Saddle Functions and Robust Sets of Equilibria

Vladyslav Nora* and Hiroshi Uno†

This version: January 28, 2013

Abstract

This paper introduces games with a saddle function. A saddle function is a real-valued function on the set of action profiles such that there is a single *minimizing* player, for whom minimizing the function implies choosing her best response, and the other players are *maximizing* players, for whom maximizing the function implies choosing their best responses. We provide a sufficient condition for the robustness of sets of equilibria to incomplete information in the sense of Kajii and Morris (1997, *Econometrica*), Morris and Ui (2005, *J. of Econ. Theory*) for games with a saddle function. Our result generalizes sufficient conditions for zero-sum and best-response potential games.

Journal of Economic Literature Classification Numbers: C72, D82

Key words: Incomplete information; robustness; potential; saddle function; team-maximin equilibrium.

1 Introduction

We often model a strategic situation as a complete information game. However, an equilibrium outcome of a complete information game may differ from the outcomes of an arbitrarily "close" incomplete information game [2, 10]. This leads Kajii and Morris [5] to introduce an equilibrium robust to incomplete information and then Morris and Ui [8] to generalize it to robust sets of equilibria. A set of equilibria of a complete information game is robust if every incomplete information game close to the original game has a Bayesian Nash equilibrium that induces an observed behavior close to some equilibrium in the set. If a robust set is a singleton, it is robust to incomplete information in the sense of Kajii and Morris [5].

*Center for Operations Research and Econometrics, Université catholique de Louvain, Voie du Roman Pays 34, L1.03.01, B-1348 Louvain-la-Neuve, Belgium. E-mail: vladyslav.nora@uclouvain.be

†Graduate School of Economics, Osaka Prefecture University, 1-1 Gakuen-cho, Nakaku, Sakai, Osaka 599-8531, Japan. E-mail: uno@eco.osakafu-u.ac.jp

We introduce games with a saddle function and provide a sufficient condition for the robustness of sets of equilibria in such games. A saddle function of a game is a real-valued function on the set of action profiles such that there is a single *minimizing* player for whom minimizing the function implies choosing her best response, and the other players are *maximizing* players, for whom maximizing the function implies choosing their best-responses. The value of a saddle function is a *maximin* attained when the saddle function is maximized over the strategy profiles of maximizing players and is minimized over the strategies of the minimizing player. We show that the set of correlated equilibria that induce an expectation of a saddle function greater or equal to the value of the saddle function is robust.

A game with a saddle function is "strategically equivalent" to a zero-sum game where a set of players with identical payoffs plays against a single adversary. We call such games, studied by von Stengel and Koller [14], *team vs. adversary games*. Consider the three-person game below.

	L	R
U	$1, \ 1, -1$	$-1, -1, \ 1$
D	$-2, -2, \ 2$	$0, \ 0, \ 0$

T

	L	R
U	$2, \ 2, -2$	$-2, -2, \ 2$
D	$-1, -1, \ 1$	$3, \ 3, -3$

B

Players 1 and 2 form the team and choose a row and a column respectively; the adversary chooses a matrix. The payoff function of player 1 is the saddle function with the value 1. Observe that a Nash equilibrium (U, L, T) is the unique correlated equilibrium that induces an expectation of the saddle function greater or equal to 1. Our result implies that (U, L, T) is robust.

We contribute to the literature on the robustness to incomplete information. Kajii and Morris [5] introduce the notion of robust equilibrium and give sufficient conditions in terms of a unique correlated equilibrium and **p**-dominant equilibrium. The former condition implies that a unique Nash equilibrium of a two-person zero-sum game is robust. Next, Ui [12] proves the robustness of a unique maximizer of a potential function defined by Monderer and Shapley [7], and Tercieux [11] obtains a condition for games with **p**-best response sets. Morris and Ui [8] provide a sufficient condition in terms of generalized potential functions which generalizes the above approaches, but is difficult to apply. Therefore they also develop tractable conditions in terms of special classes of generalized potential functions: best response, monotone and local potentials, the latter two being further generalized by Oyama and Tercieux [9] using iterated monotone potentials. However, all these conditions are not applicable to games with a saddle function. We provide a new sufficient condition for the robustness for games with a saddle function. Our condition generalizes those in terms of zero-sum and best-response potential games. Although the condition is simpler to apply than the one in terms of generalized potential functions, it can guarantee the robustness of smaller sets of equilibria.

We also contribute to the literature on team vs. adversary games, which is a special class of games with a saddle function. Von Stengel and Koller [14] characterize the maximum payoff of a team in a finite game. We generalize their result to infinite games with multilinear payoffs.

2 Robust sets of equilibria

A complete information game consists of a finite set of players N and, for each $i \in N$, a finite set of actions A_i and a payoff function $g_i : A \to \mathbb{R}$, where $A = \prod_{i \in N} A_i$. Since we fix N and A, we simply denote a complete information game by $\mathbf{g} := (g_i)_{i \in N}$.

For a profile $(X_i)_{i \in N}$ of sets, $X := \prod_{i \in N} X_i$ and $x := (x_i)_{i \in N} \in X$. Given a nonempty subset T of N, we write $X_T := \prod_{i \in T} X_i$ and $x_T := (x_i)_{i \in T} \in X_T$. For $i \in N$ we write $X_{-i} := \prod_{k \neq i} X_k$ and $x_{-i} := (x_k)_{k \neq i} \in X_{-i}$. A set of probability distributions on a set X is denoted by $\Delta(X)$.

An action distribution $\mu \in \Delta(A)$ is a *correlated equilibrium* of $\mathbf{g}$ if, for each $i \in N$ and each $a_i \in A_i$ with $\mu_i(a_i) > 0$,

$$\sum_{a_{-i} \in A_{-i}} \mu(a_{-i}|a_i) g_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \mu(a_{-i}|a_i) g_i(a_i', a_{-i})$$

for each $a_i' \in A_i$, where $\mu_i(a_i)$ is the marginal probability of a_i and $\mu(a_{-i}|a_i)$ is the conditional probability of a_{-i} given a_i. An action distribution $\mu \in \Delta(A)$ is a *Nash equilibrium* of $\mathbf{g}$ if it is a correlated equilibrium of $\mathbf{g}$ and $\mu(a) = \prod_{i \in N} \mu_i(a_i)$ for each $a \in A$.

