu \in \Delta(A) \text{ implies } \kappa \leq 0. \quad (\text{D})$$

– *Step 3.* Finally, we show that sequential consistency implies (D) and is implied by (P). Fix any sequence $\{\lambda_m^+\} \subset \mathcal{D}_+$. For every player i , let α_{im} be the deviation plan defined pointwise by $\alpha_{im}(a_i) = p_{im} \lambda_{im}^+(a_i) + (1 - p_{im}) \theta_i(a_i)$, where $p_{im} = \min\{\|\lambda_{im}^+\|, 1/\|\lambda_{im}^+\|\} \in (0, 1]$. By hypothesis, the sequence of deviation profiles $\{\alpha_m = (\alpha_{1m}, \dots, \alpha_{nm})\}$ has a subsequence $\{\alpha_{m_k}\}$ with $\hat{\mu} \in \bigcap_k \{\mu : \Pi(\mu)(\alpha_{m_k}) \leq 0\}$ for some $\hat{\mu} \in \Delta(A)$, hence $\liminf \Pi(\hat{\mu})(\alpha_m) \leq 0$. By linearity of Π^* and the fact that θ_i is obedient, $\liminf \Pi^*(\lambda_m^+)(\hat{\mu}) \leq 0$, so (D) holds. Conversely, existence of correlated equilibrium clearly implies sequential consistency. \square

Theorem 3. The proof follows similar steps to that of [Theorem 2](#). Let $U : M(X) \rightarrow M(Y)^*$ be the map defined by $U(\sigma) = \int_X u(x, \cdot) \sigma(dx)$, where, as usual, $M(X)^* = B(X)$. The adjoint map $U^* : M(Y) \rightarrow M(X)^*$ is given by $U^*(\tau) = \int_Y u(\cdot, y) \tau(dy)$.

– *Step 1.* Finding Nash equilibrium is equivalent to the following primal problem:

exists if and only if there exists $(\sigma, \tau, v) \in M(X) \times M(Y) \times \mathbb{R}$ such that $(\sigma, \tau, v) \geq 0$ and

$$\text{Find } (\sigma, \tau, v) \in M(X) \times M(Y) \times \mathbb{R} \text{ with } (\sigma, \tau, v) \geq 0 \text{ and } \begin{cases} U(\sigma) \geq v\mathbf{1}_Y \geq U^*(\tau), \\ \sigma(\mathbf{1}_X) = \tau(\mathbf{1}_Y) = 1. \end{cases} \quad (\text{P})$$

We will transform (P) into a canonical system of linear equations for which we shall seek a non-negative solution. Let $T : M(X) \times B(Y) \times M(Y) \times B(X) \times \mathbb{R} \times \mathbb{R} \rightarrow B(Y) \times B(X) \times \mathbb{R} \times \mathbb{R}$ be defined pointwise by

$$T(\sigma, g, \tau, f, c_+, c_-) = (U(\sigma) + g - (c_+ - c_-)\mathbf{1}_Y, -U^*(\tau) + f + (c_+ - c_-)\mathbf{1}_X, -\sigma(\mathbf{1}_X), -\tau(\mathbf{1}_Y)).$$

Clearly, (P) is equivalent to the following standard problem

$$\text{Find } (\sigma, g, \tau, f, c_+, c_-) \geq 0 \text{ such that } T(\sigma, g, \tau, f, c_+, c_-) = (\mathbf{0}_Y, \mathbf{0}_X, -1, -1). \quad (\text{P}^*)$$

– *Step 2.* In order to appeal to [Lemma A.1](#), define $F = \{T^*(\beta, \alpha, \gamma, \delta) : \gamma + \delta = 0\}$ and $J = (B(X) \times M(Y) \times B(Y) \times M(X) \times \mathbb{R} \times \mathbb{R})_+ - F$. By [Lemma A.1](#), there is a solution to (P*) if and only if given $(\beta, \alpha, \gamma, \delta) \in M(Y) \times M(X) \times \mathbb{R} \times \mathbb{R}$, $T^*(\beta, \alpha, \gamma, \delta) \in \bar{J}$ implies that $\gamma + \delta \leq 0$. Hence, Nash equilibrium exists if and only if given $(\beta, \alpha, \gamma, \delta)$, there is a sequence $\{(f_m^+, \beta_m, \gamma_m, \beta_m^+, g_m^+, \alpha_m, \delta_m, \alpha_m^+, \gamma_m^+, \delta_m^+)\}$ (where $^+$ means ≥ 0) such that

