

Michele Aghassi · Dimitris Bertsimas

Robust game theory

Received: June 15, 2004 / Accepted: March 16, 2005

Published online: February 2, 2006 – © Springer-Verlag 2006

Abstract. We present a distribution-free model of incomplete-information games, both with and without private information, in which the players use a robust optimization approach to contend with payoff uncertainty. Our “robust game” model relaxes the assumptions of Harsanyi’s Bayesian game model, and provides an alternative distribution-free equilibrium concept, which we call “robust-optimization equilibrium,” to that of the *ex post* equilibrium. We prove that the robust-optimization equilibria of an incomplete-information game subsume the *ex post* equilibria of the game and are, unlike the latter, guaranteed to exist when the game is finite and has bounded payoff uncertainty set. For arbitrary robust finite games with bounded polyhedral payoff uncertainty sets, we show that we can compute a robust-optimization equilibrium by methods analogous to those for identifying a Nash equilibrium of a finite game with complete information. In addition, we present computational results.

Key words. Game theory – Robust optimization – Bayesian games – *Ex post* equilibria

1. Introduction

1.1. Finite games with complete information

Game theory is a field in economics that examines multi-agent decision problems, in which the rewards to each agent, or player, can depend not only on his action, but also on the actions of the other players. In his seminal paper [39], John Nash introduced the notion of an equilibrium of a game. He defined an equilibrium as a profile of players’ strategies, such that no player has incentive to unilaterally deviate from his strategy, given the strategies of the other players.

In [39] and [40], Nash focused on non-cooperative, simultaneous-move, one-shot, finite games with complete information, a class of games encompassing various situations in economics. “Simultaneous-move” refers to the fact that the players choose strategies without knowing those selected by the other players. “One-shot” means that the game is played only once. “Finite” connotes that there are a finite number of players, each having a finite number of actions, over which mixed strategies in these actions may

M. Aghassi: Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, USA. e-mail: maghassi@alum.mit.edu.

The research of the author was partially supported by a National Science Foundation Graduate Research Fellowship and by the Singapore-MIT Alliance.

D. Bertsimas: Boeing Professor of Operations Research, Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, E53-363, Cambridge, MA 02139, USA.

e-mail: dbertsim@mit.edu.

The research of the author was partially supported by the Singapore-MIT Alliance.

Mathematics Subject Classification (2000): 90C05, 90C47, 91A06, 91A10, 91A15

be defined. Finally, “complete information” implies that all parameters of the game, including individual players’ payoff functions, are common knowledge.

In his analysis, Nash modeled each player as rational and wanting to maximize his expected payoff with respect to the probability distributions given by the mixed strategies. Nash proved that each game of the aforementioned type has an equilibrium in mixed strategies. In fact, Nash gave two existence proofs, one in [39] based on Kakutani’s Fixed Point Theorem [28] and one in [40] based on the less general Brouwer’s Fixed Point Theorem [10].

Nash’s equilibrium concept and existence theorem have become a cornerstone in the field of game theory and earned him the 1994 Nobel Prize in Economics. The concept is regarded as practically significant largely because, under the standard assumptions that players are rational and that the structure of a game is common knowledge, the concept offers a possible approach to predicting the outcome of the game. The argument is as follows. Any rational player who thinks his opponents will use certain strategies should never play anything other than a best response to those strategies. By the common knowledge assumption, the other players know this, the player knows that the other players know this, etc., *ad infinitum*. Thus, the players may be able to reach consistent predictions of what each other will play. The classical game theory literature concludes from this observation that we should expect the realized behavior in a game to belong to the set of equilibria. As discussed in the introduction of Fudenberg and Levine [19], in practice, this conclusion can prove to be unreliable. Nonetheless, the concept of Nash equilibrium remains the central idea in game theory, in part because no solution concept has been offered that overcomes these prediction issues.

1.2. Finite games with incomplete information

While the existence of an equilibrium can be asked in any game, Nash’s existence result addresses only non-cooperative, simultaneous-move, one-shot, finite games with complete information. Of course, in real-world, game-theoretic situations, players are often uncertain of some aspects of the structure of the game, such as payoff functions.

Harsanyi [24] modeled these incomplete-information games as what he called “Bayesian games.” He defined a player’s “type” as an encoding of information available to that player, including knowledge of his own payoff function and beliefs about other players’ payoff functions. In this way, he used type spaces to model incomplete-information games, in which some players may have private information. He assumed that the players share a common-knowledge prior probability distribution over the type space. Harsanyi suggested that each player would use this prior probability distribution, together with his type, to derive a conditional probability distribution on the parameter values remaining unknown to him. Furthermore, he assumed that each player’s goal would then be to, using a Bayesian approach in this way, maximize his expected payoff with respect to both the aforementioned conditional probability distribution and the mixed strategies of the players.

In this framework, Harsanyi extended Nash’s result to games with *incomplete* information. In particular, he showed that any Bayesian game is equivalent to an extensive-form game with complete, but imperfect information. This extensive-form game, in turn,

is known to have a static-form representation. Using an equilibrium existence result more general than Nash's and due to Debreu [14], Harsanyi thus proved the existence of equilibria, which he called "Bayesian equilibria," in Bayesian games. For his work on games with incomplete information, he won the 1994 Nobel Prize in Economics, alongside Nash and Selten.

Harsanyi's work relaxes the assumption that all parameters affecting the payoffs of the players are known with certainty. His model technique is essentially analogous to the stochastic programming approach to data uncertainty in mathematical optimization problems. As in stochastic programming, Harsanyi's model assumes the availability of full prior distributional information for all unknown parameters. In addition, his analysis assumes that all players use the same prior, and that this fact is common knowledge. Many, including Morris [38] and Wilson [55], have questioned the common prior and common knowledge aspects of these assumptions. Nonetheless, Harsanyi's Bayesian model remains the accepted convention for analyzing static games with incomplete information.

Some contributions to the literature have relaxed the common prior and common knowledge assumptions of Harsanyi's model. Most importantly, Mertens and Zamir [37] formalized the notion of a "universal type space," a type space large enough to capture players' higher-order beliefs, players' use of different prior probability distributions on the uncertain parameters, and the absence of common knowledge of these priors.

Taking a different approach, other contributors to the game theory literature have offered distribution-free equilibrium concepts for incomplete-information games. These analyses address the possibility that distributional information is not available to the players, or that they opt not to use potentially inaccurate distributional information. The notion of an *ex post* equilibrium is the most common distribution-free solution concept and is especially prevalent in the auction theory literature. Holmström and Myerson [25] first introduced this notion under the name "uniform incentive compatibility," and Crémer and McLean [12] first used it in the context of auctions. The *ex post* equilibrium is a refinement of the Bayesian equilibrium and is an appealing solution concept, because it is relevant even when the players lack distributional information on the uncertain parameters. However, it is a strong concept, and *ex post* equilibria need not exist in an incomplete-information game.

1.3. Contributions and structure of the paper

The contributions of this paper are as follows.

1. We propose a distribution-free, robust optimization model for incomplete-information games. Our model relaxes the assumptions of Harsanyi's Bayesian games model, and at the same time provides a more general equilibrium concept than that of the *ex post* equilibrium.

Specifically, in Section 2, we review relevant results from the literature on robust optimization. In Section 3, we formally propose the robust optimization model for non-cooperative, simultaneous-move, one-shot, finite games with incomplete information. We start by discussing precedents, from the game theory literature, for using a worst-case approach to uncertainty, in the absence of probability distributions.

We then note the novelty of our approach with respect to these works. After setting forth our model, we describe the “robust games,” analogous to Harsanyi’s Bayesian games, to which this approach gives rise. We compare the equilibrium conditions for such robust games to those for Bayesian equilibria and *ex post* equilibria. Furthermore, we note that any *ex post* equilibria of an incomplete-information game are what we call “robust-optimization equilibria,” i.e., equilibria of the corresponding robust game, just as they are Bayesian equilibria under Harsanyi’s model. We then discuss our union of the notion of equilibrium with the robust optimization paradigm, and we give interpretations of mixed strategies in the context of robust games. We relate this discussion and these interpretations to those in the literature on complete-information games.

At the end of Section 3, to concretize the idea of a robust game, we present some examples. In addition, we use one of these examples to illustrate that *ex post* equilibria need not exist, thereby motivating the need for an alternate distribution-free equilibrium concept.

Let us pause to note that, for the sake of simplicity, in Sections 3 through 6, we focus on situations of uncertainty in which no player has private information. This focus allows for a clearer and a sufficiently rich discussion of the main ideas underlying our model and results, without hindering the reader’s understanding through the use of cumbersome notation and references to results from the theory of Banach spaces. Such notation and results are required for the general case of incomplete-information games involving potentially private information. In Section 7, we extend our analysis to this general case.

2. In Section 4, we prove the existence of equilibria in robust finite games with bounded uncertainty sets and no private information.
3. In Section 5, we formulate the set of equilibria of an arbitrary robust finite game, with bounded polyhedral uncertainty set and no private information, as the dimension-reducing, component-wise projection of the solution set of a system of multilinear equalities and inequalities. We provide a general method for approximately computing sample robust-optimization equilibria, and we present numerical results from the application of this method. For a special class of such games, we show equivalence to finite games with complete payoff information and the same action spaces. As a result, in order to compute sample equilibria for robust games in this class, one need only solve for the equilibria of the corresponding complete-information game.
4. In Section 6, we compare properties of robust finite games with those of the corresponding complete-information games, in which the uncertain payoff parameters of the former are commonly known to take fixed, nominal values. In the absence of private information, these nominal games are precisely the Bayesian games arising from attributing symmetric probability distributions over the uncertainty sets of the corresponding robust games. In addition, turning our attention to a notion of symmetry unrelated to the symmetry of probability distributions, we extend the definition of a symmetric game, i.e., one in which the players are indistinguishable with respect to the structure of the game, to the robust game setting. We prove the existence of symmetric equilibria in symmetric, robust finite games with bounded uncertainty sets and no private information.

5. In Section 7, we generalize our model to the case with potentially private information. We extend our existence result to this context and generalize our computation method to such situations involving private information and finite type spaces.

1.4. Notation

We will use the following notation conventions throughout the paper. Boldface letters will denote vectors and matrices. In general, upper case letters will signify matrices, while lower case will denote vectors. To designate uncertain coefficients and their nominal counterparts, we will use the tilde (e.g., \tilde{a}) and the check (e.g., \check{a}), respectively. Lastly, $\text{vec}(\mathbf{A})$ will denote the column vector obtained by stacking the row vectors of the matrix \mathbf{A} one on top of another. Thus, if \mathbf{A} is an $m \times n$ matrix, $\text{vec}(\mathbf{A})$ will be an $mn \times 1$ vector.

2. Review of robust linear optimization

For the purpose of more precisely characterizing the robust optimization approach, let us consider the mathematical optimization problem (MP)

$$P : \max_{\mathbf{x}} f(\mathbf{x})$$

$$\text{s.t. } \mathbf{x} \in X(\check{\delta}_1, \check{\delta}_2, \dots, \check{\delta}_\omega),$$

where \mathbf{x} is the vector of decision variables, $X(\delta_1, \delta_2, \dots, \delta_\omega)$ is the feasible region defined by parameters $\delta_i, i \in \{1, \dots, \omega\}$, and $f(\mathbf{x})$ is the objective function. In the above nominal MP, we regard the parameter values as being fixed at $\delta_i = \check{\delta}_i, i \in \{1, \dots, \omega\}$. Suppose instead that we do not know the exact values of these parameters $\check{\delta}_1, \dots, \check{\delta}_\omega$, but know that $(\tilde{\delta}_1, \dots, \tilde{\delta}_\omega)$ belongs to some uncertainty set U . Under this model of uncertainty, the robust counterpart of the above nominal problem is given by

$$RP : \max_{\mathbf{x}} f(\mathbf{x})$$

$$\text{s.t. } \mathbf{x} \in X(\tilde{\delta}_1, \tilde{\delta}_2, \dots, \tilde{\delta}_\omega), \quad \forall (\tilde{\delta}_1, \tilde{\delta}_2, \dots, \tilde{\delta}_\omega) \in U.$$

Without loss of generality, we can restrict our discussion of data uncertainty to the constraints, since the objective can always be incorporated into the constraints. Indeed, the nominal problem P can be rewritten as a maximization over a new decision variable, z .

$$\max_{\mathbf{x}, z} z$$

$$\text{s.t. } z \leq f(\mathbf{x})$$

$$\mathbf{x} \in X(\check{\delta}_1, \check{\delta}_2, \dots, \check{\delta}_\omega).$$

Initial results on robust linear optimization were given by Soyster in [47]. Soyster considered linear optimization problems (LPs) subject to column-wise uncertainty in

the constraint matrix. His model is equivalent to the LP in which all uncertain parameters have been fixed at their worst-case values from the uncertainty set. Soyster’s model is overly conservative; in practice, it seems quite unlikely that the uncertain parameters would all simultaneously realize their worst-case values. In addition, his model is specific to column-wise uncertainty and does not easily generalize.

Twenty years later, Ben-Tal and Nemirovski [1–3] and, independently of them, El Ghaoui et al. [16, 17] renewed the discussion of optimization under uncertainty. They examined ellipsoidal models of uncertainty, which, for the robust LP case, are less conservative than the column-wise model considered by Soyster. They showed that the robust counterpart of an LP under such ellipsoidal uncertainty models is a second-order cone optimization problem (SOCP). Furthermore, they remarked that polyhedral uncertainty can be regarded as a special case of ellipsoidal uncertainty. As a result, LPs under polyhedral uncertainty of the coefficient matrix can be solved via SOCPs.

Ellipsoidal uncertainty formulations of robustness are attractive in that they offer a reduced level of conservatism, as compared with the Soyster model, and lead to efficient solutions, via SOCPs, of LPs under uncertainty. Unfortunately, ellipsoidal uncertainty formulations give rise to robust counterparts whose exact solution is more computationally demanding than that of the corresponding nominal problem. In response to this drawback, Bertsimas and Sim [6, 7] offered an alternative model of uncertainty, under which the robust counterpart of an LP is an LP. Their formulation yields essentially the same level of conservatism as do those of Ben-Tal and Nemirovski and El Ghaoui et al.

Bertsimas, Pachamanova, and Sim [5] further extended the results of Bertsimas and Sim [7] to the case of general polyhedral uncertainty of the coefficient matrix. In particular, where $\tilde{\mathbf{A}}$ is an $m \times n$ coefficient matrix, they modeled the uncertainty set as a polyhedron

$$U = \left\{ \text{vec}(\tilde{\mathbf{A}}) \mid \mathbf{G} \cdot \text{vec}(\tilde{\mathbf{A}}) \leq \mathbf{d} \right\}, \tag{1}$$

where $\mathbf{G} \in \mathbb{R}^{\ell \times mn}$, $\text{vec}(\tilde{\mathbf{A}}) \in \mathbb{R}^{mn \times 1}$, and $\mathbf{d} \in \mathbb{R}^{\ell}$. They considered the robust LP given by

$$\begin{aligned} \max_{\mathbf{x}} \quad & \mathbf{c}'\mathbf{x} \\ \text{s.t.} \quad & \tilde{\mathbf{A}}\mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \in S \\ & \forall \tilde{\mathbf{A}} \in U, \end{aligned} \tag{2}$$

where $\mathbf{x} \in \mathbb{R}^n$ is the decision vector and S is a polyhedron defined by constraints that are not subject to uncertainty. They showed that, if the number of constraints defining S is r , then robust LP (2) in n variables and $m + r$ constraints is equivalent to a nominal LP in $n + m\ell$ variables and $m^2n + m + m\ell + r$ constraints.

3. A robust approach to payoff uncertainty in games

As an alternative to Harsanyi’s model and the notion of the *ex post* equilibrium, we propose a new distribution-free model of and equilibrium concept for incomplete-information games. Our model is based on robust optimization, in which one takes a deterministic

approach to uncertainty and seeks to optimize worst-case performance, where the worst case is taken with respect to a set of possible values for the uncertain parameters. Let us note that, in Sections 3 through 6, we will focus on incomplete-information games without private information. In Section 7, we will extend our analysis to the general case involving potentially private information.

3.1. Precedents for a worst-case approach

In fact, the game theory literature is ripe with precedents for using a worst-case approach. The field arose in large part from von Neumann's and Morgenstern's "max-min" formulation of behavior in games [52]. More recently, for example, Goldberg et al. [22] proposed a worst-case, competitive-analysis approach to auction mechanism design. In the more general context of games in normal form, several authors, including Gilboa and Schmeidler [21], Dow and Werlang [15], Klibanoff [29], Lo [32], and Marinacci [34], have argued for a max-min-based approach to "ambiguous uncertainty," uncertainty in the absence of probabilistic information. They contend that expected utility models are well-suited for decision-theoretic situations characterized by "risk," uncertainty with distributional information, but that these models do not capture behavior observed in practice in situations of ambiguous uncertainty.¹ As Dow and Werlang [15] note, the former type of uncertainty is exemplified by the outcome of a coin toss, while the latter is typified by the outcome of a horse race.

