STRONG ROBUSTNESS TO INCOMPLETE INFORMATION AND THE UNIQUENESS OF A CORRELATED EQUILIBRIUM

Ezra Einy, Ori Haimanko, and David Lagziel

Discussion Paper No. 19-01

May 2019

Monaster Center for
Economic Research
Ben-Gurion University of the Negev
P.O. Box 653
Beer Sheva, Israel

Fax: 972-8-6472941
Tel: 972-8-6472286

Strong Robustness to Incomplete Information and the Uniqueness of a Correlated Equilibrium*

Ezra Einy[†], Ori Haimanko[‡], David Lagziel[§]

April 19, 2019

Abstract:

We define and characterize the notion of *strong robustness* to incomplete information, whereby a Nash equilibrium in a game u is strongly robust if, given that each player knows that his payoffs are those in u with high probability, *all* Bayesian-Nash equilibria in the corresponding incomplete-information game are close – in terms of action distribution – to that equilibrium of u. We prove, under some continuity requirements on payoffs, that a Nash equilibrium is strongly robust if and only if it is the unique correlated equilibrium. We then review and extend the conditions that guarantee the existence of a unique correlated equilibrium in games with a continuum of actions. The existence of a strongly robust Nash equilibrium is thereby established for several domains of games, including those that arise in economic environments as diverse as Tullock contests, Cournot and Bertrand competitions, network games, patent races, voting problems, and location games.

Journal of Economic Literature classification numbers: C62, C72, D82.

Keywords: strong robustness to incomplete information; Nash equilibrium; correlated equilibrium.

*The authors wish to thank Eddie Dekel, Ehud Lehrer, Yehuda John Levy, Daisuke Oyama, David Schmeidler, and Aner Sela for their valuable comments. Their special grattitude goes to Atsushi Kajii, whose encouragement to put the first basic ideas into writing led to the creation of this work.

[†]Department of Economics, Ben-Gurion University of the Negev, Israel. e-mail: einy@bgu.ac.il

[‡]Department of Economics, Ben-Gurion University of the Negev, Israel. e-mail: orih@bgu.ac.il.

[§]Department of Economics, Ben-Gurion University of the Negev, Israel. e-mail: davidlag@bgu.ac.il.

1 Introduction

Nash equilibrium is an immensely popular and long-established solution concept in economics. By comparison, its generalized version, the correlated equilibrium of Aumann (1974), is a far less frequent choice in economic modelling. Clearly, the two solution concepts have their merits and drawbacks: Nash equilibrium is believed to have a high predictive power and does not require a mediation or a correlation device, while the correlated equilibrium is superior in terms of computational complexity and arises naturally in a range of simple learning processes.[1]

In this paper we bring to light the synergy produced when the two concepts happen to coincide, in relation to the robustness of equilibrium outcome to the presence of incomplete information. To motivate our notion of robustness, we first take a step back to discuss an important issue in the field of economics – *the need to predict*.

1.1 The need to predict, and strong robustness to incomplete information

A major difficulty in the profession of economics is the perpetual requirement to provide accurate predictions in a realm affected by uncertainty and randomness. Similarly to weather forecasters, economists are repeatedly evaluated by their ability to produce solid assessments. When the latter deviate from the eventual outcomes to a significant degree, doubts may be cast not only over the treatment of the available empirical data and its validity, but also over the suitability of the underlying theoretical model.

The work of Kajii and Morris (1997) (henceforth KM) partially deals with this concern by introducing the notion of equilibrium *robustness to incomplete information*. Roughly speaking, an equilibrium in a complete-information scenario is robust if, when some uncertainty is introduced, there exists an equilibrium that is "sufficiently close" to the original one. Thus, existence of a robust equilibrium in a game-theoretical model reinforces that model, because allowing limited uncertainty may lead only to small changes in the predicted behavior. The practical implications are clear – an economist who advises policy makers would be rather confident in her recommendations if they are based on a robust equilibrium, even when there is some unmodelled uncertainty regarding the agents' true characteristics.

The current work continues this quest by defining and characterizing a stricter robustness notion – *strong robustness to incomplete information*. We say that a Nash equilibrium in a complete-information game u is strongly robust if, under uncertainty about the individual payoffs but with each player knowing that his payoffs are those in u with high probability, *all*

[1] See, e.g., Papadimitriou and Roughgarden (2008) for a discussion of computational complexity, and Foster and Vohra (1997) and Hart and Mas-Colell (2000) for results on convergence to correlated equilibria.

Bayesian-Nash equilibria in the incomplete-information game are sufficiently close (in terms of the induced action distribution) to the equilibrium in u. Thus, when some uncertainty is introduced, the effect on all possible equilibrium outcomes should be minor.

It is obvious that the imposition of the closeness requirement on all equilibria in nearby games makes strong robustness a hugely demanding notion. At the same time, this requirement is very beneficial for an analyst who either designs or models a strategic interaction between rational agents. Whenever a strongly robust Nash equilibrium exists, the analyst can be sure that the behavior of the agents remains sufficiently close to the designed or predicted outcome, as long as the agents' behavior in nearby incomplete-information games is consistent with *any* equilibrium.

The main result of this paper provides a characterization of a strongly robust Nash equilibrium in a framework with a continuum of individual actions. Under a set requirements (inspired by Dasgupta and Maskin (1986)) on the payoff functions that limit the extent of possible discontinuity,[2] our main theorems show that a Nash equilibrium is strongly robust if and only if it is the unique correlated equilibrium. In other words, it is precisely the coincidence of being a Nash and a correlated equilibrium that makes such an equilibrium strongly robust.[3]

In the last part of this paper we review and extend the conditions that guarantee the existence of a unique correlated equilibrium in games with a continuum of actions. This will, via our main result, imply the existence of a strongly robust Nash equilibrium in various economic environments, such as Tullock contests, certain types of Bertrand and Cournot competition, network games, patent races, the median-voter problem and pure-location Hotelling games.

1.2 The main contribution and relation to the literature

The first and main contribution of this paper is the formulation and characterization of strong robustness to incomplete information. Our notion of strong robustness preserves the spirit of informational robustness of KM,[4] but is far stricter since strong robustness requires the closeness of *all*, not just *some*, equilibria in incomplete information settings to the complete-information Nash equilibrium that is being approximated. Similarly to Proposition 3.2 of KM concerning

[2]The conditions in Dasgupta and Maskin (1986) guarantee the existence of a (mixed-action) Nash equilibrium in a game, and our set will also be sufficient for equilibrium existence. Although Dasgupta and Maskin were only concerned with equilibrium existence, variants of their conditions are useful in the proofs of our main results since, like them, we make extensive use of (weak) convergence of probability measures and of the corresponding integrals of payoff functions.

[3]The uniqueness of a correlated equilibrium has been known to imply a different type of robustness, w.r.t. payoff perturbations (see Viossat (2008)).

[4]The work of KM was preceded by the approaches of Fudenberg et al. (1988), Dekel and Fudenberg (1990), and Carlsson and van Damme (1993).

finite games, the uniqueness of a correlated equilibrium[5] implies its robustness also in our setting with a continuum of actions, albeit requiring conditions on payoffs and necessitating a non-trivial proof. Strong robustness of a Nash equilibrium is, moreover, *equivalent* to its being the unique correlated equilibrium in the game.

From a practical perspective, this equivalence result ensures that the behavior of rational economic agents is always sufficiently close to the designed outcome as long as a unique correlated equilibrium is guaranteed. From a theoretical point of view, the result highlights the role of correlated equilibria in determining the sensitivity of a Nash equilibrium to the presence of incomplete information. The main example in Section 3.1 of KM showcases what may go wrong in terms of robustness even if a game u has a unique, pure-action Nash equilibrium: when there is a correlated equilibrium that is distinct from the latter, there may be some incomplete-information games nearby in which the (unique) Bayesian-Nash equilibrium approximates the correlated, and not the unique Nash, equilibrium of u.

The motivation for allowing a continuum of actions to be available to each player in our framework[6] comes partially from the fact that the existing results on the uniqueness of a correlated equilibrium, necessary and sufficient for it being strongly robust, are mostly in the continuum setting. An earliest example is due to Milgrom and Roberts (1990), who showed, as an application of their characterization of undominated action sets in supermodular games, that a Bertrand oligopoly with differentiated products has a unique correlated equilibrium for certain families of demand functions. Also relying on supermodularity techniques, Amir (1996) proved the uniqueness of a correlated equilibrium for a Cournot duopoly with a log-concave strictly decreasing inverse demand function. Liu (1996) went beyond two firms, and showed the uniqueness for linear Cournot oligopolies. His result was generalized by Neyman (1997), who proved the existence of a unique correlated equilibrium in every potential game with a compact and convex set of actions and a strictly concave smooth potential function. The latter class of potential games partially includes network games, as shown by Bramoullé et al. (2014) and Ui (2016). Generalizing the work of Neyman (1997), Ui (2008) showed, under the condition of Rosen (1965) for Nash equilibrium uniqueness in smooth concave games, that the same equilibrium is also the unique correlated one.[7] Recently, Hart and Mas-Colell (2015) proved the uniqueness of a correlated equilibrium in social strictly concave games, without any payoff-smoothness requirements.

We also contribute to this line of work by showing that every Tullock rent-seeking game

[5]Since in our setting the existence of a (mixed-action) Nash equilibrium will be guaranteed, if a correlated equilibrium is unique then it must be a Nash equilibrium.

[6]We also concomitantly admit uncountable, measurable state-spaces in incomplete information approximations of a complete information game.

[7]Ui (2008) also generalized the original condition of Rosen.

(contest), and every equivalent patent race, have a unique correlated equilibrium.[8] Certain features of Tullock contests (discontinuity of payoffs when all efforts vanish, and the sum of payoffs not being strictly concave) make them unsuitable for the frameworks of both Ui (2008) and Hart and Mas-Colell (2015), and thus necessitate a separate approach. It will also be observed that a correlated equilibrium is unique in two-player constant-sum games whenever their Nash equilibrium is unique and consists of pure actions; this implies correlated equilibrium uniqueness in median-voter problems and pure-location Hotelling games, which are non-concave and discontinuous. We thereby expand the known part of the domain of games with a unique correlated equilibrium by adding to it important sets of non-smooth and non-concave games. In order to demonstrate the scope of our strong robustness notion, we will offer a formal survey of what is known on that domain, as its constituent games have a strongly robust NE in light of our main result.

1.3 Structure of the paper

The rest of the paper is organized as follows. In Section 2 we present the basic complete-information framework, and extend it to incomplete information. In Section 3 we define and explain the notion of strong robustness to incomplete information. In Section 4 we present our main result on the equivalence of the existence of a strongly robust Nash equilibrium and the uniqueness of a correlated one. In Section 5 we survey the games for which a correlated equilibrium is known to be unique and state uniqueness results of our own.