$$\begin{aligned} U^*(\beta) - \gamma\mathbf{1}_X &= \lim f_m^+ - U^*(\beta_m) + \gamma\mathbf{1}_X \\ \beta &= \lim \beta_m^+ - \beta_m \\ -U(\alpha) - \delta\mathbf{1}_Y &= \lim g_m^+ + U(\alpha_m) + \delta_m\mathbf{1}_Y \\ \alpha &= \lim \alpha_m^+ - \alpha_m \\ \alpha(\mathbf{1}_X) - \beta(\mathbf{1}_Y) &= \lim \gamma_m^+ + \beta_m(\mathbf{1}_Y) - \alpha_m(\mathbf{1}_X) \\ \beta(\mathbf{1}_Y) - \alpha(\mathbf{1}_X) &= \lim \delta_m^+ + \alpha_m(\mathbf{1}_X) - \beta_m(\mathbf{1}_Y) \end{aligned}$$

and $\gamma_m + \delta_m = 0$ for all m only if $\gamma + \delta \leq 0$. By continuity and linearity of U and U^* , simple calculations show that this system of equations is equivalent to (i) $\lim \alpha_m^+(\mathbf{1}_X) - \beta_m^+(\mathbf{1}_Y) = 0$, (ii) $\lim U^*(\beta_m^+) - f_m^+ - \gamma_m\mathbf{1}_X = \gamma\mathbf{1}_X$, and (iii) $\lim -U(\alpha_m^+) - g_m^+ - \delta_m\mathbf{1}_Y = \delta\mathbf{1}_Y$.

– *Step 3.* Adding (ii) and (iii) together, it follows that for every pair $(\sigma, \tau) \in \Delta(X) \times \Delta(Y)$, $\lim U^*(\beta_m^+)(\sigma) - \sigma(f_m^+) - U(\alpha_m^+)(\tau) - \tau(g_m^+) = \gamma + \delta$, so $\gamma + \delta \leq \liminf U^*(\beta_m^+)(\sigma) - U(\alpha_m^+)(\tau)$. The rest of the proof follows Step 3 of [Theorem 2](#). If every sequence of mixed-strategy profiles $\{(\sigma_m, \tau_m)\} \subset \Delta(X) \times \Delta(Y)$ has a subsequence $\{(\sigma_{m_k}, \tau_{m_k})\}$ with the property that $\bigcap_k \{(\sigma, \tau) \in \Delta(X) \times \Delta(Y) : U^*(\tau_{m_k})(\sigma) - U(\sigma_{m_k})(\tau) \leq 0\} \neq \emptyset$ then clearly $\liminf U^*(\beta_m^+)(\sigma) - U(\alpha_m^+)(\tau) \leq 0$, so Nash equilibrium exists. Conversely, if there exists Nash equilibrium then clearly sequential product consistency holds. \square

Theorem 4. The proof follows the same line as the previous ones. \square

Theorem 5. The multi-step proof below follows that of [Theorem 2](#) with minor adjustments. We first introduce some notation. Let $\mathcal{C} = \{\mu : T \rightarrow M(A) \text{ measurable and } \|\mu\| < \infty\}$ endowed with the norm $\|\mu\| = \sup_t |\mu|(A|t)$ be the vector space (indeed, Banach) where communication mechanisms live. With reference to the notation of [Definition 6](#), let us view $B(\Theta_i)$ as $M(\Theta_i)^*$ (when endowed with the weakest topology that makes this work). Write $\mathcal{B}_i = M(\Theta_i)$, $\mathcal{B}_i^* = M(\Theta_i)^*$, $\mathcal{B} = \prod_i M(\Theta_i)$ and $\mathcal{B}^* = \prod_i M(\Theta_i)^*$. Now, Φ_i can be viewed as an operator $\mathcal{C} \rightarrow \mathcal{B}^*$, and $\Phi : \mathcal{C} \rightarrow \mathcal{B}^*$.

– *Step 1.* With reference to the above notation, consider the following primal problem:

Find $\mu \in \mathcal{C}$ so that $\mu(t) \in \Delta(A)$ for all t and $\Phi(\mu)(\beta) \leq 0$ for all $\beta \in \mathcal{B}$ with $\beta \geq 0$. (P)

Clearly, problem (P) characterizes communication equilibrium, i.e., there exists a communication equilibrium if and only if (P) has a solution.

We will transform (P) into a canonical system of linear equations for which we shall seek a non-negative solution. Let $G : \mathcal{C} \times \mathcal{B}^* \rightarrow B(T) \times \mathcal{B}^*$ be defined pointwise by

$$G(\mu, \nu) = (-\mu(A|\cdot), \Phi(\mu) + \nu)$$

for every $\mu \in \mathcal{C}$ and $\nu \in \mathcal{B}^*$. Clearly, (P) is equivalent to the following version:

$$\text{Find } (\mu, \nu) \in \mathcal{C} \times \mathcal{B}^* \text{ with } (\mu, \nu) \geq 0 \text{ and } G(\mu, \nu) = (-\mathbf{1}, \mathbf{0}), \quad (\text{P}^*)$$

where $\mathbf{1} \in B(T)$ and $\mathbf{0} \in \mathcal{B}^*$. Here, ν is a slack variable.

