Economic Theory (2022) 73:91–119
https://doi.org/10.1007/s00199-020-01327-4

RESEARCH ARTICLE

Strong robustness to incomplete information and the uniqueness of a correlated equilibrium

Ezra Einy[1] · Ori Haimanko[1] · David Lagziel[1]

Received: 5 June 2020 / Accepted: 5 November 2020 / Published online: 12 November 2020

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, all-pay auctions, Cournot and Bertrand competitions, network games, patent races, voting problems and location games.

Keywords Strong robustness to incomplete information · Nash equilibrium · Correlated equilibrium

Mathematics Subject Classification C62 · C72 · D82

The authors wish to thank Eddie Dekel, Ehud Lehrer, Yehuda John Levy, Daisuke Oyama, David Schmeidler, Aner Sela, and the participants of SWET 2019 workshop held in Otaru, Japan, for their valuable comments. The authors' special grattitude goes to Atsushi Kajii, whose encouragement to put the first basic ideas into writing led to this work.

✉ David Lagziel
davidlag@bgu.ac.il

Ezra Einy
einy@bgu.ac.il

Ori Haimanko
orih@bgu.ac.il

1 Department of Economics, Ben-Gurion University of the Negev, Beer Sheva, Israel

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 effect produced when the two concepts happen to coincide on 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 $\mathbf{u}$ is strongly robust if, under uncertainty about the individual payoffs but with each player knowing that his payoffs are those in $\mathbf{u}$ with high probability, *all* Bayesian–Nash equilibria in the incomplete-information game are sufficiently close (in terms of the induced action distribution) to the equilibrium in $\mathbf{u}$. 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 strate-

[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.

gic 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 of 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, all-pay auctions, 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 finite games,[5] the uniqueness of a correlated equilibrium[6] implies its strong robustness also in our setting with a continuum of actions, albeit requiring certain continuity conditions on payoffs and necessitating a new, non-trivial proof.[7] Strong robustness of a Nash equilibrium is, moreover, *equivalent* to its being the unique correlated equilibrium in the game.

[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).

[5] Proposition 3.2 in KM establishes KM-robustness of a unique correlated equilibrium in any finite game, but its proof implicitly shows strong robustness of such an equilibrium.

[6] 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.

[7] Although our proofs in the continuum setting are quite involved, they use arguments that would have been rather straightforward had finite action sets been assumed instead (notice, e.g. that the lengthy Lemma 3 in the proof of Theorem 1 would have then become nearly trivial).

This equivalence result puts the notion of strong robustness in a proper light, *a posteriori*. Unlike KM-robustness, the equilibrium refinement based on strong robustness does not provide an equilibrium-selection tool because, as it turns out, such refinements only exist in games where the selection problem is trivial anyway (on account of there being a unique Nash equilibrium that is also the unique correlated one). Yet, strong robustness is a powerful, always-present feature of a unique correlated equilibrium. From a practical perspective, our 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 $\mathbf{u}$ 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 $\mathbf{u}$.

The motivation for allowing a continuum of actions to be available to each player in our framework[8] comes partially from the fact that the existing results on the uniqueness of a correlated equilibrium are mostly in the continuum setting.[9] 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.[10] 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 (contest), and every equivalent patent race, have a unique correlated equilibrium.[11] 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

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

[9] In the finite setting, KM mention two classes of games with a unique correlated equilibrium: two-player zero-sum games with a unique optimal strategy for each player, and dominance-solvable games.

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

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

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; this implies correlated equilibrium uniqueness in all-pay auctions, 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 Sect. 2 we present the basic complete-information framework, and extend it to incomplete information. In Sect. 3 we define and explain the notion of strong robustness to incomplete information. In Sect. 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 Sect. 5 we survey the games for which a correlated equilibrium is known to be unique and state uniqueness results of our own. All proofs are given in the "Appendix".

2 Preliminaries

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

2.1 Games with a continuum of pure actions

Fix a finite set of players $N = \{1, 2, \ldots, n\}$. The set A_i of (pure) actions of each player i is assumed to be a compact and full-dimensional.[12] 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 $\mathbf{u}$ will be identified with an n-tuple $(u_i)_{i\in N}$ where $u_i : A \to \mathbb{R}$ is the payoff function of player i.

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.[13] In general, for any product set $C = \times_{i\in N} C_i$ and any $j \in N$,

[12] 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.

[13] 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 $f : B \to \mathbb{R}$.

