Proceedings of Machine Learning Research vol 117:1–27, 2020 31st International Conference on Algorithmic Learning Theory

Finding Robust Nash equilibria

Vianney Perchet VIANNEY.PERCHET@NORMALESUP.ORG
CREST, ENSAE
& Criteo AI Lab, Paris

Editors: Aryeh Kontorovich and Gergely Neu

Abstract

When agents or decision maker have uncertainties of underlying parameters, they may want to take "robust" decisions that are optimal against the worst possible values of those parameters, leading to some max-min optimization problems. With several agents in competition, in game theory, uncertainties are even more important and robust games - or game with non-unique prior - have gained a lot of interest recently, notably in auctions.

The existence of robust equilibria in those games is guaranteed using standard fixed point theorems as in classical finite games, so we focus on the problem of finding and characterizing them. Under some linear assumption on the structure of the uncertainties, we provide a polynomial reduction of the robust Nash problem to a standard Nash problem (on some auxiliary different game). This is possible by proving the existence of a lifting transforming robust linear programs into standard linear programs. In the general case, the above direct reduction is not always possible. However, we prove how to adapt the Lemke-Howson algorithm to find robust Nash equilibria in non-degenerate games.

Keywords: Robust game; Computing Equilibria

1. Introduction.

In many decision problems, the agent - or decision maker - has some uncertainties on the ambient environment. For instance, hidden parameters of a parametric model might be a-priori unknown. Those parameters can sometimes be learned if enough data are gathered and/or Bayesian inference can be made. However, those techniques requires that the distribution of data is "constant" with time and/or that the prior of the agent on the parameters is unique and "well-chosen". Those two assumptions are often too strong and unrealistic and this has motivated a different approach to deal with uncertainties, whether it is in optimization (Ben-Tal and Nemirovski, 1998), operations research (Gilboa and Schmeidler, 1989), decision science (Paolo Crespi et al., 2017; Li et al., 2019), game theory (Costikyan, 2013), Markov Decision Processes – MDP – (White and El Deib, 1992), learning theory (Xu et al., 2009), etc.

A typical way to handle uncertainties is to aim at finding "robust" solutions that try to maximize a worst-case scenario. This concept emerged in decision theory, in the celebrated Ellsberg paradox (Ellsberg, 1961) where agents express preferences on choices that are incompatible with the existence of subjective probabilities. It is also present in OR, where it is often impossible to evaluate accurately or to have any relevant a-priori information about the time required to develop a new technology, as aircrafts or missiles, promised to a government (Hitch, 1960). Similarly, the optimal planning of a company, as a supply chain (Ben-Tal et al., 2009b), depends heavily on the future demand and it might be impossible to evaluate it precisely, or to estimate it with some adequate Bayesian prior. As mentioned above, the theory of *robust optimization* (Soyster, 1973; Ben-Tal and Nemirovski, 1998; Bertsimas and Thiele, 2006; Ben-Tal et al., 2009a,b; Mevissen et al., 2013; Bogunovic et al., 2018) and *decisions under non-unique priors* (Gilboa and Schmeidler, 1989; Klibanoff, 1996) provide some answers to these problems and have gained a lot of interest recently and have been generalized to robust learning (Xu et al., 2009; Shafieezadeh Abadeh et al., 2015, 2018), robust MDP (White and El Deib, 1992; Xu and Mannor, 2007, 2010; Osogami, 2012; Tirinzoni et al., 2018) where it is typically the transitions of the MDP that are uncertain when trying to solve it.

Uncertainties are even more important when there are more than just one agent interacting with the environment. For instance, recently, and this is due to the vast amount of data available on internet, anyone can obtain recommendations – typically when looking for a specific diagnosis or treatments – made by several different experts or charlatans, without having the possibility to know their true types. Indeed, their answers, accuracy and correctness may vary from problems (Gilboa and Marinacci, 2011), and one cannot just assess the probability of being honest. Similarly, a standard result in auction theory (Myerson, 1981) indicates that the optimal auction design depends on the valuation distributions of the bidders. More importantly, billions of auctions are run every day for the online ads market. In those auctions, sellers have uncertainties on the type of bidders they are facing, or on their distribution of valuations, leading to the concept of robust auctions (Allouah and Besbes, 2018; Agrawal et al., 2018; Roughgarden et al., 2018). Those are a specific case of a similar general concept of robustness in *games* (Hayashi et al., 2005; Aghassi and Bertsimas, 2006).

Indeed, in general games, uncertainties can emerge from the doubts players have about the type of their opponents, their numbers, the action they use, etc., and they aim at using strategies performing relatively well in any case (and, typically, are evaluated on the worst one). For instance, in online poker, the wealth and origins of players, the technology they use are hidden; it is even possible that two supposably different opponents are controlled by the same person. These games have been called in the literature as *with ambiguity* (Bade, 2010; Bose and Renou, 2013), with *uncertainty* (Klibanoff, 1996; Hitch, 1960; Renou and Schlag, 2009), *partially specified* (Lehrer, 2012), etc. (see also Gilboa and Marinacci (2011)). Yet uncertainties in game can also arise from the partial or raw information a player has about her actual payoff mapping as it may depend on some unpredictable future events, on some hidden and unknown types (Renou and Schlag, 2009), or on some exogenous state variable that a player can not see nor compute. Those kind of games where actually called *robust games* (Hayashi et al., 2005; Aghassi and Bertsimas, 2006). In adversarial multi-armed bandits, uncertainties of payoffs are referred to as *partial monitoring* (Mertens et al., 1994; Rustichini, 1999; Lugosi et al., 2008; Perchet, 2011a,b; Mannor et al., 2014; Perchet and Quincampoix, 2014; Kwon and Perchet, 2017; Lattimore and Szepesvari, 2019).

Robust strategies of players maximize the worst possible payoff compatible with their information; this is also called their *maxmin utility*. Equilibria can still be defined in those games as fixed points of some complicated – yet regular enough– set-valued mappings; their existence is then ensured, as in the original paper of Nash (Nash, 1950), by some fixed point theorem (Brower, Kakutani or Fan's version), see, e.g., Perchet (2014) and references therein for more details on the existence properties. So the question of computing (and possibly characterizing) these equilibria is the most challenging one. In classical finite games, some algorithms – such as the celebrated Lemke-Howson algorithm (Lemke and Howson, 1964) – can be used to construct Nash equilibria, although it is a difficult problem (as the Nash problem is PPAD complete (Papadimitriou, 2007; Daskalakis et al., 2009; Chen and Deng, 2006)). Some topological properties of the set of equilibria, as the oddness of its cardinality or the existence of ± 1 index of equilibria, can be proved or recovered using this algorithm.

In this paper, we obtain similar kind results – computation and characterization– of equilibria in non-degenerate robust games. This class encompasses the aforementioned frameworks behind uncertainty aversion (Klibanoff, 1996), conjectural equilibria (Battigalli and Guaitoli, 1988), partial monitoring (Mertens et al., 1994) and robust games (Hayashi et al., 2005; Aghassi and Bertsimas, 2006), under an additional mild assumption for the latter. These finite games are defined by **two** sets of mappings (instead of only one in classical games). The first set describes the actual payoffs of the players while the second set of mappings represents the uncertainties of the players.

Using topological properties of linear mappings and projection of polytopes, we recover a fundamental property of finite games. Namely, there exists a finite subset of mixed actions that contains, for every profile of actions of the opponents, a best response. This result is obvious in finite classical games: just consider the set of pure actions. It is no longer immediate in the robust framework (as players no longer maximize linear utilities but concave ones).

Main results and organization of the paper.

Our objectives are twofolds: first, we aim at constructing an algorithm that finds robust Nash equilibria in this class of games and second to deduce some properties of their set. For the sake of clarity, we consider in the paper only 2-players game as it encompasses all the main ideas (we will briefly explain though how our reductions generalise to more than 2 players).

In a first step, we shall consider the specific subclass of games with *semi-standard structure of information*. We will construct an auxiliary bi-matrix game whose set of Nash equilibria corresponds to the set of robust equilibria of the original game. This is done in Section 3

The general case, where such an easy reduction is impossible, is considered in Section 5. We shall prove that these games still satisfy another crucial property: best response areas form a polytopial complex (Shapley, 1974; Perchet, 2011a). Using this, we give some characterization of the set of Nash equilibria and an algorithm to compute them based on an adapted version of the Lemke-Howson algorithm (Lemke and Howson, 1964). For the sake of completeness, its description and main properties are recalled briefly in Section 4. Finally, since robust Nash equilibria are end-points of a special instance of the Lemke-Howson algorithm, the oddness property of their set is preserved under, as usual, some non-degeneracy assumption (this is one is investigated in more details in Appendix C).

We only consider the two players case in that paper for the sake of clarity, yet we quickly mention in the conclusion how the result extend to more than two players.

Links with the existing literature.

As already mentioned, the literature on ambiguity, uncertainties and robustness in optimization, MDP and games is growing rapidly. The main concern is often the existence of robust optima/strategy/equilibria and their implication from a strategic point of view (at least in decision theory or in economics).

In the most related papers, see Hayashi et al. (2005); Aghassi and Bertsimas (2006), a description of robust Nash equilibria is given as the solution of some huge linear or conical programs with complementarity constraints. Unfortunately, this is not a simple computing algorithm (such as the Lemke-Howson one in the classical case) and sets of robust Nash equilibria cannot be characterized this way; for instance to prove that it is generically finite and of odd cardinality, etc.

Moreover, only games with an additional assumption that we will call "semi-standard" (see Section 3) fit the setup considered in Aghassi and Bertsimas (2006). The setting of Hayashi et al.

(2005) is a bit different as it uses conical programming instead of linear programming, but this is at the cost of very-specific structure of uncertainty (to respect the conical one): roughly speaking, they have to be some ”balls” around the true unknown parameter. Quite interestingly, they also consider the semi-standard structure as special cases (yet topological structure like finiteness of the set of equilibria and oddness of its cardinality can not be recovered with conical programming either).

So our model is in some sense more general. Yet, on the other hand, degenerate games (that happens with probability 0 if all parameters are taken uniformly at random) can be considered at no cost using the techniques of Hayashi et al. (2005); Aghassi and Bertsimas (2006), while one needs to slightly perturbe the payoffs in our frameworks, following usual techniques (von Stengel, 2007).

2. Two-player finite robust games.

We consider from now on the class of two-players finite robust games. They are defined by **two pairs** of mappings that respectively describe the payoffs and the uncertainties of players. Most of the concepts of robustness, uncertainty, ambiguity mentioned above in the introduction are captured by our generic model, see Appendix A.

Payoff mappings. The finite sets of actions of player 1 and 2 are respectively denoted by $\mathcal{A}$ and $\mathcal{B}$, of respective cardinality A and B. The choices of the actions $a \in \mathcal{A}$ and $b \in \mathcal{B}$ generate a payoff $u(a, b)$ for player 1 and $v(a, b)$ for player 2. As usual, these payoff mappings are extended multi-linearly to the set of *mixed actions* $\mathcal{X} := \Delta(\mathcal{A})$ and $\mathcal{Y} := \Delta(\mathcal{B})$, where $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$ stand for the set of probability distributions on the finite sets $\mathcal{A}$ or $\mathcal{B}$, by

$$u(x, y) = \sum_{a \in \mathcal{A}} \sum_{b \in \mathcal{B}} x_a y_b u(a, b),$$

where $x_a \in [0, 1]$ is the weight put by the mixed action $x \in \mathcal{X}$ on the pure action $a \in \mathcal{A}$. For the sake of notations, $\mathcal{X}$ and $\mathcal{Y}$ are seen as subsets of, respectively, $\mathbb{R}^A$ and $\mathbb{R}^B$ so that $\sum_{a \in \mathcal{A}} x_a U_a = \mathbb{E}_x[U]$ can be rewritten in a more compact way as $\langle x, U \rangle$, for every $x \in \mathcal{X}$ and $U \in \mathbb{R}^A$.

Uncertainty mappings. We recall that the idea behind robust games is that players have doubts, partial knowledge, uncertainties or multiple priors about the strategies of their opponents and/or about their own payoff mappings. We will consider the following quite general model of uncertainties (see Appendix A for more details).