Consider an incomplete information game with the set N of players and the set A of action profiles (same as in $\mathbf{g}$). Let Θ_i be a countable set of types of $i \in N$, and let P be a prior probability distribution on Θ with $P_i(\theta_i) := \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_i, \theta_{-i}) > 0$ for all $i \in N$ and $\theta_i \in \Theta_i$. A payoff function of player $i \in N$ is a bounded function $u_i : A \times \Theta \to \mathbb{R}$. Since we fix N, A and Θ, we denote an incomplete information game by $(\mathbf{u}, P)$, where $\mathbf{u} := (u_i)_{i \in N}$.

For each player $i \in N$, a strategy is a function $\sigma_i : \Theta_i \to \Delta(A_i)$ and Σ_i is a set of strategies. We write $\sigma_i(a_i|\theta_i)$ for the probability that an action $a_i \in A_i$ is chosen given a type $\theta_i \in \Theta_i$ under a strategy $\sigma_i \in \Sigma_i$. Given a subset T of N and a strategy profile $\sigma_T \in \Sigma_T$, a probability of an action profile $a_T \in A_T$ given a type profile $\theta_T \in \Theta_T$ is $\sigma_T(a_T|\theta_T) := \prod_{i \in T} \sigma_i(a_i|\theta_i)$.

A strategy profile $\sigma \in \Sigma$ is a *Bayesian Nash equilibrium* of $(\mathbf{u}, P)$ if, for each $i \in N$,

$$\sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) \sum_{a \in A} \sigma(a|\theta) u_i(a, \theta) \geq \sum_{\theta_{-i} \in \Theta_{-i}} P(\theta_{-i}|\theta_i) \sum_{a \in A} \sigma_i'(a_i|\theta_i) \sigma_{-i}(a_{-i}|\theta_{-i}) u_i(a, \theta)$$

for each $\theta_i \in \Theta_i$ and each $\sigma_i' \in \Sigma_i$, where $P(\theta_{-i}|\theta_i) = P(\theta_i, \theta_{-i}) / P_i(\theta_i)$.

Given a complete information game $\mathbf{g}$ and an incomplete information game $(\mathbf{u}, P)$, for each $i \in N$, consider the subset $\bar{\Theta}_i$ of Θ_i such that, if $\theta_i \in \bar{\Theta}_i$ is realized, payoffs of i are given by g_i, and he knows his payoffs:

$$\bar{\Theta}_i = \{\theta_i \in \Theta_i | u_i(a, (\theta_i, \theta_{-i})) = g_i(a) \text{ for all } a \in A, \theta_{-i} \in \Theta_{-i} \text{ with } P(\theta_i, \theta_{-i}) > 0\}.$$

An incomplete information game $(\mathbf{u}, P)$ is an *ε-elaboration* of $\mathbf{g}$ if $P(\bar{\Theta}) = 1 - \varepsilon$, where $\varepsilon \in [0, 1]$. Kajii and Morris [5] prove the following lemma.

Lemma 1. *Let $\{(\mathbf{u}^m, P^m)\}$ be a sequence of ε-elaborations with $\varepsilon^m \to 0$ and let σ^m be a Bayesian Nash equilibrium of $(\mathbf{u}^m, P^m)$. Then, there exist a subsequence $\{\sigma^l\}$ of $\{\sigma^m\}$ and a correlated equilibrium $\mu \in \Delta(A)$ of $\mathbf{g}$ such that $\sum_{\theta \in \Theta} P^l(\theta)\sigma^l(a|\theta) \to \mu(a)$ for each $a \in A$.*

A type $\theta_i \in \Theta_i \backslash \bar{\Theta}_i$ of player $i \in N$ is *committed* if it has a strictly dominant action $a_i^{\theta_i} \in A_i$:

$$u_i\left((a_i^{\theta_i}, a_{-i}), (\theta_i, \theta_{-i})\right) > u_i\left((a_i, a_{-i}), (\theta_i, \theta_{-i})\right)$$

for each $a_i \in A_i \backslash \{a_i^{\theta_i}\}$, each $a_{-i} \in A_{-i}$ and each $\theta_{-i} \in \Theta_{-i}$ with $P(\theta_i, \theta_{-i}) > 0$. An ε-elaboration of $\mathbf{g}$ is *canonical* if, for each $i \in N$, each $\theta_i \in \Theta_i \backslash \bar{\Theta}_i$ is a committed type.

Morris and Ui [8] study the sets of correlated equilibria robust to canonical elaborations.

Definition 1. A set $\mathcal{E} \subseteq \Delta(A)$ of correlated equilibria is *robust to canonical elaborations* in $\mathbf{g}$ if, for each $\delta > 0$, there exists $\bar{\varepsilon} > 0$ such that, for all $\varepsilon \leq \bar{\varepsilon}$, each canonical ε-elaboration of $\mathbf{g}$ has a Bayesian Nash equilibrium $\sigma \in \Sigma$ such that $\max_{a \in A} |\mu(a) - \sum_{\theta \in \Theta} P(\theta)\sigma(a|\theta)| \leq \delta$ for some $\mu \in \mathcal{E}$.

If $\mathcal{E}$ is a singleton, then the correlated equilibrium in the set is a Nash equilibrium robust to canonical elaborations in the sense of Kajii and Morris [4].

We consider a set-valued concept because a robust equilibrium may not exist, as shown by Kajii and Morris [5].[1] On the contrary, by Lemma 1 a set of all correlated equilibria is robust to canonical elaborations.