2 Preliminaries

Our basic framework is laid out in Section 2.1, where we formally define games with a continuum of pure actions. It is then extended in Section 2.2 to accommodate incomplete information.

2.1 Games with a continuum of pure actions

Fix a finite set of players $N = \{1, 2, ..., n\}$. The set A_i of (pure) actions of each player i is assumed to be a compact and full-dimensional[9] convex subset of a Euclidean space $\mathbb{R}^{m_i}$, and $A = \times_{i \in N} A_i \subset \mathbb{R}^{\Sigma_{i \in N} m_i}$ denotes the set of players' action profiles. A game is given by an n-tuple $u = (u_i)_{i \in N}$, where $u_i : A \to \mathbb{R}$ is the payoff function of player i.

[8] For the proof of equivalence between patent races and Tullock contests see Baye and Hoppe (2003), who follow the model of Loury (1979).

[9] The assumption of full dimension entails no loss of generality, since otherwise A_i can be replaced by an equivalent strategy set of lower, full, dimension.

To formally treat mixed actions, some general notations are in order. For a positive integer m and a compact set $B \subset \mathbb{R}^m$, denote by $M(B)$ the set of Borel probability measures on B. When needed, any $b \in B$ will be identified with a Dirac measure supported on $\{b\}$, and hence B may be viewed as a subset of $M(B)$. We shall endow $M(B)$ with the topology of weak convergence of measures, in which $M(B)$ is metrizable and compact.[10] In general, for any product set $C = \times_{i \in N} C_i$ and any $j \in N$, the notation c_j will refer to a generic element of the set C_j, and c_{-j} – to a generic element of the set $C_{-j} = \times_{i \neq j} C_i$.

A *mixed action* of player i will be an element of $M(A_i)$, and an element of $M(A)$ will be referred to as an *action distribution.* Similarly to the definition used by Hart and Schmeidler (1989), an action distribution $\mu \in M(A)$ is a *correlated equilibrium* (henceforth, CE) of a game u if, for any player i and any Borel-measurable function $\psi_i : A_i \to A_i$,

$$\int_A u_i(a)\, d\mu(a) \geq \int_A u_i(\psi_i(a_i), a_{-i})\, d\mu(a). \tag{1}$$

In fact, if μ is a CE, then Ineq. (1) holds for any Borel-measurable function[11] $\psi_i : A_i \to M(A_i)$, with $u_i(\psi_i(a_i), a_{-i})$ being defined as $\int_{A_i} u_i(b_i, a_{-i})\, d\psi_i(a_i)(b_i)$ in this case. See Appendix A.1 for the proof of this claim.

Given a mixed-action profile $\nu = (\nu_i)_{i \in N}$, with each $\nu_i \in M(A_i)$ being a mixed action of player i, let $\widehat{\nu} = \times_{i \in N} \nu_i \in M(A)$ be the product action distribution that is induced by ν when the individual action choices are independent. The expected payoff $u_i(\nu)$ of player i in the latter scenario is given by

$$u_i(\nu) = \int_A u_i(a)\, d\widehat{\nu}(a). \tag{2}$$

A mixed-action profile ν is a *Nash equilibrium* (henceforth, NE) of u if $\widehat{\nu}$ is a CE. This is equivalent to the requirement that $u_i(\nu) \geq u_i(a_i, \nu_{-i})$ for every player i and $a_i \in A_i$.

2.2 Incomplete information games

In an incomplete information game, the underlying uncertainty is described by a measurable space $(\Omega, \mathcal{F})$ of states of nature and a countably additive probability measure P on Ω, which is the common prior belief of the players about the actual state. The information of player i is given by a σ-subfield $\mathcal{F}_i$ of $\mathcal{F}$; the interpretation is that given any $E \in \mathcal{F}_i$, player i knows whether the realized state of nature belongs to E. The payoffs to player i are determined by a state-dependent payoff function $U_i : A \times \Omega \to \mathbb{R}$ that is measurable w.r.t. the product σ-algebra $(Borel\ Sets) \times \mathcal{F}$. The incomplete information game with the above attributes will be denoted by $\mathcal{U} = \left\{(\Omega, \mathcal{F}), \{\mathcal{F}_i\}_{i \in N}, \{U_i\}_{i \in N}\right\}$. We shall henceforth assume that payoff functions

[10] Recall that under this topology a sequence $\{\mu_k\}_{n=1}^{\infty} \subset M(B)$ converges to $\mu \in M(B)$ if and only if $\lim_{k \to \infty} \int_B f(a)\, d\mu_k(a) = \int_B f(a)\, d\mu(a)$ for any continuous $f : B \to \mathbb{R}$.

[11] That is, $\psi_i(a_i)(B)$ is a measurable function of a_i for every Borel subset B of A_i.

in all complete and incomplete information games are bounded in absolute value by the same exogenously fixed constant.

A (*behavioral*) *strategy* of player i is a $\mathcal{F}_i$-measurable function $\sigma_i : \Omega \to M(A_i)$, i.e., $\sigma_i(\omega)(B)$ is an $\mathcal{F}_i$-measurable function of ω for every Borel set $B \subset A_i$. A *strategy profile* is an n-tuple $\sigma = (\sigma_i)_{i \in N}$, where σ_i is a strategy of player i. Given such σ, the expected payoff to player i is

$$\overline{U}_i(\sigma) := \int_\Omega U_i(\sigma(\omega), \omega)\, dP(\omega), \tag{3}$$

where $U_i(\sigma(\omega), \omega)$ denotes the extension of $U_i(\cdot, \omega)$ into mixed-action profiles, which is done by the same procedure as in Eq. (2).[12]

A strategy profile σ is a *Bayesian-Nash equilibrium* (henceforth, BNE) of $\mathcal{U}$ if $\overline{U}_i(\sigma) \geq \overline{U}_i(\tau_i, \sigma_{-i})$ for every player i and for every strategy τ_i of i. Given a BNE σ, its induced action distribution $\mu(\sigma) \in M(A)$ is given by

$$\mu(\sigma)(B) = \int_\Omega \left[\int_A \chi_B(a)\, d\left(\widehat{\sigma(\omega)}\right)(a)\right] dP(\omega)$$

for every Borel subset B of A, where χ_B denotes the indicator function of the set B.

3 Strong robustness to incomplete information

To accurately define strong robustness, we first need to make precise the sense in which an incomplete-information game $\mathcal{U}$ can approximate a (complete-information) game u. We will consider an incomplete-information game $\mathcal{U}$ as being close to u if, with high probability, each player i knows that his payoff in $\mathcal{U}$ is given by u_i. Formally, for any $\delta \geq 0$, an incomplete-information game $\mathcal{U}$ is said to be a δ-*elaboration* of u if for every player i there exists an event $\Omega_i(\mathcal{U}, u) \in \mathcal{F}_i$ such that

$$\Omega_i(\mathcal{U}, u) \subset \{\omega \mid U_i(a, \omega) = u_i(a) \text{ for all } a \in A\},$$

and $P(\cap_{i \in N} \Omega_i(\mathcal{U}, u)) = 1 - \delta$. Note that the above notion of a close incomplete-information game is in line with that of KM for finite games, with the additional possibility of an uncountable state space.

We shall use these δ-elaborations to define *strong robustness* of NE, a notion that preserves the spirit of informational robustness of KM but is far more demanding.

Definition 1 *Given a complete-information game u, its* NE ν *is* strongly robust (to incomplete information) *if, for any sequence $\{\mathcal{U}^k\}_{k=1}^\infty$ of incomplete information games where each $\mathcal{U}^k$ is a*

[12]The integrand in Eq. (3) is bounded and Borel-measurable by, e.g., Proposition 7.29 in Bertsekas and Shreve (2004), and hence $\overline{U}_i$ is well-defined.

δ_k-elaboration of u that possesses some BNE σ^k *and* $\lim_{k\to\infty}\delta_k = 0$*, the sequence* $\{\mu(\sigma^k)\}_{k=1}^{\infty}$ *of action distributions induced by* $\{\sigma^k\}_{k=1}^{\infty}$ *weakly converges to the action distribution* $\widehat{\nu}$ *of* ν.

In other words, an NE of u is strongly robust if its induced product action-distribution is close to the action distributions of BNE in every incomplete-information elaboration that is sufficiently close to u.

Notice that the definition allowsus to choose *any* BNE in an elaboration $\mathcal{U}^k$, and so strong robustness requires *all* corresponding BNE sequences to approximate ν. This is the main difference between our strong robustness and KM-robustness for finite games, as the latter notion only requires elaborations near u to have *some* BNE that approximate ν. The corresponding definition for our class of games could have been termed just *robustness*, and would require $\widehat{\nu}$ to be approximable by $\{\mu(\sigma^k)\}_{k=1}^{\infty}$ for some selection of BNE $\{\sigma^k\}_{k=1}^{\infty}$ in any sequence $\{\mathcal{U}^k\}_{k=1}^{\infty}$. But since our focus is on the farthest possible extent to which robustness can constrain equilibrium outcomes, it is the strong robustness that we need.

There are two other distinctions between strong robustness and KM-robustness that we would like to point out. Unlike KM-robustness, Definition 1 employs a sequential statement, which obviates the need to specify a particular metric that governs weak convergence of action distributions.[13] Also, the KM-robustness is defined for action distributions in general, while strong robustness only applies to NE and presupposes the existence of an NE in u. This difference is in appearance only. In the KM set-up of finite games, the existence of a mixed-action NE is guaranteed and any robust action distribution is clearly an NE. Hence, KM-robustness in actuality applies only to NE. The reason we focus on a strongly robust *NE* is to exclude cases where an equilibrium does not exist and all action distributions are strongly robust by default.

It can be readily seen that the implications of there being a strongly robust NE in u are quite stark: such an NE is necessarily unique, and its action distribution must be the only CE in the game. For the sake of completeness, we state this observation in the following proposition.

Proposition 1 *If a complete-information game u possesses a strongly robust* NE ν*, then* $\widehat{\nu}$ *is the unique* CE *of u. In particular, if a strongly robust* NE *exists, then it is unique.*

(The proofs of Proposition 1 and all subsequent results are deferred to the Appendix.)

We conclude that a *necessary* condition for an NE ν to be strongly robust is the uniqueness of a CE in the game. In the next section we establish conditions under which the uniqueness of a CE is both necessary and *sufficient* for the existence of a strongly robust NE.

[13] No specific metric on $M(A)$ that induces the topology of weak convergence, including the Lévy–Prokhorov metric, seems to be sufficiently appealing to make an ε, δ-statement preferable to our (equivalent) statement in terms of sequence convergence.

