

Equilibria in Ordinal Games: A Framework based on Possibility Theory.

Nahla Ben Amor¹ and H el ene Fargier² and R egis Sabbadin³

¹ISG-Tunis

²Universit e de Toulouse

³INRA-MIAT, Universit e de Toulouse

nahla.benamor@gmx.fr, helene.fargier@irit.fr, regis.sabbadin@inra.fr

Abstract

The present paper proposes the first definition of a mixed equilibrium in an ordinal game. This definition naturally extends possibilistic (single agent) decision theory. Our first contribution is to show that ordinal games always admit a possibilistic mixed equilibrium, which can be seen as a qualitative counterpart to mixed (probabilistic) equilibrium. Then, we show that a possibilistic mixed equilibrium can be computed in polynomial time (wrt the size of the game), which contrasts with mixed probabilistic equilibrium computation in cardinal game theory. The definition we propose is thus operational in two ways: (i) it tackles the case when no pure Nash equilibrium exists in an ordinal game; and (ii) it allows an efficient computation of a mixed equilibrium.

1 Introduction

Game theory used to be mainly a science for economic theorists, until recently when, with the expansion of internet, many practical instances have emerged. Accordingly, in the recent years there has been a growing interest in the AI community for algorithmic game theory [Nisan *et al.*, 2007]. From a practical viewside of game theory, there are instances where the assumption of cardinal payoffs for players can be questioned. For example, in a recent paper, [Ouenniche *et al.*, 2016] describe a realistic example of Public-Private Partnership project, which is dealt with using ordinal game theory.

Ordinal games [Cruz Jr. and Simaan, 2000; Xu, 2000] are thus a natural framework to consider as soon as no cardinal payoffs are available to players. The first analysis of equilibria in ordinal games dates back to [Barany *et al.*, 1992], with a classification of two-person ordinal games in terms of their Pareto sets (pairs of pure strategies which are not uniformly dominated). [Xu, 2000] defines *Ordinal Nash Equilibria (ONE)* as the natural counterpart of Nash equilibria and proposes an algorithm that computes ordinal Nash equilibria for 2 and 3-players ordinal games, in polynomial time. From the theoretical point of view, [Durieu *et al.*, 2008] study a few categories of ordinal games, among which *ordinal best-response potential games*, and show that they always

admit ordinal Nash equilibria. This property is a counterpart of the one holding for cardinal potential games, which are also known to admit a pure Nash equilibrium [Monderer and Shapley, 1996]. Finally, [De Clercq *et al.*, 2014; 2015] study “possibilistic Boolean games”. These are a concise version of ordinal games, constructed from possibilistic logic preference bases. The set of alternatives of each player is given by the instantiation of a separate set of propositional variables. A pure Nash equilibrium (i.e. ONE) is defined as usual. Then, the authors focus on the ill-known ordinal game case, where players do not precisely know others’ preferences. They propose an evaluation of ONE in terms of their possibility and necessity degrees to be an ONE for the *true* (ill-known) ordinal game.

Besides, there has been little progress, if any, in the analysis of ordinal games which do not admit pure strategy equilibria. The only proposed approach [Cruz Jr. and Simaan, 2000; Ouenniche *et al.*, 2016] consists in listing all potential combinations of pure strategies and use the list of ranks of pure strategies of each player to analyze the game’s possible pure equilibria. This kind of approach is hardly feasible in practice and, again, it is oriented toward the search of a pure equilibrium. A consequence is that there is no real solution concept for ordinal games which do not admit an ONE.

The present paper proposes the first definition of *mixed equilibria* for ordinal games. This definition naturally extends possibilistic (single agent) decision theory [Dubois and Prade, 1995]. Our first contribution is to show that ordinal games always admit a possibilistic mixed equilibrium, which can be seen as a qualitative counterpart to mixed probabilistic equilibria. Then, we show that a possibilistic equilibrium can be computed in polynomial time, which contrasts with mixed probabilistic equilibria in cardinal game theory. The definition we propose is thus *operational* in two ways: (i) it tackles the case when no ONE exists in an ordinal game and (ii) it allows an efficient computation of mixed equilibria.