We suppose that there exists an exogenous linear mapping $\mathbf{M} : \mathcal{Y} \to \mathbb{R}^m$ representing the uncertainties of player 1. The intuition is the following: an action $y \in \Delta(\mathcal{B})$ of player 2 generates a “message” (or a signal) $\mathbf{M}(y)$ to player 1. The latter has some uncertainty if two action y and y' generate the same message or are in the same equivalence class, i.e., if $\mathbf{M}(y) = \mathbf{M}(y')$. Stated otherwise, player 1 can not infer if player 2 used action y or $y' \in \Delta(\mathcal{B})$ or, in terms of payoff, if her payoff is either $u(x, y)$ or $u(x, y') \in \mathbb{R}$.

As a consequence, if player 1 plays $x \in \mathcal{X}$ and player 2 plays $y \in \mathcal{Y}$, then player 1 knows only that her payoff belongs to the set $\{u(x, y') \in \mathbb{R};\ \mathbf{M}(y') = \mathbf{M}(y)\} \subset \mathbb{R}$. Equivalently, the payoff of player 1 is of the form $\langle x, U \rangle$ where $U \in \mathbb{R}^A$ belongs to the set $\Phi(y)$ defined by

$$\Phi(y) := \Big\{ \big[u(a, y')\big]_{a \in \mathcal{A}} \in \mathbb{R}^A;\ \mathbf{M}(y') = \mathbf{M}(y) \Big\} \subset \mathbb{R}^A.$$

The set $\Phi(y)$ is called the *uncertainty set* of player 1 on her payoff.

Similarly, we denote by $\mathbf{N} : \mathcal{X} \to \mathbb{R}^{k'}$ the linear mapping that plays the same role as $\mathbf{M}$ for player 2. The uncertainty set of player 2, denoted by $\Psi(x) \subset \mathbb{R}^B$, is therefore defined by

$$\Psi(x) = \left\{ \left[v(x', b)\right]_{b \in \mathcal{B}} \in \mathbb{R}^B;\ \mathbf{N}\left(x'\right) = \mathbf{N}(x) \right\} \subset \mathbb{R}^B.$$

Maximization of payoffs under uncertainties. Following the precepts of optimization under non-unique prior (Gilboa and Schmeidler, 1989) and robust optimization (Aghassi and Bertsimas, 2006; Soyster, 1973; Ben-Tal et al., 2009a), a mixed action $x^* \in \mathcal{X}$ of player 1 is a robust best response to some uncertainty set $\mathcal{U} \subset \mathbb{R}^A$, a fact we denote by $x^* \in \mathrm{BR}_1(\mathcal{U})$, if it solves the following problem

$$x^* \in \mathrm{BR}_1(\mathcal{U}) \qquad \Longleftrightarrow \quad x^* \in \arg\max_{x \in \mathcal{X}} \inf_{U \in \mathcal{U}} \langle x, U \rangle .$$

Stated differently, x^* maximizes the worst possible compatible payoffs. No matter the set $\mathcal{U}$, this problem is well-defined since $x \mapsto \inf_{U \in \mathcal{U}} \langle x, U \rangle$ is upper semi-continuous (and concave) hence maxima are indeed attained on compact sets. Best responses of player 2 are defined similarly.

Definition of robust Nash equilibria. The concept of robust Nash equilibria (Klibanoff, 1996; Aghassi and Bertsimas, 2006) naturally follows from the aforementioned concepts.

Definition 1 *A pair $(x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ is a robust Nash equilibrium if x^* and y^* are mutually robust best responses, i.e., if $x^* \in \mathrm{BR}_1(\Phi(y^*))$ and $y^* \in \mathrm{BR}_2(\Psi(x^*))$ or, stated otherwise, if*

$$x^* \in \arg\max_{x \in \mathcal{X}} \inf_{U \in \Phi(y^*)} \langle x, U \rangle \ \text{ and } y^* \in \arg\max_{y \in \mathcal{Y}} \inf_{V \in \Psi(x^*)} \langle y, V \rangle .$$

We recall that we have defered some remarks and interpretations of robust games to the Appendix A, for the sake of the flow of presentation.

3. A warm-up: semi-standard structure

Before the general case, we first consider the specific class of games with a "*semi-standard* structure of information" where the set of robust Nash equilibria has a simpler structure.

3.1. Description of the semi-standard structure

Basically, the main insight behind a semi-standard structure is that some actions are equivalent from the point of view of the information they provide and that it is always possible to distinguish between two non-equivalent actions (but equivalent actions are undistinguishable).

Definition 2 *The structure of information of player 1 is semi-standard if there exists a partition of $\mathcal{B}$ into "cells" $\{\mathcal{B}_i;\ i \in \mathcal{I}\}$ such that*

i) If b and b' belong to the same cell $\mathcal{B}_i$ then $\mathbf{M}(b) = \mathbf{M}(b') := \mathbf{M}_i$ and

ii) The family $\{\mathbf{M}_i;\ i \in \mathcal{I}\}$ is linearly independent, i.e. if $\sum_{i \in \mathcal{I}} \lambda_i \mathbf{M}_i = \sum_{i \in \mathcal{I}} \gamma_i \mathbf{M}_i$ then $\lambda_i = \gamma_i$, for every $i \in \mathcal{I}$.

A game is semi-standard if both players have a semi-standard structure of information.

The key point of the assumption is obviously ii). The robust game introduced in the proof of Proposition 16 (in Appendix A) has a semi-standard structure; other examples are provided in Appendix B.

The key property of semi-standard structure is the linearity of the uncertainty sets.

Proposition 3 *If the structure of information of player 1 is semi-standard, then Φ is linear, in the sense that*

$$\Phi(y) = \sum_{b \in \mathcal{B}} y_b \Phi(b) = \left\{ \sum_{b \in \mathcal{B}} y_b U_b \, ; \, U_b \in \Phi(b) \right\}, \quad \forall y \in \mathcal{Y}.$$

In particular, $\Phi(y)$ is a polytope for every $y \in \mathcal{Y}$.

Under this specific assumption, robust games can be reduced to bi-matrix games, as shown in the following Lemma 4. But first, we recall the general concept of *polytopial complex* on which our results rely. A finite set $\{P_k;\ k \in \mathcal{K}\}$ is a polytopial complex of a polytope $P \subset \mathbb{R}^d$ with non-empty interior if for every $k \in \mathcal{K}$, $P_k \subset P$ is a polytope with non empty interior; the union $\bigcup_{k \in \mathcal{K}} P_k$ is equal to P; and every intersection of two different polytopes $P_k \cap P_{k'}$ has an empty interior.

Lemma 4 *There exists $L \in \mathbb{N}$, a finite subset $\{x[1], \ldots, x[L]\}$ of $\mathcal{X}$ that contains, for every $y \in \mathcal{Y}$, a maximizer of the program $\max_{x \in \mathcal{X}} \min_{U \in \Phi(y)} \langle x, U \rangle$ and such that its convex hull contains the whole set of maximizers. Moreover, there exists a polytopial complex $\{\mathcal{Y}_1, \ldots, \mathcal{Y}_L\}$ of $\mathcal{Y}$ such that, for every $\ell \in \{1, \ldots, L\}$, $x[\ell]$ is a maximizer on $\mathcal{Y}_\ell$.*

Lemma 4 might be surprising to a reader familiar with linear and non-linear programming. Indeed, it is quite clear that if $u(\cdot, y)$ is linear then it is always maximized at one of the vertices of $\mathcal{X}$. However, in our case, $\min_{U \in \Phi(y)} \langle \cdot, U \rangle$ is not linear but merely concave. So, at first sight, it could be maximized anywhere in $\mathcal{X}$, and the sets of maximizers could be different for every mixed action $y \in \mathcal{Y}$. Thus, without any regularity of Φ, the result would obviously be wrong. The key point of the proof is that, in our framework, Φ is itself induced by another linear mapping and Lemma 4 holds because $\min_{U \in \Phi(y)} \langle \cdot, U \rangle$ is not just *any* concave mapping.

Another intuition behind this result comes from the fact that, in some cases, robust linear programming can be solved by linear programming (see, e.g., Bertsimas and Sim (2004); Bertsimas et al. (2004)). Similar statement holds for player 2:

Corollary 5 *There exists a finite subset $\{y[1], \ldots, y[K]\}$ of $\mathcal{Y}$ that contains, for every $x \in \mathcal{X}$, a maximizer of the program $\max_{y \in \mathcal{Y}} \min_{V \in \Psi(x)} \langle y, V \rangle$ and such that its convex hull contains the whole set of maximizers.*

Without additional assumptions on the structure of the game, the size of the subset of possible best responses of Lemma 4 can be exponential in A, the number of actions of player 1. The additional assumption will concern the "dimension" of uncertainties, namely the dimension $d_{\mathbf{M}}$ of the kernel of $\mathbf{M}$; $d_{\mathbf{M}}$ is actually equal to B minus the number of class of undistinguishable actions.

Corollary 6 *Let $d_{\mathbf{M}}$ be the dimension of the kernel of $\mathbf{M}$ be fixed. Then the cardinality L defined in Lemma 4 is polynomial in A.*

We have decided to consider as constant (as A and B might increase to infinity) the dimension of the kernel of $\mathbf{M}$ and $\mathbf{N}$ in order to have some succinct representations of the set of possible best responses. An alternative – and actually equivalent – assumption would be that the number of actions in some uncertainty class (or, for that matter, the number of uncertainty classes) is fixed.

3.2. Computation of robust Nash equilibria in semi-standard games

This section contains the first main result that describes the set of robust Nash equilibria as the set of Nash equilibria of some associated bi-matrix game. This game is defined as follows.

Action sets of the associated bi-matrix game. The *pure action sets* of player 1 and 2 are precisely the finite subsets $\{x[1], \ldots, x[L]\}$ and $\{y[1], \ldots, y[K]\}$ given by Lemma 4 and Corollary 5. To avoid confusion between pure and mixed actions, we introduce the set of indices of pure actions $\mathcal{L} = \{1, \ldots, L\}$ and $\mathcal{K} = \{1, \ldots, K\}$ so that *mixed actions* in this bi-matrix game are elements of $\Delta(\mathcal{L})$ and $\Delta(\mathcal{K})$.

Payoffs of the associated bi-matrix game.. The payoffs associated to the choices of $\ell \in \mathcal{L}$ and $k \in \mathcal{K}$ are given by

$$\widetilde{u}(\ell, k) = \min_{U \in \Phi(y[k])} \langle x[\ell], U \rangle \quad \text{and} \quad \widetilde{v}(\ell, k) = \min_{V \in \Psi(x[\ell])} \langle y[k], V \rangle.$$

As usual, $\widetilde{u}$ (and similarly $\widetilde{v}$) is extended multi-lineary to $\Delta(\mathcal{L})$ and $\Delta(\mathcal{K})$ by

$$\forall (\mathbf{x}, \mathbf{y}) \in \Delta(\mathcal{L}) \times \Delta(\mathcal{K}), \quad \widetilde{u}(\mathbf{x}, \mathbf{y}) = \sum_{\ell \in \mathcal{L}} \sum_{k \in \mathcal{K}} \mathbf{x}_\ell \mathbf{y}_k \min_{U_{\ell,k} \in \Phi(y[k])} \langle x[\ell], U_{\ell,k} \rangle.$$

Because of the construction, a mixed action $\mathbf{x} \in \Delta(\mathcal{L})$ can be seen as a probability distribution over mixed actions in $\mathcal{X}$. We can therefore associate to any mixed action in $\Delta(\mathcal{L})$ a unique element of $\mathcal{X}$, its expectation and we say that $x \in \mathcal{X}$ is induced by $\mathbf{x} \in \Delta(\mathcal{L})$ if

$$\forall a \in \mathcal{A}, \quad x_a = \sum_{\ell=1}^{L} \mathbf{x}_\ell x[\ell]_a, \qquad \text{i.e., if } \; x = \mathbb{E}_{\mathbf{x}}\big(x[\ell]\big).$$

Similarly, we that say $y \in \mathcal{Y}$ is induced by $\mathbf{y} \in \Delta(\mathcal{K})$ if $y = \mathbb{E}_{\mathbf{y}}\big(y[k]\big)$.