Originally, Kajii and Morris [5] proposed the stronger notion of robustness to *all* elaborations. To get a corresponding set-valued notion one allows for all, not only canonical ε-elaborations in Definition 1. It is clear that a set of correlated equilibria robust to all elaborations is also robust to canonical elaborations.[2]

3 Saddle functions

In a complete information game we fix a player j in N, and let T be the set of the other players. A saddle function of a complete information game is a real-valued function f on the set of action

[1] In fact, Haimanko and Kajii [3] show that even a two-person zero-sum game may have no robust equilibrium.
[2] Whether the converse holds is an open question.

profiles such that, for each member i of T, every best response against his belief over the other players' actions in a game where i's payoff is given by f is a best response against the same belief in the original game; and j's best response against her belief over the other players' actions in a game where j's payoff is given by $-f$ is a best response against the same belief in the original game as well.

Definition 2. Let $j \in N$ and $T := N \backslash \{j\}$. A function $f : A \to \mathbb{R}$ is a *j-saddle function* of $\mathbf{g}$ if, for each $i \in T$,

$$\arg\max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \lambda_i(a_{-i}) f(a) \subseteq \arg\max_{a_i \in A_i} \sum_{a_{-i} \in A_{-i}} \lambda_i(a_{-i}) g_i(a)$$

for all $\lambda_i \in \Delta(A_{-i})$; and for j,

$$\arg\min_{a_j \in A_j} \sum_{a_T \in A_T} \lambda_j(a_T) f(a) \subseteq \arg\max_{a_j \in A_j} \sum_{a_T \in A_T} \lambda_j(a_T) g_j(a)$$

for all $\lambda_j \in \Delta(A_T)$. A real number v^* is the *value of f* if

$$v^* := \max_{\mu_T \in \prod_{i \in T} \Delta(A_i)} \min_{\mu_j \in \Delta(A_j)} \sum_{a \in A} \left(\prod_{i \in N} \mu_i(a_i) \right) f(a).$$

Since we fix a game $\mathbf{g}$, a player $j \in N$ and a j-saddle function f of $\mathbf{g}$, we simply say that f is a *saddle function* and v^* is the *value.*

Team vs. adversary games studied by von Stengel and Koller [14] form a special class of games with a saddle function. A game $\mathbf{g}$ is *a team vs. adversary game* if there exists $j \in N$ such that $g_i = f$ for each $i \in T := N \setminus \{j\}$ and $g_j = -f$, where $f : A \to \mathbb{R}$. We call j an adversary and T a team. Then the payoff function f of team members is a saddle function of the game.

Games with a saddle function generalize best-response potential games of Morris and Ui [8]: add a dummy player j with a singleton action set to a best-response potential game. Then, a best-response potential function is a saddle function and the maximum of potential is the value.[3]

Let $\mathcal{E}$ be a set of correlated equilibria of $\mathbf{g}$ inducing an expectation of a saddle function greater or equal to v^*:

$$\mathcal{E} := \{\mu \in \Delta(A) | \mu \text{ is a correlated equilibrium of } \mathbf{g} \text{ and } \textstyle\sum_{a \in A} \mu(a) f(a) \geq v^*\}. \tag{1}$$

Now we are ready to state our main result.

Theorem 1. *If* $\mathbf{g}$ *has a saddle function, then the set* $\mathcal{E}$ *is robust to canonical elaborations in* $\mathbf{g}$.

[3] Note also that games with a j-saddle function form a special class of multi-potential games introduced by Monderer [6], or games with a partition $\{\{j\}, T\}$-potential introduced by Uno [13].

If $\mathcal{E}$ is a singleton, then the unique Nash equilibrium in the set is robust to canonical elaborations in the game.

In the remainder of this section we prove Theorem 1. Suppose that $\mathbf{g}$ has a saddle function f. Let ($\mathbf{u}$,P) be a canonical ε-elaboration of $\mathbf{g}$. Define a function $V : \Sigma \to \mathbb{R}$ by

$$V(\sigma) := \sum_{\theta\in\Theta}\sum_{a\in A} P(\theta)\sigma(a|\theta)f(a).$$

For each $i \in N$, let $\bar{\Sigma}_i := \{\sigma_i \in \Sigma_i | \sigma_i(a_i^{\theta_i}|\theta_i) = 1 \text{ for all } \theta_i \in \Theta_i \setminus \bar{\Theta}_i\}$ be the set of strategies such that all committed types choose their dominant actions. A Bayesian Nash equilibrium $(\sigma_T^*, \sigma_j^*) \in \Sigma$ of a canonical ε-elaboration ($\mathbf{u}$,P) is a *saddle point* of ($\mathbf{u}$,P) if

$$\sigma_T^* \in \arg\max_{\sigma_T\in\bar{\Sigma}_T}\min_{\sigma_j\in\bar{\Sigma}_j} V(\sigma_T, \sigma_j).$$

Note that a saddle point of a 0-elaboration of $\mathbf{g}$ where Θ_i is a singleton for each $i \in N$ is a Nash equilibrium of $\mathbf{g}$ in a set $\mathcal{E}$.

To prove the existence of a saddle point in every canonical ε-elaboration of $\mathbf{g}$, we generalize von Stengel and Koller's [14] result on team vs. adversary games.[4]

Proposition 1. *Let $j \in N$ and $T := N\backslash\{j\}$. Consider a game $(N, (S_i, g_i)_{i\in N})$, where S_i for each $i \in N$ is a compact convex subset of a locally convex Hausdorff space and $g_i = f$ for each $i \in T$ and $g_j = -f$, where $f : S \to \mathbb{R}$ is a continuous multilinear function. Then, $\arg\max_{s_T\in S_T}\min_{s_j\in S_j} f(s_T, s_j)$ is nonempty and, for each $s_T^* \in \arg\max_{s_T\in S_T}\min_{s_j\in S_j} f(s_T, s_j)$, there exists $s_j^* \in S_j$ such that (s_T^*, s_j^*) is a Nash equilibrium.*

In Proposition 1, if we let $S_i = \bar{\Sigma}_i$ for each $i \in N$ and $f = V$, then we can show that (s_T^*, s_j^*) is a saddle point of a canonical ε-elaboration ($\mathbf{u}$,P).