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 $\mathbf{u}$ if, for any player i and any Borel-measurable function $\psi_i : A_i \to A_i$,

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

In fact, if μ is a CE, then Ineq. (1) holds for any Borel-measurable function[14] $\psi_i : A_i \to M(A_i)$, with $u_i\left(\psi_i\left(a_i\right), a_{-i}\right)$ being defined as $\int_{A_i} u_i\left(b_i, a_{-i}\right) d\psi_i\left(a_i\right)\left(b_i\right)$ 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

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

A mixed-action profile ν is a *Nash equilibrium* (henceforth, NE) of $\mathbf{u}$ if $\widehat{\nu}$ is a CE. This is equivalent to the requirement that $u_i\left(\nu\right) \geq u_i\left(a_i, \nu_{-i}\right)$ 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 $(\Omega, \mathcal{F})$, 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.[15] The payoffs to player i are determined by a state-dependent payoff function $U_i : A \times \Omega \to \mathbb{R}$ that is $\mathcal{B}(A) \otimes \mathcal{F}$-measurable, where $\mathcal{B}(A)$ stands for the σ-field of Borel subsets of A. 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

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

[15] If $\mathcal{F}_i$ is generated by a finite or countable partition Π_i of Ω (in which case any $E \in \mathcal{F}_i$ is a union of finitely or countably many elements of Π_i), then player i knows the exact partition element $\pi_i(\omega) \in \Pi_i$ that contains the realized state of nature ω, and thus Π_i can be viewed as the collection of *information types* of player i. In fact, the present framework of incomplete information can accommodate a general type space. Indeed, consider a state space $(\Omega, \mathcal{F}, \mu)$ in which Ω is a product set $\Theta \times (\times_{i \in N} T_i)$, $\mathcal{F}$ is a tensor-product σ-field $\Xi \otimes (\otimes_{i \in N} \Upsilon_i)$, and, for each $i \in N$, $\mathcal{F}_i$ is the inverse image of Υ_i under the projection $\left(\theta, (t_i)_{i \in N}\right) \mapsto t_i$ on Ω. This corresponds to each player i having a type, or signal, space (T_i, Υ_i), and Θ being a set of payoff-relevant parameters.

payoff functions 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).[16]

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

$$\mu(\sigma)(B) = \int_\Omega \widehat{\sigma(\omega)}(B) \, dP(\omega)$$

for every $B \in \mathcal{B}(A)$.

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 $\mathbf{u}$. We will consider an incomplete-information game $\mathcal{U}$ as being close to $\mathbf{u}$ if, with high probability, each player i knows that his payoff in $\mathcal{U}$ is given by u_i. Formally, for any $\epsilon \geq 0$, an incomplete-information game $\mathcal{U}$ is said to be an ϵ*-elaboration* of $\mathbf{u}$ if for every player i there exists an event $\Omega_i(\mathcal{U}, \mathbf{u}) \in \mathcal{F}_i$ such that

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

and $P(\cap_{i \in N} \Omega_i(\mathcal{U}, \mathbf{u})) = 1 - \epsilon$. 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 ϵ-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 $\mathbf{u}$, 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 an ϵ_k -elaboration of $\mathbf{u}$ that possesses some *BNE* σ^k and $\lim_{k \to \infty} \epsilon_k = 0$, 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.

[16] 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.

In other words, an NE of $\mathbf{u}$ 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 $\mathbf{u}$.

Notice that the definition allows us 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 $\mathbf{u}$ 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 $\{\mu\left(\sigma^k\right)\}_{k=1}^{\infty}$ 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 furthest 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.[17] 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 $\mathbf{u}$. 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 $\mathbf{u}$ 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* $\mathbf{u}$ *possesses a strongly robust* NE ν, *then* $\widehat{\nu}$ *is the unique* CE *of* $\mathbf{u}$. *In particular, if a strongly robust* NE *exists, then it is unique.*

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.

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

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

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.[18]

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 corresponding 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, as the following lemma:

Lemma 1 *If* **u** *satisfies* (**a**), (**b**), *and* (**c**), *then it possesses a mixed-action* NE.

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.

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. For this extension, we shall confine ourselves to one-dimensional action sets, i.e., assume that each A_i is a closed non-degenerate interval $\left[\underline{a}_i, \overline{a}_i\right] \subset \mathbb{R}$. 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 $A^+ = \cup_{i \in N} \left[\underline{a}_i - \delta, \overline{a}_i + \delta\right] \times A_{-i}$ 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

[18] 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 $\limsup_{k \to \infty} \sum_{i \in N} u_i\left(a^k\right) \leq \sum_{i \in N} u_i(a)$, for any sequence $\{a^k\}_{k=1}^{\infty} \subset A$ that converges to a.

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 monotone 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 $A^*(i)$, as defined above, for some given sets $D(i)$ and functions f_{ij}^d.

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 $0 < \lambda_1, \ldots, \lambda_n < 1$ such that, for every $i \in N$ and $a \in A$,

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

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.

Note that Theorems 1 and 2 complement each other in terms of generality: Theorem 1 admits higher-dimension actions sets, whereas Theorem 2 allows broader discontinuity features.