With these notations at hand, we can state our first main result, the correspondence between the set of robust Nash equilibria and the set of Nash equilibria of the associated bi-matrix game.

Theorem 7 *Every Nash equilibrium of the associated bi-matrix game induces a robust Nash equilibrium. Reciprocally, every robust Nash equilibrium is induced by a Nash equilibrium of the associated bi-matrix game.*

The main consequence of Theorem 7 is that the robust Nash problem (finding one robust Nash equilibrium) can be polynomially reduced to the standard Nash problem (finding one Nash equilibrium) - if the dimension of the kernels of $\mathbf{M}$ and $\mathbf{N}$ are bounded[1].

Corollary 8 *The robust Nash equilibria problem, with a semi-standard structure and uniformly bounded kernels of uncertainties, can be polynomially reduced to the Nash problem.*

We reming here that the Nash problem can be solved (for *non-degenerate* game, see (von Stengel, 2002) for the specific assumption and how to circumvent degeneracy issues) with the Lemke-Howson algorithm (Lemke and Howson, 1964) (see also Wilson (1971)); this is actually a key

1. To be more precise, the reduction is polynomial in the inputs of the problem but not in the dimension of the kernels, i.e., the different exponents of the polynomes involve might depend on them.

element in the proof that the Nash problem is PPAD (Papadimitriou, 2007) – the proof of its completeness even for 2 players (Daskalakis et al., 2009; Chen and Deng, 2006) being more involved.

Theorem 7 also guarantees the *semi-algebraic* structure of the set of robust Nash equilibria – we recall that a semi-algebraic set is defined by a finite number of polynomial inequalities.

Corollary 9 *The set of robust Nash equilibria of a game with a semi-standard structure is the projection of the set of equilibria of some bi-matrix game; it is composed by the union of a finite number of connected semi-algebraic and closed components.*

This result is directly induced by the fact that this property holds for the set of Nash equilibria of the associated bi-matrix game.

As before, for the sake of fluency, we postpone remarks and additional properties of the semi-standard structure to Appendix B. We now briefly recall the key elements behind the Lemke-Howson algorithm that outputs Nash equilibria as it can obviously be used to output robust Nash equilibria with a semi-standard structure, but also, as we will see later, for general information structure.

4. Quick reminder on Lemke-Howson algorithm

The Lemke-Howson algorithm (Lemke and Howson, 1964) is based on the decompositions of $\mathcal{X}$ and $\mathcal{Y}$ into *best-response areas*. Recall that, for any $a \in \mathcal{A}$, the a-th best-response area of player 1 is $\mathcal{Y}_a := \Big\{y \in \mathcal{Y} \text{ s.t. } a \in \arg\max_{a' \in \mathcal{A}} u(a', y)\Big\} \subset \mathcal{Y}$. A required *non-degeneracy* assumption is:

Assumption 1 $\{\mathcal{Y}_a;\ a \in \mathcal{A}\}$ *forms a polytopial complex of* $\mathcal{Y}$ *and any* $y \in \mathcal{Y}$ *belongs to at most* s_y *best response areas* $\mathcal{Y}_a$*, where* s_y *is the cardinality of the support of* y*. The similar condition holds for* $\{\mathcal{X}_B;\ b \in \mathcal{B}\}$*.*

Stated otherwise, Assumption 1 means that every $y \in \mathcal{Y}$ has at most s_y best responses, i.e., a pure action has a unique best response, a mixture of two pure actions, at most two pure best responses, and so on.

By linearity, it is clear that $\mathcal{Y}_a$ is a polytope. We denote by V_2 and E_2 the set of all vertices and edges of all these sets $\mathcal{Y}_a$; in particular, one necessarily has that $\mathcal{B} \subset V_2$. For technical reason, we shall also assume that V_2 contains another abstract point 0_2 and that $(0_2, b)$ belongs to E_2 for every $b \in \mathcal{B}$. This construction defines a graph $\mathcal{G}_2 = (V_2, E_2)$ over $\mathcal{Y}$.

To each vertex $v \in V_2$, and more generally to any $y \in \mathcal{Y}$ is associated the following set of labels:

$$L(v) := \Big\{a \in \mathcal{A} \text{ s.t. } v \in \mathcal{Y}_a\Big\} \bigcup \Big\{b \in \mathcal{B} \text{ s.t. } v_b = 0\Big\} \subset \mathcal{A} \cup \mathcal{B},$$

i.e., labels are the set of best responses to v_2 and the set of pure actions on which v does not put any weight. The label set of the abstract point 0_2 is $L(0_2) = \mathcal{B}$.

A similar construction defines a labeled graph $\mathcal{G}_1 = (V_1, E_1)$ over $\mathcal{X}$, with labels also in $\mathcal{A} \cup \mathcal{B}$.

The product of those graphs defines a product labeled graph $\mathcal{G}_0 = (V_0, E_0)$ over $\mathcal{X} \times \mathcal{Y}$, whose set of vertices is the cartesian product $V_0 = V_1 \times V_2$ and there exists an edge in E_0 between (v_1, v_2) and (v_1', v_2') if and only if $v_1 = v_1'$ and $(v_2, v_2') \in E_2$ or $v_2 = v_2'$ and $(v_1, v_1') \in E_1$. The set of labels of (v_1, v_2) is defined as the union $L(v_1, v_2) = L(v_1) \cup L(v_2)$.

Nash equilibria are exactly *fully labeled* pairs (v_1, v_2), i.e., points such that $L(v_1, v_2) = \mathcal{A}_1 \cup \mathcal{A}_2$; indeed, this means that an action a is either not played (if $(v_1)_a = 0$) or a best reply to v_2 (if

$v_2 \in \mathcal{Y}_a$). The LH-algorithm walks along edges of $\mathcal{G}_0$, from vertices to vertices, and stops at a one of those points. We describe quickly in the remaining of this section how this works generically, for almost all games; for more details we refer to Shapley (1974) or von Stengel (2007) and references therein.

Starting at the fully labeled point $v_0 = (0_1, 0_2)$, one label ℓ in $\mathcal{A} \cup \mathcal{B}$ is chosen and dropped arbitrarily. The LH algorithm visits sequentially *almost fully labeled* vertices $(v_t)_{t\in\mathbb{N}}$ of $\mathcal{G}_0$, i.e., points such that $L(v_t) \supset \mathcal{A} \cup \mathcal{B}\backslash\{\ell\}$ and (v_t, v_{t+1}) is an edge in E_0. Under the non-degeneracy assumption, at any v_t there exists at most one point (apart from v_{t-1}) satisfying both properties, and any end point must be fully labeled.

As a consequence, the LH algorithm follows paths whose endpoints are necessarily Nash equilibria or $(0_1, 0_2)$. This property is an elementary proof of the existence of Nash equilibria, to compute some of them and, maybe more surprisingly, to recover the generic oddness of their cardinality.

5. Computation and characterization of robust Nash eq. The general case.

Without semi-standardness, or at least the linearity of Φ, Lemma 4 and Theorem 7 might not hold (see Example 2). As a consequence, the reduction to a bi-matrix game is therefore not possible in the general case, because of the lack of linearity of Φ. However, we will show that a weaker property holds, namely that Φ is *piecewise linear*, i.e., Φ is linear on a polytopial complex of $\mathcal{Y}$.

For instance, in Example 2, the mapping Φ is linear on both of the following subsets

$$\left\{y \in \mathcal{Y} \text{ s.t. } y_L \geq 2y_C\right\} \text{ and } \left\{y \in \mathcal{Y} \text{ s.t. } y_L \leq 2y_C\right\}.$$

Using this, it will be possible to show that best-responses areas still form a polytopial complex, enabling the generalization of LH-algorithm. Similar decompositions have been recently used in related frameworks (von Stengel and Zamir, 2010).

Lemma 10 *The correspondence Φ is piecewise linear on $\mathcal{Y}$. Moreover, if $d_{\mathbf{M}}$ is constant, then the number of elements on which Φ is linear is polynomial in B.*

By considering the polytopial complex with respect to which Φ is piecewise linear and applying Lemma 4 on each element of the complex, one immediately obtains the following lemma that generalizes Lemma 4 to non-semi standard structures.

Lemma 11 *There exists a finite subset $\{x[\ell];\ \ell \in \mathcal{L}\}$ of $\mathcal{X}$ that contains, for every $y \in \mathcal{Y}$, a maximizer of the program $\max_{x\in\mathcal{X}} \min_{U\in\Phi(y)}\langle x, U\rangle$ and such that its convex hull contains the set of maximizers. Moreover, for every $\ell \in \mathcal{L}$, $x[\ell]$ is a maximizer on $\mathcal{Y}_\ell$ which is a finite union of polytopes.*

Similarly, there exists a finite subset $\{y[k];\ k \in \mathcal{K}\}$ of $\mathcal{Y}$ that contains, for every $x \in \mathcal{X}$, a maximizer of the program $\max_{y\in\mathcal{Y}} \min_{V\in\Psi(y)}\langle y, V\rangle$ and such that its convex hull contains the set of maximizers. Moreover, for every $k \in \mathcal{K}$, $y[k]$ is a maximizer on $\mathcal{X}_k$ which is a finite union of polytopes.

Both L and K, the cardinality of $\mathcal{L}$ and $\mathcal{K}$ are polynomial in A and B.

The second main result is the following characterization of robust Nash equilibria. We are going to characterize them with respect to points in $\Delta(\mathcal{L})$ and $\Delta(\mathcal{K})$ that induce them. We recall that any $\mathbf{x} \in \Delta(\mathcal{L})$ induces a mixed action $x = \mathbb{E}_{\mathbf{x}}(x[\ell]) \in \mathcal{X}$ and $\mathbf{y} \in \Delta(\mathcal{K})$ induces similarly $y \in \mathcal{Y}$.

Similarly to the labeling of $\mathcal{X}$ and $\mathcal{Y}$ used by the LH-algorithm, the labels of a point $\mathbf{y} \in \Delta(\mathcal{K})$ belongs to $\mathcal{L} \cup \mathcal{K}$: $k \in \mathcal{K}$ is a label of $\mathbf{y}$ iff $\mathbf{y}_k = 0$ and $\ell \in \mathcal{L}$ is also a label of $\mathbf{y}$ iff

$$\begin{aligned}\mathbf{y} \in \mathbf{Y}_\ell &:= \Big\{\mathbf{y}' \in \Delta(\mathcal{K}) \text{ s.t. } \mathbb{E}_{\mathbf{y}'}(y[k]) \in Y_\ell\Big\} \\ &= \left\{\mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } x_\ell \in \arg\max_{\ell' \in \mathcal{L}} \inf_{U \in \Phi(\mathbb{E}_{\mathbf{y}}(y[k]))} \langle x_{\ell'}, U\rangle\right\},\end{aligned}$$

i.e., either $\mathbf{y}$ puts no weight on $k \in \mathcal{K}$ or $x[\ell]$ is a best reply to the mixed action induced by $\mathbf{y}$.

Theorem 12 *Robust Nash equilibria are induced by points in $\Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ that are fully labeled and, reciprocally, any fully labeled point induces a robust Nash equilibrium*

It remains to describe why the LH algorithm can be used in this framework. We defined, in the proof of Lemma 11, Y_ℓ as the union of a finite number of polytopes $Y_{\ell,i}$.

Lemma 13 *Any element of the families $\{\mathbf{Y}_l;\ l \in \mathcal{L}\}$ and $\{\mathbf{X}_k;\ k \in \mathcal{K}\}$ is a polytope.*

We can now generalize the LH-algorithm to robust games satisfying some non-degeneracy assumptions.

Theorem 14 *If $\{\mathbf{Y}_\ell;\ \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k;\ k \in \mathcal{K}\}$ satisfy the non-degeneracy Assumption 1, then any fully labelled point induces a robust Nash equilibrium and reciprocally. So robust Nash equilibria are induced by a finite, odd number of points.*

If Assumption 1 is not satisfied, then the set of robust Nash equilibria is still a finite union of connected closed semi-algebraic sets.