2 Ordinal Games and Nash Equilibria in Possibility Theory

2.1 Normal Form Games and Nash Equilibria

A *strategic game*, or *normal form game* is classically defined as follows [Von Neumann and Morgenstern, 1948]:

Definition 1 (Normal form game) A normal form game is a triple $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, where:

- $N = \{1, \dots, n\}$ is a finite set of players.
- A_i is a finite set of actions available to player $i \in N$ (also called strategy space). $A = \times_{i \in N} A_i$ is the product set of available actions. $a = (a_i)_{i \in N} \in A$ is a joint action of the n players (or a strategy profile).
- $(u_i)_{i \in N}$ is a set of utility functions. For any player i , $u_i : A \rightarrow L$ is the utility function of player i , i.e. $u_i(a)$ is the utility for agent i of the joint action $a \in A$.

$L = \{\alpha_0, \dots, \alpha_k\}$ is a finite totally ordered set, equipped with preorder relation \geq , which is sufficient to define pure Nash equilibria. We will use the following notation $\forall a \in A$:

- $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in A_{-i} = \times_{j \neq i} A_j$
- $(a'_i, a_{-i}) = (a_1, \dots, a_{i-1}, a'_i, a_{i+1}, \dots, a_n) \in A$

Informally, a pure Nash equilibrium, for a game $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$, is a joint action $a^* \in A$, from which no player has interest to deviate unilaterally. More formally:

Definition 2 (Pure Nash equilibrium)

Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a normal form game. $a^* \in A$ is a pure Nash equilibrium of G , if and only if:

$$u_i(a^*) \geq u_i(a_i, a_{-i}^*), \forall i \in N, \forall a_i \in A_i.$$

One known problem with normal form games is that a pure Nash equilibrium may not exist for a given game G . However, when utility functions take real values, the following has been proved (see, e.g. [Osborne and Rubinstein, 1994]). By allowing players to adopt mixed strategies in the form of probability distributions $p_i : A_i \rightarrow [0, 1]$, one can ensure the existence of a mixed Nash equilibrium p^* , verifying:

Definition 3 (Mixed Nash equilibrium)

Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a given normal form game. Let $p^* = (p_i^*)_{i \in N}$, where p_i^* is a probability distribution over A_i . p^* is a mixed Nash equilibrium of G , if and only if it verifies: $u_i(p^*) \geq u_i(p_i, p_{-i}^*), \forall i \in N, \forall p_i$, where, $u_i(p) =_{def} \sum_{a \in A} (\prod_{k \in N} p_k(a_k)) u_i(a)$ is the expected utility to player i of the joint mixed strategy p and $p_{-i}^* = \{p_j^*, j \in N, j \neq i\}$.

The existence of a mixed Nash equilibrium is guaranteed for any game [Nash, 1950]. However, [Daskalakis et al., 2006] have shown that it is hard to compute one (PPAD-complete: a complexity class of problems which are not assumed to be polynomial time solvable).

2.2 Ordinal Games

There are many natural examples of games in which the utility functions of players take their values in an ordinal scale. Let us consider the following example:

Example 1 (Firms competition) An established firm (E) and a newcomer (N) to a market of fixed size have to choose the packaging of a similar product. Each firm can choose between two different packagings: X and Y . The established producer prefers the newcomers product to look like its own-while the newcomer prefers that the products look different (to attract customers). The option Y is more risky than X for

the newcomer (in case E also plays Y), but will generate more profit if E plays X . E , however, is indifferent between X and Y . Ordinal outcomes (u_E, u_N) , compatible with this knowledge, are listed in the following table. These outcomes are ranked from 0 (worst case) to 4 (best case):

		N	
		X	Y
E	X	(2,2)	(1,4)
	Y	(1,3)	(2,0)

Preferences are purely ordinal, so are not proportional to any expected profit. If we translate ordinal utilities to their corresponding real values, this problem admits a single mixed Nash equilibrium¹, $p_E^* = (0.6; 0.4)$ and $p_N^* = (0.5; 0.5)$.

However, such a direct translation to quantitative values raises questions. If we replace the value $u_N(X, Y) = 4$ with $u_N(X, Y) = 5$, the new mixed equilibrium becomes $p_E^* = (0.5; 0.5)$ and $p_N^* = (0.5; 0.5)$. If we switch for $u_N(X, Y) = 6$, we get $p_E^* = (0.43; 0.57)$ and $p_N^* = (0.5; 0.5)$. The three games give identical ordinal orderings of all pure joint strategies to all players. However, with three different compatible cardinal translations of the utilities, the resulting Nash equilibria verify either $p_E^*(X) > p_E^*(Y)$, $p_E^*(X) = p_E^*(Y)$ or $p_E^*(X) < p_E^*(Y)$. Thus, a game theory that directly works on the ordinal expression of utilities would be useful. We propose such an approach, based on possibility theory, which handles ordinal games and avoids the bias linked to ordinal-cardinal utility translation.