We now comment on the proofs of Theorems 1 and 2, given in the "Appendix". Because the "only if" direction of both theorems is settled by Proposition 1, it is the "if" direction that needs attention. The proofs follow an expected route, established in Proposition 3.2 of KM concerning the KM-robustness of a unique CE in finite games. It consists of showing that any limit point of action distributions of BNE in a sequence of incomplete information elaborations that approximate **u** is a CE of **u**; but, since the CE is unique, all limit action distributions must be the same, and thus the BNE action distributions *converge* to that CE, establishing strong robustness of the underlying NE. However, unlike in finite games, in games with a continuum of actions the convergence of BNE action distributions to a limit point does not necessarily imply convergence of expected payoffs (including those obtained by unilateral deviations), on account of possible discontinuity of the payoff functions and of deviation methods. This makes checking the fact that any limit point is a CE quite challenging technically, and establishing that fact is what the proofs of Theorems 1 and 2 are mainly dedicated to.

Remark 1 Occasionally, a weighted version of condition *(* **c***)* will be useful, requiring upper semi-continuity on A of the function $\sum_{i\in N} \omega_i u_i$ for some collection $(\omega_i)_{i\in N}$ of fixed, positive weights. This version can replace *(* **c***)* in both our theorems, requiring only self-evident changes in the proofs.

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 Sect. 5.1 we review known results concerning "smooth" games—essentially 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*, strictly concave potential games, a subclass of Cournot oligopolies with differentiated products, and some types of network games. Attention will also be drawn to Bertrand oligopolies with differentiated products, where the CE is unique by the supermodularity arguments of Milgrom and Roberts (1990) for some specifications of the demand functions.

Sections 5.2, 5.3 and 5.4 concern games that are not necessarily smooth. Contrary to Sect. 5.1, many games therein require varying degrees of analysis beyond applying ready-made results, in order to prove the uniqueness of a CE. Specifically, Sect. 5.2 is dedicated to Tullock contests, in which the payoff functions are not differentiable (or even continuous) 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,[19] 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 deduce CE uniqueness in socially strictly concave games from the result of Hart and Mas-Colell (2015). 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 considers some two-player games with major discontinuities and provides an argument for the uniqueness of their CE (whose strong robustness is established by appealing to a more potent Theorem 2). Among these games are two types of auctions and two classical constant-sum games (namely, the median-voter problem and a pure-location Hotelling game).

5.1 Smooth games

Following Ui (2008), a game $\mathbf{u}$ is called *smooth* if every payoff function u_i is continuously differentiable in $a_i \in A_i$; we will additionally impose the assumption that each u_i is continuous on A. Ui's work generalized the results of Rosen (1965) on NE uniqueness in smooth games to CE uniqueness, with one sufficient condition being the following. A game $\mathbf{u}$ is said to have *strictly monotone payoff gradients* (henceforth, SMPG) if, for every $a \neq a'$,[20]

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

[19] While the uniqueness of an NE is Tullock contests is well known, our result on CE uniqueness is new.

[20] In the following inequality, ∇_i denotes the gradient of u_i as a function of a_i.

An easy way to verify the SMPG condition, assuming that all payoff function are twice continuously differentiable, is to consider the matrix $\left[\partial^2 u_i\left(a\right)/\partial a_i \partial a_j\right]$; if it is negative definite then **u** has SMPG.[21]

The main result of Ui (2008) shows that any smooth game with SMPG has a unique CE (which is also the unique pure-action NE).[22] Since our continuity conditions (**a**)–(**c**) hold trivially for any smooth game, the following is a corollary of Theorem 1.

Corollary 1 *Any smooth game with* SMPG *has a unique pure-action* NE *which is also the unique* CE,[23] *and therefore is strongly robust.*

In what follows we review some examples of smooth games. The majority are known to have SMPG. One exception is a class of Bertrand oligopolies considered in Sect. 5.1.4, where the CE uniqueness is based on a log-supermodularity argument due to Milgrom and Roberts (1990) instead of the SMPG condition.

5.1.1 Strictly 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 $a_i' \neq a_i$, $u_i(a_i', a_{-i}) - u_i(a) = P(a_i', a_{-i}) - P(a)$. Among the best-known smooth potential games is a single-product Cournot oligopolies with linear demand and costs. For a smooth potential game, its potential is *strictly concave* if and only if the game has SMPG (by Lemma 4 of Ui 2008), in which case Corollary 1 applies.[24]

5.1.2 Network games

In *network games,* each player's action set is a compact interval, and the payoffs depend on the network (i.e., a graph) that links different players to one another. Ui (2016) lists in his Section 5.3 several such games, in particular those from Ballester et al. (2006) and Bramoullé et al. (2014), whose payoffs have the following quadratic form: for each $i \in N$ and $a \in A$,

$$u_i\left(a\right) = \theta_i a_i - \frac{1}{2} q_{ii} a_i^2 - a_i \sum_{j \neq i} q_{ij} a_j, \tag{5}$$

[21] See Lemma 3 and the proof of Corollary 6 in Ui (2008). A necessary condition for SMPG is strict concavity of each u_i in a_i, by Lemma 5 in Ui (2008).

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

[23] We identify a (pure) action profile a with the action distribution supported on a. Thus, in interpreting a pure-action NE as a CE there is no need to add that a as is viewed as the corresponding degenerate "action distribution".