Proof: If $\{\mathbf{Y}_\ell;\ \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k;\ k \in \mathcal{K}\}$ satisfy the non-degeneracy Assumption 1, any end point of the LH-algorithm is fully labeled, hence a robust Nash equilibrium. □

It is not compulsory to use the induced polytopial complexes of $\Delta(\mathcal{L})$ and $\Delta(\mathcal{K})$. One can work directly in $\mathcal{X}$ and $\mathcal{Y}$ by considering the projection of the skeleton of the complexes $\{\mathbf{Y}_\ell;\ \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k;\ k \in \mathcal{K}\}$ onto them. However, the graphs generated might not be planar and there are, at first glance, no guarantee that the LH-algorithm will work. The proof of Theorem 12, based on a lifting of the problem, ensures that graphs are planar.

The fact that there was an odd number of Nash equilibria in the game of Example 2, where the generalized version of the Lemke-Howson algorithm is illustrated, is therefore not surprising; in the usual bi-matrix case, this is one of the consequences of the analysis of the LH-algorithm. So, as soon as $\{\mathbf{Y}_\ell;\ \ell \in \mathcal{L}\}$ and $\{\mathbf{X}_k;\ k \in \mathcal{K}\}$ satisfy the same assumption, there will exist an odd number of fully labeled points in $\Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ inducing robust Nash equilibria.

6. Conclusion

We have considered robust games with 2 players, and we have shown that if the structure is semi-standard (and the number of undistinguishable auctions is constant), then the robust Nash problem can be polynomially reduced to the Nash problem (on some auxiliary game). If the structure is not semi-standard, then the same reduction can not be used; however, a generalized version of the Lemke-Howson algorithm can be used on some lifted action space.

Our reduction can also be translated to the case of more than 2 players. The only constraint is that the uncertainty set of a player is the product set of uncertainties upon each of her opponents (or all information structures are semi-standard) so that $\Phi(\cdot)$ is always a polytope and our machinery can be developed again in the N-persons variant of the Lemke-Howson algorithm (Rosenmuller, 1971).

Acknowledgments

This work was supported by a public grant as part of the Investissement d'avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH, in a joint call with Gaspard Monge Program for optimization, operations research and their interactions with data sciences. V. Perchet also acknowledges the support of the ANR under the grant ANR-19-CE23-0026.

References

M. Aghassi and D. Bertsimas. Robust game theory. *Math. Program., Ser. B*, 107, 2006.

S. Agrawal, C. Daskalakis, V. Mirrokni, and B. Sivan. Robust repeated auctions under heterogeneous buyer behavior. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 171–171, New York, NY, USA, 2018. ACM.

A. Allouah and O. Besbes. Prior-independent optimal auctions. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, pages 503–503, New York, NY, USA, 2018. ACM.

S. Bade. Ambiguous act equilibria. *Games. Econom. Behav.*, 71:246–260, 2010.

P. Battigalli and D. Guaitoli. Conjectural equilibrium. *mimeo*, 1988.

A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Math. Oper. Res.*, 23:769–805, 1998.

A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009a.

A. Ben-Tal, B. Golany, and S. Shtern. Robust multi echelon multi period inventory control. *European J. on operations Research*, 199:922–935, 2009b.

D. Bertsimas and M. Sim. The price of robustness. *Oper. Res.*, 52:35–53, 2004.

D. Bertsimas and A. Thiele. Robust and data-driven optimization: modern decision-making under uncertainty. *Tutorials on Operations Research.*, Chapter 4:122–195, 2006.

D. Bertsimas, D. Pachamanova, and M. Sim. Robust linear optimization under general norms. *Oper. Res. Lett.*, 32:510–516, 2004.

L. J. Billera and B. Sturmfels. Fiber polytopes. *Ann. of Math.*, 135:527–549, 1992.

I. Bogunovic, J. Scarlett, S. Jegelka, and V. Cevher. Adversarially robust optimization with gaussian processes. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 5760–5770. Curran Associates, Inc., 2018.

S. Bose and L. Renou. Mechanism design and ambiguous communication devices. *manuscript*, 2013.

X. Chen and X. Deng. Settling the complexity of two-player nash equilibrium. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pages 261–272, 2006.

G. Costikyan. *Uncertainty in Games*. The MIT Press, 2013.

C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

D. Ellsberg. Risk, ambiguity and the savage axioms. *Quarterly Journal of Economics*, 75:643–669, 1961.

L. Gilboa and M. Marinacci. Ambiguity and the bayesian paradigm. *manuscript*, 2011.

L.J. Gilboa and D. Schmeidler. Maxmin expected utility with a non-unique prior. *J. Math. Econom.*, 61:141–153, 1989.

S. Hayashi, N. Yamashita, and M. Fukushimay. Robust nash equilibria and second-order cone complementarity problems. *Journal of Nonlinear and Convex Analysis*, 6, 01 2005.

C. Hitch. Uncertainties in operations research. *Oper. Res.*, 8:437–445, 1960.

P. Klibanoff. Uncertainty, decision and normal form games. *mimeo*, 1996.

Joon Kwon and Vianney Perchet. Online Learning and Blackwell Approachability with Partial Monitoring: Optimal Convergence Rates. In Aarti Singh and Jerry Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 604–613, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR.

T. Lattimore and C. Szepesvari. Cleaning up the neighborhood: A full classification for adversarial partial monitoring. In A. Garivier and S. Kale, editors, *Algorithmic Learning Theory 2019*, pages 529–556. Proceedings of Machine Learning Research, 2019.

E. Lehrer. Partially specified probabilities: decisions and games. *Am. Econom. J.: Microeconom*, 4:70–100, 2012.

C. E. Lemke and J. T. Howson, Jr. Equilibrium points of bimatrix games. *J. Soc. Indust. Appl. Math.*, 12:413–423, 1964.

B. Li, J. Sun, H. Xu, and M. Zhang. A class of two-stage distributionally robust games. *Journal of Industrial & Management Optimization*, 15:387, 2019.

G. Lugosi, S. Mannor, and G. Stoltz. Strategies for prediction under imperfect monitoring. *Math. Oper. Res.*, 33:513–528, 2008.

S. Mannor, V. Perchet, and G. Stoltz. Approachability in unknown games: Online learning meets multi-objective optimization. *Proceedings of the Conference On Learning Theory (COLT'14)*, pages 339–355, 2014.

J.-F. Mertens, S. Sorin, and S. Zamir. *Repeated Games*. CORE discussion paper 9420–9422. 1994.

M. Mevissen, E. Ragnoli, and J. Y. Yu. Data-driven distributionally robust polynomial optimization. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 37–45. Curran Associates, Inc., 2013.

Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

J.F. Nash. Equilibrium points in n-person games. *Proc. Nat. Acad. Sci. USA.*, 36, 1950.

T. Osogami. Robustness and risk-sensitivity in markov decision processes. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 233–241. Curran Associates, Inc., 2012.

G. Paolo Crespi, D. Radi, and M. Rocca. Robust games: theory and application to a cournot duopoly model. *Decisions in Economics and Finance*, 40:1–22, 2017.

C. Papadimitriou. The complexity of finding nash equilibria. In N. Nisan, T. Roughgarden, E. Tardos, and V Vazirani, editors, *Algorithmic Game Theory*, pages 29–52. Cambridge Uiversity Press, Cambridge, 2007.

V. Perchet. Internal regret with partial monitoring, calibration based optimal algorithms. *J. Mach. Learn. Res.*, 12:1893–1921, 2011a.

V. Perchet. Approachability of convex sets in games with partial monitoring. *Journal of Optimization Theory and Applications*, 149:665–677, 2011b.

V. Perchet. A note on robust nash equilibria with uncertainties. *RAIRO - Operations Research - Recherche Opérationnelle*, 48(3):365–371, 2014. doi: 10.1051/ro/2014001.

V. Perchet and M. Quincampoix. On a unified framework for approachability with full or partial monitoring. *Mathematics of Operations Research*, 40:596–610, 2014.

J. Rambau and G. M. Ziegler. Projections of polytopes and the generalized Baues conjecture. *Discrete Comput. Geom.*, 16:215–237, 1996.

L. Renou and K. Schlag. Minimax regret and strategic uncertainty. *Journal of Economic Theory*, 145:264–286, 2009.

R. T. Rockafellar. *Convex Analysis*. Princeton Mathematical Series, No. 28. Princeton University Press, Princeton, N.J., 1970.

J. Rosenmuller. On a generalization of the lemke–howson algorithm to noncooperative $n$person games. *Siam Journal on Applied Mathematics - SIAMAM*, 21, 07 1971.

T. Roughgarden, I. Talgam-Cohen, and Q. Yan. Robust auctions for revenuevia enhanced competition. *manuscript*, 2018.

A. Rustichini. Minimizing regret: the general case. *Games. Econom. Behav.*, 29:224–243, 1999.

S. Shafieezadeh Abadeh, P. M. Mohajerin Esfahani, and D. Kuhn. Distributionally robust logistic regression. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 1576–1584. Curran Associates, Inc., 2015.

S. Shafieezadeh Abadeh, V.A. Nguyen, D. Kuhn, and P. M. Mohajerin Esfahani. Wasserstein distributionally robust kalman filtering. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8474–8483. Curran Associates, Inc., 2018.

L. S. Shapley. A note on the Lemke-Howson algorithm. *Math. Program. Stud.*, pages 175–189, 1974.

A. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.*, 21:1154–1157, 1973.

A. Tirinzoni, M. Petrik, X. Chen, and B. Ziebart. Policy-conditioned uncertainty sets for robust markov decision processes. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 8939–8949. Curran Associates, Inc., 2018.

B. von Stengel. Computing equilibria for two-person games. In R.J. Aumann and S. Hart, editors, *Hanbook of Game Theory with Economic Applications*, volume 3, pages 1723–1759. Elsevier, Amsterdam, 2002.

B. von Stengel. Equilibrium computation for two-player games in strategic and extensive form. In N. Nisan, T. Roughgarden, E. Tardos, and V Vazirani, editors, *Algorithmic Game Theory*, pages 53–78. Cambridge Uiversity Press, Cambridge, 2007.

B. von Stengel and S. Zamir. Leadership games with convex strategy sets. *Games Econom. Behav.*, 69:446–457, 2010.

C.C. White and H. K. El Deib. Markov decision processes with imprecise transition probabilities. *Operations Research*, 42:739–748, 1992.

R. Wilson. Computing equilibria of n-person games. *SIAM J. Appl. Math.*, 21:80–87, 1971.

H. Xu and S. Mannor. The robustness-performance tradeoff in markov decision processes. In B. Schölkopf, J. C. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1537–1544. MIT Press, 2007.

H. Xu and S. Mannor. Distributionally robust markov decision processes. In J. D. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, and A. Culotta, editors, *Advances in Neural Information Processing Systems 23*, pages 2505–2513. Curran Associates, Inc., 2010.

H. Xu, C. Caramanis, and S. Mannor. Robust regression and lasso. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1801–1808. Curran Associates, Inc., 2009.

G. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

Appendix A. Appendix. Remarks and interpretations of robust games.

We formulate in this section several remarks on the definition, the generalization and the interpretation of robust games that answer the following questions:

– Are usual bi-matrix games special instances of robust games?

– Does this framework embed the case where the mappings $u(\cdot,\cdot)$ and $v(\cdot,\cdot)$ are unknown?

– Could uncertainties of a player depend on the action played?

– What are the possible interpretations of the timeline of a robust game?

A.1. Bi-matrix games.

The structure of bi-matrix games can be recovered from our setup.

Proposition 15 *Bi-matrix games are special instances of robust games*

Proof: One just has to define $\mathbf{M} : \mathcal{Y} \to \mathbb{R}^m$ as the identity, i.e., to assume that $m = B$ and

$$\mathbf{M}(y) = (y_b)_{b\in\mathcal{B}} \in \mathbb{R}^B \text{ so that } \Phi(y) = \{u(\cdot,y)\} \subset \mathbb{R}^A.$$

By definition of robust best responses,

$$x^* \in \mathrm{BR}_1(\Phi(y)) \iff x^* \in \mathrm{argmax}_{x\in\mathcal{X}} \inf_{U\in\{u(\cdot,y)\}} \langle x, U\rangle \iff x^* \in \mathrm{argmax}_{x\in\mathcal{X}}\, u(x,y).$$

If $\mathbf{N}$ is also defined as the identity, then the concepts of robust Nash equilibria and Nash equilibria coincide. And this is simply because robust best responses to singleton are the usual best responses. □

A.2. Unknown games.