4 Strong robustness of a unique CE

Our main result, which identifies a strongly robust NE with a unique CE, will be proved under a set of conditions requiring partial continuity of the payoff functions. We will present two versions of the result. Theorem 1 below assumes continuity of payoffs in the interior of the action-profile set A. Theorem 2, on the other hand, allows discontinuity of payoffs along "diagonal" curves in A, when the action sets of all players are one-dimensional.

The following three conditions will be used in the statement of Theorem 1:

(a) each payoff function $u_i(a)$ is continuous in a whenever a_i is an interior point of A_i;

(b) each payoff function $u_i(a_i, a_{-i})$ is lower semi-continuous in a_i for a fixed a_{-i}; and

(c) the sum $\sum_{i \in N} u_i$ is upper semi-continuous on A.[14]

These continuity conditions resemble those that were introduced in Dasgupta and Maskin (1986), and were shown therein to be sufficient for the existence of a mixed-action NE in a complete-information game. Specifically, our condition (**b**) is a strengthened, and (**c**) an exact, version of the conditions in Theorem 5 of Dasgupta and Maskin (1986). Condition (**a**) is not, however, directly comparable with their requirement of partial continuity. As may be expected, under our conditions, u also possesses a mixed-action NE. We state this claim separately, for later reference:

Remark 1 *If u satisfies* (**a**), (**b**), *and* (**c**), *then it possesses a mixed-action* NE (*the proof is given in* Appendix A.3).

Our main result, Theorem 1, extends Proposition 1 to a full characterization of strong robustness, by showing an equivalence between the existence of a strongly robust NE and its being a unique CE under conditions (**a**)–(**c**).

Theorem 1 *Consider a complete-information game u which satisfies* (**a**), (**b**), *and* (**c**). *Then an* NE ν *is strongly robust if and only if its induced action distribution $\widehat{\nu}$ is the unique* CE.

In other words, the *existence* of a strongly robust NE is tantamount to the *uniqueness* of a CE in the game. In particular, the quest for strongly robust distributions may be reduced to finding conditions that ensure CE uniqueness. We pursue this latter goal in the next section which reviews and extends the known settings that possess a unique CE.

[14]The mentioned lower (upper) semi-continuity is respectively defined by the requirements that $\liminf_{k\to\infty} u_i\left(a_i^k, a_{-i}\right) \geq u_i\left(a_i, a_{-i}\right)$ and $\limsup_{k\to\infty} \sum_{i\in N} u_i\left(a^k\right) \leq \sum_{i\in N} u_i\left(a\right)$, for any sequence $\{a^k\}_{k=1}^{\infty} \subset A$ that converges to a.

We shall now show that the equivalence result in Theorem 1 remains in force even when some discontinuity of payoffs occurs in the interior of A, along appropriately defined "diagonal" curves. To facilitate the exposition and the proof, we will assume that all action sets are one-dimensional, i.e., that each A_i is a closed non-degenerate interval $[\underline{a}_i, \overline{a}_i] \subset \mathbb{R}$. For technical reasons (the need for which will become clear in the proof), it will be further assumed that each payoff function u_i is defined on (or can be extended to) a superset $A^+ = \cup_{i \in N} [\underline{a}_i - \delta, \overline{a}_i + \delta] \times A_{-i}$ of A, for some $\delta > 0$, in such a way that all player i's actions above $\overline{a}_i$ are weakly dominated by $\overline{a}_i$, and all actions below $\underline{a}_i$ are weakly dominated by $\underline{a}_i$, when other players are restricted to A_{-i}.

Condition (**a**) on the continuity of each individual payoff function u_i when i's own actions are in the interior of A_i will be replaced by the following assumption, based on the continuity requirement posited in Section 4 of Dasgupta and Maskin (1986). Let $\{D(i)\}_{i \in N}$ be a collection of finite sets, and let $\{f_{ij}^d : \mathbb{R} \to \mathbb{R}\}_{i \neq j \in N, d \in D(i)}$ be a collection of strictly monotonic and continuous functions. Define

$$A^*(i) := \{a \in A^+ \mid \exists j \neq i, \exists d \in D(i) \text{ s.t. } a_j = f_{ij}^d(a_i)\}.$$

We shall assume that:

(a') for every $i \in N$, the set of discontinuity points of the payoff function u_i on A^+ is a subset of $A^*(i)$.

Our earlier condition (**b**) will also be replaced by a (somewhat strengthened) version of *weak* lower semi-continuity of Dasgupta and Maskin (1986). Specifically, we will assume that:

(b') there exist $0 < \lambda_1, ..., \lambda_n < 1$ such that, for every $i \in N$ and $a \in A$,

$$\lambda_i \liminf_{\varepsilon \to 0+} u_i(a_i - \varepsilon, a_{-i}) + (1 - \lambda_i) \liminf_{\varepsilon \to 0+} u_i(a_i + \varepsilon, a_{-i}) \geq u_i(a).$$

By Theorem 5 of Dasgupta and Maskin (1986), any u that satisfies (**a'**), (**b'**), and (**c**) possesses a (mixed-action) NE. Furthermore, strong robustness of that NE is subject to the same characterization as in Theorem 1:

Theorem 2 *Consider a complete-information game u which satisfies* (**a'**)*,* (**b'**)*, and* (**c**)*. Then an* NE ν *is strongly robust if and only if its induced action distribution* $\widehat{\nu}$ *is the unique* CE.

5 Applications: survey of games with a unique CE

In this section we examine various settings in which the uniqueness of a CE is known, or can be proved, and to which our results on the existence of a strongly robust NE may therefore be applied. In Section 5.1 we consider "smooth" games – those with continuously differentiable payoff functions – for which the question of CE uniqueness was addressed, in fullest generality to date, by Ui (2008). The class of smooth games for which Ui's conditions for CE uniqueness hold includes, *inter alia*, Cournot oligopolies with linear demand, strictly concave potential games, and a subclass of network games. Attention will also be drawn to some types of smooth Bertrand oligopolies with differentiated products, where the CE is unique by the supermodularity arguments of Milgrom and Roberts (1990).

Sections 5.2, 5.3 and 5.4 concern games that are not necessarily smooth. Section 5.2 is dedicated to Tullock contests, in which the payoff functions are not differentiable at a boundary point of the set of action profiles. We prove that the unique NE of a Tullock contest is also its unique CE, and therefore strongly robust. Section 5.3 discards the differentiability assumption altogether. In that section, we recall the notion of a socially concave game that is due to Even-dar et al. (2009), and the result of Hart and Mas-Colell (2015) on CE uniqueness in socially strictly concave games. The latter class of games includes various imperfectly discriminating contests (such as those arising from patent races), Cournot oligopolies with linear demand and possibly non-differentiable costs, and equilibrium implementation games for quasi-linear exchange economies. Section 5.4 provides examples of discontinuous games with a unique CE (whose strong robustness is established by appealing to a more potent Theorem 2). They include two classical two-player constant-sum games, namely the median-voter problem and a pure-location Hotelling game.

5.1 Smooth games

Following Ui (2008), a game u is called *smooth* if every payoff function u_i is continuously differentiable on A. Ui's work generalized the results of Rosen (1965) by showing that the unique NE is also the unique CE in Rosen's games. Ui uses a condition that is weaker than Rosen's, but in the interest of brevity we will confine ourselves to the original condition of Rosen in this exposition which Ui termed *strict monotonicity of payoff gradients.*

Formally, u has strictly monotone payoff gradients (SMPG)[15] if, for every $a \neq a'$,

$$\sum_{i \in N} [\nabla_i u_i(a) - \nabla_i u_i(a')] (a_i - a'_i) < 0. \tag{4}$$

An easy way to verify the SMPG condition, assuming that all payoff function are twice continuously differentiable, is to consider the matrix $[\partial^2 u_i (a) / \partial a_i \partial a_j]$; if it is negative definite then u has SMPG.[16] For any game u with SMPG, every payoff function u_i is strictly concave in a_i.[17]

[15] Ui (2008) allows weighted sums in Ineq. (4). By dividing each player's payoff by a positive constant, if needed, an equivalent game for which Ineq. (4) holds can always be obtained.

[16] See Lemma 3 and the proof of Corollary 6 in Ui (2008).

[17] See Lemma 5 in Ui (2008).

The main result of Ui (2008) shows that any smooth game with SMPG has a unique CE which is also the unique NE.[18] Since our continuity conditions (**a**)–(**c**) hold trivially for any smooth game, the unique NE is strongly robust by Theorem 1.

In what follows we review examples of smooth games with SMPG: strictly concave potential games (in Section 5.1.1); a subclass of network games (in Section 5.1.2); and differentiable Cournot oligopolies with linear demand (in Section 5.1.3). In Section 5.1.4 we point out that some types of smooth Bertrand oligopolies also have a unique CE (which is then a strongly robust NE), but in this case the argument, put forward in Milgrom and Roberts (1990), exploits log-supermodularity of these oligopolies instead of the SMPG condition.

5.1.1 Concave potential games

An important class of games which is relevant to our context is that of potential games, defined by Monderer and Shapley (1996). A game u is a *potential* game if there exists a function $P : A \to \mathbb{R}$ such that for every player i, every action profile $a = (a_i, a_{-i})$, and every action $a'_i \neq a_i$, $u_i(a'_i, a_{-i}) - u_i(a) = P(a'_i, a_{-i}) - P(a)$. In words, the potential function P mimics player i's payoff changes obtained by his unilateral deviations. Monderer and Shapley (1996) showed that potential games include congestion games as well as Cournot oligopolies with linear demand.

Neyman (1997) showed that for any game u with a continuously differentiable and strictly concave potential function, its unique NE is also a unique CE. By our Theorem 1, the unique NE is strongly robust.

Two additional comments are in order. First, although u with a continuously differentiable potential need not be smooth, it is (trivially) strategically equivalent to a smooth potential game. For the latter, a potential is strictly concave if and only if the game has SMPG (by Lemma 4 of Ui (2008)), which is an alternative way to deduce strong robustness of the unique NE of u. Based on the discussion following Corollary 6 in Ui (2008), we further remark that a twice continuously differentiable game u is a potential game if and only if the matrix $[\partial^2 u_i(a) / \partial a_i \partial a_j]$ is symmetric for each $a \in A$; in order for the potential to be strictly concave, the matrix must be negative definite. As has been mentioned already, this negative definiteness alone (without the symmetry) suffices for SMPG, and hence for strong robustness of the unique NE.