[24] One may, alternatively, use the result of Neyman (1997), who was the first to observe CE uniqueness in games with a continuously differentiable and strictly concave potential, and then deduce strong robustness by applying Theorem 1.

where $\theta_i > 0$ and $Q = [q_{ij}]$ is an $n \times n$-matrix accounting for the network links. Ui (2016) states conditions on the corresponding models that ensure positive definiteness of Q.[25] For such a Q, by an observation following (4) the game $\mathbf{u}$ has SMPG. It is clearly smooth, and hence Corollary 1 applies.

5.1.3 Cournot oligopoly with differentiated products

In a Cournot oligopoly model with differentiated products, each firm $i \in N$ chooses a non-negative output level a_i of its product from some compact interval $A_i \subset \mathbb{R}_+$, expecting to obtain the price $P_i(a) \geq 0$ given any output-profile a, and incurring a production cost $c_i(a_i) \geq 0$. Assume further that the two functions P_i and c_i are twice continuously differentiable and that c_i is convex. As shown in Example 1 of Ui (2008), the game $\mathbf{u}$ has SMPG whenever the matrix $\left[\frac{\partial^2 P_i(a)a_i}{\partial a_i \partial a_j}\right]$ is negative-definite for each $a \in A$, in which case Corollary 1 applies.

5.1.4 Supermodularity and Bertrand oligopoly with differentiated products

In their study of supermodular games, Milgrom and Roberts (1990) showed that pure-action NE provide bounds on the sets of rationalizable actions. Specifically, their Theorem 5 states that in a supermodular game the set of serially undominated actions[26] of each player is bounded from above (and, respectively, from below) by the largest (respectively, the smallest) pure NE action of that player. Since no CE can, with positive probability, prescribe strictly serially dominated actions to any player, an obvious corollary is that uniqueness of a pure-action NE in a supermodular game implies its being a unique CE. The same claims also apply to games that become supermodular after a strictly increasing transformation of payoffs, such as log-supermodular games.[27]

In their Section 4(2), Milgrom and Roberts (1990) considered smooth Bertrand oligopolies with differentiated products that are substitutes and constant marginal costs, providing a condition (the elasticity of each demand being a non-increasing function of the other firms' prices) for its log-supermodularity. They noted that the condition holds, and that the pure-action NE is unique, for several demand types (including linear, CES and logit demand functions). With the oligopoly being smooth by assumption, our continuity conditions (**a**)–(**c**) hold trivially, and the unique NE is strongly robust by Theorem 1.

---

25 For a non-symmetric Q, positive definiteness is defined as that of $Q + Q^T$. If Q is symmetric, then the game $\mathbf{u}$ has a smooth strictly concave potential, but it is not a potential game otherise.

26 Serially undominated actions are those that survive the iterative process of eliminating strongly dominated actions.

27 That is because the sets of serially undominated actions and of pure-action NE are detemined only by ordinal comparisons.

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.[28] The Tullock rent-seeking game (see Tullock 2001), or Tullock contest, is a complete-information game $\mathbf{u}$ 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 contestant 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 $\mathbf{0}$,

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

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,[29] which implies its strong robustness via Theorem 1.

Theorem 3 *The pure-action* 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

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

[29] A recent work by Ewerhart and Quartieri (2020) establishes the uniqueness of a BNE in (a generalization of) the Szidarovszky and Okuguchi model with incomplete information, where each player has finitely many information types. This implies that, among finitely-supported action distributions in a complete-information Tullock contest, its NE is the unique CE, but does not settle the question of CE uniqueness among all action distributions.

properties. In addition, each u_i is continuously differentiable on $\mathbb{R}^n_+ \setminus \{\mathbf{0}\}$. Since the sum $\sum_{i \in N} u_i(a) = 1 - \sum_{i \in N} a_i$ 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 relies on the fact of there being a pure-action NE. Its main step is to show that every player is indifferent between following a CE recommended-action and deviating to his NE action; this part of the proof uses a method akin to the one employed by Liu (1996) in showing CE uniqueness for linear Cournot oligopolies. This indifference, combined with strict concavity of each player's payoff in his own action, then pinpoints the CE as a Dirac measure concentrated on the pure-action NE.

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 Sect. 5.3.2), single product Cournot oligopolies with linear demand (see Sect. 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 $\mathbf{u}$ 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 every u_i is (strictly) 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 Lemma 1, conditions (**a**)–(**c**) imply the existence of an NE, which is in particular a CE. (The NE is evidently in pure actions, due to the strict concavity of each player's payoffs in his own action.) This leads to the following Corollary:

Corollary 2 *Any socially strictly concave game* $\mathbf{u}$ *that satisfies* (**a**)*,* (**b**)*, and* (**c**) *has a pure-action* 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 $\mathbf{u}$ 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 2 since they do not have a strictly concave sum of payoffs, and we must rely on Theorem 3 for the CE uniqueness result. However, the following family of imperfectly discriminating contests does satisfy the conditions of Corollary 2.