While the max-min approaches of Gilboa and Schmeidler, Dow and Werlang, Klibanoff, Lo, and Marinacci share the worst-case perspective of our model, their approaches are fundamentally different from ours for at least three reasons. First, these authors consider complete-information games, whereas we address incomplete-information games. In their models, players know, with certainty, the payoffs under given tuples of actions, but do not know which tuple of actions will be played. In our model, the players may be uncertain of the payoffs, even under given tuples of actions. Accordingly, the aforementioned authors use a pessimistic approach to model each player's uncertainty of the other players' behaviors, whereas we use a worst-case approach to model each player's uncertainty of the payoff functions themselves.

Second, although these authors take a worst-case approach to some extent, their models are nonetheless inherently probabilistic, unlike our approach, which is fundamentally deterministic. Klibanoff [29] and Lo [32] model each player's uncertainty of the other players' behaviors using the notion of multiple prior probability distributions. They characterize each player as believing his counterparts' actions are a realization from some unknown probability distribution, belonging to a family of known multiple priors. Each player then seeks to maximize his minimum expected utility, where the minimum is taken with respect to this set of multiple priors. Gilboa and Schmeidler [21], Dow and Werlang [15], and Marinacci [34] propose a related approach using non-additive probability distributions in place of sets of multiple priors. Unlike these authors, we offer a model in which the players give no consideration whatsoever to probability distributions over the uncertain values. Under our approach, the players regard the uncertain values as simply unknown and not as realizations from some probability distribution, even a

¹ Knight [30] was one of the first to draw a distinction between these two forms of uncertainty.

distribution that is itself not exactly known. Consequently, our model of the players' responses to uncertainty is distribution-free and deterministic in nature.

Third, these authors offer no guidelines for equilibria computation in the context of their models. In contrast, in Section 5, we propose such a computation method. Despite these differences, the aforementioned authors' contributions provide ample support for the robust optimization model we propose for games with incomplete information.

In addition, since the submission of our paper, Hyafil and Boutilier [27] have recently offered a worst-case approach for incomplete-information games, based on the distribution-free decision criterion of minimax regret, popular in the online optimization literature [9]. Their approach is in contrast to our framework of modeling the players as each seeking to maximize his worst-case expected payoff. Hyafil and Boutilier provide an existence result for a very restricted, special case of incomplete-information games, involving private information, but finite type spaces. They offer no ideas on computation of their equilibria.

Having discussed some precedents for taking a worst-case approach to analyzing game-theoretic situations, let us now formalize our robust games model.

3.2. Formalization of the robust game model

In our robust optimization model of incomplete-information games, we assume that the players commonly know only an uncertainty set of possible values of the uncertain payoff function parameters.² They need not, as Harsanyi's model additionally assumes, have distributional information for this uncertainty set. In addition, we suppose that each player uses a robust optimization, and therefore a worst-case, approach to the uncertainty, rather than seeking, as in Harsanyi's model, to optimize "average" performance with respect to a distribution over the uncertainty set. In the game theory literature, the "performance" of a player's mixed strategy is measured by his expected payoff. Accordingly, in our model, given the other players' strategies, each player seeks to maximize his worst-case expected payoff. The worst-case is taken with respect to the uncertainty set, and the expectation is taken, as in complete-information games, over the mixed strategies of the players. Analogous to Harsanyi's "Bayesian game" terminology, we call the resulting games "robust games," and we refer to their equilibria as "robust-optimization equilibria" of the corresponding incomplete-information games.

In this section we will formalize our robust game model and its relation to Nash's and Harsanyi's models for the complete- and incomplete-information settings, respectively. We will also compare the notion of *ex post* equilibrium with the concept of robust-optimization equilibrium.

Let us first define some terms and establish some notation. Suppose there are N players and that player $i \in \{1, \dots, N\}$ has $a_i > 1$ possible actions.

Definition 1. A game is said to be *finite* if the number of players N and the number of actions a_i available to each player $i \in \{1, \dots, N\}$ are all finite.

² Incomplete-information games in the absence of distributional information are sometimes called "games in informational form" [26].

So, we will use the term “robust finite game” to refer to robust games that have finitely many players with finitely many actions each, even when the uncertainty sets are not finite.

In the complete-information game setting, a multi-dimensional payoff matrix $\check{\mathbf{P}}$, indexed over $\{1, \dots, N\} \times \prod_{i=1}^N \{1, \dots, a_i\}$, records the payoffs to the players under all possible action profiles for the players. In particular, for $i \in \{1, \dots, N\}$, $(j_1, \dots, j_N) \in \prod_{i=1}^N \{1, \dots, a_i\}$, let $\check{P}_{(j_1, \dots, j_N)}^i$ denote the payoff to player i when player $i' \in \{1, \dots, N\}$ plays action $j_{i'} \in \{1, \dots, a_{i'}\}$. Let

$$S_{a_i} = \left\{ \mathbf{x}^i \in \mathbb{R}^{a_i} \mid \mathbf{x}^i \geq 0, \sum_{j_i=1}^{a_i} x_{j_i}^i = 1 \right\}.$$

That is, S_{a_i} is the set of mixed strategies over action space $\{1, \dots, a_i\}$. Let us define $\boldsymbol{\pi} : U \times \prod_{i=1}^N S_{a_i} \rightarrow \mathbb{R}^N$ as the vector function mapping a payoff matrix and the mixed strategies of N players to a vector of expected payoffs to the N players. In particular, $\pi_i(\mathbf{P}; \mathbf{x}^1, \dots, \mathbf{x}^N)$ will denote the expected payoff to player i when the payoff matrix is given by \mathbf{P} and player $i' \in \{1, \dots, N\}$ plays mixed strategy $\mathbf{x}^{i'} \in S_{a_{i'}}$. That is,

$$\pi_i(\mathbf{P}; \mathbf{x}^1, \dots, \mathbf{x}^N) = \sum_{j_1=1}^{a_1} \dots \sum_{j_i=1}^{a_i} \dots \sum_{j_N=1}^{a_N} P_{(j_1, \dots, j_N)}^i \prod_{i=1}^N x_{j_i}^i.$$

Now that we have established some notation, we can formulate the best response correspondence in our robust optimization model for games with incomplete payoff information. We will compare this correspondence with those in Nash’s and Harsanyi’s models for games with complete and incomplete information, respectively. In the remainder of the paper, we will use the following shorthands.

$$\begin{aligned} \mathbf{x}^{-i} &\triangleq (\mathbf{x}^1, \dots, \mathbf{x}^{i-1}, \mathbf{x}^{i+1}, \dots, \mathbf{x}^N) \\ (\mathbf{x}^{-i}, \mathbf{u}^i) &\triangleq (\mathbf{x}^1, \dots, \mathbf{x}^{i-1}, \mathbf{u}^i, \mathbf{x}^{i+1}, \dots, \mathbf{x}^N) \\ S &\triangleq \prod_{i=1}^N S_{a_i} \\ S_{-i} &\triangleq \prod_{\substack{i'=1 \\ i' \neq i}}^N S_{a_{i'}}. \end{aligned} \tag{3}$$

Every model of a game attributes some objective to each player. A player’s objective in turn determines the set of best responses to the other players’ strategies.

Definition 2. *A player’s strategy is called a **best response** to the other players’ strategies if, given the latter, he has no incentive to unilaterally deviate from his aforementioned strategy.*

In the complete-information game setting, with payoff matrix $\check{\mathbf{P}}$, the classical model assumes that each player seeks to maximize his expected payoff. So, player i 's best response to the other players' strategies $\mathbf{x}^{-i} \in S_{-i}$ belongs, by definition, to

$$\arg \max_{\mathbf{u}^i \in S_{a_i}} \pi_i \left(\check{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right).$$

In games with incomplete payoff information, the payoff matrix $\tilde{\mathbf{P}}$ is subject to uncertainty. In Harsanyi's Bayesian model, in the context of games without private information, in which the type spaces are singletons, player i 's best response to the other players' strategies $\mathbf{x}^{-i} \in S_{-i}$ must belong to

$$\arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\mathbb{E}_{\tilde{\mathbf{P}}} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right) \right].$$

In our robust model, for the case without private information, player i 's best response to the other players' strategies $\mathbf{x}^{-i} \in S_{-i}$ must belong to

$$\arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right) \right].$$

Thus, in moving from Harsanyi's Bayesian approach to our robust optimization model, we replace the expectation in the definition of the best response correspondence with an infimum operator.

Note that $\forall i \in \{1, \dots, N\}$ and $\forall (\mathbf{x}^{-i}, \mathbf{u}^i) \in S$, by the linearity of π_i over U and by the linearity of the expectation operator,

$$\mathbb{E}_{\tilde{\mathbf{P}}} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right) = \pi_i \left(\mathbb{E}_{\tilde{\mathbf{P}}} [\tilde{\mathbf{P}}]; \mathbf{x}^{-i}, \mathbf{u}^i \right), \quad (4)$$

where $\mathbb{E}_{\tilde{\mathbf{P}}} [\tilde{\mathbf{P}}]$ is the component-wise expected value of $\tilde{\mathbf{P}}$. Hence, in the Bayesian game setting, the average expected payoffs and expected average payoffs are in fact equivalent.³ Recall, from Harsanyi [24], that any Bayesian game with incomplete information is equivalent to a static game with complete but imperfect information. As indicated by Equation (4), in the absence of private information, a Bayesian game is equivalent to a finite game with complete and *perfect* information, with the same action spaces and with payoff matrix $\mathbb{E}_{\tilde{\mathbf{P}}} [\tilde{\mathbf{P}}]$.

In contrast, under the robust model, the worst-case expected payoff expressed above is no less than, and is generally strictly greater than, the expected worst-case payoff. That is,

$$\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right) \geq \pi_i \left(\inf_{\tilde{\mathbf{P}} \in U} [\tilde{\mathbf{P}}]; \mathbf{x}^{-i}, \mathbf{u}^i \right),$$

³ We use the terms "average" and "expected" in an effort to distinguish between two different types of expectations, namely, the expectation ("average") taken with respect to the distribution over the uncertainty set of payoff parameter values and the expectation ("expected payoff") taken with respect to the distributions induced by the players' mixed strategies.

where $\inf_{\tilde{\mathbf{P}} \in U} [\tilde{\mathbf{P}}]$ denotes the component-wise infimum of $\tilde{\mathbf{P}}$. Thus, a robust finite game without private information is, in general, not equivalent to the complete-information, finite game with the same action spaces and with payoff matrix commonly known to be $\inf_{\tilde{\mathbf{P}} \in U} [\tilde{\mathbf{P}}]$. We will see in Section 5 that this equivalence does, however, hold for certain classes of robust games.⁴

We are now ready to apply the concept of equilibrium to robust finite games.

Definition 3. A tuple of strategies is said to be an **equilibrium** if each player's strategy is a best response to the other players' strategies.

Accordingly, the criterion for an equilibrium is completely determined by the best response correspondence, which in turn is completely determined by the players' objectives. For example, in the complete-information game setting, $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is said to be a Nash equilibrium iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \pi_i(\check{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i). \quad (5)$$

Similarly, under Harsanyi's model for finite games with incomplete payoff information and with no private information, $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is said to be an equilibrium iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\mathbb{E}_{\tilde{\mathbf{P}}} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) \right]. \quad (6)$$

Finally, under our robust model for finite games with incomplete payoff information and with no private information, $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is said to be an equilibrium, i.e., a robust-optimization equilibrium of the corresponding game with incomplete information, iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\inf_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) \right]. \quad (7)$$

Let us contrast the equilibrium concepts arising from Harsanyi's Bayesian game model and our robust game model with the notion of the *ex post* equilibrium, defined as follows.

⁴ One could model each player as wishing to maximize his expected worst-case payoff, rather than, as we have done, his worst-case expected payoff. We chose the latter over the former for two reasons. First, the former model is *not*, while the latter model is, in the spirit of robust optimization, in which we seek to optimize a worst-case version of the nominal objective, i.e., the expected payoff. Second, while a robust approach is by its nature pessimistic, the former model is even more, and perhaps excessively, pessimistic. In it, each player assumes that the uncertain data realization will be maximally hostile with respect to the *action outcomes* of the randomizations yielded by all the players' mixed strategies. In contrast, in the robust model we propose, the maximal hostility is assumed by the players to be with respect to the mixed strategy *probability distributions* themselves; i.e., the "adversary" does not have the benefit of seeing the action outcomes of the randomizations, before he is forced to choose values of the uncertain data.

If one nonetheless opts, despite these drawbacks, to model each player as seeking to maximize his expected worst-case payoff, rather than his worst-case expected payoff, the game with incomplete information will be equivalent to one with complete information and with payoff matrix $\inf_{\tilde{\mathbf{P}} \in U} [\tilde{\mathbf{P}}]$. Accordingly, the existence and computation results that we will present in this paper for the robust model follow trivially for this excessively pessimistic model.

Definition 4. A tuple of strategies is said to be an **ex post equilibrium** if each player's strategy is a best response to the other players' strategies, under all possible realizations of the uncertain data.

More precisely, in the absence of private information, $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is said to be an *ex post* equilibrium iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right), \quad \forall \tilde{\mathbf{P}} \in U. \tag{8}$$

By definition, an *ex post* equilibrium must be an equilibrium of every nominal game in the family of nominal games arising from U . This condition is quite strong. In fact, it is easy to prove the well-known result that every *ex post* equilibrium of an incomplete-information game is an equilibrium of any corresponding Bayesian game arising from the assignment of a distribution over the set U . Similarly, we have the following lemma, establishing an analogous result for the set of robust-optimization equilibria.

Lemma 1. Any *ex post* equilibrium of an incomplete-information game, without private information, is a robust-optimization equilibrium of the game.

Proof. Suppose $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is an *ex post* equilibrium of the game with incomplete information and uncertainty set U . Suppose, $\exists i \in \{1, \dots, N\}$ and $\exists \mathbf{u}^i \in S_{a_i}$, such that

$$\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i \right) < \inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right).$$

By the definition of *ex post* equilibrium,

$$\pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right) \leq \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i \right), \quad \forall \tilde{\mathbf{P}} \in U,$$

yielding a contradiction of the fact that $\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i \right)$ is the greatest lower bound on $\pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i \right)$ over $\tilde{\mathbf{P}} \in U$. Therefore, $\forall i \in \{1, \dots, N\}$, and $\forall \mathbf{u}^i \in S_{a_i}$,

$$\inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i \right) \geq \inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i \right),$$

establishing that $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is an equilibrium of the corresponding robust game. □

In Section 3.5, we will illustrate, with examples, our robust optimization model for games with incomplete information. Using one of these examples, in Section 3.6, we will demonstrate that, in general, *ex post* equilibria do not exist in incomplete-information games. Before giving these examples, we wish to address two questions that the reader may have regarding our approach. In Section 3.3, we will discuss why, in the context of distribution-free, incomplete-information games, it is reasonable, and in fact natural, to combine the notion of equilibrium with a worst-case viewpoint. In Section 3.4, we will discuss our motivation for considering mixed strategies.

3.3. Why combine equilibrium and worst-case notions?

Recall that, with the exception of two-person, zero-sum games with complete information, mixed strategy equilibria do not generally consist of max-min strategies. That is, a player's strategy in a mixed-strategy equilibrium is not generally the one guaranteeing him the best possible expected payoff when his counterparts collude to minimize this quantity. The reason is that a player's counterparts generally have incentive to deviate from such collusive behavior, in order to try to individually maximize their own payoffs. In turn, the player himself therefore generally has incentive to deviate from the aforementioned max-min strategy.

In contrast, the robust optimization paradigm is fundamentally such a max-min, or a worst-case, approach. In our robust games model, given his counterparts' strategies, each player formulates a best response as the solution of a robust optimization problem. Based on the discussion in the preceding paragraph, one may worry that, by analogy, best responses based on robust optimization are not conducive to equilibrium. This analogy fails, and this worry is therefore unfounded, for the following reason. In our model, a player's counterparts are outside the scope of that player's pessimistic viewpoint. In particular, each player takes a worst-case view *only* of the uncertain parameters that define his payoff function, under a *given* tuple of his counterparts' strategies. Each player does *not* take a worst-case approach to his uncertainty with respect to this tuple itself, as is done in classical max-min strategies. Indeed, "nature," rather than any of the players themselves, selects these unknown payoff parameter values. Accordingly, in order for the analogy to hold, nature must be a participant, on the same footing as the other players, in the game. However, nature receives no payoff in the game, and therefore cannot be characterized as a player itself.