5.1.2 Network games

The class of smooth potential games includes a subclass of *network games*, in which the players' payoffs depend on the realized action profile $a \in \mathbb{R}^N_+$ and on the network (i.e., a graph) that

[18] See Proposition 5 in Ui (2008).

links different players to one another.[19] Bramoullé et al. (2014) follow the model of Ballester et al. (2006) in their study of network games with linear best-response functions. They consider a network game u in which player i's payoff is

$$u_i(a_i, a_{-i}) = a_i - \tfrac{1}{2}a_i^2 - \delta \sum_{j \neq i} g_{ij} a_j a_i,$$

where $\delta > 0$, the values $g_{ij} \in \{0, 1\}$ indicate whether players i and j are linked or not, and $g_{ij} = g_{ji}$ for every $i \neq j$; w.l.o.g., each player i can be constrained to use actions from an undominated set $A_i = [0, 2]$. Bramoullé et al. (2014) constructed a potential function for u, which is strictly concave precisely when the matrix $[1 + \delta g_{ij}]$ is positive definite. In the latter case, the NE of the game is also its unique CE by the result of Neyman (1997). While their interest lied primarily in NE stability w.r.t. continuous Nash tâtonnement (and, indeed, a unique NE turns out to be stable), that NE is also strongly robust by Theorem 1.

5.1.3 Cournot oligopoly with linear demand and differentiated products

In a Cournot oligopoly model with linear demand for differentiated products, each firm $i \in N$ chooses a non-negative output level a_i of its product from some compact $A_i \subset \mathbb{R}_+$ and incurs a production cost of $c_i(a_i)$, where each function c_i is assumed to be continuously differentiable, strictly increasing and convex. For each firm i, the inverse demand function for its product is given by $P_i(a) = \mathbf{B}_i + \sum_{j \in N} \mathbf{A}_{ij} a_j$, and the net-profit function of firm i is therefore

$$u_i(a) = \left(\mathbf{B}_i + \sum_{j \in N} \mathbf{A}_{ij} a_j \right) a_i - c_i(a_i). \tag{5}$$

As observed in Example 1 of Ui (2008), if the matrix $[(1 + \delta_{ij})\mathbf{A}_{ij}]$ is negatively definite (where δ_{ij} denotes the Kronecker delta), then the game has SMPG, and hence its NE is a unique CE that is strongly robust by Theorem 1.

When $\mathbf{A}_{ij} = -1$ and $\mathbf{B}_i = \mathbf{B}_j =: \mathbf{B} > 0$ for any i and j, the standard single-product oligopoly model with linear demand is obtained. The matrix $[-(1 + \delta_{ij})]$ is negative definite, and the uniqueness of CE follows. This fact is also a direct corollary of the main result of Neyman (1997), who observed in Section 2 of his paper that such an oligopoly has a continuously differentiable and strictly concave potential function.

5.1.4 Bertrand oligopoly with differentiated products

In their study of supermodular games, Milgrom and Roberts (1990) showed that NE provide bounds on the set of rationalizable actions of each player. Specifically, they proved that the

[19] For a broader review of these games see, for example, Bramoullé et al. (2014), Blume et al. (2015), and Section 5.3 of Ui (2016).

set of serially undominated actions[20] of each player i in a supermodular game has a largest element $\overline{a}_i$ and a smallest element $\underline{a}_i$,[21] and that the corresponding action profiles $\overline{a} = (\overline{a}_i)_{i \in N}$ and $\underline{a} = (\underline{a}_i)_{i \in N}$ are the largest, and respectively the smallest, NE in the game. As no CE can, with positive probability, prescribe strictly dominant actions to any player i, an obvious corollary is that uniqueness of a pure-action NE in the game implies its being a unique CE.

Following Milgrom and Roberts (1990),[22] we apply the above corollary to Bertrand competition. Consider a Bertrand oligopoly consisting of the set N of firms, each selling a different product. Each firm i has a constant unit cost $\overline{c}_i > 0$, and faces a twice continuously differentiable demand $D_i(a)$, where $a \in \times_{i \in N}[\overline{c}_i, \overline{p}]$ is a price vector. It is assumed that the products are substitutes and the elasticity of demand of firm i is a non-increasing function of the other firms' prices. The latter condition is equivalent requiring that $\partial^2 \log(D_i(a))/\partial a_i \partial a_j \geq 0$, for every profile a and every $j \neq i$, and it holds for various demand types (including, e.g., linear, CES, and logit demand functions). The payoff of every firm i is given by $u_i(a) = (a_i - \overline{c}_i) D_i(a)$.

Under the above conditions, the log-transformed game $\log u = (\log u_i)_{i \in N}$ is supermodular, and hence the result on the existence and the role of the NE $\overline{a}$ and $\underline{a}$ applies to both $\log u$ and u. For several families of demand functions, including linear, CES and logit, Milgrom and Roberts (1990) furthermore show that the game $\log u$ (and thus u) has a unique NE. This, by the arguments above, implies that $\overline{a} = \underline{a}$ a unique CE of u. Since u is smooth, our continuity conditions (**a**)–(**c**) hold trivially, and that NE is strongly robust by Theorem 1.

5.2 The Tullock rent-seeking game

Many economic settings, ranging from political races to investments in R&D, can be modelled as contests where players exert effort to win the competition and the winner receives a reward.[23] The Tullock rent-seeking game (see Tullock (2001)), or Tullock contest, is a complete-information game $u = (u_i)_{i \in N}$ of this type.

In a Tullock contest with $n \geq 2$ players, every player i exerts an effort $a_i \in \mathbb{R}_+$ for a chance to win a single prize, e.g., an economic rent. The *success function* $p = (p_i)_{i \in N}$ specifies the probability of each contester to receive the prize based on the realized effort profile a, and is assumed to have the following form: for each player i and any profile $a = (a_i)_{i \in N}$ that is distinct from the zero-effort profile $\mathbf{0}$,

$$p_i(a) = \frac{f_i(a_i)}{\sum_{j \in N} f_j(a_j)},$$

[20] Undominated actions are those that survive the iterative process of eliminating strongly dominated actions.

[21] In a supermodular game the action sets A_i are lattices, and largest and smallest elements of a subset of A_i are defined w.r.t. the lattice order on A_i.

[22] See Section 4, Example (2), of Milgrom and Roberts (1990).

[23] See, for example, Dasgupta and Stiglitz (1980), Dixit (1987) and Skaperdas (1996) among many others.

where $\{f_j : \mathbb{R}_+ \to \mathbb{R}_+\}_{j \in N}$ are *effort impact* functions of the players. These functions are assumed to be twice differentiable, strictly increasing, concave, and vanishing at 0. If a is the zero-effort profile $\mathbf{0}$, then $p(\mathbf{0})$ can be an arbitrary strictly positive probability vector.

All efforts are costly. Each effort is identified with its cost, and the value of the prize to each player is normalized to 1, giving rise to net utilities $(u_i)_{i \in N}$ where

$$u_i(a_i, a_{-i}) = p_i(a) - a_i \tag{6}$$

for every $i \in N$ and $a \in \mathbb{R}^n_+$. Note that this formulation also allows for the case of player-specific general costs of effort. Namely, given a twice-differentiable, strictly increasing and convex cost function $c_i : \mathbb{R}_+ \to \mathbb{R}_+$ with $c_i(0) = 0$, one can obtain an equivalent game with payoffs given by (6) by redefining each player i's effort impact function to be equal to $f_i \circ c_i^{-1}$.

Since all efforts above 1 are strictly dominated by effort 0 for all players, it can be assumed w.l.o.g. that player i's action set is $A_i = [0, 1]$. Tullock contests may thereby be viewed as belonging to our basic framework, and one can easily verify that they meet requirements (**a**)–(**c**). Moreover, it has already been established by Szidarovszky and Okuguchi (1997) that a Tullock contest has a unique (pure-action) NE. In the following theorem we prove that the same equilibrium is also the unique CE, which implies its strong robustness via Theorem 1.

Theorem 3 *The* NE *of a Tullock contest is also its unique* CE*, and therefore is strongly robust.*

In a Tullock contest, each p_i is *strictly* concave in a_i for a fixed $a_{-i} \neq \mathbf{0}_{-i}$, and convex in a_{-i} for a fixed a_i; this implies that the payoff function u_i has the same properties. In addition, each u_i is continuously differentiable on $\mathbb{R}^n_+ \backslash \{\mathbf{0}\}$. Since the sum $\sum_{i \in N} u_i(a) = \sum_{i \in N} a_i$ is linear in a, Tullock contests have SMPG on $\mathbb{R}^n_{++}$ by an observation made in Goodman (1980). The lack of differentiability of the payoff functions at $\mathbf{0}$ prevents, however, a direct application of Proposition 5 of Ui (2008) on CE uniqueness, and necessitates the separate approach that we have taken. The proof of Theorem 3 has an additional merit of being quite short and straightforward, with the main auxiliary function H inspired by the method of proof in Liu (1996).

5.3 Socially concave games

The class of *socially concave* games, introduced in Even-dar et al. (2009) and, with greater generality, in Hart and Mas-Colell (2015) contains games that arise in several widely used models such as Tullock contests, patent races (see Section 5.3.2), Cournot oligopolies with linear demand (see Section 5.3.3), and quasi-linear exchange economies (see Hart and Mas-Colell (2015)). This class of games is also important in our context since they tend to have a unique CE.

Let us now formally define socially concave games. A game u is socially (strictly) concave if the sum of payoffs $\sum_{i\in N} u_i(a)$ is (strictly) concave in a, and every payoff function $u_i(a_i, a_{-i})$ is convex in a_{-i}. Note that the combination of these two conditions immediately implies that player i's utility function is concave in a_i for a fixed a_{-i}.

Hart and Mas-Colell (2015) showed that socially strictly concave games have at most one CE. In addition, by Remark 1, conditions (**a**)–(**c**) imply the existence of an NE, which is in particular a CE. This leads to the following Corollary:

Corollary 1 *Any socially strictly concave game u that satisfies* (**a**), (**b**), *and* (**c**) *has an* NE *which is also the unique* CE, *and therefore is strongly robust. If all actions sets are intervals in $\mathbb{R}$,* (**a**) *and* (**b**) *above can be replaced by* (**a'**) *and* (**b'**).

The proof is omitted since it follows directly from our theorems 1, 2 and Proposition 10 of Hart and Mas-Colell (2015).

Remark 2 *It is well-known that any concave function on a convex polytope is lower semi-continuous* (*see, e.g.,* Gale et al. (1968))*. Hence, if u is a socially concave game in which each action set A_i is a polytope, then condition* (**b**) [*or* (**b'**), *when $A_i \subset \mathbb{R}$*] *holds trivially, and condition* (**c**) *is equivalent to*

(**c'**) the sum $\sum_{i\in N} u_i(a)$ is a continuous function.

5.3.1 Imperfectly discriminating contests

Tullock contests do not fall under the purview of Corollary 1 since they do not have a strictly convex sum of payoffs, and we must rely on Theorem 3 for the CE uniqueness result. However, the following (rather big) family of imperfectly discriminating contests does satisfy the conditions of Corollary 1.