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[30] 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 $\mathbf{u}$ 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 2 applies to imperfectly discriminating contests when they are strictly socially concave:

Corollary 3 *An imperfectly discriminating contest* $\mathbf{u}$ *has a pure-action* 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 3.

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 make a discovery by the 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 $\mathbf{u}$ is included in the scope of Corollary 3 even if h is merely continuous.

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

[30] 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).

5.3.3 Single-product Cournot oligopoly

Consider a single-product Cournot oligopoly model with linear demand, which is given by the description in Sect. 5.1.3 using a (common) inverse demand function $P(a) = \mathbf{B} - \mathbf{A}\left(\sum_{i \in N} a_i\right)$, where $\mathbf{A}, \mathbf{B} > 0$. Now discard the assumption of continuous differentiability of costs; assume instead that each $c_i : A_i \to \mathbb{R}_+$ is merely continuous and convex. 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 $\mathbf{u}$ is clearly socially strictly concave. Since it trivially satisfies (**a**) and (**c'**), Corollary 2 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).[31] 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),[32] that the unique pure-action Cournot equilibrium is also a unique CE.

5.4 Some two-player games with major discontinuities

In this section we will show the existence of a unique, possibly strictly mixed, CE in three classes of discontinuous and zero-sum-equivalent games with two players: all-pay auctions (see Sect. 5.4.1), median-voter problems (see Sect. 5.4.2) and pure-location Hotelling games (see Sect. 5.4.3). In these games, the discontinuity in payoffs occurs along an entire diagonal line, 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 non-zero-sum-equivalent game (specifically, a common-value first-price auction) where the payoffs are discontinuous along the diagonal but a unique CE is known to exist.

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

Lemma 2 *If a two-player zero-sum game has a unique* NE*, then its induced action distribution is the unique* CE *in the game.*

If $\mathbf{u}$ is strategically equivalent (in mixed actions) to a zero-sum game, Lemma 2 is also applicable to such $\mathbf{u}$, i.e., the induced action distribution of a unique NE is the unique CE of the game. In particular, Lemma 2 is applicable to two-player constant-sum games.

5.4.1 Two auction types

Consider an *all-pay* auction with two players, in which player 1 values the auctioned object at $V_1 > 0$ and player 2 values it at $V_2 > 0$, where (w.l.o.g.) $V_2 \leq V_1$. Both

[31] 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 $QP(Q) - \min_{i=1,2} c_i(Q) < 0$ for every $Q > \overline{Q}$.

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

players submit bids in $A_1 = A_2 = [0, V_1]$. Each bid is paid, and the object is awarded to the highest bidder (with a symmetric tie-breaking rule). Each payoff function u_i is given, for every $a \in A$, by

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

It is well known that this auction has a unique NE in which player 1's bids have the uniform distribution on $[0, V_2]$, whereas player 2 bids uniformly on $[0, V_2]$ with probability $\frac{V_2}{V_1}$ and submits the zero bid with the complementary probability (see, e.g., Proposition 2 in Hillman and Riley 1989). Also, **u** is strategically equivalent (in mixed actions) to a zero-sum game **w** with $w_i(a) = \frac{1}{V_i} u_i(a) + \frac{1}{V_{-i}} a_{-i} - \frac{1}{2}$ for every $i = 1, 2$ and $a \in A$.[33] By Lemma 2 and the comment following it, the unique mixed NE is the unique CE of **u**. Conditions (**a'**) and (**b'**) obviously hold for the game **u**,[34] while a weighted version of (**c**) is satisfied: $\frac{1}{V_1} u_1(a) + \frac{1}{V_2} u_2(a) = 1 - \frac{1}{V_1} a_1 - \frac{1}{V_1} a_2$ is a continuous (and in particular upper semi-continuous) function on A. Thus, the unique NE is strongly robust by Theorem 2 and Remark 1.

The above class of auctions is important because in all of them the strongly robust NE is comprised of *strictly mixed* actions. That stands in contrast to games in Sects. 5.1, 5.2 and 5.3, where the strongly robust NE are in pure actions; the latter is, however, primarily an artifact of the strict concavity of each player's payoff in his own action in most of those games.

We end this discussion by considering a *first-price* auction with *common values*, in which $V_1 = V_2 = 1$ (the second equality is assumed w.l.o.g.) and only the highest bidder pays his bid. Accordingly, the modified payoffs are given, for every $i = 1, 2$ and $a \in A$, by

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

While this auction is not strategically equivalent to any zero-sum game and thus Lemma 2 does not apply, it was shown in Section 3.1 of Dütting et al. (2014) by using direct arguments that both players bid 1 in a unique (pure-action) NE and CE. As before, it is easy to check conditions (**a**') and (**b** '); the sum $u_1(a) + u_2(a) = 1 - \min\{a_1, a_2\}$ is obviously continuous and so (**c**) is satisfied as well. The unique NE is therefore strongly robust by Theorem 2.