Thus, it is indeed reasonable to combine, as we have done, the notion of equilibrium with the robust optimization paradigm. Let us now explain why this union is in fact natural, in the context of incomplete-information games. If the players commonly know that they all take a robust optimization approach to the payoff uncertainty, then they would all commonly know each others' best response correspondences. Armed with this common knowledge, the players could then attempt to mutually predict each other's behavior, just as they could in a complete-information game, as discussed in Section 1.1. Recall from this discussion that the set of Nash equilibria are the set of *consistent* such mutual predictions in a finite, complete-information game. Analogously, the set of equilibria of a robust finite game are the set of consistent such mutual predictions in the corresponding finite, distribution-free, incomplete-information game. As such, our notion of equilibrium in a robust game offers a natural approach to attempting to predict the outcomes of such incomplete-information games.

3.4. Interpretation of mixed strategies

We will now explain our motivation for considering mixed strategies, and we will relate this discussion to interpretations of mixed strategies in the context of complete-information games (see, for example, Chapter 3 of Osborne and Rubinstein [41]). In the case of finite games with complete information, some game theorists support the literal

interpretation of mixed strategies as actual randomizations by the players over their action spaces. Others are dissatisfied with this viewpoint. The latter group note the following property of mixed strategy equilibria, in finite, complete-information games. In response to his counterparts' behaviors in any such equilibrium, each player's mixed strategy does as well, but no better than the actions contained in its support. The opponents of the literal interpretation therefore argue that this lack of strict preference for randomization undermines the belief that players randomize in reality.

In the case of robust finite games, this argument against the literal interpretation does not hold. In particular, because of the infimum in the worst-case expected payoff function, this function is nonlinear. Consequently, for any mixed strategy equilibrium in such a game, in response to his counterparts' behavior in this equilibrium, each player will, in general, strictly prefer his mixed strategy over the actions in its support. Accordingly, one may argue that the literal viewpoint of mixed strategies is more justified in the context of robust games than it is in the context of complete-information games.

One may nonetheless remain dissatisfied with this belief that players randomize in real-world, game-theoretic settings, even those involving payoff uncertainty. Let us then consider an alternative interpretation of mixed strategies. In the literature on finite, complete-information games, some have advocated the viewpoint of mixed strategy equilibria as limiting, empirical frequencies of actions played, when the game is repeated.

The same empirical frequency interpretation extends to robust finite games. Imagine that the players engage concurrently in many instances of the same game, with the same, unknown payoff matrix $\tilde{\mathbf{P}}$. Suppose the players know that $\tilde{\mathbf{P}}$ is constant across all instances, but are uncertain of its true value. As before, suppose each player knows only an uncertainty set to which $\tilde{\mathbf{P}}$ belongs, has no distributional information with respect to this set, and takes a worst-case approach to this uncertainty. Lastly, suppose that, in each instance of the game, each agent may play a different action. Each player thus builds, in essence, a "portfolio" of actions. The payoff from each action in the portfolio is determined by the other players' actions in the corresponding instance of the game and by the single unknown value of $\tilde{\mathbf{P}}$. Accordingly, we may view the mixed strategy equilibria as the limiting, empirical frequencies describing each player's level of diversification within his portfolio of actions. Note that this portfolio interpretation can be recast in terms of sequentially repeated games, in which the players know that the uncertain payoff matrix is constant over all rounds, and in which they do not receive their payoffs until the final round is played. That is, the players do not know, until at least after play has terminated, the true value of $\tilde{\mathbf{P}}$.

3.5. Examples of robust finite games

Having presented our robust games model and addressed some interpretation issues, we will now illustrate our approach with a few examples.

Example 1. Robust Inspection Game

Consider the classical inspection game discussed in [20]. The row player, the employee, can either shirk or work (actions 1 and 2, respectively). The column player, his employer, can either inspect or not inspect (actions 1 and 2, respectively). The purpose of

inspecting is to learn whether the employee is working. The two players simultaneously select their actions. When the employee works, he suffers an opportunity cost \tilde{g} , and his employer enjoys a value of work output of \tilde{v} . When the employer inspects, she suffers an opportunity cost of \tilde{h} . If she inspects and finds the employee shirking, she need not pay him his wage w . Otherwise, she must pay him w . In the nominal version of the game, \tilde{v} , w , \tilde{g} , and \tilde{h} are commonly known with certainty by the players. In practice, it seems reasonable that the opportunity costs and the value of work output (e.g., subject to unpredictable defects) would be subject to uncertainty. To this end, suppose that \tilde{v} , \tilde{g} , and \tilde{h} are subject to independent uncertainty, the nature of which is common knowledge between the two players. For example, we may consider the robust game in which the payoff uncertainty set is given by

$$U = \left\{ \left(\begin{array}{cc} (0, -\tilde{h}) & (w, -w) \\ (w - \tilde{g}, \tilde{v} - w - \tilde{h}) & (w - \tilde{g}, \tilde{v} - w) \end{array} \right) \mid (\tilde{g}, \tilde{v}, \tilde{h}) \in [\underline{g}, \bar{g}] \times [\underline{v}, \bar{v}] \times [\underline{h}, \bar{h}] \right\}.$$

Example 2. Robust Free-Rider Problem

Consider the symmetric version of the classical, 2-player, free-rider problem discussed in [20]. Each player must make a binary decision of whether or not (actions 1 and 2, respectively) to contribute to the construction of a public good. The players make their decisions simultaneously. If a player contributes, the player incurs some cost \tilde{c} , which is subject to minor uncertainty (e.g., because projected costs are rarely accurate), in a way that is common knowledge to the two players. If the public good is built, each player enjoys a payoff of 1. The good will not be built unless at least one player contributes. So, we may consider the resulting robust game with payoff uncertainty set

$$U = \left\{ \left(\begin{array}{cc} (1 - \tilde{c}, 1 - \tilde{c}) & (1 - \tilde{c}, 1) \\ (1, 1 - \tilde{c}) & (0, 0) \end{array} \right) \mid \tilde{c} \in [\check{c} - \Delta, \check{c} + \Delta] \right\},$$

for some fixed $\Delta > 0$.

Example 3. Robust Network Routing

Network routing games, formulated as early as 1952 by Wardrop [54], have become an increasingly popular topic in the game theory literature. Within the last five years, Papadimitriou [42] and others have studied the so-called “price of anarchy,” or the ratio of total payoffs at equilibrium to those at Pareto optimality.

Consider a network routing game, in which N internet service providers must each contract for the use of a single “path” in a network of a paths (e.g., servers, wiring, etc.). The providers must make these arrangements simultaneously and prior to knowing the demand to be faced (i.e., the amount of data their customers will want to route). So, each provider’s action space is the set of paths in the network. Suppose edge latencies in the network are linear and additive, and that the payoff to provider i when he uses path j_i is given by the negative of total latency experienced on edge j_i . That is, higher latencies yield lower payoffs. Specifically, we can express the uncertain payoff matrix $\tilde{\mathbf{P}}$ as a function, \mathbf{P} , of the uncertain demands to be faced. Let \tilde{d}_i denote the uncertain demand to be faced by provider i . $\forall i \in \{1, \dots, N\}, \forall (j_1, \dots, j_N) \in \{1, \dots, a\}^N$, let

$$P_{(j_1, \dots, j_N)}^i(\tilde{d}_1, \dots, \tilde{d}_N) = - \sum_{\{i' \mid j_{i'}=j_i\}} \lambda_{(i', j_{i'})} \tilde{d}_{i'},$$

where $\lambda_{(i,j_i)}$ are nonnegative coefficients that account for the fact that the marginal latencies may differ by provider and path. The demand uncertainty may arise from the fact that the providers commonly know the total demand D to be faced by all of them, but do not know how this demand will be distributed among them (e.g., uncertainty of projected subscribership for a future year). For example, the uncertainty set may be given by

$$U = \left\{ \mathbf{P} \left(\tilde{d}_1, \dots, \tilde{d}_N \right) \mid \sum_{i=1}^N \tilde{d}_i = D, \tilde{d}_i \geq \underline{d}_i, i = 1, \dots, N \right\},$$

where

$$D > \sum_{i=1}^N \underline{d}_i$$

$$\underline{d}_i > 0, \quad i = 1, \dots, N,$$

are commonly known by the players.

3.6. Nonexistence of *ex post* equilibria

We will now use the incomplete-information inspection game presented in Example 1 to illustrate that not all incomplete-information games have an *ex post* equilibrium. Each possible realization of $(\tilde{g}, \tilde{v}, \tilde{h}) \in [\underline{g}, \bar{g}] \times [\underline{v}, \bar{v}] \times [\underline{h}, \bar{h}]$ gives rise to a nominal game, i.e., a game with complete information. It is easy to show that each such game has a unique equilibrium in which the employee shirks (action 1) with probability \tilde{h}/w and the employer inspects (action 1) with probability \tilde{g}/w . So, unless $\underline{g} = \bar{g}$ and $\underline{h} = \bar{h}$, this incomplete-information game has no *ex post* equilibria.

Accordingly, the *ex post* equilibrium concept cannot be applied to all games with incomplete information, because such equilibria need not exist. In contrast, in the next section, we will prove that any robust finite game with bounded uncertainty set has an equilibrium. In this way, robust games offer an alternative distribution-free notion of equilibrium, whose existence is guaranteed.

4. Existence of equilibria in robust finite games

Having formalized and given examples illustrating our robust optimization model of games with incomplete payoff information, let us now establish the existence of equilibria in the resulting robust games, when these games are finite and have bounded uncertainty sets. Our proof of existence directly uses Kakutani’s Fixed Point Theorem [28] and parallels Nash’s first existence proof in [39]. As already mentioned, we focus in this section on incomplete-information games not involving private information. In Section 7, we extend our existence result to the general case involving potentially private information.

To begin, let us state Kakutani’s theorem and a relevant definition. Kakutani’s definition of upper semi-continuity relates to mappings from a closed, bounded, convex set S in a Euclidean space into the family of all closed, convex subsets of S . 2^S will denote the power set of S .

Definition 5 (Kakutani [28]). A point-to-set mapping $\Psi : S \rightarrow 2^S$ is called **upper semi-continuous** if

$$\begin{aligned} \mathbf{y}^n &\in \Psi(\mathbf{x}^n), & n = 1, 2, 3, \dots \\ \lim_{n \rightarrow \infty} \mathbf{x}^n &= \mathbf{x} \\ \lim_{n \rightarrow \infty} \mathbf{y}^n &= \mathbf{y} \end{aligned}$$

imply that $\mathbf{y} \in \Psi(\mathbf{x})$. In other words, the graph of $\Psi(\mathbf{x})$ must be closed.

Theorem 1 (Kakutani’s Fixed Point Theorem [28]). If S is a closed, bounded, and convex set in a Euclidean space, and Φ is an upper semi-continuous point-to-set mapping of S into the family of closed, convex subsets of S , then $\exists \mathbf{x} \in S$ s.t. $\mathbf{x} \in \Phi(\mathbf{x})$.

To use Kakutani’s Fixed Point Theorem, we must first establish some properties of the worst-case expected payoff functions, given by

$$\rho_i \left(\mathbf{x}^1, \dots, \mathbf{x}^N \right) \triangleq \inf_{\tilde{\mathbf{P}} \in U} \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^1, \dots, \mathbf{x}^N \right), \tag{9}$$

$i \in \{1, \dots, N\}$. In an N -person, robust finite game, let $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ be the uncertainty set of possible payoff matrices $\tilde{\mathbf{P}}$.

Lemma 2. Let $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ be bounded. Then, $\forall (\mathbf{x}^1, \dots, \mathbf{x}^N) \in \mathbb{R}^{a_1 + \dots + a_N}$ and $\forall \varepsilon > 0$, $\exists \delta (\varepsilon, \mathbf{x}^1, \dots, \mathbf{x}^N) > 0$ such that, $\forall \tilde{\mathbf{P}} \in U$ and $\forall i \in \{1, \dots, N\}$,

$$\left\| (\mathbf{y}^1, \dots, \mathbf{y}^N) - (\mathbf{x}^1, \dots, \mathbf{x}^N) \right\|_{\infty} < \delta (\varepsilon, \mathbf{x}^1, \dots, \mathbf{x}^N)$$

implies

$$\left| \pi_i \left(\tilde{\mathbf{P}}; \mathbf{y}^1, \dots, \mathbf{y}^N \right) - \pi_i \left(\tilde{\mathbf{P}}; \mathbf{x}^1, \dots, \mathbf{x}^N \right) \right| < \varepsilon.$$

Proof. $\forall (\mathbf{x}^1, \dots, \mathbf{x}^N) \in \mathbb{R}^{a_1 + \dots + a_N}$, and $\forall \varepsilon > 0$, consider $\delta (\varepsilon, \mathbf{x}^1, \dots, \mathbf{x}^N)$ given by

$$\delta (\varepsilon, \mathbf{x}^1, \dots, \mathbf{x}^N) = \frac{\min\{\varepsilon, 1\}}{2(2^N - 1)M \cdot \prod_{i=1}^N \left(a_i \max \left\{ \max_{j_i \in \{1, \dots, a_i\}} |x_{j_i}^i|, 1 \right\} \right)},$$

where $1 < M < \infty$ satisfies

$$\left| \tilde{P}_{(j_1, \dots, j_N)}^i \right| \leq M, \quad \forall i \in \{1, \dots, N\}, \quad \forall (j_1, \dots, j_N) \in \prod_{i=1}^N \{1, \dots, a_i\}, \quad \forall \tilde{\mathbf{P}} \in U.$$

The result follows from algebraic manipulation. □

Lemma 2 immediately gives rise to the following continuity result, which we therefore state without proof.

Lemma 3. Let $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ be bounded. Then, $\forall i \in \{1, \dots, N\}$, $\rho_i (\mathbf{x}^1, \dots, \mathbf{x}^N)$ is continuous on $\mathbb{R}^{a_1 + \dots + a_N}$.

Similarly, it is trivial to prove the following lemma.

Lemma 4. $\forall i \in \{1, \dots, N\}$ and $\forall \mathbf{x}^{-i} \in S_{-i}$ fixed, $\rho_i(\mathbf{x}^{-i}, \mathbf{x}^i)$ is concave in \mathbf{x}^i .

We are now ready to prove the existence of equilibria in robust finite games with bounded uncertainty sets.

Theorem 2 (Existence of Equilibria in Robust Finite Games). *Any N -person, non-cooperative, simultaneous-move, one-shot robust game, in which $N < \infty$, in which player $i \in \{1, \dots, N\}$ has $1 < a_i < \infty$ possible actions, in which the uncertainty set of payoff matrices $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ is bounded, and in which there is no private information, has an equilibrium.*

Proof. We will proceed by constructing a point-to-set mapping that satisfies the conditions of Kakutani’s Fixed Point Theorem [28], and whose fixed points are precisely the equilibria of the robust game. To begin, clearly, S is closed, bounded, and convex, since S_{a_i} is, $\forall i \in \{1, \dots, N\}$. Define $\Phi : S \rightarrow 2^S$ as

$$\begin{aligned} &\Phi(\mathbf{x}^1, \dots, \mathbf{x}^N) \\ &= \left\{ (\mathbf{y}^1, \dots, \mathbf{y}^N) \in S \mid \mathbf{y}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \rho_i(\mathbf{x}^{-i}, \mathbf{u}^i), i = 1, \dots, N \right\}. \end{aligned} \quad (10)$$

Let us show that $\Phi(\mathbf{x}^1, \dots, \mathbf{x}^N) \neq \emptyset, \forall (\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$. By Lemma 3, $\forall i, \forall \mathbf{x}^{-i} \in S_{-i}$ fixed, $\rho_i(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is continuous on S_{a_i} , a nonempty, closed, and bounded subset of \mathbb{R}^{a_i} . Thus, by Weierstrass’ Theorem,

$$\arg \max_{\mathbf{u}^i \in S_{a_i}} \rho_i(\mathbf{x}^{-i}, \mathbf{u}^i) \neq \emptyset.$$

Accordingly, $\forall (\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$,

$$\Phi(\mathbf{x}^1, \dots, \mathbf{x}^N) \neq \emptyset.$$

It is obvious from the definition of Φ , that $\forall (\mathbf{x}^1, \dots, \mathbf{x}^N) \in S, \Phi(\mathbf{x}^1, \dots, \mathbf{x}^N) \subseteq S$, and that $(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is an equilibrium of the robust game iff it is a fixed point of Φ . Thus, we need only prove the existence of a fixed point of Φ . Let us therefore establish that Φ satisfies the remaining conditions of Kakutani’s Fixed Point Theorem; that is, we must show that Φ maps S into a family of closed, convex sets, and that Φ is upper semi-continuous.