Lemma 2. *If $\mathbf{g}$ has a saddle function, then every canonical ε-elaboration of $\mathbf{g}$ has a saddle point.*

Proof. Suppose that $\mathbf{g}$ has a saddle function f. Fix a canonical ε-elaboration $(\mathbf{u}, P)$ of $\mathbf{g}$. Consider a game $\mathbf{V} := (N, (\bar{\Sigma}_i, V_i)_{i\in N})$, where $V_i = V$ for $i \in T$ and $V_j = -V$. For each $i \in N$, a strategy set $\bar{\Sigma}_i$ is a convex subset of a locally convex Hausdorff space and is compact in a product topology by Tyhonoff's theorem. The function V is continuous and linear in a strategy of each player. By Proposition 1 there exists $\sigma_T^* \in \arg\max_{\sigma_T\in\bar{\Sigma}_T}\min_{\sigma_j\in\bar{\Sigma}_j} V(\sigma_T, \sigma_j)$ and $\sigma_j^* \in \bar{\Sigma}_j$ such that (σ_T^*, σ_j^*) is a Nash equilibrium of $\mathbf{V}$. We shall show that (σ_T^*, σ_j^*) is also a saddle point of $(\mathbf{u}, P)$. It suffices to show that (σ_T^*, σ_j^*) is a Bayesian Nash equilibrium of $(\mathbf{u}, P)$, that is for each $i \in N$ and for each $\theta_i \in \Theta_i$ we have

[4] The proof is referred to Appendix.

$$\sigma_i^*(\theta_i) \in \arg\max_{\sigma_i(\theta_i)\in\Delta(A_i)} \sum_{\theta_{-i}\in\Theta_{-i}} P(\theta_{-i}|\theta_i) \sum_{a\in A} \sigma_i(a_i|\theta_i)\sigma_{-i}^*(a_{-i}|\theta_{-i})u_i(a,\theta). \tag{2}$$

First, (2) holds for each $i \in N$ and each committed type $\theta_i \in \Theta_i \setminus \bar{\Theta}_i$ by definition of $\bar{\Sigma}_i$.

Next, we show that (2) holds for each $i \in N$ and for each $\theta_i \in \bar{\Theta}_i$. Fix $i \in T$. We have $V(\sigma_i^*, \sigma_{-i}^*) \geq V(\sigma_i, \sigma_{-i}^*)$ for all $\sigma_i \in \bar{\Sigma}_i$. We can rewrite it as

$$\sum_{\theta_i\in\bar{\Theta}_i} P_i(\theta_i) \sum_{\theta_{-i}\in\Theta_{-i}} P(\theta_{-i}|\theta_i) \sum_{a\in A} \sigma_i^*(a_i|\theta_i)\sigma_{-i}^*(a_{-i}|\theta_{-i})f(a) \geq$$
$$\sum_{\theta_i\in\bar{\Theta}_i} P_i(\theta_i) \sum_{\theta_{-i}\in\Theta_{-i}} P(\theta_{-i}|\theta_i) \sum_{a\in A} \sigma_i(a_i|\theta_i)\sigma_{-i}^*(a_{-i}|\theta_{-i})f(a)$$

for all $\sigma_i \in \bar{\Sigma}_i$. Now (2) follows for each $i \in T$ since f is a saddle function of $\mathbf{g}$. A symmetric argument for j concludes the proof. $\square$

Finally, we are ready to prove Theorem 1.

Proof of Theorem 1. Let $(\mu_i^*)_{i\in T} \in \prod_{i\in T} \Delta(A_i)$ be such that

$$(\mu_i^*)_{i\in T} \in \arg\max_{\mu_T\in\prod_{i\in T}\Delta(A_i)} \min_{\mu_j\in\Delta(A_j)} \sum_{a\in A} \left(\prod_{i\in N} \mu_i(a_i)\right) f(a). \tag{3}$$

Fix a canonical ε-elaboration $(\mathbf{u}, P)$ of $\mathbf{g}$. By Lemma 2 there exists a saddle point $\sigma \in \Sigma$ of $(\mathbf{u}, P)$. First, we shall find a lower bound on $V(\sigma)$. For each $i \in T$ and each $\theta_i \in \bar{\Theta}_i$, let $\sigma_T^* \in \bar{\Sigma}_T$ be such that $\sigma_i^*(a_i|\theta_i) = \mu_i^*(a_i)$ for each $a_i \in A_i$. Since $\sigma_T \in \arg\max_{\sigma_T'\in\bar{\Sigma}_T} \min_{\sigma_j'\in\bar{\Sigma}_j} V(\sigma_T', \sigma_j')$, we have $V(\sigma_T, \sigma_j) \geq \min_{\sigma_j'\in\bar{\Sigma}_j} V(\sigma_T^*, \sigma_j')$.

Let $\varepsilon_T \geq 0$ be a marginal probability such that there exists a player in T of a committed type, i.e., $\varepsilon_T := P\left((\Theta_T \setminus \bar{\Theta}_T) \times \Theta_j\right) \leq \varepsilon$. By definition of V, for each $\sigma_j' \in \bar{\Sigma}_j$, we have

$$\begin{aligned} V(\sigma_T^*, \sigma_j') &= \sum_{\theta\in(\Theta_T\setminus\bar{\Theta}_T)\times\Theta_j} P(\theta) \sum_{a\in A} \sigma_T^*(a_T|\theta_T)\sigma_j'(a_j|\theta_j)f(a) \\ &\quad + \sum_{\theta\in\bar{\Theta}_T\times\Theta_j} P(\theta) \sum_{a\in A} \left(\prod_{i\in T} \mu_i^*(a_i)\right) \sigma_j'(a_j|\theta_j)f(a) \\ &\geq \varepsilon_T f_{min} + \sum_{\theta\in\bar{\Theta}_T\times\Theta_j} P(\theta) \sum_{a\in A} \left(\prod_{i\in T} \mu_i^*(a_i)\right) \sigma_j'(a_j|\theta_j)f(a), \end{aligned}$$

where $f_{min} := \min_{a\in A} f(a)$. Observe that $\sum_{a\in A} \left(\prod_{i\in T} \mu_i^*(a_i)\right) \mu_j(a_j)f(a) \geq v^*$ for all $\mu_j \in \Delta(A_j)$

by (3). It follows that

$$\varepsilon_T f_{min} + \sum_{\theta \in \bar{\Theta}_T \times \Theta_j} P(\theta) \sum_{a \in A} \left(\prod_{i \in T} \mu_i^*(a_i) \right) \sigma_j(a_j|\theta_j) f(a) \geq \varepsilon_T f_{min} + (1 - \varepsilon_T) v^*.$$