Consider a contest where the action (effort) set of every player i is a closed bounded interval A_i and player i's payoff is given by $u_i(a) = p_i(a) - c_i(a_i)$ for every action profile a. Let us assume that $(p_i(a))_{i\in N}$ is a concave-sum sub-probability vector[24] such that $p_i(a_i, a_{-i})$ is a convex function of a_{-i} for a fixed a_i, and the cost function $c_i : A_i \to \mathbb{R}_+$ is continuous and convex. Notice that u is a socially concave game which satisfies (**a**) [or (**a'**)] and (**c'**) whenever $p = (p_i)_{i\in N}$ does so. Thus, in light of Remark 2, Corollary 1 applies to imperfectly discriminating contests when they are strictly socially concave:

[24]That is, $\sum_{i\in N} p_i(a)$ is a concave function and $\sum_{i\in N} p_i(a) \leq 1$ (the latter means that the prize may be withheld with positive probability).

Corollary 2 *An imperfectly discriminating contest u has an* NE *which is also the unique* CE*, and therefore is strongly robust, provided that $p = (p_i)_{i \in N}$ satisfies* (**a**) *[or* (**a'**)*] and* (**c'**)*, and that either $\sum_{i \in N} p_i$ is* strictly *concave or c_i is* strictly *convex for every $i \in N$.*

5.3.2 Patent races

Baye and Hoppe (2003) consider the patent-race model of Loury (1979), where firms compete over an infinite-life patent. Baye and Hoppe prove that the patent race is strategically equivalent to an imperfectly discriminating contest, which is a variant of the Tullock competition. We will now show that this contest also meets the requirement of Corollary 2.

A patent race is an n-firms game where each firm i chooses to invest $a_i \in \mathbb{R}_+$ in R&D for a patent of value $v > 0$. Given a_i, the probability of firm i to reach a discovery until time $t \geq 0$ is $1 - e^{-h(a_i)t}$, where $h : \mathbb{R}_+ \to \mathbb{R}_+$ is a strictly increasing, concave, twice-differentiable function. Taking a positive interest rate r, the payoff of firm i is given by

$$\begin{aligned} u_i(a_i, a_{-i}) &= \int_{\mathbb{R}_+} h(a_i) v e^{-t\left[r + \sum_{j \in N} h(a_j)\right]} dt - a_i \\ &= v \frac{h(a_i)}{r + \sum_{j \in N} h(a_j)} - a_i. \end{aligned}$$

Since all sufficiently high investments are strictly dominated by the null investment (i.e., $a_i = 0$), it can be assumed w.l.o.g. that every player i's action set is some bounded closed interval $A_i \subseteq \mathbb{R}_+$.

Evidently, if r tends to zero, the patent race is strategically equivalent to a Tullock contest, as noted in Theorem 3 of Baye and Hoppe (2003). Moreover, when $r > 0$ and after a division of all payoffs by v, the game u is included in the scope of Corollary 2 even if h is merely continuous.

Claim 4 *The patent-race game has an* NE *which is also the unique* CE*, and therefore is strongly robust.*

5.3.3 Cournot oligopoly

Consider the standard (single-product) Cournot oligopoly model with linear demand, given by the description in Section 5.1.3 with $\mathbf{A}_{ij} = -\mathbf{A} < 0$, $\mathbf{B}_i = \mathbf{B} > 0$ for any i and j, but discard the assumption that the cost functions are continuously differentiable. Each cost function $c_i : A_i \to \mathbb{R}_+$ is now assumed to be continuous, convex, and strictly increasing.

The results on smooth games are not applicable to such an oligopoly, but, with payoff functions given by $u_i(a) = \left(\mathbf{B} - \mathbf{A}\left(\sum_{i \in N} a_i\right)\right) a_i - c_i(a_i)$, the game u is clearly socially strictly concave. Since it trivially satisfies (**a**) and (**c'**), Corollary 1 applies.

When a *duopoly* is considered, CE is unique under much more general conditions on the inverse demand that the firms face: it only needs to be strictly decreasing and log-concave (see Theorem 2.3 and Corollary 2.4 in Amir (1996)).[25] Amir showed that such a duopoly can be viewed as an ordinally supermodular game, which implies, via the previously mentioned method of Milgrom and Roberts (1990),[26] that the unique Cournot equilibrium is also a unique CE.

5.4 Discontinuous zero-sum games

In this section we will show the existence of a unique CE in two classical, discontinuous and zero-sum-equivalent games: the median-voter problem (see Section 5.4.1) and a pure-location Hotelling game (see Section 5.4.2). In these games, the discontinuity in payoffs occurs at interior action profiles that lie along strictly increasing curves, which necessitates the use of Theorem 2, instead of Theorem 1 as in the previous sections. To further demonstrate the applicability of Theorem 2, we will also consider a two-bidder first-price auction with common values and complete information, where the payoffs are discontinuous along the diagonal but a unique CE is known to exist, and hence an appeal to Theorem 2 is required (see Section 5.4.3).[27]

The main tool in establishing the uniqueness of a CE will be the following observation:

Remark 3 *Consider a two-player zero-sum game u. If μ is a* CE *of u, it is easy to see that μ_{A_i} (the marginal distribution induced by μ on A_i) is an optimal strategy of each player i, and the action-profile (μ_{A_1}, μ_{A_2}) is an* NE. *Thus, if u has a unique* NE, *which is moreover in pure actions, then that* NE *is the unique* CE *of u.*

Since any two-player constant-sum game u is obviously strategically equivalent to a zero-sum game, the last conclusion in Remark 3 also applies to such u: if u has a unique, *pure-action* NE, then it is the unique CE of the game.

5.4.1 The median-voter problem

The median voter problem, also known as a Hotelling-Downs game, is a simple model of bipartisan political competition with a one-dimensional policy space. Following Persson and Tabellini

[25]For the results of Amir (1996) to hold, it must be further assumed that there exists $\overline{Q} > 0$ such that the inverse demand function P satisfies $QP(Q) - c_i(Q) < 0$ for every $Q > \overline{Q}$ and $i = 1, 2$. This assumption is satisfied if, e.g., for each firm i, $A_i = [0, q_i]$ for some $q_i > 0$, and all output levels above q_i are strictly dominated by output 0.

[26]In fact, an extension due to Milgrom and Shannon (1994) of this method is required to deal with ordinally supermodular games.

[27]Although not zero-sum, such an auction is a *continuous-sum* game on account of the equal values that bidders have for the auctioned object.

(2000) (Section 3.2, p. 49 – 51), we assume that there are two players (i.e., candidates), and that their action sets, representing possible policy promises, are the interval $[0,1]$. Voters, of which there is a continuum, have single-peaked preferences over the policy space, and their ideal points are continuously distributed on $[0,1]$ with a strictly positive density function f.

The game begins by each player $i = 1,2$ choosing an action in $A_i = [0,1]$, in a possibly mixed fashion. Given a realized action profile $a = (a_1, a_2)$, every voter with an ideal point $x \in [0,1]$ votes for player i whose action a_i is the closest to x, with a symmetric tie-breaking rule. For every profile a, denote by $W_i(a)$ the mass of voters who vote for i; that is,

$$W_i(a) = \begin{cases} \int_0^{\frac{a_1+a_2}{2}} f(x)dx, & \text{if } a_i < a_{-i}, \\ \frac{1}{2}, & \text{if } a_i = a_{-i}, \\ \int_{\frac{a_1+a_2}{2}}^1 f(x)dx, & \text{if } a_i > a_{-i}. \end{cases}$$

The payoff function of every player i is then given by

$$u_i(a) = \begin{cases} -1, & \text{if } W_i(a) < W_{-i}(a), \\ 0, & \text{if } W_i(a) = W_{-i}(a), \\ 1, & \text{if } W_i(a) > W_{-i}(a), \end{cases}$$

which defines a two-player zero-sum game u.

Note that the action $a_i = m$, where m is the median voter (characterized by the equation $\int_0^m f(x)dx = \frac{1}{2}$), guarantees player i the payoff of 0, and leads to a strictly positive payoff if his opponent uses any (mixed) action that is different from m. It follows that m is the unique optimal strategy for each player in the zero-sum game u, and therefore (m,m) is its unique, pure-action NE. By Remark 3, that NE is also the unique CE. Moreover, one can easily verify that conditions (**a'**), (**b'**), and (**c**) hold in this framework,[28] and so the unique pure NE is strongly robust by Theorem 2.

5.4.2 Hotelling model of pure location

A pure-location Hotelling game is a classical motivating scenario for a more general Hotelling (1929) duopoly model of spatial competition. In a location game, each firm $i = 1,2$ chooses a location (sale point) in the interval $[0,1]$, which may represent the main street in a town, and hence $A_1 = A_2 = [0,1]$. Both firms offer for sale the same product, and charge the same mill price for each unit of the good and at each sale point. Unit-demand customers are located

[28]Notice that (**a'**) holds for $f_{ij}^d(a_i) = 2m - a_i$ with $|D(i)| = 1$; (**b'**) holds for $\lambda_1 = \lambda_2 = \frac{1}{2}$, and (**c**) is satisfied trivially because u is constant-sum.

along $[0, 1]$; the continuous distribution of their locations has a positive density function f. Each customer patronizes the closest seller, with a symmetric tie-breaking rule, and both firms' objective is to maximize their respective market shares. The corresponding constant-sum game u may thus be described in terms of the functions W_i from Section 5.4.1: $u_i(a) = W_i(a)$ for each $a \in A$ and $i = 1, 2$.

Just as in, e.g., Corollary 1 of Ben-Porat and Tennenholtz (2016) (taken for $k = 1$), it can be seen that the location game u has a unique, pure-action NE, in which both firms choose the location of the median customer m. By Remark 3, that NE is also the unique CE of the game. As conditions (**a'**), (**b'**), and (**c**) hold for the game u, the unique NE is strongly robust by Theorem 2.

5.4.3 Two-bidder first-price auction with common values

Consider an auction in which two bidders with the common value of 1 for the auctioned object submit bids in $A_1 = A_2 = [0, 1]$, of which the highest wins and is paid (with a symmetric tie-breaking rule). Each payoff function u_i is therefore given by

$$u_i(a) = \begin{cases} 1 - a_i, & \text{if } a_i > a_{-i}, \\ \frac{1}{2}\left(1 - a_i\right), & \text{if } a_i = a_{-i}, \\ 0, & \text{if } a_i < a_{-i}. \end{cases}$$

The restriction to two bidders and equal value is undoubtedly severe, but that is the only context in which such an auction is guaranteed to have a unique NE and CE (see Section 3.1 in Dütting et al. (2014)). Conditions (**a'**), (**b'**), and (**c**) obviously hold for the game u,[29] and hence the unique NE is strongly robust by Theorem 2.

References

Amir, R. (1996): "Cournot Oligopoly and the Theory of Supermodular Games," *Games and Economic Behavior*, 15, 132–148.

Aumann, R. J. (1974): "Subjectivity and Correlation in Randomized Strategies," *Journal of Mathematical Economics*, 1, 67–96.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74, 1403–1417.