Let us first prove that, $\forall (\mathbf{x}^1, \dots, \mathbf{x}^N) \in S, \Phi(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is a convex set. Suppose

$$(\mathbf{u}^1, \dots, \mathbf{u}^N), (\mathbf{v}^1, \dots, \mathbf{v}^N) \in \Phi(\mathbf{x}^1, \dots, \mathbf{x}^N).$$

Then, by the definition of $\Phi, \forall i \in \{1, \dots, N\}, \forall \mathbf{y}^i \in S_{a_i}$,

$$\rho_i(\mathbf{x}^{-i}, \mathbf{u}^i) = \rho_i(\mathbf{x}^{-i}, \mathbf{v}^i) \geq \rho_i(\mathbf{x}^{-i}, \mathbf{y}^i).$$

It follows that, $\forall \lambda \in [0, 1], \forall \mathbf{y}^i \in S_{a_i}$,

$$\lambda \rho_i \left(\mathbf{x}^{-i}, \mathbf{u}^i \right) + (1 - \lambda) \rho_i \left(\mathbf{x}^{-i}, \mathbf{v}^i \right) \geq \rho_i \left(\mathbf{x}^{-i}, \mathbf{y}^i \right).$$

By the concavity result of Lemma 4,

$$\lambda \left(\mathbf{u}^1, \dots, \mathbf{u}^N \right) + (1 - \lambda) \left(\mathbf{v}^1, \dots, \mathbf{v}^N \right) \in \Phi \left(\mathbf{x}^1, \dots, \mathbf{x}^N \right).$$

Let us now show that Φ is upper semi-continuous, per Kakutani's definition. Suppose that, for $n = 1, 2, 3, \dots$,

$$\begin{aligned} \left(\mathbf{x}^{1,n}, \dots, \mathbf{x}^{N,n} \right) &\in S \\ \left(\mathbf{y}^{1,n}, \dots, \mathbf{y}^{N,n} \right) &\in \Phi \left(\mathbf{x}^{1,n}, \dots, \mathbf{x}^{N,n} \right) \\ \lim_{n \rightarrow \infty} \left(\mathbf{x}^{1,n}, \dots, \mathbf{x}^{N,n} \right) &= \left(\mathbf{u}^1, \dots, \mathbf{u}^N \right) \in S \\ \lim_{n \rightarrow \infty} \left(\mathbf{y}^{1,n}, \dots, \mathbf{y}^{N,n} \right) &= \left(\mathbf{v}^1, \dots, \mathbf{v}^N \right) \in S. \end{aligned}$$

By the definition of Φ , we know that, $\forall n = 1, 2, 3, \dots, \forall i \in \{1, \dots, N\}$ and $\forall \mathbf{w}^i \in S_{a_i}$,

$$\rho_i \left(\mathbf{x}^{-i,n}, \mathbf{y}^{i,n} \right) \geq \rho_i \left(\mathbf{x}^{-i,n}, \mathbf{w}^i \right).$$

Taking the limit of both sides, and using Lemma 3 (continuity of ρ_i), we obtain that, $\forall i \in \{1, \dots, N\}$ and $\forall \mathbf{w}^i \in S_{a_i}$,

$$\rho_i \left(\mathbf{u}^{-i}, \mathbf{v}^i \right) \geq \rho_i \left(\mathbf{u}^{-i}, \mathbf{w}^i \right).$$

Hence,

$$\left(\mathbf{v}^1, \dots, \mathbf{v}^N \right) \in \Phi \left(\mathbf{u}^1, \dots, \mathbf{u}^N \right),$$

and Φ is upper semi-continuous. Note that the fact that $\Phi \left(\mathbf{x}^1, \dots, \mathbf{x}^N \right)$ is closed follows from the fact that Φ is upper semi-continuous (take $\left(\mathbf{x}^{1,n}, \dots, \mathbf{x}^{N,n} \right) = \left(\mathbf{u}^1, \dots, \mathbf{u}^N \right) = \left(\mathbf{x}^1, \dots, \mathbf{x}^N \right), \forall n = 1, 2, 3, \dots$).

This completes the proof that Φ satisfies the conditions of Kakutani's Fixed Point Theorem, and thereby establishes the existence of an equilibrium in the robust game. \square

5. Computing sample equilibria of robust finite games

In Section 4, we established the existence of robust-optimization equilibria in any finite, incomplete-information game with bounded uncertainty set and no private information. In this section, for any resulting robust game with bounded polyhedral uncertainty set, we show that the set of equilibria is a projection of the set of solutions to a system of multilinear equalities and inequalities. This projection is a simple component-wise one, into a space of lower dimension. Based on this formulation, we present an approximate

computation method for finding a sample equilibrium of such a robust game. We provide numerical results from the application of our method. Finally, we describe a class of robust finite games whose set of equilibria are precisely the set of equilibria in a related complete-information, finite game in the same action spaces. As noted before, in this section, we focus on robust games not involving private information. In Section 7, we provide a more general result on the computation of robust-optimization equilibria in games with private information.

5.1. Review for complete-information, finite games

Before describing our technique for finding robust-optimization equilibria, let us review the state of the art for complete-information, finite games.

Solving for an equilibrium of a general complete-information, finite game is regarded as a difficult task [42]. Two-person, zero-sum games are the exception. As noted in von Stengel [53], in any such game, the set of Nash equilibria is precisely the set of maximinimizers, as defined by von Neumann and Morgenstern [52]. Accordingly, the equilibria are pairs of solutions of two separate LPs, one for each player, and the set of equilibria is therefore convex. For non-fixed-sum games, solving for Nash equilibria is more computationally demanding, and the set of equilibria is generally nonconvex. As discussed in McKelvey and McLennan [35], the set of Nash equilibria can be cast as the solution set of several well-known problems in the optimization literature: fixed point problems, nonlinear complementarity problems (linear in the case of two-player games), stationary point problems, systems of multilinear equalities and inequalities, and unconstrained penalty function minimization problems, in which a penalty is incurred for violations of the multilinear constraints.

Algorithms for finding sample Nash equilibria exploit special properties of these formulations. Traditionally, the favored algorithm for two-player, non-fixed-sum, finite games with complete information has been the Lemke-Howson path-following algorithm [31] for linear complementarity problems. For this more general class of problems, the algorithm's worst-case runtime is exponential. The worst-case runtime in the specific application context of two-person games is unknown. For N -player complete-information, finite games with $N > 2$, the traditionally favored algorithms have been different versions of path-following methods based on Scarf's simplicial subdivision approach [44, 45] to computing fixed points of a continuous function on a compact set. These simplicial subdivision algorithms include that of van der Laan and Talman [49, 50], and their worst-case runtimes are also exponential. The Lemke-Howson and simplicial subdivision algorithms form the backbone of the well-known game theory software Gambit [36].

More recent approaches to solving for sample Nash equilibria have exploited the multilinear system formulation, and have applied general root-finding methods to the complementarity conditions that arise from this system. For an overview of the formulation, we refer the reader to Chapter 6 of [48]. For a comparison of these Gröbner basis and homotopy continuation methods of computation with the more traditional Gambit software, we refer the reader to Datta [13]. Govindan's and Wilson's [23] global Newton method similarly uses the multilinear system formulation. Finally, Porter, Nudelman, and Shoham [43] offer a potential shortcut, which exploits the fact that, for

complete-information games, it is easier to solve for a Nash equilibrium with a fixed support, and that smaller supports yield lower runtimes.

These more recent numerical techniques are more powerful in terms of their aptitudes at computing all equilibria of a complete-information, finite game, a more difficult task than computing a single, sample equilibrium. For example, PHCpack [51] can find all isolated roots of a system of polynomials.

5.2. Robust finite games

5.2.1. Multilinear system formulation for equilibria In this subsection, we show that the set of equilibria of a robust finite game, with bounded polyhedral uncertainty set and no private information, is the projection of the solution set of a system of multilinear equalities and inequalities. The projection is a simple component-wise projection into a space of lower dimension.

As a basis of comparison, for an N -player, complete-information, finite game, in which player i has action space $\{1, \dots, a_i\}$, and in which the payoff matrix is $\tilde{\mathbf{P}}$, let us formulate the multilinear system whose solutions are the set of Nash equilibria. From Condition (5), we see that $(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is a Nash equilibrium iff it satisfies the following system.

$$\begin{aligned} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i - \mathbf{e}_{j_i}^i) &\geq 0, & i = 1, \dots, N; j_i = 1, \dots, a_i \\ \mathbf{e}'\mathbf{x}^i &= 1, & i = 1, \dots, N \\ \mathbf{x}^i &\geq \mathbf{0}, & i = 1, \dots, N, \end{aligned}$$

where \mathbf{e} is the vector, of appropriate dimension, of all ones, and where $\mathbf{e}_{j_i}^i$ denotes the j_i^{th} unit vector in \mathbb{R}^{a_i} .

Analogously, from Condition (7), $(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is an equilibrium of the robust finite game with closed and bounded uncertainty set $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$, and with no private information, iff

$$\begin{aligned} \min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{x}^i) - \max_{\mathbf{u}^i \in S_{a_i}} \min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) &\geq 0, & i = 1, \dots, N \\ \mathbf{e}'\mathbf{x}^i &= 1, & i = 1, \dots, N \\ \mathbf{x}^i &\geq \mathbf{0}, & i = 1, \dots, N. \end{aligned}$$

Stated in another way, $(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is an equilibrium of the robust finite game iff, for each $i \in \{1, \dots, N\}$, \mathbf{x}^i is a max-min strategy in a two-person, zero-sum game between player i and an adversary. In this two-person, zero-sum game, the payoff matrix is determined by \mathbf{x}^{-i} , and the adversary's strategy space is U .

Although the above system may not seem amenable to reformulation as a system of multilinear equalities and inequalities, we establish in Theorem 3 that, when U is a bounded polyhedron, the system can, in fact, be reformulated in this way. Before stating and proving this theorem, let us state and prove the following lemma, inspired by the LP duality proof techniques used in [7].

Lemma 5. Let $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ be a bounded polyhedral set, given by

$$U = \left\{ \tilde{\mathbf{P}} \mid \mathbf{F} \cdot \text{vec}(\tilde{\mathbf{P}}) \geq \mathbf{d} \right\} \neq \emptyset, \tag{11}$$

where

$$\text{vec}(\mathbf{P}) \triangleq \left(P^i_{(j_1, \dots, j_N)} \right)_{i=1, \dots, N; (j_1, \dots, j_N) \in \prod_{i=1}^N \{1, \dots, a_i\}}.$$

Let $\mathbf{G}(\ell)$, $\ell \in \{1, \dots, k\}$, denote the extreme points of U . $\forall i, \forall (\mathbf{x}^{-i}, \mathbf{u}^i) \in S$, the following three conditions are equivalent.

Condition 1) $z_i \leq \min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i)$

Condition 2) $z_i \leq \pi_i(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{u}^i)$, $\ell = 1, \dots, k$.

Condition 3) $\exists \boldsymbol{\eta}^i \in \mathbb{R}^m$ such that

$$\begin{aligned} \mathbf{d}' \boldsymbol{\eta}^i &\geq z_i \\ \mathbf{F}' \boldsymbol{\eta}^i &= \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{u}^i \\ \boldsymbol{\eta}^i &\geq \mathbf{0}, \end{aligned}$$

where $\mathbf{Y}^i(\mathbf{x}^{-i}) \in \mathbb{R}^{(N \prod_{i=1}^N a_i) \times a_i}$ denotes the matrix such that

$$\text{vec}(\mathbf{P})' \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{u}^i = \pi_i(\mathbf{P}; \mathbf{x}^{-i}, \mathbf{u}^i). \tag{12}$$

Proof. Conditions 1 and 2 are equivalent, since by the linearity of π_i in $\tilde{\mathbf{P}}$,

$$\min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) = \min_{\ell \in \{1, \dots, k\}} \pi_i(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{u}^i).$$

To prove the equivalence of Conditions 1 and 3, consider the following primal-dual pair, in which $(\mathbf{x}^{-i}, \mathbf{u}^i)$ is treated as data.

$$\begin{aligned} \min_{\text{vec}(\mathbf{P})} \quad & \text{vec}(\mathbf{P})' \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{u}^i \\ \text{s.t.} \quad & \mathbf{F} \cdot \text{vec}(\mathbf{P}) \geq \mathbf{d} \end{aligned} \tag{13}$$

$$\begin{aligned} \max_{\boldsymbol{\eta}^i} \quad & \mathbf{d}' \boldsymbol{\eta}^i \\ \text{s.t.} \quad & \mathbf{F}' \boldsymbol{\eta}^i = \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{u}^i \\ & \boldsymbol{\eta}^i \geq \mathbf{0}. \end{aligned} \tag{14}$$

Since $U \neq \emptyset$, Problem (13) is feasible. Suppose $(\mathbf{x}^{-i}, \mathbf{u}^i)$ satisfies Condition 1. Then, Problem (13) is also bounded. By strong duality, Problem (14) is feasible and bounded with optimal value equal to that of Problem (13). Thus, Condition 3 is satisfied. For the other direction, suppose Condition 3 is satisfied. Then, Problem (14) is feasible. By weak duality, Condition 1 must hold. \square

Theorem 3 (Computation of Equilibria in Robust Finite Games). *Consider the N -player robust game, in which player $i \in \{1, \dots, N\}$ has action set $\{1, \dots, a_i\}$, $1 < a_i < \infty$, in which the payoff uncertainty set $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ is polyhedral, bounded, and given by (11), and in which there is no private information. Let $\mathbf{G}(\ell)$, $\ell \in \{1, \dots, k\}$, denote the extreme points of U . The following three conditions are equivalent.*

Condition 1 $(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is an equilibrium of the robust game.

Condition 2 For all $i \in \{1, \dots, N\}$, there exists $z_i \in \mathbb{R}$, $\theta^i \in \mathbb{R}^k$, $\phi_i \in \mathbb{R}$ such that $(\mathbf{x}^1, \dots, \mathbf{x}^N, z_i, \theta^i, \phi_i)$ satisfies

$$\begin{aligned}
 z_i &= \phi_i \\
 z_i - \pi_i(\mathbf{G}(\ell); \mathbf{x}^1, \dots, \mathbf{x}^N) &\leq 0, & \ell = 1, \dots, k \\
 \mathbf{e}'\mathbf{x}^i &= 1 \\
 \mathbf{x}^i &\geq \mathbf{0} \\
 \mathbf{e}'\theta^i &= 1 \\
 \sum_{\ell=1}^k \theta_\ell^i \pi_i(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{e}_{j_i}^i) - \phi_i &\leq 0, & j_i = 1, \dots, a_i \\
 \theta^i &\geq \mathbf{0},
 \end{aligned} \tag{15}$$

where \mathbf{e} is the vector, of appropriate dimension, of all ones, and where $\mathbf{e}_{j_i}^i$ is the j_i^{th} unit vector in \mathbb{R}^{a_i} .

Condition 3 For all $i \in \{1, \dots, N\}$, there exists $\boldsymbol{\eta}^i \in \mathbb{R}^m$ and $\boldsymbol{\xi}^i \in \mathbb{R}^N \prod_{i=1}^N a_i$ such that $(\mathbf{x}^1, \dots, \mathbf{x}^N, \boldsymbol{\eta}^i, \boldsymbol{\xi}^i)$ satisfies

$$\begin{aligned}
 (\boldsymbol{\xi}^i)' \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{e}_{j_i}^i &\leq \mathbf{d}'\boldsymbol{\eta}^i, & j_i = 1, \dots, a_i \\
 \mathbf{F}'\boldsymbol{\eta}^i - \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{x}^i &= \mathbf{0} \\
 \mathbf{e}'\mathbf{x}^i &= 1 \\
 \mathbf{x}^i &\geq \mathbf{0} \\
 \boldsymbol{\eta}^i &\geq \mathbf{0} \\
 \mathbf{F}\boldsymbol{\xi}^i &\geq \mathbf{d},
 \end{aligned} \tag{16}$$

where $\mathbf{Y}^i(\mathbf{x}^{-i}) \in \mathbb{R}^{(N \prod_{i=1}^N a_i) \times a_i}$ is as defined in (12).