Thus we obtain a lower bound as a function of $\varepsilon_T \leq \varepsilon$:

$$V(\sigma) \geq v^* + \varepsilon_T (f_{min} - v^*). \tag{4}$$

To complete the proof we show that for each $\delta > 0$, there exists $\bar{\varepsilon} > 0$ such that, for all $\varepsilon \leq \bar{\varepsilon}$, each canonical ε-elaboration of $\mathbf{g}$ has a Bayesian Nash equilibrium $\sigma \in \Sigma$ such that $\max_{a \in A} |\mu(a) - \sum_{\theta \in \Theta} P(\theta) \sigma(a|\theta)| \leq \delta$ for some $\mu \in \mathcal{E}$.

To get a contradiction suppose that for some $\delta > 0$ there exists a sequence $\{(\mathbf{u}^m, P^m)\}$ of canonical ε-elaboration of $\mathbf{g}$ with $\varepsilon^m \to 0$ such that $\max_{a \in A} | \sum_{\theta \in \Theta} P^m(\theta) \sigma^m(a|\theta) - \mu(a)| > \delta$ for all $\mu \in \mathcal{E}$, where $\sigma^m \in \Sigma$ is a saddle point of $(\mathbf{u}^m, P^m)$. By Lemma 1 there exist a subsequence $\{\sigma^l\}$ of $\{\sigma^m\}$ and a correlated equilibrium $\nu \in \Delta(A)$ of $\mathbf{g}$ such that $\sum_{\theta \in \Theta} P^l(\theta) \sigma^l(a|\theta) \to \nu(a)$ for all $a \in A$. Since $\sigma^l \in \Sigma$ is a saddle point of $(\mathbf{u}^l, P^l)$, by (4) we have $v^* \leq \sum_{a \in A} \nu(a) f(a)$. Therefore ν belongs to $\mathcal{E}$. The contradiction completes the proof. $\square$

4 Discussion

4.1 Generalized potential maximizers

Morris and Ui [8] introduce generalized potential functions and provide a sufficient condition for a set of equilibria to be robust. Generalized potentials are real-valued functions on a domain $\mathcal{A} = \prod_{i \in N} \mathcal{A}_i$, where $\mathcal{A}_i \subseteq 2^{A_i} \backslash \{\emptyset\}$ is a collection of nonempty subsets of A_i such that $\bigcup_{X_i \in \mathcal{A}_i} X_i = A_i$ for each $i \in N$. Given $i \in N$ and a distribution $\Lambda_i \in \Delta(\mathcal{A}_i)$ let

$$\Delta_{\Lambda_i}(A_i) := \{\lambda \in \Delta(A_i) | \lambda(a_i) = \sum_{X_i \in \mathcal{A}_i} \Lambda_i(X_i) \lambda^{X_i}(a_i) \text{ for each } a_i \in A_i \text{ and}$$
$$\lambda^{X_i} \in \Delta(A_i) \text{ with } \sum_{a_i \in X_i} \lambda^{X_i}(a_i) = 1 \text{ for each } X_i \in \mathcal{A}_i\}$$

be the set of distributions over A_i induced by Λ_i.

A function $F : \mathcal{A} \to \mathbb{R}$ is a *generalized potential function* of $\mathbf{g}$ if, for each $i \in N$, all

$Q_i \in \Delta(\mathcal{A}_{-i})$ and all $q_i \in \Delta_{Q_i}(A_{-i})$,

$$X_i \cap \arg\max_{a'_i \in A_i} \sum_{a_{-i} \in A_{-i}} q_i(a_{-i}) g_i(a'_i, a_{-i}) \neq \emptyset$$
$$\text{for every } X_i \in \arg\max_{X'_i \in \mathcal{A}_i} \sum_{X_{-i} \in \mathcal{A}_{-i}} Q_i(X_{-i}) F(X'_i, X_{-i})$$

such that X_i is maximal in the argmax set ordered by the set inclusion relation. An action subset X^* is a *generalized potential maximizer (GP maximizer)* if $F(X^*) > F(X)$ for each $X \in \mathcal{A} \backslash \{X^*\}$.

Let $\mathcal{E}_{X^*}$ be a set of correlated equilibria of $\mathbf{g}$ that assign probability 1 to X^*:

$$\mathcal{E}_{X^*} := \{\mu \in \Delta(A) | \mu \text{ is a correlated equilibrium of } \mathbf{g} \text{ such that } \textstyle\sum_{a \in X^*} \mu(a) = 1\}. \tag{5}$$

Morris and Ui [8] show the following result.

Theorem 2. *If* $\mathbf{g}$ *has a generalized potential function with a generalized potential maximizer* X^**, then* $\mathcal{E}_{X^*}$ *is robust to canonical elaborations in* $\mathbf{g}$*.*

Theorem 1 can identify a smaller set of robust equilibria than Theorem 2 as we show in the following example.

T:

	L	R
U	$1, \; 1, -1$	$-1, -1, \; 1$
D	$-1, -1, \; 1$	$0, \; 0, \; 0$

B:

	L	R
U	$0, \; 0, \; 0$	$-1, -1, \; 1$
D	$-1, -1, \; 1$	$1, \; 1, -1$

Example 1. Consider the three-person game above where player 1, 2 and 3 are choosing a row, a column and a matrix respectively. Since the payoff function of player 1 (or 2) is a saddle function with the value 0, the set $\mathcal{E}$, defined by (1), does not include a completely mixed Nash equilibrium. We shall show that, if there exists a GP maximizer X^*, then $X_i^* = A_i$ for $i = 1, 2, 3$. Therefore $\mathcal{E}_{X^*}$ of (5) is the set of all correlated equilibria. Thus we have $\mathcal{E}_{X^*} \subset \mathcal{E}$.