[29] Condition (**a'**) holds for $f^d_{ij}(a_i) = a_i$ with $|D(i)| = 1$, $i = 1, 2$; (**b'**) holds for $\lambda_1 = \lambda_2 = 0$, and (**c**) holds because $u_1(a) + u_2(a) = 1 - \max(a_1, a_2)$.

Baye, M. R. and H. C. Hoppe (2003): "The Strategic Equivalence of Rent-Seeking, Innovation, and Patent-Race Games," *Games and Economic Behavior*, 44, 217–226.

Ben-Porat, O. and M. Tennenholtz (2016): "Multi-Unit Facility Location Games," in *WINE 2016: The 12th Conference on Web and Internet Economics*.

Bertsekas, D. P. and S. E. Shreve (2004): *Stochastic optimal control : the discrete time case*, Athena Scientific.

Blume, L. E., W. A. Brock, S. N. Durlauf, and R. Jayaraman (2015): "Linear Social Interactions Models," *Journal of Political Economy*, 123, 444–496.

Borkar, V. S. (1995): *Probability theory : an advanced course*, Springer.

Bramoullé, Y., R. Kranton, and M. D'Amours (2014): "Strategic Interaction and Networks," *The American Economic Review*, 104, 898–930.

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

Dasgupta, P. and E. Maskin (1986): "The Existence of Equilibrium in Discontinuous Economic Games, Part I (Theory)," *Review of Economic Studies*, 53.

Dasgupta, P. and J. Stiglitz (1980): "Uncertainty, Industrial Structure, and the Speed of R&D," *The Bell Journal of Economics*, 11, 1.

Dekel, E. and D. Fudenberg (1990): "Rational Behavior with Payoff Uncertainty," *Journal of Economic Theory*, 52, 243–267.

Dixit, A. (1987): "Strategic Behavior in Contests," *The American Economic Review*, 77, 891–898.

Dütting, P., T. Kesselheim, and É. Tardos (2014): "Mechanism with Unique Learnable Equilibria," in *Proceedings of the Fifteenth ACM Conference on Economics and Computation - EC '14*, New York, New York, USA: ACM Press, 877–894.

Even-dar, E., Y. Mansour, and U. Nadav (2009): "On the Convergence of Regret Minimization Dynamics in Concave Games," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing - STOC '09*, New York, New York, USA: ACM Press, 523–532.

Foster, D. P. and R. V. Vohra (1997): "Calibrated Learning and Correlated Equilibrium," *Games and Economic Behavior*, 21, 40–55.

Fudenberg, D., D. M. Kreps, and D. K. Levine (1988): "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44, 354–380.

Gale, D., V. Klee, and R. T. Rockafellar (1968): "Convex Functions on Convex Polytopes," *Proceedings of the American Mathematical Society*, 19, 867–873.

Goodman, J. C. (1980): "Note on Existence and Uniqueness of Equilibrium Points for Concave N-Person Games," *Econometrica*, 48, 251–251.

Hart, S. and A. Mas-Colell (2000): "A Simple Adaptive Procedure Leading to Correlated Equilibrium," *Econometrica*, 68, 1127–1150.

——— (2015): "Markets, Correlation, and Regret-Matching," *Games and Economic Behavior*, 93, 42–58.

Hart, S. and D. Schmeidler (1989): "Existence of Correlated Equilibria," *Mathematics of Operations Research*, 14, 18–25.

Hotelling, H. (1929): "Stability in Competition," *The Economic Journal*, 39, 41.

Kajii, A. and S. Morris (1997): "The Robustness of Equilibria to Incomplete Information," *Econometrica*, 65, 1283–1309.

Liu, L. (1996): "Correlated Equilibrium of Cournot Oligopoly Competition," *Journal of Economic Theory*, 68, 544–548.

Loury, G. C. (1979): "Market Structure and Innovation," *The Quarterly Journal of Economics*, 93, 395.

Milgrom, P. and J. Roberts (1990): "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 58, 1255–1277.

Milgrom, P. and C. Shannon (1994): "Monotone Comparative Statics," *Econometrica*, 62, 157.

Monderer, D. and L. S. Shapley (1996): "Potential Games," *Games and Economic Behavior*, 14, 124–143.

Neyman, A. (1997): "Correlated Equilibrium and Potential Games," *International Journal of Game Theory*, 26, 223–227.

Papadimitriou, C. H. and T. Roughgarden (2008): "Computing Correlated Equilibria in Multi-Player Games," *Journal of the ACM*, 55, 1–29.

Persson, T. and G. E. Tabellini (2000): *Political economics : explaining economic policy*, MIT Press.

Reny, P. J. (1999): "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games," *Econometrica*, 67, 1029–1056.

Rosen, J. B. (1965): "Existence and Uniqueness of Equilibrium Points for Concave N-Person Games," *Econometrica*, 33, 520–534.

Skaperdas, S. (1996): "Contest Success Functions," *Economic Theory*, 7, 283–290.

Szidarovszky, F. and K. Okuguchi (1997): "On the Existence and Uniqueness of Pure Nash Equilibrium in Rent-Seeking Games," *Games and Economic Behavior*, 18, 135–140.

Tullock, G. (2001): "Efficient Rent Seeking," in *Efficient Rent-Seeking*, Boston, MA: Springer US, 3–16.

Ui, T. (2008): "Correlated Equilibrium and Concave Games," *International Journal of Game Theory*, 37, 1–13.

——— (2016): "Bayesian Nash Equilibrium and Variational Inequalities," *Journal of Mathematical Economics*, 63, 139–146.

Viossat, Y. (2008): "Is Having a Unique Equilibrium Robust?" *Journal of Mathematical Economics*, 44, 1152–1160.

A Appendices

A.1 Extending CE to mixed-action deviations

Proposition 2 *If μ is a* CE *then* Ineq. (1) *holds for any Borel-measurable function $\psi_i : A_i \to M(A_i)$, with $u_i(\psi_i(a_i), a_{-i})$ in* Ineq. (1) *being defined as $\int_{A_i} u_i(b_i, a_{-i})\, d\psi_i(a_i)(b_i)$.*

Proof. Suppose that Ineq. (1) does not hold for some $i \in N$ and some measurable $\psi_i' : A_i \to M(A_i)$. It is well known (see, e.g., Corollary 3.1.2 of Borkar (1995)) that conditional distribution $\mu(\cdot \mid a_i) \in M(A_{-i})$, induced by μ on A_{-i} given a_i, can be defined for every $a_i \in A_i$ in such a way that the stochastic kernel $(a_i, B) \mapsto \mu(B \mid a_i)$ is Borel-measurable in a_i for any Borel subset B of A_{-i}. By assumption, the stochastic kernel $(a_i, B) \mapsto \psi_i'(a_i)(B)$ is also Borel-measurable in a_i for any Borel subset B of A_i. By Proposition 7.29 of Bertsekas and Shreve (2004) on integration involving Borel-measurable stochastic kernels, the functions $(a_i, b_i) \mapsto \int_{A_{-i}} u_i(b_i, a_{-i})\, d\mu(a_{-i} \mid a_i)$ and $a_i \mapsto \int_{A_{-i}} u_i(\psi_i'(a_i), a_{-i})\, d\mu(a_{-i} \mid a_i)$ are Borel-measurable. Hence the graph of the (non-empty-valued) correspondence

$$\Psi_i(a_i) := \{b_i \in A_i \mid \int_{A_{-i}} u_i(b_i, a_{-i})\, d\mu(a_{-i} \mid a_i) \geq \int_{A_{-i}} u_i(\psi_i'(a_i), a_{-i})\, d\mu(a_{-i} \mid a_i)\}$$

is also Borel-measurable. By the measurable choice theorem, the correspondence Ψ_i possesses a single-valued measurable selection ψ_i.

Clearly,

$$\int_{A_{-i}} u_i\left(\psi_i(a_i), a_{-i}\right) d\mu\left(a_{-i} \mid a_i\right) \geq \int_{A_{-i}} u_i\left(\psi_i'\left(a_i\right), a_{-i}\right) d\mu\left(a_{-i} \mid a_i\right)$$

for every $a_i \in A_i$, and integrating both terms w.r.t. μ_{A_i} (the marginal distribution induced by μ on A_i) yields

$$\int_{A_{-i}} u_i\left(\psi_i\left(a_i\right), a_{-i}\right) d\mu(a) \geq \int_{A_{-i}} u_i\left(\psi_i'\left(a_i\right), a_{-i}\right) d\mu\left(a\right).$$

Therefore, ψ_i violates Ineq. (1) because ψ_i' does so, contradicting the assumption that μ is a CE. ■

A.2 Proof of Proposition 1

Proof. Let μ' be any CE of u. Consider a 0-elaboration $\mathcal{U}_{0,\mu'}$ of u in which $(\Omega, \mathcal{F})$ is the set of action profiles A with the Borel σ-algebra on it, $P = \mu'$ and, for each player i, $\mathcal{F}_i = \{B_i \times A_{-i} \mid B_i \subset A_i \text{ is a Borel set}\}$ and $U_i \equiv u_i$ in a state-independent fashion. It follows from Proposition 2 that a strategy profile τ in which $\tau_i(a) = a_i$ for every $i \in N$ and $a \in A$ is a pure-action BNE of $\mathcal{U}_{0,\mu'}$, with $\mu\left(\tau\right) = \mu'$. It therefore follows from Definition 1 that $\widehat{\nu}$, the product action distribution of the strongly robust ν, must coincide with the CE μ'. Thus $\widehat{\nu}$ must coincide with *any* CE of u, and hence it is the unique CE. ■

A.3 Proof of Remark 1

Proof. Assume that u satisfies (**a**), (**b**), and (**c**). As in the proof of Proposition 5.1 in Reny (1999), it can be seen that (**b**) implies lower semi-continuity of each $u_i\left(\nu_i, \nu_{-i}\right)$ in ν_i when players use mixed strategies. Furthermore, it follows from (**a**) and the Portmanteau theorem that $u_i\left(\nu_i, \nu_{-i}\right)$ is continuous at any point ν as long as $\nu_i \in M(A_i)$ satisfies $\nu_i(\partial(A_i)) = 0$. These two observations, together with the fact that any $\nu_i \in M(A_i)$ can be approximated by probability measures on A_i for which $\partial(A_i)$ is a zero-measure set, imply that the payoffs in mixed strategies are payoff-secure. That is, for every $\nu \in \times_{i\in N} M(A_i)$ and $\varepsilon > 0$, each player i can secure a payoff of at least $u_i\left(\nu\right) - \varepsilon$. (The latter means that there exists $\overline{\nu}_i \in M(A_i)$ such that $u_i\left(\overline{\nu}_i, \nu_{-i}'\right) \geq u_i\left(\nu\right) - \varepsilon$ for any ν_{-i}' in some open neighborhood of ν_{-i}.) Given the payoff-security of the mixed-strategy extension of u, and condition (**c**) on pure-strategy payoffs, the existence of a mixed-strategy NE in u follows from Proposition 5.1 and Corollary 5.2 of Reny (1999). ■

A.4 Proof of Theorem 1

The "only if" direction of the theorem is given by Proposition 1. As for the "if" direction, consider a sequence $\{\mathcal{U}^k\}_{k=1}^{\infty}$ of incomplete information games and a sequence of corresponding BNE $\{\sigma^k\}_{k=1}^{\infty}$ such that each $\mathcal{U}^k = \left\{\Omega^k, P^k, \{F_i^k\}_{i\in N}, \{U_i^k\}_{i\in N}\right\}$ is a δ_k-elaboration of u and $\lim_{k\to\infty}\delta_k = 0$. We will show that, for any subsequence of $\{\mu(\sigma^k)\}_{k=1}^{\infty} \subset M(A)$ that converges to some $\mu' \in M(A)$, the limit μ' is a CE of u. W.l.o.g., we will take such a subsequence to be $\{\mu(\sigma^k)\}_{k=1}^{\infty}$ itself in our forthcoming considerations.