Proof. Since U is closed and bounded, Condition 1 is equivalent, by Relation (7), to

$$\begin{aligned}
 (\mathbf{x}^1, \dots, \mathbf{x}^N) &\in S \\
 \mathbf{x}^i &\in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) \right], & i = 1, \dots, N.
 \end{aligned}$$

In turn, these constraints are equivalent to the requirement that, $\forall i \in \{1, \dots, N\}, \exists z_i \in \mathbb{R}$ such that (\mathbf{x}^i, z_i) is a maximizer of the following robust LP.

$$\begin{aligned} & \max_{\mathbf{u}^i, z_i} z_i \\ & \text{s.t. } z_i \leq \min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) \\ & \mathbf{e}'\mathbf{u}^i = 1 \\ & \mathbf{u}^i \geq \mathbf{0}. \end{aligned} \tag{17}$$

In this robust LP, \mathbf{x}^{-i} is regarded as given data, and \mathbf{e} denotes the vector, of appropriate dimension, of all ones.

Suppose Condition 1 is satisfied. Then, by Lemma 5, $\exists z_i \in \mathbb{R}$ and $\boldsymbol{\eta}^i \in \mathbb{R}^m$ such that (\mathbf{x}^i, z_i) is a maximizer of

$$\begin{aligned} & \max_{\mathbf{u}^i, z_i} z_i \\ & \text{s.t. } z_i \leq \pi_i(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{u}^i), \quad \ell = 1, \dots, k \\ & \mathbf{e}'\mathbf{u}^i = 1 \\ & \mathbf{u}^i \geq \mathbf{0}, \end{aligned} \tag{18}$$

whose dual is

$$\begin{aligned} & \min_{\boldsymbol{\theta}^i, \phi_i} \phi_i \\ & \text{s.t. } \mathbf{e}'\boldsymbol{\theta}^i = 1 \\ & \sum_{\ell=1}^k \theta_\ell^i \pi_i(\mathbf{G}(\ell); \mathbf{x}^{-i}, \mathbf{e}_{j_i}^i) - \phi_i \leq 0, \quad j_i = 1, \dots, a_i \\ & \boldsymbol{\theta}^i \geq \mathbf{0}, \end{aligned} \tag{19}$$

and such that $(\mathbf{x}^i, \boldsymbol{\eta}^i, z_i)$ is a maximizer of

$$\begin{aligned} & \max_{\mathbf{u}^i, \boldsymbol{\eta}^i, z_i} z_i \\ & \text{s.t. } z_i - \mathbf{d}'\boldsymbol{\eta}^i \leq 0 \\ & \mathbf{F}'\boldsymbol{\eta}^i - \mathbf{Y}^i(\mathbf{x}^{-i})\mathbf{u}^i = 0 \\ & \mathbf{e}'\mathbf{u}^i = 1 \\ & \mathbf{u}^i \geq \mathbf{0} \\ & \boldsymbol{\eta}^i \geq \mathbf{0}, \end{aligned} \tag{20}$$

whose dual is

$$\begin{aligned} & \min_{\boldsymbol{\xi}^i, v_i} v_i \\ & \text{s.t. } \mathbf{F}\boldsymbol{\xi}^i \geq \mathbf{d} \\ & v_i \geq (\boldsymbol{\xi}^i)' \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{e}_{j_i}^i, \quad j_i = 1, \dots, a_i. \end{aligned} \tag{21}$$

Conditions 2 and 3 follow from LP strong duality.

For the reverse direction, suppose that Condition 2 holds. Then, for all $i \in \{1, \dots, N\}$, and for \mathbf{x}^{-i} , (\mathbf{x}^i, z_i) is a feasible solution of (18), and $(\boldsymbol{\theta}^i, \phi_i)$ is a feasible solution of (19), such that $z_i = \phi_i$. By LP weak duality, (\mathbf{x}^i, z_i) is an optimizer of (18). Equivalently, by Lemma 5, (\mathbf{x}^i, z_i) is an optimizer of (17), and Condition 1 follows.

Similarly, suppose that Condition 3 holds. $\forall i \in \{1, \dots, N\}$, let

$$z_i = \mathbf{d}'\boldsymbol{\eta}^i$$

$$v_i = \max_{j_i \in \{1, \dots, a_i\}} (\boldsymbol{\xi}^i)' \mathbf{Y}^i(\mathbf{x}^{-i}) \mathbf{e}_{j_i}^i.$$

Then, for \mathbf{x}^{-i} , $(\mathbf{x}^i, \boldsymbol{\eta}^i, z_i)$ is a feasible solution of (20) and $(\boldsymbol{\xi}^i, v_i)$ is a feasible solution of (21) such that $z_i \geq v_i$. By LP weak duality, $(\mathbf{x}^i, \boldsymbol{\eta}^i, z_i)$ is an optimizer of (20). Equivalently, by Lemma 5, (\mathbf{x}^i, z_i) is an optimizer of (17), and Condition 1 follows. \square

Remark. Note that Systems (15) and (16) are derived using the extreme-point and constraint representations of the polyhedral set U , respectively. These systems are very sparse as a result of their multilinearity. In addition, it is possible to formulate System (16) more compactly if U can be described by m constraints and only v variables, with $v < N \prod_{i=1}^N a_i$. Let $a_{\text{tot}} = \sum_{i=1}^N a_i$, v and m be the number of variables and constraints, respectively, needed to define U , and k be the number of extreme points of U . Table 1 summarizes the sizes of the different multilinear systems of equalities and inequalities whose solution sets are precisely the set of equilibria of an N -player game in which player $i \in \{1, \dots, N\}$ has action space $\{1, \dots, a_i\}$.

5.2.2. Computation method For robust finite games with bounded polyhedral uncertainty sets and no private information, we showed in Section 5.2.1 that the set of equilibria is a projection of the solution set of a system of multilinear equalities and inequalities. This projection is a simple component-wise projection into a space of lower dimension. Currently available and computationally effective solvers for large polynomial systems tend to be specific to systems of *equations* and not inequalities. Accordingly, we propose to solve the multilinear systems for the robust-optimization equilibria by converting any such system into a corresponding penalty function, and then solving the resulting unconstrained minimization problem. The penalty method we use is based on Courant’s quadratic loss technique [11], which Fiacco and McCormick later more fully developed in [18].

To more concretely describe our approach, consider any system

$$\begin{aligned} g_n(\mathbf{y}) &= 0, & n &\in E \\ g_n(\mathbf{y}) &\leq 0, & n &\in I, \end{aligned} \tag{22}$$

Table 1. Sizes of multilinear systems for equilibria

	Robust game using constraints	Robust game using extreme points	Complete-info game
Variables	$a_{\text{tot}} + N(m + v)$	$a_{\text{tot}} + N(k + 2)$	a_{tot}
Constraints	$2a_{\text{tot}} + N(2m + v + 1)$	$2a_{\text{tot}} + N(2k + 3)$	$2a_{\text{tot}} + N$
Maximum degree	N	N	N

with $\mathbf{y} \in \mathbb{R}^V$, $|I| < \infty$, and $|E| < \infty$. Let

$$h(\mathbf{y}) = \frac{1}{2} \sum_{n \in E} [g_n(\mathbf{y})]^2 + \frac{1}{2} \sum_{n \in I} [\max \{g_n(\mathbf{y}), 0\}]^2.$$

Since $h(\mathbf{y}) \geq 0, \forall \mathbf{y} \in \mathbb{R}^V$, it is easy to see that \mathbf{y} satisfies System (22) iff

$$h(\mathbf{y}) = \min_{\mathbf{u} \in \mathbb{R}^V} h(\mathbf{u}) = 0.$$

So, we can solve System (22) by solving the unconstrained minimization problem

$$\min_{\mathbf{u} \in \mathbb{R}^V} h(\mathbf{u}).$$

For the unconstrained minimization problem, we propose the use of a pseudo-Newton method using the Armijo rule (see, for example, [4]) for determining step size at each iteration. Each pseudo-Newton method run attempts to find a single point satisfying the constraints. It is possible, though not guaranteed, that, when the system of constraints has more than one solution, multiple pseudo-Newton method runs may identify multiple, distinct approximate solutions of the system. Furthermore, in contrast to most other state-of-the-art polynomial system solvers, this method is capable of finding non-isolated, as well as isolated solutions.

In the next subsection, we present numerical results from the implementation of this technique for approximately computing a sample robust-optimization equilibrium.

5.2.3. Numerical results For each problem instance, we formulated the set of equilibria using System (15). We executed all computations in MATLAB 6.5.0 R13, running on the Red Hat Linux 7.2-1 operating system, on a Dell with a Pentium IV processor, at 1.7 GHz with 512 MB RAM. To encourage the numerical method to find points satisfying the nonnegativity and normalization constraints on \mathbf{x}^i and $\theta^i, i \in \{1, \dots, N\}$, we multiplied the amount of violation of each such constraint by $M = 100$, before halving the square of this violation. We initialized all runs of the pseudo-Newton method by, for each $i \in \{1, \dots, N\}$, randomly generating \mathbf{x}^i and θ^i , satisfying the aforementioned nonnegativity and normalization constraints. We initialized z_i to be the maximum possible value satisfying the upper-bound constraint on z_i , and we set ϕ_i either equal to z_i or to the minimum possible value allowed by the lower-bound constraint on ϕ_i . For each pseudo-Newton method run, we terminated the run if the current and previous iterate were too close, if the norm of the direction was too small, if the penalty was already sufficiently small, or if the number of iterations already executed was too large.

We executed the method on the robust inspection game, described in Example 1 in Section 3.5, with

$$\begin{array}{llll} \underline{g} = 8, & \underline{v} = 16, & \underline{h} = 4, & w = 15, \\ \bar{g} = 12, & \bar{v} = 24, & \bar{h} = 6. & \end{array}$$

The multilinear system for the equilibria of this robust game has 22 constraints in 10 variables, after elimination of some redundant variables. We terminated the pseudo-Newton method run once the penalty function dipped below 10^{-8} . As will follow from

Theorem 4, in the unique equilibrium of this robust game, the employee (row player) shirks (plays action 1) with probability $x_1^1 = 2/5$, and the employer (column player) inspects (plays action 1) with probability $x_1^2 = 4/5$. Our numerical method terminates at $(x_1^1, x_2^2) = (0.4000, 0.8000)$ after 0.5000 seconds of one pseudo-Newton run, requiring 71 iterations.

In addition, we executed the method on the robust free-rider game, described in Example 2 in Section 3.5, with

$$\underline{c} = 1/4, \quad \bar{c} = 5/8.$$

The multilinear system for the equilibria of this robust game has 18 constraints in 8 variables, after elimination of some redundant variables. We terminated each pseudo-Newton method run once the penalty function dipped below 10^{-10} or the number of iterations reached 2000. We used $M = 1$, since the method did not seem to be attracted to strategy profiles outside of the simplex. Let x_1^i denote the probability with which player $i \in \{1, 2\}$ contributes. As will follow from Theorem 4, this robust game has 3 equilibria (x_1^1, x_1^2) : $(1, 0)$, $(0, 1)$, and $(1 - \bar{c}, 1 - \bar{c}) = (3/8, 3/8)$. We made 15 sequential runs of the pseudo-Newton method, each initialized at a randomly generated point. These 15 runs required 1.8458 minutes, with each run executing an average of 1,652.1 iterations in an average of 7.3827 seconds. Terminal points with penalty function less than 10^{-10} included $(0.0000, 0.9999)$, $(0.9999, 0.0000)$, $(0.3750, 0.3751)$, and $(0.3751, 0.3750)$. This example demonstrates that the method is capable of finding multiple equilibria, and possibly all equilibria, of a robust game.

Lastly, we executed the method on several instances of the robust network routing game, described in Example 3 in Section 3.5. The instances differ in terms of their values of N , the number of players, and a , the number of paths available. The resulting versions of System (15) consist of $2N^2 + N(2a + 3)$ constraints in $N^2 + N(2 + a)$ variables.

For all the instances, we used the same values for D and λ . In particular, we set $D = 5$ and $\lambda_{(i,j_i)}$ to be a realization of the uniform distribution on $[0, 4]$. The computational results for these robust network routing games are summarized in Table 2. For each instance, we made only one run of the pseudo-Newton method, and terminated it after the lesser of 50 iterations or the minimum number of iterations required to produce an iterate with associated penalty less than 10^{-5} . The “vars” and “constr’s” columns in

Table 2. Numerical results for instances of robust network routing game

	Vars	Constr's	Cpu time (mins)	Iters	Penalty	Proportional error
$N = 2, a = 2$	12	22	0.0612	37	7.9811×10^{-6}	1.8850×10^{-7}
$N = 3, a = 2$	21	39	0.3887	17	5.5489×10^{-8}	8.0825×10^{-10}
$N = 3, a = 3$	24	45	1.4572	50	9.7746×10^{-3}	2.2952×10^{-4}
$N = 4, a = 2$	32	60	3.4895	50	9.2659×10^{-3}	1.3192×10^{-4}
$N = 4, a = 3$	36	68	3.8935	50	1.2910×10^{-1}	3.0427×10^{-3}
$N = 4, a = 4$	40	76	4.7893	50	3.8569	1.4566×10^{-1}
$N = 5, a = 2$	45	85	7.3268	50	5.7834×10^{-1}	7.7508×10^{-3}
$N = 5, a = 3$	50	95	9.2945	50	1.3551	2.5495×10^{-2}
$N = 5, a = 4$	55	105	12.6322	50	3.4239	8.7083×10^{-2}
$N = 5, a = 5$	60	115	17.6880	50	15.3203	5.9287×10^{-1}

Table 2 give the number of variables and constraints, respectively, in System (15) for each problem instance. The “iters” column gives the number of iterations executed. The “penalty” column gives the penalty value of the final iterate. Finally, the “proportional error” column gives

$$\frac{\text{penalty}}{\min_{i \in \{1, \dots, N\}} \min \left\{ |\hat{z}_i|, |\hat{\phi}_i| \right\}},$$

where \hat{z}_i and $\hat{\phi}_i$ denote the values of z_i and ϕ_i in the final iterate. We could obviously achieve better speed or accuracy by varying the cap on the number of iterations and the penalty threshold used to decide whether to terminate the pseudo-Newton method run.

These numerical results demonstrate that a practical method, simple in nature and general in its applicability, exists for approximately solving, with considerable accuracy and speed, for sample equilibria of robust games of small size. Furthermore, with longer runtimes and lower accuracy, this method may be capable of finding solutions for robust finite games of larger size.

5.2.4. A special class of Robust finite games Under certain conditions, the set of equilibria of a robust finite game is equivalent to that of a related finite game with complete payoff information, with the same number of players, and with the same action spaces. In these cases, equilibria computation for the robust game will reduce to computation in the context of the related complete-information finite game. As shown in Table 1, the multilinear systems arising from robust finite games are larger than those arising from complete-information, finite games with the same number of players and with the same action spaces. Thus, it will be computationally beneficial to take advantage of this equivalence when it holds. As we will discuss in Section 6, the complete-information equivalent of the robust finite game will generally not be the nominal (i.e., average) version of the robust game.

The following theorem establishes sufficient conditions for the equivalence of a robust finite game with a complete-information finite game having the same number of players and the same action spaces.