We mentioned in the introduction that our framework encompasses the *unknown games* where the mappings u and v are partially unknown to, respectively, player 1 or 2.

More precisely, in a unknown game, player 1 only knows that $\big[u(a,b)\big]_{a\in\mathcal{A}}$ belongs, for every action $b \in \mathcal{B}$, to some polytope[2] $\mathcal{U}_b \subset \mathbb{R}^A$ generated by at most $s_\mathcal{B}$ vertices. As a consequence, when $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ are played, although y is observed by player 1, only information is that her payoff is of the form $\langle x, U\rangle$ where $U \in \mathbb{R}^A$ belongs to the set $\mathcal{U}(y)$ defined by

$$\mathcal{U}(y) := \Big\{ \sum_{b\in\mathcal{B}} y_b U_b\, ;\; U_b \in \mathcal{U}_b \Big\} \subset \mathbb{R}^A. \tag{1}$$

Similarly player 2 only knows that, for each action $a \in \mathcal{A}$, $\big[v(a,b)\big]_{b\in\mathcal{B}}$ belongs to some polytope $\mathcal{V}_a \subset \mathbb{R}^B$; $\mathcal{V}(x) \subset \mathbb{R}^B$ is defined as in Equation (1) and is generated by at most $s_\mathcal{A}$ vertices.

In this setup, a robust Nash equilibria is a pair $(x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ such that $x^* \in \mathrm{BR}_1(\mathcal{U}(y^*))$ and $y^* \in \mathrm{BR}_2(\mathcal{V}(x^*))$.

2. We recall that a *polytope* is the convex hull of a finite number of points.

Proposition 16 *Finding equilibria in unknown games can be polynomially (in $Bs_{\mathcal{B}}+As_{\mathcal{A}}$) reduced to finding equilibria in robust games.*

Proof: Consider any unknown game defined by the polytopes $\mathcal{U}_b$ for player 1 and $\mathcal{V}_a$ for player 2. As $\mathcal{U}_b$ is a polytope, it is the convex hull of a finite number of points. We denote them by $u_{b,j} \in \mathbb{R}^A$ where j is indexed over some finite set $\mathcal{J}_b$. Similarly, we represent $\mathcal{V}_a$ as the convex hull of some $v_{a,i} \in \mathbb{R}^B$ for i in the finite index set $\mathcal{I}_a$. Finally, we choose arbitrarily one specific index $i^*_a \in \mathcal{I}_a$ and $j^*_b \in \mathcal{J}_b$ in every such set, and we introduce $\mathcal{I}^* = \{i^*_a, a \in \mathcal{A}\}$ and $\mathcal{J}^* = \{j^*_b, b \in \mathcal{B}\}$. Since all sets are compact, we can assume that all payoffs are bounded by some $M > 0$.

We introduce the auxiliary robust game where

i) Action sets of player 1 and 2 are respectively $\mathcal{I} = \bigcup_{a\in\mathcal{A}} \mathcal{I}_a$ and $\mathcal{J} = \bigcup_{b\in\mathcal{B}} \mathcal{J}_b$.

ii) For every $j \in \mathcal{J}_b$, $\mathbf{M}(j) = \mathbf{e}_b$, where $\mathbf{e}_b$ is the unit vector of $\mathbb{R}^B$ whose only non-zero coordinate is the b-th. Similarly, for every $i \in \mathcal{I}_a$, $\mathbf{N}(i) = \mathbf{e}_a$.

iii) Payoff mappings $\widehat{u}$ and $\widehat{v}$ are defined by

$$\widehat{u}(i,j) = \begin{cases} \langle \mathbf{e}_a, u_{b,j}\rangle & \text{if} \quad i = i^*_a \in \mathcal{I}^* \\ \\ -2M & \text{otherwise} \end{cases}, \quad \text{for all } j \in \mathcal{J}_b \text{ and}$$

$$\widehat{v}(i,j) = \begin{cases} \langle \mathbf{e}_b, v_{a,i}\rangle & \text{if} \quad j = j^*_b \in \mathcal{J}^* \\ -2M & \text{otherwise} \end{cases}, \quad \text{for all } i \in \mathcal{I}_a .$$

We now show how this auxiliary game is equivalent to the original one. To any $\widehat{y} = (\widehat{y}_{b,j}) \in \Delta(\mathcal{J})$ and any $\widehat{x} = (\widehat{x}_{a,i}) \in \Delta(\mathcal{I})$, we associate $y \in \Delta(\mathcal{B})$ and $x \in \Delta(\mathcal{A})$ defined by

$$y_b := \sum_{j\in\mathcal{J}_b} \widehat{y}_{b,j}, \quad \forall b \in \mathcal{B} \quad \text{and} \quad x_a := \sum_{i\in\mathcal{I}_a} \widehat{x}_{a,i}, \quad \forall a \in \mathcal{A}$$

The information available to player 1 when $\widehat{y}$ is played in the robust game is the same as when y is played in the unknown game since, by construction,

$$\Big\{\langle \widehat{x}, \widehat{U}\rangle\,;\, \widehat{U} \in \Phi(\widehat{y})\Big\} = \Big\{\langle x, U\rangle\,;\, U \in \mathcal{U}(y)\Big\}.$$

Moreover, because of the choices of $\widehat{u}$ and $\widehat{v}$, the pure actions in $\mathcal{I}$ that are not in $\mathcal{I}^*$ are strictly dominated by any action in $\mathcal{I}^*$. More generally, any mixed action whose support is not included in $\mathcal{I}^*$ and $\mathcal{J}^*$ is strictly dominated.

As a consequence, when looking for equilibria, we can only consider mixed actions supported on $\mathcal{I}^*$ or $\mathcal{J}^*$. Since $|\mathcal{I}^*| = |\mathcal{A}| = A$, we identify $\mathcal{X} = \Delta(\mathcal{A}) = \Delta(\mathcal{I}^*)$ and $\mathcal{Y} = \Delta(\mathcal{B}) = \Delta(\mathcal{J}^*)$. The uncertainty set of player 1 should be a subset of $\mathbb{R}^{|\mathcal{I}|}$ but since we only consider mixed action with support on $\mathcal{I}^*$, we can consider its restriction to the coordinates in $\mathcal{I}^*$ by defining

$$\Phi^*(y) = \Big\{\big[\widehat{u}(i^*_a, y')\big]_{a\in\mathcal{A}} \in \mathbb{R}^A;\ \mathbf{M}\left(y'\right) = \mathbf{M}\left(y\right)\Big\} \subset \mathbb{R}^A.$$

As a consequence, $\Phi^*(y) = \mathcal{U}(y)$ for all $y \in \mathcal{Y}$ and a pair (x, y), seen as an element of $\Delta(\mathcal{I}^*) \times \Delta(\mathcal{J}^*)$, is a robust Nash equilibrium of the auxiliary robust game if and only if it is a robust Nash equilibrium of the original unknown game, when seen as an element of $\Delta(\mathcal{A}) \times \Delta(\mathcal{B})$. □

A.3. Action-dependent uncertainties.

In the underlying framework of games with partial monitoring (Mertens et al., 1994) and conjectural equilibria (Battigalli and Guaitoli, 1988), a player can receive signals depending on both $a \in \mathcal{A}$ and $b \in \mathcal{B}$. Mathematically, there exists a mapping $M : \mathcal{A} \times \mathcal{B} \to \mathbb{R}^{m'}$ and another mapping $N : \mathcal{A} \times \mathcal{B} \to \mathbb{R}^{n'}$ such that, if actions chosen are $a \in \mathcal{A}$ and $b \in \mathcal{B}$, player 1 gets the signal $M(a, b)$ and player 2 the signal $N(a, b)$.

Within this framework, one shall define $\mathbf{M} : \mathcal{Y} \to (\mathbb{R}^m)^A$ by

$$\mathbf{M}(y) = \Big[M(a, y)\Big]_{a\in\mathcal{A}} = \Bigg[\sum_{b\in\mathcal{B}} y_b M(a, b)\Bigg]_{a\in\mathcal{A}} \in \Big(\mathbb{R}^{m'}\Big)^A.$$

The motivating reasons behind this definition are given in Perchet (2014). In a few words, player 1 cannot distinguish between two mixed actions $y, y' \in \mathcal{Y}$ such that $\mathbf{M}(y) = \mathbf{M}(y')$, no matter her choice of actions. Moreover, the vector $\mathbf{M}(y)$ can be estimated at almost no cost by player 1 (Lugosi et al., 2008; Perchet, 2011a). Finally, this definition guarantees the existence of equilibria, which is not necessarily the case of alternative propositions found in the literature; we refer to Perchet (2014) for more details.

A.4. Different game theoretic interpretations of robust Nash equilibria.

There are different classical alternative ways to apprehend robust games and more specifically robust Nash equilibria.

1. When the true payoff mapping is unknown, it can be seen as the simultaneous and independent choices of mixed actions that maximize the worst possible payoffs. An equilibrium is a point such that none of the players has incentives to deviate.

   In other terms, it is a usual two-player game where the payoff mapping is no longer $u(x, y)$ which is replaced by $\mathbf{u}(x, y) := \min_{u\in\mathbf{U}} u(x, y)$, where $\mathbf{U}$ is some set of equicontinuous mappings.

   This mapping $\mathbf{u}$ is clearly non-linear but concave and continuous. Hence existence of Nash equilibria is guaranteed. But as claimed in the introduction, it is the computation and characterization of those equilibria that should matter in the present.

2. When payoff mappings are known, a robust game describes an interaction where player 1 makes decision on the basis of the value $\Phi(y)$, representing her information, rather than on the underlying decision y. The latter cannot be seen explicitly but only through the information window $\Phi(\cdot)$.

   For instance, consider the interpretation of Nash equilibria where players follow instructions of a referee, telling secretly to player 1 that her recommendation is to play x^* and that she will recommend player 2 to play y^*; the same construction can be made for correlated equilibrium. In a robust game, the referee would say to player 1 that her recommendation is x^* and that she will recommend to player 2 an action compatible with $\mathbf{M}(y)$, i.e., some $y' \in \mathcal{Y}$ such that $\mathbf{M}(y') = \mathbf{M}(y)$ without specifying the true value.

   Instead of having partial information on the actions of player 2, uncertainty represent the fact that player 1 has several priors on the behaviour of her opponent. She then optimizes her payoff with respect to the worst possible prior.

3. Robust Nash equilibria can represent stable situations in population games when the fitness cannot be evaluated precisely because of lack of information. They can also be seen as static and stable strategies in repeated games with partial monitoring.

Appendix B. Remarks on the semi-standard assumption

B.1. Example and counter-examples of games with semi-standard structure

Example 1 *A semi-standard structure of information for player 1 can be constructed as follows: for every $b \in \mathcal{B}$, $\mathbf{M}(b)$ belongs to a fixed basis of $\mathbb{R}^m$. It is trivial to show that $\mathbf{M}$ satisfies all axioms.*

Notice, on the other hand, that this is no longer true if the mapping $\mathbf{M}$ is defined by $\mathbf{M}(y) = \Big(M(a,y)\Big)_{a\in\mathcal{A}}$ where $M(a,b)$ is some basis vector of $\mathbb{R}^m$ (see Appendix A for more details). Consider for instance the case where $\mathcal{A} = \{a_1, a_2\}$, $\mathcal{B} = \{b_1, b_2, b_3, b_4\}$ and M is defined as

M:	b_1	b_2	b_3	b_4
a_1	$\mathbf{e}_1$	$\mathbf{e}_2$	$\mathbf{e}_1$	$\mathbf{e}_2$
a_2	$\mathbf{e}_1$	$\mathbf{e}_2$	$\mathbf{e}_2$	$\mathbf{e}_1$

with $\mathbf{e}_1 = (1,0)$ and $\mathbf{e}_2 = (0,1)$.