First we show that in every correlated equilibrium players 1 and 2 choose each action with positive probabilities, unless the correlated equilibrium is a pure strategy Nash equilibrium (D, R, T) or (U, L, B). By symmetry of the game it's enough to consider a correlated equilibrium $\mu \in \Delta(A_1 \times A_2 \times A_3)$ such that player 1 chooses U with probability 1. We have

$$\begin{aligned} -\mu(U, R, T) - \mu(U, R, B) &\geq \mu(U, R, T), \\ -\mu(U, L, T) + \mu(U, R, T) &\geq \mu(U, R, B), \end{aligned}$$

which implies that $\mu(U, L, T) = \mu(U, R, T) = \mu(U, R, B) = 0$.

Let X^* be a GP maximizer of the game. By Theorem 2, the set $\mathcal{E}_{X^*}$ is robust to canonical elaborations. Suppose that there exists $i \in \{1, 2\}$ such that $X_i^* \neq A_i$. Then, as we have established before, $\mathcal{E}_{X^*}$ contains a unique distribution assigning probability 1 to either (D, R, T) or (U, L, B). Now we show that (D, R, T) and (U, L, B) are not robust to canonical elaborations which is a contradiction. Therefore $X_i^* = A_i$ for $i \in \{1, 2\}$. By symmetry it suffices to show only that (U, L, B) is not robust. Consider the following ε-elaboration of the game. Sets of types of player 1 and 2 are $\Theta_1 = \{0, 1, 2, \dots\}$ and $\Theta_2 = \{1, 2, \dots\}$ respectively; player 3 has a single type θ_3. We represent the probability distribution on type profiles in the table below, where rows and columns are players' 1 and 2 types respectively:

	1	2	3	$\cdots$	m	$m+1$	$\cdots$
0	ε	0	0	$\cdots$			
1	$\varepsilon(1-\varepsilon)$	$\varepsilon(1-\varepsilon)^2$	0	$\cdots$			
2	0	$\varepsilon(1-\varepsilon)^3$	$\varepsilon(1-\varepsilon)^4$				
$\vdots$	$\vdots$	$\vdots$		$\ddots$			
m					$\varepsilon(1-\varepsilon)^{2m-1}$	$\varepsilon(1-\varepsilon)^{2m}$	$\cdots$
$m+1$					0	$\varepsilon(1-\varepsilon)^{2m+1}$	$\cdots$
$\vdots$					$\vdots$	$\vdots$	$\ddots$

Formally, the probability distribution on type profiles is given by

$$P(\theta_1, \theta_2, \theta_3) = \begin{cases} \varepsilon(1-\varepsilon)^{2m-1} & \text{if } \theta_1 = \theta_2 = m, \\ \varepsilon(1-\varepsilon)^{2m} & \text{if } \theta_1 = m,\ \theta_2 = m+1, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\theta_1 = 0$ be a committed type of player 1 with the strictly dominant action D. The other types of player 1 as well as all types of player 2 and 3 have payoffs as in the complete information game. Suppose that there exists a Bayesian Nash equilibrium such that players choose (U, L, B) with a probability close to 1. Observe that $\frac{\varepsilon(1-\varepsilon)^{2m-1}}{\varepsilon(1-\varepsilon)^{2m}+\varepsilon(1-\varepsilon)^{2m-1}} > \frac{1}{2}$ for all $\varepsilon > 0$. Hence, by induction, if player 3 chooses B with a probability greater or equal to $1/2$, for all types of player 1 and 2 dominant actions are respectively D and R, a contradiction. Thus, there exists a sequence of canonical ε-elaborations with all equilibrium distributions bounded away from (U, L, B).

Finally, suppose that $X_3^* \neq A_3$ and $X_i^* = A_i$ for $i = 1, 2$. But then X_3^* does not contain a best response to either profile (U, L) or (D, R), a contradiction. Therefore $X_3^* = A_3$.

4.2 Team-maximin equilibria

Consider a team vs. adversary game with a team payoff function f. Let $S_i := \Delta(A_i)$ for $i \in N$. Von Stengel and Koller [14] introduce team-maximin equilibria. A *team-maximin* strategy profile is $s_T^* \in S_T$ such that $s_T^* \in \arg\max_{s_T \in S_T} \min_{s_j \in S_j} \sum_{a \in A} \left(\prod s_i(a_i)\right) f(a)$. A *team-maxmin equilibrium* is a Nash equilibrium $s \in S$ such that $s_T \in S_T$ is a team-maximin strategy profile. The following result of von Stengel and Koller [14] is a special case of Proposition 1.

Theorem 3. *In a finite team vs. adversary game any team-maximin strategy profile is a part of a team-maximin equilibrium.*

A team-maximin equilibrium may not be robust to canonical elaborations. Indeed, it is easy to see that in Example 1 (U, L, B) is a team-maximin equilibrium. However, we show that (U, L, B) is not robust to canonical elaborations.[5]

In a case of team vs. adversary games Theorem 1 has an intuitive interpretation. The set of correlated equilibria that induce a payoff to the team greater or equal to the best Nash equilibrium payoff of the team, is robust to canonical elaborations.

Appendix

Proof of Proposition 1. Following von Stengel and Koller [14] we call T *a team* and j *an adversary.* A profile $s_T^* \in S_T$ is a *team-maximin profile* if $s_T^* \in \arg\max_{s_T \in S_T} \min_{s_j \in S_j} f(s_T, s_j)$. A profile (s_T^*, s_j^*) is a *team-maximin equilibrium* if it is a Nash equilibrium such that s_T^* is a team-maximin profile. We write $v^* := \max_{s_T \in S_T} \min_{s_j \in S_j} f(s_T, s_j)$.

The existence of a team-maximin profile is guaranteed by compactness and continuity assumptions.