Theorem 4. *Consider the robust finite game, without private information and in which the payoff uncertainty set is*

$$U = \left\{ \mathbf{P} \left(\tilde{f}_1, \dots, \tilde{f}_v \right) \mid \left(\tilde{f}_1, \dots, \tilde{f}_v \right) \in U_f \right\}, \tag{23}$$

where

$$U_f = \left\{ \left(\tilde{f}_1, \dots, \tilde{f}_v \right) \mid \tilde{f}_\ell \in \left[\underline{f}_\ell, \overline{f}_\ell \right], \ell \in \{1, \dots, v\} \right\}, \tag{24}$$

and \mathbf{P} is a continuous and differentiable vector function. Suppose that, for all $i \in \{1, \dots, N\}$ and $\forall \ell \in \{1, \dots, v\}$, $\exists \kappa(i, \ell) \in \{-1, 0, 1\}$ such that, $\forall (j_1, \dots, j_N) \in \prod_{i=1}^N \{1, \dots, a_i\}$ and $\forall \left(\tilde{f}_1, \dots, \tilde{f}_v \right) \in U_f$,

$$\text{sign} \left(\frac{\partial}{\partial f_\ell} \left[P_{(j_1, \dots, j_N)}^i \left(f_1, \dots, f_v \right) \right]_{(f_1, \dots, f_v) = (\tilde{f}_1, \dots, \tilde{f}_v)} \right) = \kappa(i, \ell).$$

Then $(\mathbf{x}^1, \dots, \mathbf{x}^N)$ is an equilibrium of this robust game iff it is a Nash equilibrium of the complete-information, finite game, with the same number of players and the same action spaces, and in which the payoff matrix is \mathbf{Q} , defined by

$$\mathbf{Q}_{(j_1, \dots, j_N)}^i = P_{(j_1, \dots, j_N)}^i(h_1^i, \dots, h_v^i)$$

$$h_\ell^i = \begin{cases} \bar{f}_\ell, & \kappa(i, \ell) < 0 \\ \underline{f}_\ell, & \kappa(i, \ell) \geq 0. \end{cases}$$

Proof. Let

$$A_i \triangleq \{1, \dots, a_i\}$$

$$A \triangleq \prod_{i=1}^N A_i.$$

$\mathbf{Q} \in U$ implies that, $\forall i \in \{1, \dots, N\}, \forall (\mathbf{x}^{-i}, \mathbf{u}^i) \in S$,

$$\pi_i(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i) \geq \min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i).$$

Conversely, by the definition of \mathbf{h} , $\forall i \in \{1, \dots, N\}, \forall (j_1, \dots, j_N) \in A$,

$$P_{(j_1, \dots, j_N)}^i(h_1^i, \dots, h_v^i) \leq P_{(j_1, \dots, j_N)}^i(f_1, \dots, f_v), \quad \forall (f_1, \dots, f_v) \in U_f.$$

Thus, $\forall i \in \{1, \dots, N\}, \forall (\mathbf{x}^{-i}, \mathbf{u}^i) \in S$,

$$\begin{aligned} \min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) &= \min_{\tilde{\mathbf{f}} \in U_f} \pi_i(\mathbf{P}(\tilde{f}_1, \dots, \tilde{f}_v); \mathbf{x}^{-i}, \mathbf{u}^i) \\ &\geq \sum_{j_1=1}^{a_1} \cdots \sum_{j_i=1}^{a_i} \cdots \sum_{j_N=1}^{a_N} \left(\prod_{\substack{i'=1 \\ i' \neq i}}^N x_{j_{i'}}^{i'} \right) u_{j_i}^i \min_{\tilde{\mathbf{f}} \in U_f} P_{(j_1, \dots, j_N)}^i(\tilde{f}_1, \dots, \tilde{f}_v) \\ &= \sum_{j_1=1}^{a_1} \cdots \sum_{j_i=1}^{a_i} \cdots \sum_{j_N=1}^{a_N} \left(\prod_{\substack{i'=1 \\ i' \neq i}}^N x_{j_{i'}}^{i'} \right) u_{j_i}^i P_{(j_1, \dots, j_N)}^i(h_1^i, \dots, h_v^i) \\ &= \pi_i(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i). \end{aligned}$$

Therefore, $\forall i \in \{1, \dots, N\}, \forall (\mathbf{x}^{-i}, \mathbf{u}^i) \in S$,

$$\min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) = \pi_i(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i).$$

By Relation (7), $(\mathbf{x}^1, \dots, \mathbf{x}^N) \in S$ is an equilibrium of the robust finite game iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{x}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\min_{\tilde{\mathbf{P}} \in U} \pi_i(\tilde{\mathbf{P}}; \mathbf{x}^{-i}, \mathbf{u}^i) \right] = \arg \max_{\mathbf{u}^i \in S_{a_i}} \left[\pi_i(\mathbf{Q}; \mathbf{x}^{-i}, \mathbf{u}^i) \right].$$

□

Let us give an example of an application of Theorem 4. For $i \in \{1, \dots, N\}$, let $I(i)^+$ and $I(i)^-$ form a partition of $\{1, \dots, v\}$. Consider U given as in Theorem 4, with the function \mathbf{P} defined as follows. $\forall i \in \{1, \dots, N\}$ and $\forall (j_1, \dots, j_N) \in A$,

$$P_{(j_1, \dots, j_N)}^i(f_1, \dots, f_v) = \sum_{\ell \in I(i)^+} \gamma_{(j_1, \dots, j_N)}^{i, \ell} f_\ell - \sum_{\ell \in I(i)^-} \gamma_{(j_1, \dots, j_N)}^{i, \ell} f_\ell$$

$$\gamma_{(j_1, \dots, j_N)}^{i, \ell} \geq 0, \quad \ell = 1, \dots, v.$$

Then

$$Q_{(j_1, \dots, j_N)}^i = \sum_{\ell \in I(i)^+} \gamma_{(j_1, \dots, j_N)}^{i, \ell} \underline{f}_\ell - \sum_{\ell \in I(i)^-} \gamma_{(j_1, \dots, j_N)}^{i, \ell} \overline{f}_\ell.$$

6. Comparison of robust and Bayesian finite games

Having established, in Section 5, a computation method for identifying equilibria of robust finite games without private information, in this section, using illustrative examples, we compare properties of these robust games with those of their nominal-game counterparts. By Equation (4), each nominal game we present is in fact equivalent to the Bayesian game that assigns a symmetric distribution to the uncertainty set in the corresponding robust game. Thus, our comparisons can be said to be between robust games and these corresponding Bayesian games.

In this same vein of comparison, turning our attention to a notion of symmetry unrelated to the symmetry of probability distributions, we end this section by discussing symmetric robust games, i.e., those in which the players are indistinguishable with respect to the game structure. We prove the existence of symmetric, robust-optimization equilibria in these games, thereby establishing a result analogous to Nash’s existence theorem for symmetric Nash equilibria of symmetric, complete-information, finite games [40].

6.1. Equilibria sets are generally not equivalent

The set of equilibria of a robust finite game and that of its nominal counterpart, e.g., the Bayesian game which assigns a symmetric distribution to the uncertainty set, may be disjoint. For example, consider the two-player inspection game presented in Example 1 in Section 3.5, with

$$\begin{aligned} \tilde{g} &\in [8, 12], & \tilde{v} &\in [16, 24], & \tilde{h} &\in [4, 6], & w &= 15, \\ \check{g} &= 10, & \check{v} &\in 20, & \check{h} &= 5. \end{aligned}$$

The nominal version of the game has payoff matrix

$$\left(\begin{array}{cc} (0, -\check{h}) & (w, -w) \\ (w - \check{g}, \check{v} - w - \check{h}) & (w - \check{g}, \check{v} - w) \end{array} \right) = \left(\begin{array}{cc} (0, -5) & (15, -15) \\ (5, 0) & (5, 5) \end{array} \right).$$

For the values given above, the nominal game has a unique equilibrium, in which the employee shirks with probability 1/3 and the employer inspects with probability 2/3.

In contrast, by Theorem 4, the robust game is equivalent to the complete-information inspection game with payoff matrix

$$\begin{pmatrix} (0, -\bar{h}) & (w, -w) \\ (w - \bar{g}, \underline{v} - w - \bar{h}) & (w - \bar{g}, \underline{v} - w) \end{pmatrix} = \begin{pmatrix} (0, -6) & (15, -15) \\ (3, -5) & (3, 1) \end{pmatrix}.$$

Thus, the robust game has a different, unique equilibrium, in which the employee shirks with probability $2/5$ and the employer inspects with probability $4/5$.

It is not surprising that the worker would shirk with higher probability and the employer would inspect with higher probability in the robust game than in the nominal game (i.e., in the Bayesian game assigning a symmetric distribution over the uncertainty set). Indeed, in moving from the average parameter values, as used in the nominal game, to the worst-case parameter values, as used in the robust game, the employee’s opportunity cost of working increases, and the employer’s cost of inspecting increases. As the employee’s opportunity cost of working increases, the employer expects that the employee will be less willing to work. In order to make the employee indifferent between shirking and working, the employer must therefore be more prone to inspect, despite her higher inspection cost. Conversely, as the employer’s cost of inspecting increases, the employee expects that the employer will be less willing to inspect. In order to make the employer indifferent between inspecting and not inspecting, the employee must therefore be more prone to shirk.

6.2. *Sizes of sets of equilibria*

The set of equilibria of a robust finite game may be smaller or larger than that of the corresponding Bayesian game assigning a symmetric distribution over the uncertainty set. For an extreme example in which the set of equilibria of a robust finite game is smaller than that of the nominal-game counterpart, consider the robust game without private information and with payoff uncertainty set

$$\left\{ \begin{pmatrix} (2, \tilde{f}) & (\tilde{f}, 2) \\ (\tilde{f}, 2) & (2, \tilde{f}) \end{pmatrix} \mid \tilde{f} \in [0, 4] \right\}.$$

Consider the nominal version of the game in which $\tilde{f} = \check{f} = 2$ is commonly known with certainty by the players. In this game, all pairs of mixed strategies for the two players are Nash equilibria. In contrast, by Theorem 4, the robust game is equivalent to the complete-information game with payoff matrix

$$\begin{pmatrix} (2, 0) & (0, 2) \\ (0, 2) & (2, 0) \end{pmatrix},$$

i.e., is equivalent to the classical, complete-information game of matching pennies (see, for example, [20]), and therefore has a unique equilibrium. In moving from the robust game to its Bayesian counterpart, the set of equilibria shrinks, because the payoff uncertainty results in reduced indifference, by each player, between his two actions.

Conversely, for an equally extreme example in which the set of equilibria of a robust finite game is larger than that of the corresponding nominal game, consider the robust game without private information and with payoff uncertainty set

$$\left\{ \left(\begin{array}{cc} (\tilde{f}_1, \tilde{f}_2) & (\tilde{f}_2, \tilde{f}_1) \\ (\tilde{f}_2, \tilde{f}_1) & (\tilde{f}_1, \tilde{f}_2) \end{array} \right) \mid (\tilde{f}_1, \tilde{f}_2) \in [0, 8] \times [0, 4] \right\}.$$

Consider the nominal version of the game in which $(\check{f}_1, \check{f}_2) = (\check{f}_1, \check{f}_2) = (4, 2)$ is commonly known with certainty by the players. This nominal game is now equivalent to the complete-information game of matching pennies and therefore has a unique equilibrium. In contrast, by Theorem 4, the robust game is equivalent to the complete-information game with payoff matrix

$$\left(\begin{array}{cc} (0, 0) & (0, 0) \\ (0, 0) & (0, 0) \end{array} \right).$$

Thus, all pairs of mixed strategies for the two players are equilibria of the robust game. In moving from the robust game to its Bayesian counterpart, the set of equilibria expands, because the payoff uncertainty results in increased indifference, by each player, between his two actions.

6.3. Zero-sum becomes non-fixed-sum under uncertainty

In general, if we subject to uncertainty the payoff matrix in a zero-sum game, the resulting robust game will not be a fixed-sum game. For example, consider the payoff uncertainty set

$$\left\{ \left(\begin{array}{cc} (\tilde{f}_1, -\tilde{f}_1) & (\tilde{f}_2, -\tilde{f}_2) \\ (\tilde{f}_3, -\tilde{f}_3) & (\tilde{f}_4, -\tilde{f}_4) \end{array} \right) \mid (\tilde{f}_1, \tilde{f}_2, \tilde{f}_3, \tilde{f}_4) \in \prod_{\ell=1}^4 [f_{\underline{\ell}}, \overline{f_{\ell}}] \right\}.$$

In the nominal version of this game, the players commonly know with certainty that

$$(\tilde{f}_1, \tilde{f}_2, \tilde{f}_3, \tilde{f}_4) = (\check{f}_1, \check{f}_2, \check{f}_3, \check{f}_4),$$

for some $(\check{f}_1, \check{f}_2, \check{f}_3, \check{f}_4) \in \prod_{\ell=1}^4 [f_{\underline{\ell}}, \overline{f_{\ell}}]$. In contrast, by Theorem 4, the robust game is equivalent to the complete-information game, with payoff matrix

$$\left(\begin{array}{cc} (\underline{f}_1, -\overline{f}_1) & (\underline{f}_2, -\overline{f}_2) \\ (\underline{f}_3, -\overline{f}_3) & (\underline{f}_4, -\overline{f}_4) \end{array} \right),$$

which is not fixed-sum unless $f_{\underline{\ell}} - \overline{f_{\ell}}$ is constant for $\ell \in \{1, 2, 3, 4\}$. This result is not surprising, since the two players' worst-case perspectives need not agree.

6.4. Symmetric robust games and symmetric equilibria

Let us turn our attention to symmetric games and their symmetric equilibria, which comprise an important topic in the game theory literature. We end this section by showing that symmetric equilibria are guaranteed to exist in symmetric, robust finite games, just as they are in symmetric, complete-information, finite games.

Stated very generally, a symmetric game is one in which the players are indistinguishable with respect to the game’s structure (action and strategy spaces, payoff functions, information, etc.). More formally, we have the following definition.

Definition 6. A finite game with complete information is said to be **symmetric** if all players have the same action space, all players’ payoff functions are invariant under permutations of the other players’ actions, and all players’ payoff functions are equivalent. That is, a complete-information game is symmetric if

$$a_i = a, \quad i = 1, \dots, N$$

$$P^i_{(j_{-i}, j_i)} = P^{i'}_{(j_{\sigma(-i)}, j_i)}, \quad i, i' = 1, \dots, N; \forall (j_{-i}, j_i) \in \{1, \dots, a\}^N; \forall \sigma \in \Sigma_{N-1},$$

where

$$(j_{-i}, j) \triangleq (j_1, \dots, j_{i-1}, j, j_{i+1}, \dots, j_N)$$

$$(j_{\sigma(-i)}, j) \triangleq (j_{\sigma(1)}, \dots, j_{\sigma(i-1)}, j, j_{\sigma(i+1)}, \dots, j_{\sigma(N)}),$$

and Σ_{N-1} denotes the set of permutations of $N - 1$ elements.

A tuple of players’ strategies will be said to be **symmetric** if all players’ strategies in the tuple are identical. In particular, a **symmetric equilibrium** refers to an equilibrium in which all players play the same strategy.

Similarly, this definition extends, as follows, to robust finite games.

Definition 7. A robust finite game with uncertainty set $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ and no private information is said to be **symmetric** if

$$a_i = a, \quad i = 1, \dots, N$$

$$\rho_i(\mathbf{x}^{-i}, \mathbf{x}^i) = \rho_{i'}(\mathbf{x}^{\sigma(-i)}, \mathbf{x}^i), \quad i, i' = 1, \dots, N; \forall (\mathbf{x}^{-i}, \mathbf{x}^i) \in S; \forall \sigma \in \Sigma_{N-1},$$

where $(\mathbf{x}^{\sigma(-i)}, \mathbf{x}^i)$ denotes $(\mathbf{x}^{\sigma(1)}, \dots, \mathbf{x}^{\sigma(i-1)}, \mathbf{x}^i, \mathbf{x}^{\sigma(i+1)}, \dots, \mathbf{x}^{\sigma(N)})$.

Accordingly, for example, the robust game presented in Example 2 of Section 3.5 is symmetric.

In [40], Nash proved the existence of symmetric equilibria in symmetric, finite games with complete information. We state and prove the following analogous existence result for robust games.

Theorem 5 (Existence of Symmetric Equilibria in Symmetric Robust Finite Games). Any N -person, non-cooperative, simultaneous-move, one-shot, symmetric robust game, in which $N < \infty$, in which each player $i \in \{1, \dots, N\}$ has $1 < a < \infty$ possible actions, in which the uncertainty set of payoff matrices $U \subseteq \mathbb{R}^{Na^N}$ is bounded, and in which there is no private information, has a symmetric equilibrium.

Proof. By the definition of symmetry of a robust game, there exists a function $\rho : S \rightarrow \mathbb{R}$ such that $\rho \equiv \rho_i, \forall i \in \{1, \dots, N\}$. Now define $\Phi : S \rightarrow 2^S$ as

$$\Phi(\mathbf{x}) = \left\{ \mathbf{y} \in S_a \mid \mathbf{y} \in \arg \max_{\mathbf{u} \in S_a} \rho(\mathbf{x}^{-i}, \mathbf{u}) \right\},$$

where \mathbf{x}^{-i} denotes the $(N - 1)$ -tuple $(\mathbf{x}, \mathbf{x}, \dots, \mathbf{x})$. The N -tuple $(\mathbf{x}, \mathbf{x}, \dots, \mathbf{x}) \in S$ is a symmetric equilibrium of the robust game iff \mathbf{x} is a fixed point of Φ . From an argument paralleling that given in the proof of Theorem 2, it follows that Φ satisfies Kakutani's Fixed Point Theorem. □

Symmetric games with incomplete information may be of particular interest for two reasons. First, incomplete-information games, in which the players are indistinguishable with respect to the game structure, may be especially amenable to the common prior assumption in Harsanyi's model and to its analog, the assumption of a common uncertainty set, in our robust game model. Second, the multilinear system formulations for symmetric equilibria of symmetric, robust finite games are smaller, by a factor of N , than those for the general equilibria of these games. Indeed, in systems (15) and (16), if we replace $\mathbf{x}^i, i \in \{1, \dots, N\}$, by the single $a \times 1$ vector variable \mathbf{x} , subsequent elimination of redundancies then reduces the number of variables and constraints in these systems by a factor of N . Thus, we may be able to compute symmetric equilibria of symmetric, robust finite games more quickly and accurately, and with less computational effort, than we can compute the general equilibria of these games.