[33] This was observed already by Pavlov (2013), but he stopped short of showing CE uniqueness, claiming instead CE's payoff-equivalence to the unique NE.

[34] Condition (**a'**) holds for $f_{ij}^d(a_i) = a_i$ with $|D(i)| = 1$, $i = 1, 2$, and (**b'**) holds for $\lambda_1 = \lambda_2 = 0$.

5.4.2 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 (2000) (Section 3.2, pp. 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 $\mathbf{u}$.

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 $\mathbf{u}$, and therefore (m, m) is its unique, pure-action NE. By Lemma 2, that NE is also the unique CE. Moreover, one can easily verify that conditions (**a'**), (**b'**), and (**c**) hold in this framework,[35] and so the unique pure NE is strongly robust by Theorem 2.

5.4.3 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 along [0, 1]; the continuous distribution of

[35] 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.

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 $\mathbf{u}$ may thus be described in terms of the functions W_i from Sect. 5.4.2 : $u_i(a) = W_i(a)$ 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 $\mathbf{u}$ has a unique, pure-action NE, in which both firms choose the location of the median customer m. By Lemma 2 and the comment following it, that NE is also the unique CE of the game. As conditions (**a'**), (**b'**), and (**c**) hold for the game $\mathbf{u}$, the unique NE is strongly robust by Theorem 2.

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 $a_i \mapsto \int_{A_{-i}} u_i\left(\psi_i'(a_i), a_{-i}\right) d\mu(a_{-i} \mid a_i)$ are Borel-measurable. Hence the graph of the (non-empty-valued) correspondence

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

is also Borel-measurable. By the measurable choice theorem, there exists a measurable $\psi_i : A_i \to A_i$ such that $\psi_i(a_i) \in \Psi_i(a_i)$ for μ_i-almost every $a_i \in A_i$.

Clearly,

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

for μ_i-almost 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

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

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 $\mathbf{u}$. Consider a 0-elaboration $\mathcal{U}_{0,\mu'}$ of $\mathbf{u}$ 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 $\mathbf{u}$, and hence it is the unique CE. □

A.3 Proof of Lemma 1

Proof Assume that $\mathbf{u}$ 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 neighbourhood of ν_{-i}.) Given the payoff-security of the mixed-strategy extension of $\mathbf{u}$, and condition (**c**) on pure-strategy payoffs, the existence of a mixed-strategy NE in $\mathbf{u}$ 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 $\{\mathcal{U}^k\}_{k=1}^{\infty}$ off incomplete information games and a sequence of corresponding BNE $\{\sigma^k\}_{k=1}^{\infty}$ such that each $\mathcal{U}^k = \left\{\Omega^k, P^k, \left\{\mathcal{F}_i^k\right\}_{i\in N}, \left\{U_i^k\right\}_{i\in N}\right\}$ is an ϵ_k-elaboration of $\mathbf{u}$ and $\lim_{k\to\infty} \epsilon_k = 0$. We will show that, for any subsequence of $\{\mu\left(\sigma^k\right)\}_{k=1}^{\infty} \subset M(A)$ that converges to some $\mu' \in M(A)$, the limit μ' is a CE of $\mathbf{u}$. W.l.o.g., we will take such a subsequence to be $\{\mu\left(\sigma^k\right)\}_{k=1}^{\infty}$ itself in our forthcoming considerations.

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

Lemma 3 *For any $i \in N$ and any measurable function $\psi_i : A_i \to A_i$,*

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

Proof of Lemma 3 Suppose to the contrary that, for some $i \in N$ and some measurable $\psi_i : A_i \to A_i$,

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

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 $\psi_i \mid_{E_\varepsilon}$, 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 $\partial\left(A_i\right)$, i.e., that $\psi_i : A_i \to A_i \backslash \partial\left(A_i\right)$. 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 $\partial\left(A_i\right)$ 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

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

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

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

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.[36] 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 3 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 $\mathbf{u}$ on a set with a $\mu\left(\sigma^k\right)$-

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

measure tending to 1) we obtain

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

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

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

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 $\int_A u_i(a)\,d\mu'(a)$ for every $i\in N$. Therefore, according to Lemma 3, for any $i\in N$ and any measurable $\psi_i : A_i\to A_i$, the inequality $\int_A u_i(a)\,d\mu'(a)\ge\int_A u_i(\psi_i(a_i),a_{-i})\,d\mu'(a)$ holds, which shows that μ' is indeed a CE of $\mathbf{u}$.

We have thereby shown that any accumulation point of $\{\mu\left(\sigma^k\right)\}_{k=1}^{\infty}$ is a CE of $\mathbf{u}$. 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

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

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,

$$\begin{aligned}&\lambda_i \lim \inf_{\varepsilon\to 0+} \int_A u_i\left(\psi_i\left(a_i\right)-\varepsilon, a_{-i}\right) d\mu'(a)\\&\quad+\left(1-\lambda_i\right) \lim \inf_{\varepsilon\to 0+} \int_A u_i\left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right) d\mu'(a)\\&\geq \int_A u_i\left(\psi_i\left(a_i\right), a_{-i}\right) d\mu'(a).\end{aligned}$$

Assume, e.g., that

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

(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).