We shall prove the existence of a team-maximin equilibrium. Let s_T^* be a team-maximin profile. We want to find a best response of the adversary such that no team member $i \in T$ has an incentive to deviate from his team-maximin strategy s_i^* when the other team members choose their team-maximin strategies $s_{T\setminus\{i\}}^*$. For each $i \in T$ and $s_i \in S_i$, let $H_i(s_i) := \{s_j \in S_j | f(s_i, s_{T\setminus\{i\}}^*, s_j) \leq v^*\}$ be a set of strategies of the adversary such that given $s_j \in H_i(s_i)$, payoff of i from s_i when others choose $s_{T\setminus\{i\}}^*$ is lower or equal to v^*. For each $i \in T$ and $s_i \in S_i$, it is clear that $H_i(s_i)$ is nonempty, convex and compact (since f is continuous). We construct a correspondence whose fixed point gives a desired best response of the adversary.

[5]In Example 1 a team-maximin equilibrium is not unique. A reader might wonder whether a unique team-maximin equilibrium is robust. In fact slightly modifying the payoffs of a game in Example 1 we can show that a unique team-maximin equilibrium may not be robust.

Define maps $\psi_i : S_j \rightrightarrows S_i$ for each $i \in T$, $\psi_j : S_T \rightrightarrows S_j$ and $\psi : S \rightrightarrows S$ by

$$\begin{aligned} \psi_i(s_j) &:= \arg\max_{s_i \in S_i} f(s^*_{T\setminus\{i\}}, s_j) \text{ for each } i \in T, \\ \psi_j(s_T) &:= \bigcap_{i \in T} H_i(s_i), \\ \psi(s) &:= \prod_{i \in N} \psi_i(s). \end{aligned}$$

Suppose there exists a fixed point $(\tilde{s}_T, s^*_j)$ of ψ. We shall show that (s^*_T, s^*_j) is a Nash equilibrium. First, we show that $s^*_i \in \arg\max_{s_i \in S_i} f(s_i, s^*_{T\setminus\{i\}}, s^*_j)$ for each $i \in T$. Observe that $f(\tilde{s}_i, s^*_{T\setminus\{i\}}, s^*_j) = f(s^*_T, s^*_j) = v^*$ for each $i \in T$. Indeed, if there exists $i \in T$ such that $\tilde{s}_i \neq s^*_i$ and $f(\tilde{s}_i, s^*_{T\setminus\{i\}}, s^*_j) < v^*$, then $i \in T$ prefers s^*_i to $\tilde{s}_i$ given s^*_j and $s^*_{T\setminus\{i\}}$, a contradiction. And, if there exists $i \in T$ such that $\tilde{s}_i \neq s^*_i$ and $f(\tilde{s}_i, s^*_{T\setminus\{i\}}, s^*_j) > v^*$ then $s^*_j \notin H_i(\tilde{s}_i)$, a contradiction. Thus $s^*_i \in \arg\max_{s_i \in S_i} f(s_i, s^*_{T\setminus\{i\}}, s^*_j)$ for each $i \in T$. Next, since s^*_T is a team-maximin profile and $f(s^*_T, s^*_j) = v^*$ we have $s^*_j \in \arg\min_{s_j \in S_j} f(s^*_T, s_j)$. Hence (s^*_T, s^*_j) is a team-maximin equilibrium.

It remains to show that ψ has a fixed point. A set S is a nonempty compact convex subset of a locally convex Hausdorff space. It is clear that ψ is upper hemicontinuous as a product of upper hemicontinuous best reply correspondences and a correspondence ψ_j which is an intersection of upper hemicontinuous correspondences $H_i : S_i \rightrightarrows S_j$. Moreover, ψ has convex values as a product of convex valued correspondences. Thus, if ψ has nonempty values, then by Kakutani-Fan-Glicksberg Theorem [1, p. 583] it has a fixed point.

So, to prove that ψ has a fixed point, it suffices to show that it has nonempty values. For each $i \in T$ the set $\psi_i(s_j)$ is nonempty for all $s_j \in S_j$ by Weierstrass' Theorem. We assert that $\psi_j(s_T) \neq \emptyset$ for all $s_T \in S_T$. For the sake of contradiction suppose that there exists $\bar{s}_T \in S_T$ such that $\bigcap_{i \in T} H_i(\bar{s}_i) = \emptyset$. For each $s_j \in S_j$, define a vector $\mathbf{f}(s_j) := \Big(f(\bar{s}_i, s^*_{T\setminus\{i\}}, s_j)\Big)_{i \in T} \in \mathbb{R}^{n-1}$. Let $K := \{\mathbf{f}(\bar{s}_T, s_j) \in \mathbb{R}^{n-1} | s_j \in S_j\}$ and $D := \{\mathbf{y} \in \mathbb{R}^{n-1} | y_i \leq v^*\}$. Note that K is a convex and compact subset of $\mathbb{R}^{n-1}$ since S_j is compact and convex and f is linear in $s_j \in S_j$. A set D is a convex and closed subset of $\mathbb{R}^{n-1}$. Since there does not exist $s_j \in S_j$ such that $s_j \in H_i(\bar{s}_i)$ for each $i \in T$, we have $K \cap D = \emptyset$. By Separating Hyperplane Theorem there exists a linear functional $\pi := (\pi_i)_{i \in T}$ on $\mathbb{R}^{n-1}$ strongly separating K and D, which clearly can be taken to satisfy $\sum \pi_i = 1$ and $\pi_i \geq 0$. Define $\hat{v} := \min_{\mathbf{y} \in K} \pi \mathbf{y} > v^*$.