7. Robust games with private information

In the preceding sections, we proposed a robust optimization approach and a corresponding distribution-free equilibrium concept for modeling games with incomplete information. We proved existence and computation results. Until now, we have focused on incomplete-information games without private information. In this section, we extend our discussion to the general case, involving potentially private information.

7.1. Extension of the model

As in the preceding sections of the paper, consider an N -person, incomplete-information game, in which player $i \in \{1, \dots, N\}$ has $a_i < \infty$ possible actions, and in which each player is in some way uncertain of the multi-dimensional payoff matrix $\tilde{\mathbf{P}}$ that parameterizes the expected payoff vector function $\boldsymbol{\pi}$. Suppose that each player may have private information about $\tilde{\mathbf{P}}$ or about the other players' beliefs. For each player $i \in \{1, \dots, N\}$, his potentially private information may be encoded in his "type" θ_i . Since the information is potentially private, player i may be uncertain of the type $\theta_{i'}$ of player $i', i' \neq i$. Let U denote, as before, a set of possible payoff matrices \mathbf{P} . Let Θ_i denote the set of possible types of player $i \in \{1, \dots, N\}$, and $\Theta = \prod_{i=1}^N \Theta_i$.

In using separate notation for the unknown payoff parameters $\tilde{\mathbf{P}}$ and the players' types $\boldsymbol{\theta}$, we make explicit the difference between the actual payoff parameters and the

players’ beliefs about these parameters and about the other players’ convictions. In addition, this notation allows us to very clearly address the situation in which players may both possess private information and yet still be uncertain of the parameters affecting their own payoffs. In fact, the model we propose in this section is sufficiently flexible to simultaneously capture the case of no private information (the Θ_i are singletons, $\forall i \in \{1, \dots, N\}$), the differential information setting involving all-but-self uncertainty (each $\theta_i \in \Theta_i$ is consistent with only a single $\tilde{\mathbf{P}}^i$), and the aforementioned differential information case in which agents may possess private information, while also being uncertain of both their own and others’ payoff functions.

In the same spirit as does Harsanyi, we assume that the players commonly know a “prior” set $V \subseteq U \times \Theta$ of realizable tuples of payoff parameters and type vectors. While Harsanyi furthermore assumes, in terms of our notation, that the players commonly know a distribution over this set V , we assume that the players lack such distributional information or have chosen not to use it. Player i ’s type θ_i induces the subset $V_i(\theta_i)$ of V consistent with θ_i ,

$$V_i(\theta_i) = \{(\mathbf{P}, \boldsymbol{\theta}_{-i}, \theta_i) \in V\}.$$

That is, $V_i(\theta_i)$ gives the set of tuples of payoff matrices and type vectors that player i , when he is of type θ_i , believes are possible. As does Harsanyi, throughout the remainder of this section, we require that $\bigcap_{i=1}^N V_i(\theta_i) \neq \emptyset, \forall \boldsymbol{\theta} \in \Theta$ such that $\{(\mathbf{P}, \boldsymbol{\theta}) \in V\} \neq \emptyset$, and that the true payoff matrix $\tilde{\mathbf{P}}$ belongs to the projection of $\bigcap_{i=1}^N V_i(\theta_i)$ onto U . The first requirement ensures that the players’ beliefs are consistent, and implies that $\boldsymbol{\theta}$ belongs to the projection of $\bigcap_{i=1}^N V_i(\theta_i)$ onto Θ , i.e., that the players believe that the true type vector is possible. The second requirement ensures that the players believe that the true payoff matrix is possible.

In the private information setting, for $i \in \{1, \dots, N\}$, player i ’s pure strategies are mappings from his type θ_i to his action space $\{1, \dots, a_i\}$. His so-called behavioral strategies (see, for example, Chapter 3 of [20] for an introduction to behavioral strategies) are mappings from his type θ_i to probability distributions over his action space $\{1, \dots, a_i\}$. More formally, we denote a behavioral strategy for player i by $\mathbf{b}^i : \Theta_i \rightarrow S_{a_i}$. That is, under behavioral strategy \mathbf{b}^i , if player i is of type θ_i , then he plays action $j_i \in \{1, \dots, a_i\}$ with probability $\mathbf{b}_{j_i}^i(\theta_i)$. Let us define the notation

$$\begin{aligned} B_{a_i} &\triangleq \left\{ \mathbf{b}^i : \Theta_i \rightarrow S_{a_i} \right\} \\ B &\triangleq \prod_{i=1}^N B_{a_i} \\ B_{-i} &\triangleq \prod_{\substack{i'=1 \\ i' \neq i}}^N B_{a_{i'}} \\ \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}) &\triangleq \left(\mathbf{b}^1(\theta_1), \dots, \mathbf{b}^{i-1}(\theta_{i-1}), \mathbf{b}^{i+1}(\theta_{i+1}), \dots, \mathbf{b}^N(\theta_N) \right) \\ \left(\mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{b}^i(\theta_i) \right) &\triangleq \left(\mathbf{b}^1(\theta_1), \dots, \mathbf{b}^N(\theta_N) \right). \end{aligned}$$

Recall that, in Harsanyi’s model, each player seeks to optimize his average performance, i.e., his average expected payoff, where the average is taken with respect to a probability distribution over $V_i(\theta_i)$. That is, in terms of our notation, in Harsanyi’s model, the set of best responses by player $i \in \{1, \dots, N\}$, when he is of type $\theta_i \in \Theta_i$, to $\mathbf{b}^{-i}(\cdot)$ is given by

$$\arg \max_{\mathbf{u}^i \in S_{a_i}} \left(E_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \mid \theta_i \right] \right),$$

where the expectation is taken with respect to the conditional probability distribution induced by θ_i over V . We use the notation $\mathbf{b}^{-i}(\cdot)$ to highlight the fact that \mathbf{b}^{-i} is a function. Since the best response correspondence completely determines the criterion for equilibrium, it follows that the tuple of behavioral strategies $(\mathbf{b}^1(\cdot), \dots, \mathbf{b}^N(\cdot)) \in B$ is a Bayesian equilibrium in Harsanyi’s model iff, $\forall i \in \{1, \dots, N\}$,

$$b^i(\theta_i) \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left(E_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \mid \theta_i \right] \right), \quad \forall \theta_i \in \Theta_i.$$

In contrast to Harsanyi, we assume that each player $i \in \{1, \dots, N\}$ lacks distributional information over V and $V_i(\theta_i)$ and therefore seeks to optimize his worst-case performance, i.e., his worst-case expected payoff, where the worst case is taken with respect to $V_i(\theta_i)$. Therefore, in a robust game involving private information, the set of best responses by player $i \in \{1, \dots, N\}$, when he is of type $\theta_i \in \Theta_i$, to $\mathbf{b}^{-i}(\cdot)$ is given by the set

$$\arg \max_{\mathbf{u}^i \in S_{a_i}} \left(\inf_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \right] \right).$$

Accordingly, the tuple of behavioral strategies $(\mathbf{b}^1(\cdot), \dots, \mathbf{b}^N(\cdot)) \in B$ is an equilibrium of the robust game with private information, i.e., is a robust-optimization equilibrium of the corresponding game with incomplete information, iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{b}^i(\theta_i) \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left(\inf_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \right] \right), \quad \forall \theta_i \in \Theta_i.$$

Before turning to the issue of equilibria existence, let us revisit the relation of the *ex post* equilibria of an incomplete-information game to the corresponding robust-optimization equilibria, this time in the context involving potentially private information. In any such game, the tuple $(\mathbf{b}^1(\cdot), \dots, \mathbf{b}^N(\cdot)) \in B$ is an *ex post* equilibrium of the incomplete-information game, iff, $\forall i \in \{1, \dots, N\}$,

$$\mathbf{b}^i(\theta_i) \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \left(\left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{u}^i \right) \right] \right), \quad \forall \theta_i \in \Theta_i; \forall (\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i).$$

By a proof analogous to that of Lemma 1, we may extend the result of that lemma to the general case involving potentially private information.

Lemma 6. *The set of ex post equilibria of an incomplete-information game is contained in the corresponding set of robust-optimization equilibria.*

7.2. Existence of equilibria

We will now extend our existence result from Section 4, in which we considered robust finite games without private information, to general robust finite games. Let us start by considering such games in which all of the players' type spaces are finite, i.e., $\forall i \in \{1, \dots, N\}$, $\Theta_i = \{1, \dots, t_i\}$, where $t_i < \infty$. Recall that player i 's pure strategies are mappings from Θ_i to $\{1, \dots, a_i\}$. Then, the set of player i 's pure strategies is simply $\{1, \dots, a_i\}^{t_i}$. Similarly, player i 's behavioral strategies can be encoded as $a_i \times t_i$ matrices, where column $\ell \in \Theta_i$ gives player i 's randomization over his action space when he is of type $\theta_i = \ell$. More precisely,

$$B_{a_i} = \{ \mathbf{X} \in \mathbb{R}^{a_i \times t_i} \mid \mathbf{X}_\ell \in S_{a_i}, \ell \in \Theta_i \},$$

where \mathbf{X}_ℓ denotes the ℓ^{th} column of the matrix \mathbf{X} . Let us define the additional shorthands

$$\mathbf{X}_{\theta_{-i}}^{-i} \triangleq \left(\mathbf{X}_{\theta_1}^1, \dots, \mathbf{X}_{\theta_{i-1}}^{i-1}, \mathbf{X}_{\theta_{i+1}}^{i+1}, \dots, \mathbf{X}_{\theta_N}^N \right)$$

$$\tau_i \left(\theta_i; \mathbf{X}^{-i}, \mathbf{x}^i \right) \triangleq \inf_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{X}_{\theta_{-i}}^{-i}, \mathbf{x}^i \right) \right],$$

where $\mathbf{X}_{\theta_i}^i$ denotes the θ_i^{th} column of the matrix \mathbf{X}^i . That is, τ_i denotes player i 's worst-case expected payoff function.

Theorem 6. *Consider a robust game that is non-cooperative, simultaneous-move, and played in one shot. Suppose there are $N < \infty$ players, that player $i \in \{1, \dots, N\}$ has $1 < a_i < \infty$ possible actions, and that the prior uncertainty set of payoff matrices $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ is bounded. Suppose that, $\forall i \in \{1, \dots, N\}$, player i 's type space is given by $\Theta_i = \{1, \dots, t_i\}$, where $t_i < \infty$. Then the robust game has an equilibrium in B .*

Proof. Let us define the point-to-set mapping $\Psi : B \rightarrow 2^B$, where 2^B is the power set of B , as

$$\Psi \left(\mathbf{X}^1, \dots, \mathbf{X}^N \right) = \left\{ \left(\mathbf{Y}^1, \dots, \mathbf{Y}^N \right) \in B \mid \mathbf{Y}_{\theta_i}^i \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \tau_i \left(\theta_i; \mathbf{X}^{-i}, \mathbf{u}^i \right), \right.$$

$$\left. \forall i \in \{1, \dots, N\}, \forall \theta_i \in \Theta_i \right\}.$$

It is obvious that $(\mathbf{X}^1, \dots, \mathbf{X}^N)$ is a behavioral strategy equilibrium of the robust game iff it is a fixed point of Ψ . That Ψ satisfies the conditions of Kakutani's Fixed Point Theorem [28] follows from the facts that, $\forall i \in \{1, \dots, N\}$ and $\forall \theta_i \in \Theta_i$, $\tau_i \left(\theta_i; \mathbf{X}^{-i}, \mathbf{x}^i \right)$ is continuous on $B_{-i} \times S_{a_i}$ and is concave in \mathbf{x}^i over S_{a_i} for fixed $\mathbf{X}^{-i} \in B_{-i}$. The details of the proof are analogous to those in our proof of Theorem 2, and we therefore omit them. □

Having treated the case in which all of the players' type spaces are finite, let us now consider the more general case in which there may exist an $i \in \{1, \dots, N\}$ such that $|\Theta_i| = \infty$. If player i has infinitely many types, his behavioral strategies

$$B_{a_i} = \left\{ \mathbf{b}^i : \Theta_i \rightarrow S_{a_i} \right\}$$

cannot be encoded as finite matrices but are functions with infinite domains, and therefore belong to an infinite dimensional space. Kakutani's Fixed Point Theorem applies to correspondences defined over Euclidean spaces, which are, by definition, finite dimensional. Accordingly, we cannot use Kakutani's theorem to prove the existence of behavioral strategy equilibria in robust finite games in which at least one player's type space is infinite. Instead, we need a fixed point theorem that applies to Banach spaces. The following fixed point result of Bohnenblust and Karlin [8] generalizes Kakutani's theorem to Banach spaces. Before stating it, we first recall a relevant definition.

Definition 8 (as stated in Smart [46]). *Let \mathcal{S} and \mathcal{T} be subsets of a normed space. Ψ is called a **K-mapping** of \mathcal{S} into \mathcal{T} if the following two conditions hold.*

1. $\forall s \in \mathcal{S}, \Psi(s) \subseteq \mathcal{T}, \Psi(s) \neq \emptyset$, and $\Psi(s)$ is compact and convex.
2. The graph $\{(s, t) \mid t \in \Psi(s)\}$ is closed in $\mathcal{S} \times \mathcal{T}$.

Theorem 7 (Bohnenblust and Karlin [8], as restated in Smart [46]). *Let \mathcal{M} be a closed, convex subset of a Banach space, and let Ψ be a K-mapping of \mathcal{M} into a compact subset \mathcal{M}' of \mathcal{M} . Then $\exists x \in \mathcal{M}$ such that $x \in \Psi(x)$.*

In order to apply this theorem to prove the existence of behavioral strategy equilibria in robust finite games with private information and potentially infinite type spaces,⁵ we must first establish some preliminary results. In the next two lemmas, $\forall i \in \{1, \dots, N\}$, we consider the metric space $(B_{-i} \times S_{a_i}) [d]$, with metric d defined as follows. $\forall (\mathbf{b}^{-i}(\cdot), \mathbf{x}^i), (\mathbf{f}^{-i}(\cdot), \mathbf{y}^i) \in B_{-i} \times S_{a_i}$,

$$d \left(\left(\mathbf{b}^{-i}(\cdot), \mathbf{x}^i \right), \left(\mathbf{f}^{-i}(\cdot), \mathbf{y}^i \right) \right) \triangleq \max \left\{ \|\mathbf{y}^i - \mathbf{x}^i\|_\infty, \max_{\substack{i' \in \{1, \dots, N\} \setminus \{i\} \\ j_{i'} \in \{1, \dots, a_{i'}\}}} \left[\sup_{\theta_{i'} \in \Theta_{i'}} \left| f_{j_{i'}}^{i'}(\theta_{i'}) - b_{j_{i'}}^{i'}(\theta_{i'}) \right| \right] \right\}.$$

Lemma 7. *Let $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ be bounded. Then $\forall \epsilon > 0, \exists \delta(\epsilon) > 0$ such that, $\forall i \in \{1, \dots, N\}, \forall \theta_i \in \Theta_i$, and $\forall (\mathbf{b}^{-i}(\cdot), \mathbf{x}^i), (\mathbf{f}^{-i}(\cdot), \mathbf{y}^i) \in B_{-i} \times S_{a_i}$,*

$$d \left(\left(\mathbf{b}^{-i}(\cdot), \mathbf{x}^i \right), \left(\mathbf{f}^{-i}(\cdot), \mathbf{y}^i \right) \right) < \delta(\epsilon)$$

implies that, $\forall (\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)$,

$$\left| \pi_i \left(\tilde{\mathbf{P}}; \mathbf{f}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{y}^i \right) - \pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{x}^i \right) \right| < \epsilon.$$

⁵ Recall that a game is said to be finite if the number of players and the number of actions available to each player are all finite. Accordingly, it is possible for a finite game with incomplete information to involve infinite type spaces.

Proof. $\forall \varepsilon > 0$, consider

$$\delta(\varepsilon) = \frac{\min \{ \varepsilon, 1 \}}{2 (2^N - 1) M \prod_{i=1}^N a_i},$$

where $1 < M < \infty$ satisfies

$$\left| \tilde{P}_{(j_1, \dots, j_N)}^i \right| \leq M, \quad \forall i \in \{1, \dots, N\}, \forall (j_1, \dots, j_N) \in \prod_{i=1}^N \{1, \dots, a_i\}, \forall \tilde{\mathbf{P}} \in U.$$

The result follows from algebraic manipulation. □

Lemma 7 immediately gives rise to the following continuity result.