In that case, $\mathbf{M}(b_1) = (\mathbf{e}_1, \mathbf{e}_1) \in \mathbb{R}^4$, $\mathbf{M}(b_2) = (\mathbf{e}_2, \mathbf{e}_2) \in \mathbb{R}^4$, $\mathbf{M}(b_3) = (\mathbf{e}_1, \mathbf{e}_2)$ and $\mathbf{M}(b_4) = (\mathbf{e}_2, \mathbf{e}_1)$. The point ii) of Definition 2 is not satisfied since

$$\mathbf{M}\left(\frac{b_1+b_2}{2}\right) = \left(\frac{\mathbf{e}_1+\mathbf{e}_2}{2}, \frac{\mathbf{e}_1+\mathbf{e}_2}{2}\right) = \mathbf{M}\left(\frac{b_3+b_4}{2}\right).$$

There are many other examples of games without the semi-standard structure of information, as illustrated in the following example.

Example 2 *Assume that $\mathcal{A} = \{T; B\}$, $\mathcal{B} = \{L, C, R\}$, that payoffs of both players are represented by the matrix on the left, that $\mathbf{N}$ is the identity matrix, i.e., player 2 has no uncertainties, and that*

$(u,v) =$	L	C	R
T	$(1,1)$	$(0,0)$	$(0,0)$
B	$(0,0)$	$(1,2)$	$(1,0)$

$\mathbf{M}(L) = 0,\ \mathbf{M}(C) = 1$ and $\mathbf{M}(R) = \frac{1}{3}$.

In particular, the mixed action $2/3L + 1/3C$ and the pure action R give the same information to player 1. This shows that Φ is not linear. Indeed, $\Phi(L) = \{(1,0)\}$ and $\Phi(C) = \{(0,1)\}$ but

$$\Phi\left(\frac{2}{3}L + \frac{1}{3}C\right) = \left\{\left(\frac{2}{3}\lambda, 1-\frac{2}{3}\lambda\right) ;\ \lambda \in [0,1]\right\} \neq \frac{2}{3}\Phi(L) + \frac{1}{3}\Phi(C) = \left\{\left(\frac{2}{3}, \frac{1}{3}\right)\right\}.$$

Following notations of Lemma 4, and after some cumbersome computations, one can define the set of best responses $\{x[\ell];\ \ell \in \mathcal{L}\}$ as $\{T, B, M\}$ where $M = 1/2T + 1/2B$ and similarly $\{y[k];\ k \in \mathcal{K}\} = \{L, C, R\}$. The associated bi-matrix game is therefore defined by the following matrix:

$(\tilde{u},\tilde{v}) =$	L	C	R
T	$(1,1)$	$(0,0)$	$(0,0)$
B	$(0,0)$	$(1,2)$	$(1/3,0)$
M	$(1/2,1/2)$	$(1/2,1)$	$(1/2,0)$

This game has three Nash equilibria: (T, L), (B, C) *and* $(2/3T + 1/3B, 1/2L + 1/2C)$. *Although the first two are indeed robust Nash equilibria, this is not the case for the last one. Indeed,* $\Phi(1/2L + 1/2C) = \{(\lambda/2; 1 - \lambda);\ \lambda \in [0, 1]\}$ *and its set of best responses is the singleton* $\{T\}$.

There are actually three robust Nash equilibria which are (T, L), (B, C) *and* $(2/3T + 1/3B, 3/4L + 1/4C)$. *Indeed, the polytopial complexes* $\{Y_\ell;\ \ell \in \mathcal{L}\} := \{Y_B, Y_M, Y_T\}$ *and* $\{X_k;\ k \in \mathcal{K}\} := \{X_C, X_L\}$ *are represented in the following figure 1.*

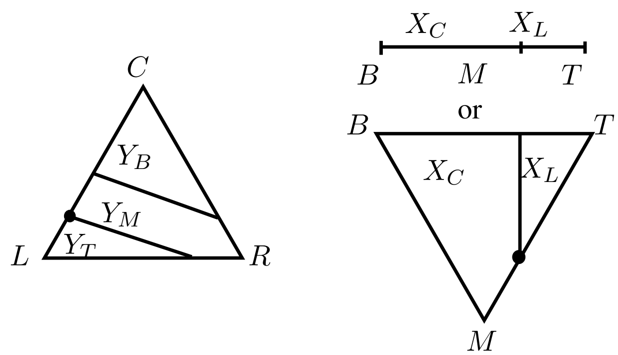

Figure 1: On the left $\mathcal{X}$ and on the right $\mathcal{Y}$ and $\Delta(\mathcal{K})$ with their complexes.

In order to describe how the LH-algorithm works, we will denote a vertex of the product graph by the cartesian product of its labels; in this example the set of labels is $\{T, B, M, L, R, C\}$. *For instance, the vertex represented with a black dot in Figure 1 is denoted by* $\{R, T, M\} \times \{B, C, L\}$.

The first step in the LH-algorithm is to drop one label arbitrarily; If the label M *is dropped then the first vertex visited by the algorithm is* $\{L, R, C\} \times \{B, T, C\}$. *The label* C *appears twice, so in order to get rid of one of them, the algorithm chooses at the next step the vertex* $\{L, R, B\} \times \{B, T, C\}$ *and the following vertex is* $\{L, R, B\} \times \{M, T, C\}$. *It's a fully labeled end point of the algorithm, hence* $(B, C) \in \mathcal{X} \times \mathcal{Y}$ *is a pure Nash equilibrium of* Γ.

Similarly, If T *is dropped at the first stage, then the first vertex is* $\{L, R, C\} \times \{B, M, L\}$ *and the second* $\{R, C, T\} \times \{B, M, L\}$. *So* (T, L) *is also a pure Nash equilibrium of* Γ.

Starting again from this point and dropping the label C *makes the LH-algorithm visit* $\{R, T, M\} \times \{B, M, L\}$, *and then* $\{R, T, M\} \times \{B, L, C\}$ *which is also a Nash equilibrium. It corresponds to* $(\mathbf{x}, \mathbf{y}) = (1/3T + 2/3M, 3/4L + 1/4C) \in \Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ *which induces* $(x, y) = (2/3T + 1/3B, 3/4L + 1/4C)$ *which is a (mixed) Nash equilibrium of* Γ.

One can check the remaining vertices of the product graph to be convinced that there does not exist any more equilibrium.

B.2. Variants of the definition of semi-standardness

We recall that $\lambda = (\lambda_1, \ldots, \lambda_k) \in \mathbb{R}^m$ are convex weights if $\lambda_i \geq 0$ for every $i \in \{1, \ldots, k\}$ and $\sum_{i=1}^k \lambda_i = 1$.

Proposition 17 *Point ii) of Definition 2 can be weakened without changing the conclusions of Theorem 7 into:*

ii') The family $\{\mathbf{M}_i;\ i \in \mathcal{I}\}$ *is convex independent, i.e. if* $\sum_{i\in\mathcal{I}} \lambda_i \mathbf{M}_i = \sum_{i\in\mathcal{I}} \gamma_i \mathbf{M}_i$ *for some convex weights* λ *and* γ *then* $\lambda = \gamma$.

Proof: The proof of Proposition 3 relies on the fact that $\mathbf{M}(y) = \mathbf{M}(y')$ if and only if $y[i] = y'[i]$ for all $i \in \mathcal{I}$. This holds true with respect to point ii') instead of point ii). □

We introduced this slight weakening of assumption in order to have a complete characterization of semi-standard structure:

Proposition 18 *The structure of information of player 1 is semi-standard (in the weak sense) if and only if the mapping $\mathbf{M}^{-1}[\mathbf{M}](\cdot) : \mathcal{Y} \rightrightarrows \mathcal{Y}$ defined by*

$$\mathbf{M}^{-1}[\mathbf{M}](y) = \Big\{ y' \in \mathcal{Y} \text{ s.t. } \mathbf{M}(y') = \mathbf{M}(y) \Big\} \subset \mathcal{Y}$$

is linear, i.e., if for every $y = (y_b)_{b \in \mathcal{B}} \in \mathcal{Y}$

$$\Big\{ y' \in \mathcal{Y} \text{ s.t. } \mathbf{M}(y') = \mathbf{M}(y) \Big\} = \sum_{b \in \mathcal{B}} y_b \Big\{ y' \in \mathcal{Y} \text{ s.t. } \mathbf{M}(y') = \mathbf{M}(b) \Big\}.$$

Proof: Assume that the structure of information is semi-standard, then $\mathbf{M}(y') = \mathbf{M}(y)$ if and only if $y'[i] = y[i]$ for all $i \in \mathcal{I}$. This yields the linearity of $\mathbf{M}^{-1}[\mathbf{M}](\cdot)$ since

$$\Big\{ y' \in \mathcal{Y} \text{ s.t. } \mathbf{M}(y') = \mathbf{M}(y) \Big\} = \Big\{ y' \in \mathcal{Y} \text{ s.t. } y'[i] = y[i],\ \forall i \in \mathcal{I} \Big\} = \Big\{ y' \in \sum_{i \in \mathcal{I}} y[i] \Delta(\mathcal{B}_i) \Big\}$$
$$= \sum_{i \in \mathcal{I}} \sum_{b \in \mathcal{B}_i} y_b \Delta(\mathcal{B}_i) = \sum_{b \in \mathcal{B}} y_b \Big\{ y' \in \mathcal{Y} \text{ s.t. } \mathbf{M}(y') = \mathbf{M}(b) \Big\} .$$

Reciprocally, assume that there exists $y', y \in \mathcal{Y}$ such that $y'[i^*] \neq y[i^*]$ for some $i^* \in \mathcal{I}$, then

$$\sum_{i \in \mathcal{I}} y'[i] \Delta(\mathcal{B}_i) \subset \mathbf{M}^{-1}[\mathbf{M}](y) \ \text{ but } \ \sum_{i \in \mathcal{I}} y'[i] \Delta(\mathcal{B}_i) \not\subset \sum_{i \in \mathcal{I}} \sum_{b \in \mathcal{B}_i} y_b \Delta(\mathcal{B}_i) = \sum_{b \in \mathcal{B}} y_b \mathbf{M}^{-1}[\mathbf{M}](b) ,$$

hence the result. □

Although the previous result characterizes completely semi-standard structure of information via the linearity of some mapping, the key argument in the proof of Lemma 4 is the linearity of $\Phi(\cdot)$ and not the linearity of $\mathbf{M}^{-1}[\mathbf{M}](\cdot)$. Although the former is implied by the latter, they are not equivalent as proved in the following non-trivial example.

Example 3 *Let $\mathcal{B} = \{1, 2, \ldots, 6\}$ and $\mathcal{A} = \{0\}$. Player 1 is a* dummy player *as she has a unique action, but the purpose of the example is only to deal with structures of information.*

Assume that the mapping $\mathbf{M}$ is defined by

$$\mathbf{M}(1) = \mathbf{M}(2) = -1, \quad \mathbf{M}(3) = \mathbf{M}(4) = 0 \ \text{ and } \ \mathbf{M}(5) = \mathbf{M}(6) = 1 .$$

Clearly, player 1 does not have a semi-standard structure of information, nonetheless, the mapping Φ is linear with respect to the following mapping u defined by

$$u(0,1) = -1,\ \ u(0,2) = 0,\ \ u(0,3) = -\frac{1}{2},\ \ u(0,4) = \frac{1}{2},\ \ u(0,5) = 0,\ \ u(0,6) = 1 .$$

Indeed, let $y = (y_1, \ldots, y_6) \in \mathcal{Y}$ then it immediately reads that

$$\sum_{b \in \mathcal{B}} y_b \Phi(b) = (y_1 + y_2).\Big[-1, 0 \Big] + (y_3 + y_4).\Big[-\frac{1}{2}, \frac{1}{2} \Big] + (y_5 + y_6).\Big[0, 1 \Big] = \Big[-\frac{1}{2} + \mu, \frac{1}{2} + \mu \Big],$$

where $\mu = [(y_5 + y_6) - (y_1 + y_2)]/2$*. On the other hand, it holds that*

$$\mathbf{M}^{-1}[\mathbf{M}](y) = \Big\{ y' \in \mathcal{Y} \text{ s.t. } (y'_5 + y'_6) = (y'_1 + y'_2) + \mu \text{ and } (y'_3 + y'_4) = 1 - \mu - 2(y'_1 + y'_2) \Big\}$$

so one obtains the following characterization of $\Phi(y)$*:*

$$\Phi(y) = \bigcup_{\lambda \in \left[0, \frac{1-\mu}{2}\right]} \lambda.\Big[-1, 0\Big] + (1 - \mu - 2\lambda).\Big[-\frac{1}{2}, \frac{1}{2}\Big] + (\lambda + \mu).\Big[0, 1\Big]$$
$$= \bigcup_{\lambda \in \left[0, \frac{1-\mu}{2}\right]} \Big[-\frac{1}{2} + \mu, \frac{1}{2} + \mu\Big] = \Big[-\frac{1}{2} + \mu, \frac{1}{2} + \mu\Big].$$

Therefore the multi-valued mapping Φ *is linear.*

The above example shows that although $\mathbf{M}^{-1}[\mathbf{M}]$ is not linear, Φ might be. Since only the linearity of the latter is used in the proofs, Theorem 7 actually holds for any structure of information implying the linearity of Φ.