The following lemma will be instrumental in the rest of the proof.

Lemma 1 *For any $i \in N$ and any measurable function $\psi_i : A_i \to A_i$,*

$$\liminf_{k\to\infty} \overline{U}_i^k\left(\sigma^k\right) \geq \int_A u_i\left(\psi_i\left(a_i\right), a_{-i}\right) d\mu'\left(a\right).$$

Proof of Lemma 1. Suppose to the contrary that, for some $i \in N$ and some measurable $\psi_i : A_i \to A_i$,

$$\liminf_{k\to\infty} \overline{U}_i^k\left(\sigma^k\right) < \int_A u_i\left(\psi_i\left(a_i\right), a_{-i}\right) d\mu'\left(a\right). \tag{7}$$

We will first show that such ψ_i can, w.l.o.g., be assumed to be continuous. Indeed, for any $\varepsilon > 0$, by Lusin's theorem, the given ψ_i is continuous on a compact subset E_ε of A_i with $\mu'(E_\varepsilon \times A_{-i}) > 1-\varepsilon$. By applying the Tietze extension theorem to each coordinate of $\psi_i|_{E_\varepsilon}$, the restriction of ψ_i to E_ε, this function may be extended to a continuous $\psi_i^\varepsilon : A_i \to \mathbb{R}^{m_i}$. If $proj_{A_i} : \mathbb{R}^{m_i} \to A_i$ is the projection onto A_i, which sends any $a_i \in \mathbb{R}^{m_i}$ into the point in A_i with the shortest Euclidean distance from a_i, then the composite function $\overline{\psi}_i^\varepsilon = proj_{A_i} \circ \psi_i^\varepsilon : A_i \to A_i$ is continuous, and is identical to ψ_i on E_ε. Since u_i is bounded and $\lim_{\varepsilon\to 0+}\mu'(E_\varepsilon \times A_{-i}) = 1$, clearly

$$\lim_{\varepsilon\to 0+} \int_A u_i\left(\overline{\psi}_i^\varepsilon\left(a_i\right), a_{-i}\right) d\mu'\left(a\right) = \int_A u_i\left(\psi_i\left(a_i\right), a_{-i}\right) d\mu'\left(a\right),$$

and so ψ_i can be replaced in Ineq. (7) by $\overline{\psi}_i^\varepsilon$ for some sufficiently small ε without affecting that inequality. Thus, it can be assumed w.l.o.g. that ψ_i for which Ineq. (7) holds is *continuous*.

Next, we will show that, w.l.o.g., it can be assumed that the values of the continuous ψ_i in Ineq. (7) avoid the boundary $\partial(A_i)$, i.e., that $\psi_i : A_i \to A_i \backslash \partial(A_i)$. To this end, for any $\varepsilon > 0$ consider the closed and convex set A_i^ε that consists of all points in A_i whose Euclidean distance from $\partial(A_i)$ is at least ε. As A_i has full dimension, A_i^ε is non-empty for all sufficiently small ε, and the projection onto A_i^ε, $proj_{A_i^\varepsilon} : \mathbb{R}^{m_i} \to A_i^\varepsilon$, is well-defined. Since the function $\overline{\overline{\psi}}_i^\varepsilon = proj_{A_i^\varepsilon} \circ \psi_i$ converges to ψ_i pointwise on A_i as $\varepsilon \to 0$, by assumption (**b**) on u_i

$$\liminf_{\varepsilon\to 0+} u_i\left(\overline{\overline{\psi}}_i^\varepsilon\left(a_i\right), a_{-i}\right) \geq u_i\left(\psi_i\left(a_i\right), a_{-i}\right)$$

for every $a \in A$. Hence, by Fatou's lemma,

$$\liminf_{\varepsilon \to 0+} \int_A u_i \left(\overline{\overline{\psi}}_i^{\varepsilon} (a_i), a_{-i} \right) d\mu' (a) \geq \int_A u_i (\psi_i (a_i), a_{-i}) \, d\mu' (a) .$$

It follows that the continuous function ψ_i can be replaced in (7) by another continuous function $\overline{\overline{\psi}}_i^{\varepsilon}$, for some sufficiently small ε, without affecting the inequality. Thus, it can be assumed w.l.o.g. that the values of the continuous ψ_i in Ineq. (7) avoid the boundary $\partial (A_i)$, i.e., that $\psi_i : A_i \to A_i \backslash \partial (A_i)$. Consequently, by assumption (**a**) on u_i, the function $u_i (\psi_i (a_i), a_{-i})$ is continuous on A.

For any $\nu_i \in M(A_i)$, let $\psi_i (\nu_i) \in M(A_i)$ be the probability measure given by $\psi_i (\nu_i) (B) = \nu_i \left(\psi_i^{-1} (B) \right)$ for every Borel set B in A_i.[30] Note that ψ_i can thus be applied to any $M(A_i)$-valued strategy σ_i^k, thereby producing a new strategy, $\psi_i \left(\sigma_i^k \right)$, for player i in the game $\mathcal{U}^k$. The uniform boundedness of U_i^k (together with the fact that $U_i^k = u_i$ on a set with a $\mu \left(\sigma^k \right)$-measure tending to 1) now implies that

$$\begin{aligned}
\lim_{k \to \infty} \overline{U}_i^k \left(\psi_i \left(\sigma_i^k \right), \sigma_{-i}^k \right) &= \lim_{k \to \infty} \int_\Omega U_i^k \left(\psi_i \left(\sigma_i^k (\omega) \right), \sigma_{-i}^k (\omega), \omega \right) dP^k (\omega) \\
&= \lim_{k \to \infty} \int_\Omega u_i \left(\psi_i \left(\sigma_i^k (\omega) \right), \sigma_{-i}^k (\omega) \right) dP^k (\omega) \\
&= \lim_{k \to \infty} \int_A u_i (\psi_i (a_i), a_{-i}) \, d\mu \left(\sigma^k \right) (a) = \int_A u_i (\psi_i (a_i), a_{-i}) \, d\mu' (a) ,
\end{aligned}$$

where the last equality follows from the weak convergence of $\{ \mu \left(\sigma^k \right) \}_{k=1}^{\infty}$ to μ' and the continuity of $u_i (\psi_i (a_i), a_{-i})$. Thus,

$$\lim_{k \to \infty} \overline{U}_i^k \left(\psi_i \left(\sigma_i^k \right), \sigma_{-i}^k \right) = \int_A u_i (\psi_i (a_i), a_{-i}) \, d\mu' (a) .$$

Combining this with Ineq. (7) shows that, for some k, $\overline{U}_i^k \left(\psi_i \left(\sigma_i^k \right), \sigma_{-i}^k \right) > \overline{U}_i^k \left(\sigma^k \right)$, in contradiction to the assumption that σ^k is a BNE of $\mathcal{U}^k$. ■

Proof of Theorem 1. By taking ψ_i to be the identity function, Lemma 1 implies that

$$\liminf_{k \to \infty} \overline{U}_i^k \left(\sigma^k \right) \geq \int_A u_i (a) \, d\mu' (a) \tag{8}$$

for every $i \in N$. On the other hand, by using the uniform boundedness of all payoff functions (together with the fact that the payoffs are given by u on a set with a $\mu \left(\sigma^k \right)$-measure tending

[30] In other words, if ν_i is the probability distribution of a random variable X, then $\psi_i (\nu_i)$ is the distribution of $\psi_i (X)$.

to 1) we obtain

$$\begin{aligned}
\limsup_{k\to\infty}\sum_{i\in N}\overline{U}_i^k\left(\sigma^k\right) &= \limsup_{k\to\infty}\sum_{i\in N}\int_\Omega U_i^k\left(\sigma^k\left(\omega\right),\omega\right)dP^k\left(\omega\right)\\
&= \limsup_{k\to\infty}\sum_{i\in N}\int_\Omega u_i\left(\sigma^k\left(\omega\right)\right)dP^k\left(\omega\right)\\
&= \limsup_{k\to\infty}\int_A\left(\sum_{i\in N}u_i\left(a\right)\right)d\mu\left(\sigma^k\right)\left(a\right)\\
&\leq \int_A\left(\sum_{i\in N}u_i\left(a\right)\right)d\mu'\left(a\right) = \sum_{i\in N}\int_A u_i\left(a\right)d\mu'\left(a\right),
\end{aligned}$$

where the inequality follows from the Portmanteau theorem and the assumption (**c**) that $\sum_{i\in N}u_i\left(a\right)$ is upper semi-continuous. Thus,

$$\limsup_{k\to\infty}\sum_{i\in N}\overline{U}_i^k\left(\sigma^k\right)\leq\sum_{i\in N}\int_A u_i\left(a\right)d\mu'\left(a\right).$$

Combined with Ineq. (8), this leads to the conclusion that $\lim_{k\to\infty}\overline{U}_i^k\left(\sigma^k\right)$ exists and is equal to $\int_A u_i\left(a\right)d\mu'\left(a\right)$ for every $i\in N$. Therefore, according to Lemma 1, for any $i\in N$ and any measurable $\psi_i:A_i\to A_i$, the inequality $\int_A u_i\left(a\right)d\mu'\left(a\right)\geq\int_A u_i\left(\psi_i\left(a_i\right),a_{-i}\right)d\mu'\left(a\right)$ holds, which shows that μ' is indeed a CE of u.