Lemma 8. *Let $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ be bounded. Then $\forall i \in \{1, \dots, N\}, \forall \theta_i \in \Theta_i$,*

$$\tau_i \left(\theta_i; \mathbf{b}^{-i}(\cdot), \mathbf{x}^i \right) \triangleq \inf_{(\tilde{\mathbf{P}}, \boldsymbol{\theta}) \in V_i(\theta_i)} \left[\pi_i \left(\tilde{\mathbf{P}}; \mathbf{b}^{-i}(\boldsymbol{\theta}_{-i}), \mathbf{x}^i \right) \right]$$

is continuous on $B_{-i} \times S_{a_i}$.

In addition, it is trivial to prove the following lemma.

Lemma 9. *$\forall i \in \{1, \dots, N\}, \forall \theta_i \in \Theta_i$, and $\forall \mathbf{b}^{-i}(\cdot) \in B_{-i}$ fixed, $\tau_i \left(\theta_i; \mathbf{b}^{-i}(\cdot), \mathbf{x}^i \right)$ is concave in \mathbf{x}^i over S_{a_i} .*

We may now apply Bohnenblust’s and Karlin’s fixed point theorem to prove the existence of behavioral strategy equilibria in robust finite games with potentially infinite type spaces.

Theorem 8 (Existence of Equilibria in Robust Finite Games). *Consider an N -person, non-cooperative, simultaneous-move, one-shot robust game, in which $N < \infty$, in which player $i \in \{1, \dots, N\}$ has $1 < a_i < \infty$ possible actions, in which player i ’s type space is given by Θ_i , and in which the prior uncertainty set of payoff matrices $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ is bounded. This robust game has an equilibrium in B .*

Proof. We will proceed by constructing a point-to-set mapping that satisfies the conditions of Bohnenblust’s and Karlin’s fixed point theorem, and whose fixed points are precisely the behavioral strategy equilibria of the robust game with private information. Recall that, for a non-empty set Θ_i , the vector space of all bounded functions defined on Θ_i is a Banach space under the supremum norm (Theorem 3-2.4 of [33]). Furthermore, the direct product of finitely many Banach spaces is a Banach space (Theorem 2-4.6 of [33]). Accordingly,

$$F = \prod_{i=1}^N \left\{ \mathbf{f}^i : \Theta_i \rightarrow \mathbb{R}^{a_i} \mid \mathbf{f}^i \text{ is bounded} \right\}$$

is a Banach space. In the notation we used to state Bohnenblust’s and Karlin’s fixed point theorem, take $\mathcal{M} = \mathcal{M}' = B$. B is a convex, closed, and compact subset of F .

Let us define the point-to-set mapping $\Psi : B \rightarrow 2^B$ as

$$\Psi \left(\mathbf{b}^1(\cdot), \dots, \mathbf{b}^N(\cdot) \right) = \left\{ \left(\mathbf{y}^1(\cdot), \dots, \mathbf{y}^N(\cdot) \right) \in B \mid \mathbf{y}^i(\theta_i) \in \arg \max_{\mathbf{u}^i \in S_{a_i}} \tau_i \left(\theta_i; \mathbf{b}^{-i}(\cdot), \mathbf{u}^i \right), \forall i \in \{1, \dots, N\}, \forall \theta_i \in \Theta_i \right\}.$$

The rest of the proof follows similarly to that of Theorem 2. \square

7.3. Computation of equilibria

Having extended our equilibria existence result to incomplete-information games involving private information, let us now establish that one may compute these robust-optimization equilibria, when the players' type spaces are finite, via a formulation analogous to the one we gave in Section 5 for the case without private information.

Theorem 9. *Consider a robust game that is non-cooperative, simultaneous-move, and played in one shot. Suppose there are $N < \infty$ players, that player $i \in \{1, \dots, N\}$ has $1 < a_i < \infty$ possible actions, and that the prior uncertainty set of payoff matrices $U \subseteq \mathbb{R}^N \prod_{i=1}^N a_i$ is bounded. Suppose $\forall i \in \{1, \dots, N\}$, player i 's type space is given by $\Theta_i = \{1, \dots, t_i\}$, where $t_i < \infty$. Let*

$$V(\boldsymbol{\theta}) = \{(\mathbf{P}, \boldsymbol{\theta}) \in V\}$$

$$T_i(\theta_i) = \{(\boldsymbol{\theta}_{-i}, \theta_i) \in \Theta \mid V(\boldsymbol{\theta}_{-i}, \theta_i) \neq \emptyset\},$$

and $\text{Proj}(A, A')$ denote the projection of a set A onto a set A' . In addition, suppose that, $\forall i \in \{1, \dots, N\}$, $\forall \theta_i \in \Theta_i$, $\forall \boldsymbol{\theta} \in T_i(\theta_i)$, there exists a polyhedron $U(\boldsymbol{\theta}) = \text{Proj}(V(\boldsymbol{\theta}), U)$. Then, the set of behavioral strategy equilibria of the robust game is the component-wise projection of the solution set of a system of multilinear equalities and inequalities.

Proof. Since $\forall i \in \{1, \dots, N\}$, $t_i < \infty$,

$$B_{a_i} = \{\mathbf{X} \in \mathbb{R}^{a_i \times t_i} \mid \mathbf{X}_\ell \in S_{a_i}, \ell \in \Theta_i\}.$$

$(\mathbf{X}^1, \dots, \mathbf{X}^N) \in B$ is an equilibrium of this robust game iff, $\forall i \in \{1, \dots, N\}$, $\forall \theta_i \in \Theta_i$, $\exists z_{\theta_i}^i \in \mathbb{R}$ such that $(\mathbf{X}_{\theta_i}^i, z_{\theta_i}^i)$ is a maximizer of the following robust LP, in which \mathbf{X}^{-i} is regarded as data.

$$\begin{aligned} & \max_{\mathbf{X}_{\theta_i}^i, z_{\theta_i}^i} z_{\theta_i}^i \\ & \text{s.t. } z_{\theta_i}^i \leq \pi_i \left(\tilde{\mathbf{P}}; \mathbf{X}_{\boldsymbol{\theta}_{-i}}^{-i}, \mathbf{X}_{\theta_i}^i \right), \quad \forall \boldsymbol{\theta} \in T_i(\theta_i); \forall \tilde{\mathbf{P}} \in U(\boldsymbol{\theta}) \\ & \quad \mathbf{e}' \mathbf{X}_{\theta_i}^i = 1 \\ & \quad \mathbf{X}_{\theta_i}^i \geq \mathbf{0}, \end{aligned}$$

where $\mathbf{e} \in \mathbb{R}^{a_i}$ is the vector of all ones. The proof follows analogously to that of Theorem 3, since $|T_i(\theta_i)| < \infty$ and $|\Theta_i| < \infty$. \square

8. Conclusions

We make several contributions in this paper. We propose a novel, distribution-free model, based on robust optimization, of games with incomplete information, and we offer a corresponding distribution-free, robust-optimization equilibrium concept. We address incomplete-information games without private information as well as those involving potentially private information. Our robust optimization model of such games relaxes the assumptions of Harsanyi's Bayesian games model and simultaneously gives a notion of equilibrium that subsumes the *ex post* equilibrium concept. In addition, we prove the existence of equilibria in any such robust finite game, when the payoff uncertainty set is bounded. This existence result is in contrast to the fact that incomplete-information games need not have any *ex post* equilibria. For any robust finite game with bounded polyhedral payoff uncertainty set and finite type spaces, we formulate the set of equilibria as the dimension-reducing, component-wise projection of the solution set of a system of multilinear equations and inequalities. We suggest a computational method for approximately solving such systems and give numerical results of the implementation of this method. Furthermore, we describe a special class of robust finite games, whose equilibria are precisely those of a related complete-information game with the same number of players and the same action spaces. Using illustrative examples of robust games from this special class, we compare properties of robust finite games with those of their Bayesian-game counterparts. Moreover, we prove that symmetric equilibria exist in symmetric, robust finite games with bounded uncertainty sets.

We hope that these contributions will provide a new perspective on games with incomplete information.

Acknowledgements. The authors would like to thank Rakesh Vohra, Sergei Izmalkov, and Drew Fudenberg for suggesting readings on games with incomplete information, and Mike Todd for suggesting readings on fixed point theorems. We would also like to thank the reviewers of the paper for their insightful comments.

References

1. Ben-Tal, A., Nemirovski, A.: Robust convex optimization. *Math. Oper. Res.* **23** (4), 769–805 (1998)
2. Ben-Tal, A., Nemirovski, A.: Robust solutions of uncertain linear programs. *Oper. Res. Lett.* **25** (1), 1–13 (1999)
3. Ben-Tal, A., Nemirovski, A.: Robust solutions of linear programming problems contaminated with uncertain data. *Math. Prog., Ser. A* **88**, 411–424 (2000)
4. Bertsekas, D.: *Nonlinear Programming*, 2nd edn. Athena Scientific, Belmont, MA, 1999
5. Bertsimas, D., Pachamanova, D., Sim, M.: Robust linear optimization under general norms. *Oper. Res. Lett.* **32**, 510–516 (2004)
6. Bertsimas, D., Sim, M.: Robust discrete optimization and network flows. *Math. Prog.* **98**, 49–71 (2003)
7. Bertsimas, D., Sim, M.: The price of robustness. *Oper. Res.* **52** (1), 35–53 (2004)
8. Bohnenblust, H., Karlin, S.: On a theorem of Ville. In: H. Kuhn, A. Tucker (eds.), *Contributions to the theory of games*, vol. **1**, Princeton UP, Princeton, 1950, pp. 155–160
9. Borodin, A., El-Yaniv, R.: *Online computation and competitive analysis*. Cambridge UP, New York, 1998
10. Brouwer, L.: Über abbildung von mannigfaltigkeiten. *Mathematische Annalen* **71**, 97–115 (1912)
11. Courant, R.: Variational methods for the solution of problems of equilibrium and vibrations. *Bull. Am. Math. Soc.* **49**, 1–23 (1943)
12. Crémer, J., McLean, R.: Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. *Econometrica* **53** (2), 345–362 (1985)
13. Datta, R.: Using computer algebra to find Nash equilibria. In: J. Senda (ed.), *Proc. 2003 Intl. Symp. on Symb. and Alg. Comp.*, ACM Press, New York, 2003, pp. 74–79

14. Debreu, G.: A social equilibrium existence theorem. *Proc. Nat. Acad. Sci., USA* **38**, 886–893 (1952)
15. Dow, J., Werlang, S.: Nash equilibrium under Knightian uncertainty: Breaking down backward induction. *J. Econ. Theory* **64**, 305–324 (1994)
16. El Ghaoui, L., Lebret, H.: Robust solutions to least-squares problems with uncertain data. *SIAM J. Matrix Analysis and Applications* **18** (4), 1035–1064 (1997)
17. El Ghaoui, L., Oustry, F., Lebret, H.: Robust solutions to uncertain semidefinite programs. *SIAM J. Optim.* **9** (1), 33–52 (1999)
18. Fiacco, A., McCormick, G.: *Nonlinear programming: Sequential unconstrained minimization techniques*. John Wiley & Sons, New York, 1968
19. Fudenberg, D., Levine, D.: *The Theory of Learning in Games*. Series on Economic Learning and Social Evolution. MIT Press, Cambridge, MA, 1998
20. Fudenberg, D., Tirole, J.: *Game Theory*. MIT Press, Cambridge, MA, 1991
21. Gilboa, I., Schmeidler, D.: Maxmin expected utility with a non-unique prior. *J. Math. Econ.* **18**, 141–153 (1989)
22. Goldberg, A., Wright, A., Karlin, A., Hartline, J., Saks, M.: Competitive auctions, 2002. Submitted to *Games and Economic Behavior*
23. Govindan, S., Wilson, R.: A global Newton method to compute Nash equilibria. *J. Econ. Theory* **110**, 65–86 (2003)
24. Harsanyi, J.: Games with incomplete information played by ‘Bayesian’ players, parts I–III. *Mgmt. Sci.* **14**, 159–182, 320–334, 486–502 (1967, 1968)
25. Holmström, B., Myerson, R.: Efficient and durable decision rules with incomplete information. *Econometrica* **51**, 1799–1820 (1983)
26. Holzman, R., Kfir-Dahav, N., Monderer, D., Tennenholtz, M.: Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior* **47** (1), 104–123 (2004)
27. Hyafil, N., Boutilier, C.: Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In: *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence (AUAI '04)*, AUAI Press, Arlington, VA, USA, 2004, pp. 268–277
28. Kakutani, S.: A generalization of Brouwer’s fixed point theorem. *Duke Math. J.* **8**, 457–459 (1941)
29. Klibanoff, P.: Uncertainty, decision, and normal form games, 1993. Manuscript, MIT
30. Knight, F.: *Risk, Uncertainty and Profit*. Houghton Mifflin, Boston, 1921
31. Lemke, C., Howson, J.: Equilibrium points of bimatrix games. *SIAM J. Appl. Math.* **12** (2), 413–423 (1964)
32. Lo, K.: Equilibrium in beliefs under uncertainty. *J. Econ. Theory* **71** (2), 443–484 (1996)
33. Ma, T.: *Banach-Hilbert spaces, vector measures and group representations*. World Scientific, New Jersey, 2002
34. Marinacci, M.: Ambiguous games. *Games and Economic Behavior* **31** (2), 191–219 (2000)
35. McKelvey, R., McLennan, A.: Computation of equilibria in finite games. In: H. Amman, D. Kendrick, J. Rust (eds.), *Handbook of computational economics*, vol. **I**. Elsevier, 1996, pp. 87–142
36. McKelvey, R., McLennan, A., Turocy, T.: *Gambit: Software tools for game theory*, version 0.97.0.6, 2004. URL <http://econweb.tamu.edu/gambit/>
37. Mertens, J., Zamir, S.: Formulation of Bayesian analysis for games with incomplete information. *Intl. J. Game Theory* **14** (1), 1–29 (1985)
38. Morris, S.: The common prior assumption in economic theory. *Econ. Philosophy* **11**, 227–253 (1995)
39. Nash, J.: Equilibrium points in N -person games. *Proc. Nat. Acad. Sci., USA* **36**, 48–49 (1950)
40. Nash, J.: Non-cooperative games. *Ann. Math.* **54** (2), 286–295 (1951)
41. Osborne, M., Rubenstein, A.: *A Course in Game Theory*. MIT Press, Cambridge, MA, 1994
42. Papadimitriou, C.: Algorithms, games, and the internet. In: *STOC '01: Proc. 33rd Annual ACM Symposium on the Theory of Computing*, ACM Press, New York, 2001, pp. 749–753
43. Porter, R., Nudelman, E., Shoham, Y.: Simple search methods for finding a Nash equilibrium. In: *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI-2004)*, 2004, pp. 664–669
44. Scarf, H.: The approximation of fixed points of a continuous mapping. *SIAM J. Appl. Math.* **15** (5), 1328–1343 (1967)
45. Scarf, H.: *The Computation of Economic Equilibria*. Yale UP, New Haven, CT, 1973. In collaboration with T. Hansen
46. Smart, D.: *Fixed Point Theorems*. Cambridge UP, New York, 1974
47. Soyster, A.: Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.* **21**, 1154–1157 (1973)
48. Sturmfels, B.: *Solving systems of polynomial equations*. American Mathematical Society, Providence, RI, 2002
49. van der Laan, G., Talman, A.: On the computation of fixed points in the product space of unit simplices and an application to noncooperative N person games. *Math. Oper. Res.* **7**, 1–13 (1982)

50. van der Laan, G., Talman, A., van der Heyden, L.: Simplicial variable dimension algorithms for solving the nonlinear complementarity problem on a product of unit simplices using a general labeling. *Math. Oper. Res.* **12**, 377–397 (1987)
51. Verschelde, J.: Algorithm 795: PHCpack: a general-purpose solver for polynomial systems by homotopy continuation. *ACM Trans. Math. Softw.* **25** (2), 251–276 (1999)
52. von Neumann, J., Morgenstern, O.: *Theory of Games and Economic Behaviour*. Princeton UP, 1944
53. von Stengel, B.: Computing equilibria for two-person games. In: R. Aumann, S. Hart (eds.), *Handbook of game theory with economic applications*, vol. **3**, chap. 45, Elsevier, 2002, pp. 1723–1759
54. Wardrop, J.: Some theoretical aspects of road traffic research. In: *Proceedings of the Institute of Civil Engineers*, vol. **1**, part II, 1952, pp. 325–78
55. Wilson, R.: Game-theoretic approaches to trading processes. In: T. Bewley (ed.), *Advances in Economic Theory: Fifth World Congress*, chap. 2, Cambridge UP, New York, 1987, pp. 33–77