In particular, this is a reason why we consider uncertainty sets defined in terms of payoffs, i.e., by $\Phi(y)$, and not in terms of actions, i.e., by $\mathbf{M}^{-1}[\mathbf{M}](y)$. The former contains actually less information, as it is as a projection of the later, but it is more relevant to the model. From a game theoretic point of view, this is related to the fact that a player is not really concerned about the actions played by her opponents but about the effects they have on her own payoff.

Another trivial example where the fact that $\mathbf{M}^{-1}[\mathbf{M}]$ is not linear is irrelevant is when the payoff mapping u is constant, say, if $u(a, b) = 0$ for all $a \in \mathcal{A}$ and $b \in \mathcal{B}$.

B.3. Semi-standard structure and unknown payoffs.

Proposition 19 *When players have only partial knowledge of their payoff mappings, the game has a semi-standard structure in the sense that the mapping* $\mathcal{U}$ *defined by Equation (1) is linear. The robust game introduced in the proof of Proposition 16 also has a semi-standard structure.*

So the conclusions of Theorem 7 hold in this framework.

This proposition derives immediately from Equation (1) and the proof of Proposition 16.

We recall that we defined partial knowledge as the fact that player 1 knows only that $\big[u(a, b)\big]_{a \in \mathcal{A}}$ belongs to some polytope $\mathcal{U}_b \subset \mathbb{R}^A$, independently of $a \in \mathcal{A}$. If we make the stronger assumption that player 1 knows that $\big[u(a, b)\big]_{a \in \mathcal{A}, b \in \mathcal{B}}$ belongs to some polytope in $\mathbb{R}^{A \times B}$, then the semi-standard structure may fails.

Example 4 *Assume that the payoff matrix of player 1 belongs to the convex hull of the following two matrices* $u_1 \in \mathbb{R}^{A \times B}$ *and* $u_2 \in \mathbb{R}^{A \times B}$ *with*

$u_1 =$

	L	R
T_1	*1*	*0*
T_2	*0*	*0*
B_1	*0*	*1*
B_2	*0*	*0*

and

$u_2 =$

	L	R
T_1	*0*	*0*
T_2	*1*	*0*
B_1	*0*	*0*
B_2	*0*	*1*

Then for any $y \in [0,1]$,

$$\Phi\Big(yL + (1-y)R\Big) = \left\{ \begin{pmatrix} y\lambda \\ y(1-\lambda) \\ (1-y)\lambda \\ (1-y)(1-\lambda) \end{pmatrix} \; ; \; \lambda \in [0,1] \right\}$$

which is clearly not linear in Y *as*

$$\Phi(L) = \left\{ \begin{pmatrix} \lambda \\ 1-\lambda \\ 0 \\ 0 \end{pmatrix} \; ; \; \lambda \in [0,1] \right\} \text{ and } \quad \Phi(R) = \left\{ \begin{pmatrix} 0 \\ 0 \\ \mu \\ 1-\mu \end{pmatrix} \; ; \; \mu \in [0,1] \right\}$$

so that we immediately obtains

$$y.\Phi(L) + (1-y).\Phi(R) = \left\{ \begin{pmatrix} y\lambda \\ y(1-\lambda) \\ (1-y)\mu \\ (1-y)(1-\mu) \end{pmatrix} \; ; \; (\lambda,\mu) \in [0,1]^2 \right\} \neq \Phi\Big(yL + (1-y)R\Big).$$

Maybe even more surprisingly, Φ *is not piece-wise linear in y, i.e., it is not linear on any polytopial complex of* $\mathcal{Y}$*. Indeed* $\Phi(yL + (1-y)R)$ *can be seen as the set of product probability distributions over* $\{T, B\} \times \{1, 2\}$ *with first marginal* $yT + (1-y)B$ *while* $y.\Phi(L) + (1-y).\Phi(R)$ *is the set of all probability distributions (non necessarily product) with the same first marginal.*

To conclude with the links between the semi-standard structure and unknown payoffs, we mention that if both players have only 2 actions, then Φ and Ψ are always linear, thus our results extend immediately to that case.

Appendix C. About the non-degeneracy assumption

C.1. On the non-degeneracy assumption

Notice that if a player has 2 actions, then the non-degeneracy assumption on the structure of information of her opponent is satisfied in *almost all game*, i.e., if payoffs are chosen at random with a density, say uniformly in $[0,1]$. This is simply due to the fact that the opponent is either in the dark or with full monitoring.

This can also be proved easily for 3 actions.

Proposition 20 *If both players have 3 actions and payoffs are chosen uniformly at random in* $[0,1]$*, then the non-degeneracy assumption is satisfied with probability 1.*

Proof: Since player 2 has 3 actions, there are only 4 possibilities for the structure of information of player 1:

i) Player 1 has full monitoring, i.e., Φ is always a singleton;

ii) Player 1 is in the dark, i.e., $\mathbf{M}^{-1}[\mathbf{M}](y) = \mathcal{Y}$ for every $y \in \mathcal{Y}$;

iii) Player 1 has a semi-standard structure of information, if $\mathbf{M}(b) = \mathbf{M}(b')$ for two different action $b, b' \in \mathcal{B}$;

iv) Player 1 has not a semi-standard structure of information.

In case i), the non-degeneracy assumption is satisfied in almost all game as in bi-matrix games. In case ii), best responses are maxmin strategies of player 1, which are generically unique.

Case iii) can also be solved by hand. Denote by $\mathcal{A} = \{T, M, B\}$ and $\mathcal{B} = \{L, C, R\}$ the action sets of both players. We assume for the moment that $\mathbf{N}$ is the identity and we define the payoff matrix of player 1 by

$$u = \begin{array}{c|c|c|c|} & L & C & R \\ \hline T & a & b & c \\ \hline M & d & e & f \\ \hline B & g & h & i \\ \hline \end{array} \qquad \text{and } \mathbf{M}(L) = 0, \mathbf{M}(C) = \mathbf{M}(R) = 1 \ ,$$

where $a, b, c, \dots, i$ are chosen uniformly independently at random in $[0, 1]$. Without loss of generality, we can assume that $b > c$, $h > i$ and $e < f$, otherwise either action C or R could be discarded. Simple computations show that the set $\{x[1], x[2], \dots, x[L]\}$ of best responses of player 1 is equal (or included in) to $\{T, B, M, \lambda T + (1 - \lambda)M, \gamma B + (1 - \gamma).M\}$ with

$$\lambda = \frac{f - e}{f - e + b - c} \in (0, 1) \ , \gamma = \frac{f - e}{f - e + h - i} \in (0, 1) \ .$$

As a consequence, the mapping $\widetilde{u}$ of player 1 is defined by

$$\widetilde{u} = \begin{array}{c|c|c|c|} & L & C & R \\ \hline T & a & c & c \\ \hline M & d & e & e \\ \hline B & g & i & i \\ \hline \lambda T + (1-\lambda)M & \lambda a + (1-\lambda)d & \lambda b + (1-\lambda)e & \lambda b + (1-\lambda)e \\ \hline \gamma B + (1-\gamma)M & \gamma g + (1-\gamma)d & \gamma h + (1-\gamma)e & \gamma h + (1-\gamma)e \\ \hline \end{array}$$

It remains to show that this game is non-degenerate for player 1, as it is immediate for player 2. We mention here that the fact the the payoffs associated to C and R are identical is not an issue for non-degeneracy. It would be an issue if payoffs on different lines were equal. We have to show the functions $y \mapsto \widetilde{u}(x[\ell], y)$ are in non-degenerate positions, and we can actually prove that it is impossible that three of them intersect at the same point.

Indeed, it is quite immediate to see that if $\{T, B, M\}$ is in non-degenerate position, which happens with probability 1, then $\{T, M, \gamma B + (1 - \gamma)M\}$, $\{B, M, \lambda T + (1 - \lambda)M\}$, $\{T, B, \lambda T + (1 - \lambda)M\}$ and $\{M, \gamma B + (1 - \gamma)M, \lambda T + (1 - \lambda)M\}$, etc. are in non-degenerate positions.

Finally, simple computations show that $\{T, M, \lambda T + (1 - \lambda)M\}$ is in non-degenerate position in the bi-matrix game with payoff $\widetilde{u}$, unless $b = c$, which happens with probability 0; the same holds for $\{B, M, \gamma B + (1 - \gamma)M\}$.

Case iv) can be seen as a consequence of case iii). Indeed, Φ is piecewise-linear with respect to either one, see Example 3, or two sub-polytopes of $\mathcal{Y}$ as in Example 2. As a consequence, case iv) can be reduced to two instances of case iii). □

This approach can be generalized to higher numbers of actions as soon as the range of $\mathbf{M}$ and $\mathbf{N}$ are of dimension at most 1 (which is always the case with 3 actions).

Appendix D. Deleted Proofs

D.1. Proof of Proposition 3

Define, for every $i \in \mathcal{I}$, the set of compatible outcomes with $\mathbf{M}_i$ by:

$$\mathcal{U}_i := \left\{[u(a,y)]_{a\in\mathcal{A}} \in \mathbb{R}^A;\ y \text{ s.t. } \mathbf{M}(y) = \mathbf{M}_i\right\} = \left\{[u(a,y)]_{a\in\mathcal{A}} \in \mathbb{R}^A;\ y \in \Delta(\mathcal{B}_i)\right\}$$
$$= \operatorname{co}\left\{[u(a,b)]_{a\in\mathcal{A}};\ b \in \mathcal{B}_i\right\},$$

where co stands for the convex hull. In particular, $\mathcal{U}_i$ is a polytope and it is equal to $\Phi(b)$, for all $b \in \mathcal{B}_i$. Since $\mathbf{M}$ is linear, we obtain that

$$\mathbf{M}(y) = \sum_{b\,\in\mathcal{B}} y_b \mathbf{M}(b) = \sum_{i\in\mathcal{I}} \sum_{b\in\mathcal{B}_i} y_b \mathbf{M}_i = \sum_{i\in\mathcal{I}} y[i]\mathbf{M}_i,$$

where $y[i]$ is the weight put by y on $\mathcal{B}_i$, i.e., $y[i] = \sum_{b\in\mathcal{B}_i} y_b$. The semi-standard assumption implies that if $\mathbf{M}(y) = \mathbf{M}(y')$ then $y[i] = y'[i]$ for all $i \in \mathcal{I}$. As a consequence,

$$\Phi(y) = \sum_{i\in\mathcal{I}} y[i]\mathcal{U}_i = \sum_{i\in\mathcal{I}} \sum_{b\in\mathcal{B}_i} y_b \mathcal{U}_i = \sum_{b\in\mathcal{B}} y_b \Phi(b).$$

$\Phi(y)$ is a polytope since finite sums of polytopes are polytopes.

D.2. Proof of Lemma 4

Fix for the moment $y \in \mathcal{Y}$ and $x^* \in \operatorname{argmax}_{x\in\mathcal{X}} \min_{U\in\Phi(y)}\langle x, U\rangle$.