We have thereby shown that any accumulation point of $\{\mu\left(\sigma^k\right)\}_{k=1}^\infty$ is a CE of u. Since $\widehat{\nu}$ has a unique CE and $M(A)$ is compact, the sequence $\{\mu\left(\sigma^k\right)\}_{k=1}^\infty$ in fact converges to $\widehat{\nu}$. As the latter is true for any such sequence, ν is strongly robust by Definition 1. ■

A.5 Proof of Theorem 2

Proof. The proof proceeds in the same way as the proof of Theorem 1. The only exception that needs to be made is in the proof of Lemma 1, the first paragraph of which we follow verbatim, establishing the fact that the inequality

$$\liminf_{k\to\infty}\overline{U}_i^k\left(\sigma^k\right)<\int_A u_i\left(\psi_i\left(a_i\right),a_{-i}\right)d\mu'\left(a\right). \tag{9}$$

holds for a continuous function $\psi_i:A_i\to A_i$. In what follows we will show that ψ_i can be modified in a way that the integrand in the right-hand term in Ineq. (9) is continuous μ'-almost everywhere.

By (**b**') and Fatou's lemma,

$$\begin{aligned}
&\lambda_i\liminf_{\varepsilon\to0+}\int_A u_i\left(\psi_i\left(a_i\right)-\varepsilon,a_{-i}\right)d\mu'\left(a\right)+\left(1-\lambda_i\right)\liminf_{\varepsilon\to0+}\int_A u_i\left(\psi_i\left(a_i\right)+\varepsilon,a_{-i}\right)d\mu'\left(a\right)\\
\geq\ &\int_A u_i\left(\psi_i\left(a_i\right),a_{-i}\right)d\mu'\left(a\right).
\end{aligned}$$

Assume, e.g., that

$$\liminf_{\varepsilon \to 0+} \int_A u_i \left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right) d\mu'\left(a\right) \geq \int_A u_i \left(\psi_i\left(a_i\right), a_{-i}\right) d\mu'\left(a\right) \tag{10}$$

(the arguments in the case where the inequality holds for $\psi_i\left(a_i\right)-\varepsilon$ instead of $\psi_i\left(a_i\right)+\varepsilon$ are symmetric). For any $j \neq i$, $d \in D(i)$, and $0 < \varepsilon < \delta$, the sets $A_{i,j,d}(\varepsilon) := \{a \in A \mid a_j = f_{ij}^d(\psi_i\left(a_i\right)+\varepsilon)\}$ are disjoint for different values of ε, and hence $\mu'(A_{i,j,d}(\varepsilon)) = 0$ for any ε outside some countable set. Since, as follows from (**a**'), the function $u_i\left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right)$ is continuous in a outside $\cup_{j\neq i, d\in D(i)} A_{i,j,d}(\varepsilon)$, this function is in fact μ'-almost everywhere continuous in a for any ε belonging to some vanishing sequence in $(0,\delta)$. By Ineq. (10), the function ψ_i can therefore be replaced in Ineq. (9) by some $\psi_i'(=\psi_i+\varepsilon) : A_i \to [\underline{a}_i, \overline{a}_i+\delta]$, for which $u_i\left(\psi_i'\left(a_i\right), a_{-i}\right)$ is μ'-almost everywhere continuous in a, and the inequality in (9) is preserved.

Now let $\psi_i'' := \min(\psi_i', \overline{a}_i)$. As in the proof of Lemma 1, we obtain

$$\lim_{k\to\infty} \overline{U}_i^k\left(\psi_i''\left(\sigma_i^k\right), \sigma_{-i}^k\right) = \lim_{k\to\infty} \int_A u_i\left(\psi_i''\left(a_i\right), a_{-i}\right) d\mu\left(\sigma^k\right)(a),$$

and, since $\overline{a}_i$ dominates all actions higher than $\overline{a}_i$ by assumption,

$$\lim_{k\to\infty} \overline{U}_i^k\left(\psi_i''\left(\sigma_i^k\right), \sigma_{-i}^k\right) \geq \lim_{k\to\infty} \int_A u_i\left(\psi_i'\left(a_i\right), a_{-i}\right) d\mu\left(\sigma^k\right)(a). \tag{11}$$

As $u_i\left(\psi_i'\left(a_i\right), a_{-i}\right)$ is μ'-almost everywhere continuous in a, the right-hand side in (11) is equal to $\int_A u_i\left(\psi_i'\left(a_i\right), a_{-i}\right) d\mu'\left(a\right)$ by the Portmanteau theorem, and so

$$\lim_{k\to\infty} \overline{U}_i^k\left(\psi_i''\left(\sigma_i^k\right), \sigma_{-i}^k\right) \geq \int_A u_i\left(\psi_i'\left(a_i\right), a_{-i}\right) d\mu'\left(a\right). \tag{12}$$

Ineq. (9) – which holds for ψ_i' – and Ineq. (12) imply that $\overline{U}_i^k\left(\psi_i''\left(\sigma_i^k\right), \sigma_{-i}^k\right) > \overline{U}_i^k\left(\sigma^k\right)$ for some k, in contradiction to the assumption that σ^k is a BNE of $\mathcal{U}^k$. This establishes the claim in Lemma 1 under conditions (**a**') and (**b**'), and the proof proceeds as that of Theorem 1 from this point onward. ■

A.6 Proof of Theorem 3

Proof. Denote by a^* the pure-action NE of the contest, whose existence and uniqueness was established in Szidarovszky and Okuguchi (1997). For any $a \in [0,1]^n$, define

$$H(a) := \sum_{i\in N}\left[u_i(a) - u_i(a_i^*, a_{-i})\right] = 1 - \sum_{i\in N}^{n} a_i - \sum_{i\in N}^{n} u_i(a_i^*, a_{-i}).$$

Clearly, $H(a^*) = 0$. As has been observed in Section 5.2, each $u_i(a_i^*, a_{-i})$ is a convex function of a_{-i}, which is also continuously differentiable whenever $a_{-i} \neq \mathbf{0}_{-i}$. It follows that H is concave on $[0,1]^n$ and continuously differentiable on $[0,1]^n \backslash \cup_{i\in N}\left([0,1]_i \times \{\mathbf{0}_{-i}\}\right)$.

Observe that at least two players exert positive effort in a^*, i.e., $a^* \notin \cup_{i\in N}\left([0,1]_i \times \{\mathbf{0}_{-i}\}\right)$, since otherwise player i, for whom $a^*_{-i} = \mathbf{0}_{-i}$, would have no best response against a^*_{-i}. As a consequence, H is differentiable at a^*.

We shall now prove that H is non-positive. For every player j and every action $a_j \in [0,1]$, we can evaluate $H(a_j, a^*_{-j})$ and get

$$\begin{aligned} H(a_j, a^*_{-j}) &= u_j(a_j, a^*_{-j}) - u_j(a^*_j, a^*_{-j}) + \sum_{i\in N\setminus\{j\}} \left[u_i(a_j, a^*_{-j}) - u_i(a^*_i, a_j, a^*_{-i,-j})\right] = \\ &= u_j(a_j, a^*_{-j}) - u_j(a^*_j, a^*_{-j}) \le 0, \end{aligned}$$

where the last inequality follows from the fact the a^* is an NE. Therefore a^* is a critical point of H, and, as the latter is differentiable at a^* and concave on $[0,1]^n$, the profile a^* is also a global maximizer of H, which implies that $H(a) \le H(a^*) = 0$ for every $a \in [0,1]^n$. Because H is non-positive, for every $a \in [0,1]^n$

$$\sum_{i\in N} u_i(a) \le \sum_{i\in N} u_i(a^*_i, a_{-i}). \tag{13}$$

Now consider any CE μ in the contest. The condition given in Ineq. (1) holds, in particular, for each $i \in N$ and the constant function $\psi(a_i) \equiv a^*_i$, i.e.,

$$\int u_i\left(a_i, a_{-i}\right) d\mu(a) \ge \int u_i\left(a^*_i, a_{-i}\right) d\mu(a). \tag{14}$$

The combination of Ineq. (13) and Ineq. (14) shows that, for every $i \in N$,

$$\int u_i\left(a_i, a_{-i}\right) d\mu(a) = \int u_i\left(a^*_i, a_{-i}\right) d\mu(a). \tag{15}$$

In words, every player i is indifferent between following the realized suggestion a_i of the CE μ and deviating to the pure NE action a^*_i.

Now assume that $\mu([0,1]^n \setminus \{a^*\}) > 0$. Then there exists $i \in N$ such that $\mu(\{a \mid a_i \ne a^*_i\}) > 0$. It cannot be that, conditional on $a_i \ne a^*_i$, the CE μ puts weight 1 on a set with $a_{-i} = \mathbf{0}_{-i}$, since otherwise

$$\psi^\varepsilon_i\left(a_i\right) = \begin{cases} a^*_i, & \text{if } a_i = a^*_i, \\ \varepsilon, & \text{otherwise,} \end{cases}$$

would violate Ineq. (1) for every sufficiently small $\varepsilon > 0$. It follows that

$$\mu(\{a \mid a_i \ne a^*_i \text{ and } a_{-i} \ne \mathbf{0}_{-i}\}) > 0. \tag{16}$$

Finally, consider a function $\psi_i : [0,1] \to [0,1]$ given by $\psi_i(a_i) = \frac{a_i + a^*_i}{2}$. It follows from Ineq.

(1) that

$$\begin{aligned}
\int u_i(a_i, a_{-i}) d\mu(a) &\geq \int u_i(\psi_i(a_i), a_{-i}) d\mu(a) \\
&= \int u_i\left(\tfrac{a_i + a_i^*}{2}, a_{-i}\right) d\mu(a) \\
&> \frac{1}{2}\int u_i(a_i, a_{-i}) d\mu(a) + \frac{1}{2}\int u_i(a_i^*, a_{-i}) d\mu(a) \\
&= \int u_i(a_i, a_{-i}) d\mu(a),
\end{aligned}$$

where the strict inequality follows from the strict concavity of u_i in a_i when $a_{-i} \neq \mathbf{0}_{-i}$ and Ineq. (16), and the last equality follows from (15). We have reached a contradiction, and therefore must conclude that any CE μ of the contest is a Dirac measure concentrated on the pure-action NE a^*. ■

A.7 Proof of Claim 4

Proof. Each function p_i is clearly continuous, and so $p = (p_i)_{i \in N}$ trivially satisfies (**a**) and (**c'**). Next, $p_i(a_i, a_{-i}) = \frac{h(a_i)}{r + \sum_{j \in N} h(a_j)}$ is convex in a_{-i} since h is concave and the function $\frac{1}{r+x}$ is decreasing and convex in $x \geq 0$ for any $r > 0$. Similarly, the sum

$$\sum_{i \in N} p_i(a) = \frac{\sum_{i \in N} h(a_i)}{\sum_{i \in N} h(a_i) + r}$$

is strictly concave since h is strictly increasing and concave, and the function $\frac{x}{r+x}$ is increasing and strictly concave in $x \geq 0$ for any $r > 0$. Thus u is an imperfectly discriminating contest that satisfies the assumptions of Corollary 2. ■