Given the set $D(i)$ and functions $\left\{f_{ij}^d\right\}_{j\neq i, j\in N}$ whose existence is postulated in condition (**a**'), for any $j \neq i$, $d \in D(i)$ and $0<\varepsilon<\delta$ consider the set $A_{i,j,d}(\varepsilon) := \{a \in A \mid a_j = f_{ij}^d(\psi_i\left(a_i\right)+\varepsilon)\}$. As each f_{ij}^d is strictly monotone, the sets $A_{i,j,d}(\varepsilon)$ are disjoint for different values of ε, and hence $\mu'(A_{i,j,d}(\varepsilon)) = 0$ for any ε outside some countable set. It follows from (**a**') that the function $u_i\left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right)$ is continuous in a outside $\cup_{j\neq i, d\in D(i)} A_{i,j,d}(\varepsilon)$.[37] Thus, $u_i\left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right)$ 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 $\psi_i'(=\psi_i+\varepsilon) : A_i \to \left[\underline{a}_i, \overline{a}_i+\delta\right]$, 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 Theorem 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,

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

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 $\int_A u_i\left(\psi_i'\left(a_i\right), a_{-i}\right) d\mu'(a)$ by the Portmanteau theorem, and so

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

[37] Indeed, given any $a^0 \notin \cup_{j\neq i, d\in D(i)} A_{i,j,d}(\varepsilon)$, consider $a' = \left(\psi_i\left(a_i^0\right)+\varepsilon, a_{-i}^0\right) \in A$. It follows from the definition of each $A_{i,j,d}(\varepsilon)$ that $a_j' \neq f_{ij}^d(a_i')$ for any $j \neq i$ and $d \in D(i)$, and hence u_i is continuous at a' by condition (**a**'). The function $u_i\left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right)$ is therefore continuous at a^0 as a composition of the continuous function $a \mapsto \left(\psi_i\left(a_i\right)+\varepsilon, a_{-i}\right)$ and u_i that is continuous at a'.

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 Sect. 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 $[0, 1]^n \setminus \cup_{i \in N}$ $([0, 1]_i \times \{\mathbf{0}_{-i}\})$.

Observe that at least two players exert positive effort in a^*, i.e., $a^* \notin \cup_{i \in N}([0, 1]_i \times \{\mathbf{0}_{-i}\})$, 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

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

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.,

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

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

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

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 $\mu(\{a \mid a_i \neq a_i^*\}) > 0$. It cannot be that, conditional on $a_i \neq a_i^*$, the CE μ puts weight 1 on a set with $a_{-i} = \mathbf{0}_{-i}$, since otherwise

$$\psi_i^{\varepsilon}(a_i) = \begin{cases} a_i^*, & \text{if } a_i = a_i^*, \\ \varepsilon, & \text{otherwise,} \end{cases}$$

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

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

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

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

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 3. □

A.8 Proof of Lemma 2

Proof Consider a two-player zero-sum game **u** with a unique NE, which is then also the unique pair of optimal actions. If μ is a CE of **u** then it is easy to see that, for almost every action recommendation a_i to player i, the conditional distribution $\mu(\cdot \mid a_i)$ on A_{-i} is the optimal action of the player $-i$. Indeed, had the conditional distribution not been almost always optimal, player i would have had profitable deviations from his recommendation on a positive-probability set of actions.[38] Thus, (almost) all conditional distributions of μ on A_{-i} are identical, implying that μ is the product distribution of the (pure or mixed) actions comprising the NE. This shows that the CE is unique. □

References

Amir, R.: Cournot oligopoly and the theory of supermodular games. Games Econ. Behav. **15**, 132–148 (1996)

Aumann, R.J.: Subjectivity and correlation in randomized strategies. J. Math. Econ. **1**, 67–96 (1974)

Ballester, C., Calvó-Armengol, A., Zenou, Y.: Who's who in networks. Wanted: the key player. Econometrica **74**, 1403–1417 (2006)

Baye, M.R., Hoppe, H.C.: The strategic equivalence of rent-seeking, innovation, and patent-race games. Games Econ. Behav. **44**, 217–226 (2003)

Ben-Porat, O., Tennenholtz, M.: Multi-unit facility location games. In: WINE 2016: The 12th Conference on Web and Internet Economics (2016)

Bertsekas, D.P., Shreve, S.E.: Stochastic Optimal Control: The Discrete Time Case. Athena Scientific, New York (2004)

Borkar, V.S.: Probability Theory: An Advanced Course. Springer, Berlin (1995)

Bramoullé, Y., Kranton, R., D'Amours, M.: Strategic interaction and networks. Am. Econ. Rev. **104**, 898–930 (2014)

Carlsson, H., van Damme, E.: Global games and equilibrium selection. Econometrica **61**, 989–1018 (1993)