If $U_* \in \Phi(y)$ is a minimizer of the program $\min_{U\in\Phi(y)}\langle x^*, U\rangle$ then it can be assumed that U_* is a vertex of $\Phi(y)$, because a linear program is always minimized on a vertex of the admissible polytope. As a consequence, $-x^*$ belongs necessarily to the normal cone at U_* to $\Phi(y)$ (Rockafellar, 1970, Theorem 27.4, page 270). So $-x^*$ belongs to the polytopial intersection of $-\mathcal{X}$ and a normal cone at one of the vertices of $\Phi(y)$. More precisely, since $\langle x, U_*\rangle$ is linear, $-x^*$ must be one of the vertices of this intersection, or a convex combination of them.

However, $\Phi(\cdot)$ is linear on $\mathcal{Y}$, so normal cones at vertices – their set is called "normal fan" – are constant, see (Ziegler, 1995, Example 7.3, page 193) and (Billera and Sturmfels, 1992, page 530). Hence there exists a finite number of intersections between $-\mathcal{X}$ and normal cones and they all have a finite number of vertices. So the set of every possible vertices is finite and denoted by $-\{x[1], \ldots, x[L]\}$: it always contains a maximizer and any maximizer belongs to its convex hull.

Since Φ is linear, $y \mapsto \min_{U\in\Phi(y)}\langle x[\ell], U\rangle$ is also linear, for every $\ell \in \{1, \ldots, L\}$. This implies that $x[\ell]$ is a maximizer on a polytopial subset of $\mathcal{Y}$.

D.3. Proof of Corollary 6

Notice that since $d_{\mathbf{M}}$, the dimension of the kernel of $\mathbf{M}$, is fixed then the number of actions of player 2 that are in some equivalent class for player 1 (i.e., for each of those actions, there exists another undistinguishable action in $\mathcal{B}$) is upper-bound by $d_{\mathbf{M}}^{d_{\mathbf{M}}}$. And thus, so is bounded the fixed cardinality of the set of polyhedra defining the normal fan and thus the total number of defining hyperplanes is also uniformly bounded by, say, M.

Recall that the different $x[\ell]$ are constructed as the intersection of M hyperplanes defining the normal fan and the A hyperplanes defining $-\mathcal{X}$. The total number of possible intersection is therefore polynomial in A.

D.4. Proof of Theorem 7

Let $(\mathbf{x}, \mathbf{y}) \in \Delta(\mathcal{L} \times \mathcal{K})$ be any Nash equilibrium of the associated bi-matrix game and let $(x, y) \in \mathcal{X} \times \mathcal{Y}$ be the induced mixed actions. Since Φ is linear and $y = \sum_{k \in \mathcal{K}} \mathbf{y}_k y[k]$, we obtain that $\sum_{k \in \mathcal{K}} \mathbf{y}_k \Phi(y[k]) = \Phi(y)$. As a consequence,

$$\begin{aligned}
\widetilde{u}(\mathbf{x}, \mathbf{y}) &= \sum_{\ell \in \mathcal{L}} \sum_{k \in \mathcal{K}} \mathbf{x}_\ell \mathbf{y}_k \widetilde{u}(\ell, k) && (\text{ by linearity of } \widetilde{u}\,) \\
&= \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell \sum_{k \in \mathcal{K}} \mathbf{y}_k \min_{U_{\ell,k} \in \Phi(y[k])} \langle x[\ell], U_{\ell,k} \rangle && (\text{ by definition of } \widetilde{u}\,) \\
&= \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell \min_{U_\ell \in \sum_{k \in \mathcal{K}} \mathbf{y}_k \Phi(y[k])} \langle x[\ell], U_\ell \rangle && (\text{ by definition of sums of sets}\,) \\
&= \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell \min_{U_\ell \in \Phi(y)} \langle x[\ell], U_\ell \rangle && (\text{ by linearity of } \Phi\,) \\
&\leq \min_{U \in \Phi(y)} \left\langle \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell x[\ell], U \right\rangle && (\text{ by concavity of } \min\,) \\
&= \min_{U \in \Phi(y)} \langle x, U \rangle. && (\text{ since } x \text{ is induced by } \mathbf{x}\,).
\end{aligned}$$

Using the same arguments, we obtain that

$$\begin{aligned}
\min_{U \in \Phi(y)} \langle x, U \rangle &\geq \widetilde{u}(\mathbf{x}, \mathbf{y}) && (\text{ previous displayed equation}\,) \\
&\geq \max_{\ell \in \mathcal{L}} \widetilde{u}(\ell, \mathbf{y}) && (\text{ because } (\mathbf{x}, \mathbf{y}) \text{ is a Nash equilibrium}\,) \\
&= \max_{\ell \in \mathcal{L}} \min_{U \in \Phi(y)} \langle x[\ell], U \rangle && (\text{ by linearity of } \Phi\,) \\
&= \max_{x' \in \mathcal{X}} \min_{U \in \Phi(\mathbf{y})} \langle x', U \rangle && (\text{ by Lemma 4}\,).
\end{aligned}$$

We have just proved that x is a best response to $\Phi(y)$; similarly y is a best response to $\Psi(x)$. Therefore (x, y) is a robust Nash equilibrium of Γ.

Reciprocally, let $(x, y) \in \mathcal{X} \times \mathcal{Y}$ be a robust Nash equilibrium. Lemma 4 implies that x is a convex combinations of elements of $\{x[1], \ldots, x[L]\}$ that maximize $\min_{U \in \Phi(y)} \langle x[\ell], U \rangle$. Denote by $\mathbf{x} \in \Delta(\mathcal{L})$ this convex combination and define $\mathbf{y}$ in a dual way.

It is clear that $\mathbf{x} \in \Delta(\mathcal{L})$ induces x and that $\mathbf{y} \in \Delta(\mathcal{K})$ induces y. As a consequence, we obtain that, for every $\ell' \in \mathcal{L}$:

$$\begin{aligned}
\widetilde{u}(\ell', \mathbf{y}) &\leq \max_{\ell \in \mathcal{L}} \widetilde{u}(\ell, \mathbf{y}) = \max_{\ell \in \mathcal{L}} \min_{U \in \Phi(y)} \langle x[\ell], U \rangle && (\text{ by linearity of } \Phi\,) \\
&= \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell \min_{U \in \Phi(y)} \langle x[\ell], U \rangle && (\text{ because } \mathbf{x}_\ell > 0 \text{ if it is a maximizer}\,) \\
&= \widetilde{u}(\mathbf{x}, \mathbf{y}) && (\text{ by linearity of } \Phi\,).
\end{aligned}$$

Therefore $\mathbf{x}$ is a best response to $\mathbf{y}$ in the bi-matrix game and the converse is true by symmetry; $(\mathbf{x}, \mathbf{y})$ is a Nash equilibrium of the bi-matrix game that induces (x, y).

D.5. Proof of Lemma 10

Since $\mathbf{M}$ is linear from $\mathcal{Y}$ into $\mathbb{R}^k$, then $\mathbf{M}^{-1}(\cdot)$ is piecewise linear on $\mathbb{R}^k$, see (Billera and Sturmfels, 1992, page 530) and (Rambau and Ziegler, 1996, Proposition 2.4, page 221). Therefore, by composition, $\mathbf{M}^{-1}[\mathbf{M}]$ is piecewise linear on $\mathcal{Y}$ and as a consequence Φ is also piecewise linear on $\mathcal{Y}$.

The number of elements on which Φ is linear can be enumerated as in the proof of Corollary 6. Indeed, uncertainty sets are defined by the $(B-1-d_{\mathbf{M}})$-faces its extreme points belong to. The result is just a matter of enumeration of those faces.

D.6. Proof of Lemma 11

Since Φ is piece-wise linear, it is linear on every element of some polytopial complex $\{\mathcal{Y}_1, \ldots, \mathcal{Y}_I\}$ of $\mathcal{Y}$. As a consequence, on each polytope $\mathcal{Y}_i$, there exists a finite set of mixed action $\{x[1,i], x[2,i], \ldots, x[L_i, i]\}$ of best-responses, that are each best response on a polytopial subset $\mathcal{Y}_{\ell,i}$ of Y_i.

Define $\{x[\ell];\ \ell \in \mathcal{L}\} = \cup_{i=1}^{I}\{x[1,i], x[2,i], \ldots, x[L_i, i]\}$ and $\mathcal{Y}_\ell$ as the union of $\mathcal{Y}_{\ell_i,i}$ such that $x[\ell_i, i] = x[\ell]$. In particular, $\mathcal{Y}_\ell$ is the union of at most I polytopes.

The construction of $y[k]$ and $\mathcal{X}_k$ is similar.

Since the total number of elements on which Φ is linear is polynomial in B and, on each element, the number of mixed actions is polynomial in A, the last statement follows.

D.7. Proof of Theorem 12

Proof: Let $(\mathbf{x}, \mathbf{y}) \in \Delta(\mathcal{L}) \times \Delta(\mathcal{K})$ be a fully labeled point and $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ be the induced mixed actions. Then x is a best response to y and reciprocally since:

$$\begin{aligned}
\min_{U \in \Phi(y)} \langle x, U\rangle &= \min_{U \in \Phi(y)} \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell \langle x[\ell], U\rangle && (\text{ by definition of induced action})\\
&\geq \sum_{\ell \in \mathcal{L}} \mathbf{x}_\ell \min_{U \in \Phi(y)} \langle x[\ell], U\rangle && (\text{ by concavity of the min})\\
&\geq \max_{x' \in \mathcal{X}} \min_{U \in \Phi(y)} \langle x', U\rangle && (\text{because } \ell \text{ is a label if } \mathbf{x}_\ell = 0\\
& && \text{or } x[\ell] \text{ is a maximiser}).
\end{aligned}$$

Therefore, any fully labeled point induces a robust Nash equilibrium.

If (x, y) is a robust Nash equilibrium then Lemma 11 implies that x and y belong to the convex hull of $\{x[\ell];\ \ell \in \mathcal{L}\}$ and $\{y[k];\ k \in \mathcal{K}\}$. More precisely, x is a convex combination of the maximizers of $\min_{U \in \Phi(y)}\langle x[\ell], U\rangle$, i.e. precisely those $x[\ell]$ such that $y \in Y[\ell]$. If we denote this convex combination as $\mathbf{x} \in \Delta(\mathcal{L})$, then necessarily either $\mathbf{x}_\ell = 0$ or y belongs to Y_ℓ, and thus $\mathbf{y} \in \mathbf{Y}_\ell$. Therefore $(\mathbf{x}, \mathbf{y})$ is fully labeled.

D.8. Proof of Lemma 13

Since, by definition,

$$Y_\ell = \left\{ y \in \mathcal{Y} \text{ s.t. } x[\ell] \in \arg\max_{\ell' \in \mathcal{L}} \inf_{U \in \Phi(y)} \langle x[\ell'], U\rangle \right\}$$

is a polytope of $\mathbb{R}^B$, there exists a finite family $\left\{b_t \in \mathbb{R}^B, c_t \in \mathbb{R}; t \in \mathcal{T}_\ell\right\}$ such that

$$Y_\ell = \bigcap_{t \in T_\ell} \left\{ y \in \mathcal{Y} \text{ s.t. } \langle y, b_t \rangle \leq c_t \right\} .$$

Therefore, $\mathbf{Y}_\ell$ is also a polytope of $\mathbb{R}^m$ as it can be written as

$$\begin{aligned} \mathbf{Y}_\ell &= \bigcap_{t \in T_\ell} \left\{ \mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } \langle \mathbb{E}_{\mathbf{y}}[y], b_t \rangle \leq c_t \right\} \\ &= \bigcap_{t \in T_\ell} \left\{ \mathbf{y} \in \Delta(\mathcal{K}) \text{ s.t. } \left\langle \mathbf{y} \,\middle|\, \big(\langle y[k], b_t \rangle \big)_{k \in \mathcal{K}} \right\rangle \leq c_t \right\}, \end{aligned}$$

where the inner product $\langle \cdot | \cdot \rangle$ in the last equation stands for the usual inner product of $\mathbb{R}^{|\mathcal{K}|}$.

Similar arguments hold for $\{\mathbf{X}_k;\ k \in \mathcal{K}\}$.