– *Step 2.* Let $F = \{G^*(0, f) : f \in \mathcal{B}\}$ and $J = (\mathcal{C}^* \times \mathcal{B})_+ - F$. By [Lemma A.1](#), there exists (μ, ν) to solve (P*) if and only if given any $q \in M(T)$ and $\lambda \in \mathcal{B}$, $G^*(q, \lambda) \in \bar{J}$ implies that $q(T) \leq 0$. Hence, there exists communication equilibrium if and only if for every (q, λ) , there is a sequence $\{(\eta_m^+, \lambda_m^+, \lambda_m)\} \subset \mathcal{C}_+^* \times \mathcal{B}_+ \times \mathcal{B}$ with $\lim \eta_m^+ - \Phi^*(\lambda_m) = \Phi^*(\lambda) - q\mathbf{1}$ and $\lim \lambda_m^+ - \lambda_m = \lambda$ (where $\mathbf{1} \in B(A)$ is the function always equal to 1) only if $q(T) \leq 0$. Plugging the second limit into the first (using continuity and linearity of Φ^* below) yields

$$\begin{aligned} \lim \Phi^*(\lambda_m) - \eta_m^+ + \Phi^*(\lambda) - q\mathbf{1} &= \\ \lim \Phi^*(\lambda_m) - \eta_m^+ + \Phi^*(\lambda_m^+ - \lambda_m) - q\mathbf{1} &= \\ \lim \Phi^*(\lambda_m^+) - q\mathbf{1} - \eta_m^+ &= 0. \end{aligned}$$

Therefore, equilibrium exists if and only if:

$$\liminf \Phi^*(\lambda_m^+)(\mu) \geq \int_T \mu(A|t)q(dt) \text{ for all } \mu \in \mathcal{C} \text{ with } \mu \geq 0 \text{ implies } q(T) \leq 0. \quad (\text{D})$$

– *Step 3.* It is easy to see that (D) holds if and only if given $q \in M(T)$ and any sequence $\{\alpha_m = (\alpha_{1m}, \dots, \alpha_{nm})\} \subset \prod_i \Delta(\Theta_i)$, if $\Phi^*(\alpha_m)(\mu) \geq \int_T \mu(A|t)q(dt)$ for every $\mu \in \mathcal{C}_+$ and all except at most finitely many m then $q(T) \leq 0$. Finally, this is clearly equivalent to $\Phi^*(\alpha_m)(\mu) \geq q(S)$ for every measurable $S \subset T$ and every measurable map $\mu : S \rightarrow \Delta(A)$ and all except at most finitely many m implying that $q(T) \leq 0$. \square

References

- ATHEY, S. (2001): “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” *Econometrica*, 69, 861–889. 2
- BAYE, M. R., G. TIAN, AND J. ZHOU (1993): “Characterizations of the Existence of Equilibria in Games with Discontinuous and Non-Quasiconcave Payoffs,” *Review of Economic Studies*, 60, 935–948. 10
- CLARK, S. A. (2006): “Necessary and Sufficient Conditions for Infinite-Dimensional Linear Inequalities,” *Positivity*, 10, 475–489. 2, 15, 17, 18
- DASGUPTA, P. AND E. MASKIN (1986): “The Existence of Equilibrium in Discontinuous Economic Games I: Theory,” *Review of Economic Studies*, 53, 1–27. 2, 10
- FOLLAND, G. (1999): *Real Analysis: Modern Techniques and their Applications*, John Wiley & Sons, second ed. 9, 11
- FORGES, F. (1993): “Five Legitimate Definitions of Correlated Equilibrium in Games with Incomplete Information,” *Theory and Decision*, 35, 277–310. 14
- HART, S. AND D. SCHMEIDLER (1989): “Existence of Correlated Equilibria,” *Mathematics of Operations Research*, 14, 18–25. 1, 2, 4, 6
- MYERSON, R. (1997): “Dual Reduction and Elementary Games,” *Games and Economic Behavior*, 21, 183–202. 6
- NIKAIDO, H. AND K. ISODA (1955): “Note on Noncooperative Convex Games,” *Pacific Journal of Mathematics*, 5, 807–815. 10
- RENY, P. J. (1999): “On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games,” *Econometrica*, 67, 1029–1056. 2, 10
- SCHAEFER, H. H. (1970): *Topological Vector Spaces*, New York: Springer-Verlag. 15
- SIMON, L. K. (1987): “Games with Discontinuous Payoffs,” *Review of Economic Studies*, 54, 569–97. 2, 10
- SIMON, L. K. AND W. ZAME (1990): “Discontinuous Games and Endogenous Sharing Rules,” *Econometrica*, 58, 861–872. 2
- SION, M. AND P. WOLFE (1957): “On a Game with No Value,” *Annals of Mathematical Studies*, 39, 299–306. 6
- YANNELIS, N. AND A. RUSTICHINI (1992): “On the Existence of Correlated Equilibria,” in *Equilibrium Theory with Infinitely Many Commodities*, Springer-Verlag, 249–265. 1