Dasgupta, P., Maskin, E.: The existence of equilibrium in discontinuous economic games, Part I (theory). Rev. Econ. Stud. **53**, 1–26 (1986)

Dasgupta, P., Stiglitz, J.: Uncertainty, industrial structure, and the speed of R&D. Bell J. Econ. **11**, 1–28 (1980)

Dekel, E., Fudenberg, D.: Rational behavior with payoff uncertainty. J. Econ. Theory **52**, 243–267 (1990)

Dixit, A.: Strategic behavior in contests. Am. Econ. Rev. **77**, 891–898 (1987)

Dütting, P., Kesselheim, T., Tardos, É.: Mechanism with unique learnable equilibria. In: Proceedings of the Fifteenth ACM Conference on Economics and Computation—EC'14, pp. 877–894. ACM Press, New York (2014)

Even-dar, E., Mansour, Y., Nadav, U.: On the convergence of regret minimization dynamics in concave games. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing—STOC'09, pp. 523–532. ACM Press, New York (2009)

Ewerhart, C., Quartieri, F.: Unique equilibrium in contests with incomplete information. Econ. Theory **70**, 243–271 (2020)

Foster, D.P., Vohra, R.V.: Calibrated learning and correlated equilibrium. Games Econ. Behav. **21**, 40–55 (1997)

[38] A single-valued *measurable* deviation function ψ_i that violates Ineq. (1) can then be chosen using a method indicated in the proof of Proposition 2.

Fudenberg, D., Kreps, D.M., Levine, D.K.: On the robustness of equilibrium refinements. J. Econ. Theory **44**, 354–380 (1988)
Gale, D., Klee, V., Rockafellar, R.T.: Convex functions on convex polytopes. Proc. Am. Math. Soc. **19**, 867–873 (1968)
Goodman, J.C.: Note on existence and uniqueness of equilibrium points for concave N-Person games. Econometrica **48**, 251–251 (1980)
Hart, S., Mas-Colell, A.: A simple adaptive procedure leading to correlated equilibrium. Econometrica **68**, 1127–1150 (2000)
Hart, S., Mas-Colell, A.: Markets, correlation, and regret-matching. Games Econ. Behav. **93**, 42–58 (2015)
Hart, S., Schmeidler, D.: Existence of correlated equilibria. Math. Oper. Res. **14**, 18–25 (1989)
Hillman, A.L., Riley, J.G.: Politically contestable rents and transfers. Econ. Polit. **1**, 17–39 (1989)
Hotelling, H.: Stability in competition. Econ. J. **39**, 41 (1929)
Kajii, A., Morris, S.: The robustness of equilibria to incomplete information. Econometrica **65**, 1283–1309 (1997)
Liu, L.: Correlated equilibrium of Cournot oligopoly competition. J. Econ. Theory **68**, 544–548 (1996)
Loury, G.C.: Market structure and innovation. Q. J. Econ. **93**, 395 (1979)
Milgrom, P.R., Roberts, J.: Rationalizability, learning, and equilibrium in games with strategic complementarities. Econometrica **58**, 1255–1277 (1990)
Milgrom, P., Shannon, C.: Monotone comparative statics. Econometrica **62**, 157 (1994)
Monderer, D., Shapley, L.S.: Potential games. Games Econ. Behav. **14**, 124–143 (1996)
Neyman, A.: Correlated equilibrium and potential games. Int. J. Game Theory **26**, 223–227 (1997)
Papadimitriou, C.H., Roughgarden, T.: Computing correlated equilibria in multi-player games. J. ACM **55**, 1–29 (2008)
Pavlov, G.: Correlated Equilibria and Communication Equilibria in All-Pay Auctions. Research Report # 2013-2 of Western University (2013). https://ir.lib.uwo.ca/cgi/viewcontent.cgi?article=1255&context=economicsresrpt
Persson, T., Tabellini, G.E.: Political Economics: Explaining Economic Policy. MIT Press, Cambridge (2000)
Reny, P.J.: On the existence of pure and mixed strategy Nash equilibria in discontinuous games. Econometrica **67**, 1029–1056 (1999)
Rosen, J.B.: Existence and uniqueness of equilibrium points for concave N-Person games. Econometrica **33**, 520–534 (1965)
Skaperdas, S.: Contest success functions. Econ Theory **7**, 283–290 (1996)
Szidarovszky, F., Okuguchi, K.: On the existence and uniqueness of pure Nash equilibrium in rent-seeking games. Games Econ. Behav. **18**, 135–140 (1997)
Tullock, G.: Efficient rent seeking. In: Efficient Rent-Seeking, pp. 3–16. Springer, Boston (2001)
Ui, T.: Correlated equilibrium and concave games. Int. J. Game Theory **37**, 1–13 (2008)
Ui, T.: Bayesian Nash equilibrium and variational inequalities. J. Math. Econ. **63**, 139–146 (2016)
Viossat, Y.: Is having a unique equilibrium robust? J. Math. Econ. **44**, 1152–1160 (2008)