For $\delta > 0$, define $s^\delta_T := [(1 - \delta\pi_i)s^*_i + \delta\pi_i\bar{s}_i]_{i \in T}$. We shall show that if $\delta > 0$ is sufficiently small, then $v^* < f(s^\delta_T, s_j)$ for all $s_j \in S_j$. Let $\hat{S}_T := \{s_T \in S_T |$there exist $i, k \in T$ such that $s_i \neq$

s_i^* and $s_k \neq s_k^*\}$. Repeatedly using the multilinearity of f we can write

$$\begin{aligned}
f(s_T^\delta, s_j) = &\left(\prod_{i\in T}(1-\delta\pi_i)\right) f(s_T^*, s_j) + \delta^2 \sum_{s_T\in \hat{S}_T} \frac{\lambda(s_T,\delta)}{\delta^2} f(s_T, s_j) + \dots \\
&+ \sum_{i\in T}\left(\delta\pi_i \prod_{k\neq i}(1-\delta\pi_k)\right) f(\bar{s}_i, s_{T\setminus\{i\}}^*, s_j) \\
= &\left(\prod_{i\in T}(1-\delta\pi_i)\right) f(s_T^*, s_j) + \delta^2 \sum_{s_T\in \hat{S}_T} \frac{q_{s_T}(\delta)}{\delta^2} f(s_T, s_j) + \dots \\
&+ \left(\sum_{i\in T}\delta\pi_i \prod_{k\neq i}(1-\delta\pi_k)\right) \sum_{i\in T}\left(\frac{\pi_i \prod_{k\neq i}(1-\delta\pi_k)}{\sum_{i\in T}\pi_i\prod_{k\neq i}(1-\delta\pi_k)}\right) f(\bar{s}_i, s_{T\setminus\{i\}}^*, s_j),
\end{aligned}$$

where q_{s_T} is a bounded function of $\delta \in [0,1]$ for each $s_T \in S_T$. Observe that $\left(\frac{\pi_i \prod_{k\neq i}(1-\delta\pi_k)}{\sum_{i\in T}\pi_i\prod_{k\neq i}(1-\delta\pi_k)}\right)_{i\in T} \to \pi$ as $\delta \to 0$. So, there exists $\delta' > 0$ such that, for all $\delta < \delta'$,

$$\sum_{i\in T}\left(\frac{\pi_i \prod_{k\neq i}(1-\delta\pi_k)}{\sum_{i\in T}\pi_i\prod_{k\neq i}(1-\delta\pi_k)}\right) f(\bar{s}_i, s_{T\setminus\{i\}}^*, s_j) \geq \frac{\hat{v}-v^*}{2}.$$

Thus for all $\delta < \delta'$ we obtain the inequality

$$f(s_T^\delta, s_j) \geq v^* \prod_{i\in T}(1-\delta\pi_i) + \check{v}\delta^2 \sum_{s_T\in\hat{S}_T}\frac{q_{s_T}(\delta)}{\delta^2} + \left(\frac{\hat{v}-v^*}{2}\right)\delta\sum_{i\in T}\left(\pi_i\prod_{k\neq i}(1-\delta\pi_k)\right)$$

for all $s_j \in S_j$, where $\check{v} := \min_{s\in S} f(s)$ and so $\check{v} \leq v^* < \frac{\hat{v}-v^*}{2}$. Since $q_{s_T}(\delta)$ is a bounded function for each $s_T \in S_T$, it follows that $\frac{\delta\sum_{i\in T}\left(\pi_i\prod_{k\neq i}(1-\delta\pi_k)\right)}{\delta^2\sum_{s_T\in\hat{S}_T}\frac{q_{s_T}(\delta)}{\delta^2}} \to \infty$ as $\delta \to 0$. Therefore, there exists $\bar{\delta} > 0$ such that, for all $\delta < \min\{\bar{\delta}, \delta'\}$, we have $f(s_T^\delta, s_j) > v^*$ for all $s_j \in S_j$, which contradicts to s_T^* being a team-maximin profile. Thus $\bigcap_{i\in T} H_i(s_i) \neq \emptyset$ for each $s_T \in S_T$. □

Acknowledgements

We thank Julio Dávila, Atsushi Kajii, Mamoru Kaneko, Daisuke Oyama, Takashi Ui, Vincent Vannetelbosch and the seminar participants at CORE, Seventh EBIM Doctoral Workshop on Economic Theory at Bielefeld University, Summer Workshop on Economic Theory at Otaru University of Commerce, University of Tsukuba and Shiga University for their helpful comments and suggestions. Nora gratefully acknowledges Financial support from the Belgian FNRS through a PhD fellowship from a M.I.S.—Mobilité 'Ulysse' F.R.S.-FNRS.

References

[1] Aliprantis, C., Border, K., 2007. Infinite Dimensional Analysis: A Hitchhiker's Guide, third edition Edition. Springer.

[2] Carlsson, H., van Damme, E., 1993. Global games and equilibrium selection. Econometrica 61 (5), pp. 989–1018.

[3] Haimanko, O., Kajii, A., 2012. Approximate robustness of equilibrium to incomplete information. Working Papers 1209, Ben-Gurion University of the Negev, Department of Economics.

[4] Kajii, A., Morris, S., 1997. Refinements and higher order beliefs: A unified survey. Northwestern Math Center Discussion Paper No. 1197, 1997 (available at http://ssrn.com/abstract=283477).

[5] Kajii, A., Morris, S., 1997. The robustness of equilibria to incomplete information. Econometrica 65 (6), pp. 1283–1309.

[6] Monderer, D., 2007. Multipotential games. In: Twentieth International Joint Conference on Artificial Intelligence (IJCAI-07).

[7] Monderer, D., Shapley, L. S., 1996. Potential games. Games and Economic Behavior 14, 124–143.

[8] Morris, S., Ui, T., 2005. Generalized potentials and robust sets of equilibria. Journal of Economic Theory 124 (1), 45 – 78.

[9] Oyama, D., Tercieux, O., 2009. Iterated potential and robustness of equilibria. Journal of Economic Theory 144 (4), 1726–1769.

[10] Rubinstein, A., 1989. The electronic mail game: Strategic behavior under "almost common knowledge". The American Economic Review 79 (3), pp. 385–391.

[11] Tercieux, O., 2006. p-best response set and the robustness of equilibria to incomplete information. Games and Economic Behavior 56 (2), 371–384.

[12] Ui, T., 2001. Robust equilibria of potential games. Econometrica 69 (5), 1373–80.

[13] Uno, H., 2007. Nested potential games. Economics Bulletin 3 (19), 1–8.

[14] von Stengel, B., Koller, D., 1997. Team-maxmin equilibria. Games and Economic Behavior 21, 309 – 321.