2.3 Possibilistic Mixed Strategies in Ordinal Games

In this paper, an ordinal game is a triple $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$. The utility functions, μ_i , take their values in a finite ordinal scale, L , contrarily to utility functions u_i which can take cardinal values.

In this section, we define a notion of ordinal mixed strategy for such games, based on possibility theory [Dubois and Prade, 1988]. A possibility distribution over the alternatives of A_i of player i is a function $\pi_i : A_i \rightarrow L$ that captures a ranking over the elements of A_i . π_i can have a dual interpretation, either in terms of preference and in terms of likelihood. More precisely:

- For player i distribution π_i models the ranking of alternatives in terms of preference or comitment. Under this interpretation, $\pi_i(a_i) = 1_L$ means that a_i is satisfactory, fully conceivable to player i to play, while $\pi_i(a_i) = 0_L$ means that it is absolutely not an option for i ².
- When interpreted by other players, the ranking π_i can be seen as a measure of likelihood of play. In this case, $\pi_i(a_i) = 1_L$ means that a_i is a completely plausible play of player i . On the contrary, $\pi_i(a_i) = 0_L$ is interpreted as a_i is an impossible play of player i .

This dual preference/likelihood interpretation is natural in game theory since, according to the other players, the most

¹Which was computed using <http://gambit.sourceforge.net/>.

² 0_L and 1_L are the lowest and highest elements of L .

preferred alternatives of player i should be the most likely to be played. Possibilistic mixed strategies are defined as *normalized* possibility distributions $\pi_i : A_i \rightarrow L$:

$$\max_{a_i \in A_i} \pi_i(a_i) = 1_L, \text{ the top element of } L.$$

The joint possibility distribution over joint strategies $a = (a_1, \dots, a_n)$, is defined as the *minimum* of the individual players possibilistic mixed strategies:

$$\pi(a) = \min_{i \in N} \pi_i(a_i), \forall a = (a_1, \dots, a_n) \in A.$$

2.4 Evaluating Possibilistic Mixed Strategies

Let an ordinal game $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$ be given and assume that a set of mixed possibilistic strategies π_i over A_i have been defined, π being the corresponding joint mixed possibilistic strategy. Now, the question is to measure the utility of π to each player, or, in other terms, the ranking of joint mixed possibilistic strategies by each player.

The two following *optimistic* and *pessimistic* criteria have been proposed and axiomatized for decision under uncertainty in the possibilistic framework [Dubois and Prade, 1995; Dubois *et al.*, 2001].

$$\begin{aligned} \mu_i^{OPT}(\pi) &= \max_{a \in A} \min(\pi(a), \mu_i(a)). \\ \mu_i^{PES}(\pi) &= \min_{a \in A} \max(n(\pi(a)), \mu_i(a)). \end{aligned} \quad (1)$$

Where n is the order-reversing map of L . For instance, if $L = \{\alpha_0, \dots, \alpha_K\}$, $n(\alpha_k) = \alpha_{K-k}$, $\forall \alpha_k \in L$.

$\mu_i^{OPT}(\pi)$ and $\mu_i^{PES}(\pi)$ can be readily adapted to the evaluation of joint mixed possibilistic strategies, when these are interpreted as measuring the likelihood of joint strategies. They are *soft* versions of the *maximax* and *maximin* criteria. When there is no knowledge about what other players will play, $\pi(a) = 1_L, \forall a \in A_1 \times \dots \times A_n$. Then:

- According to $\mu_i^{OPT}(\pi)$, each player assumes that all other players will be benevolent to him, i.e. act for the best, since $\mu_i^{OPT}(\pi) = \max_a \mu_i(a)$.
- On the contrary, according to $\mu_i^{PES}(\pi)$, each player assumes that all other players will act in the worst way for him, since $\mu_i^{PES}(\pi) = \min_a \mu_i(a)$.

Note that, as for expected utility theory, possibilistic decision criteria assume a form of *commensurateness* between uncertainty and preference levels [Dubois *et al.*, 2001]. In the framework of *competitive game theory*, μ^{OPT} is rather unnatural, so we will focus on μ^{PES} .

2.5 Possibilistic Mixed Equilibria

A possibilistic mixed equilibrium (PME) can be informally defined, as in the stochastic case, as a set $\pi^* = (\pi_1^*, \dots, \pi_n^*)$ of (normalized) preference rankings from which no player has interest to deviate unilaterally. Formally, we define:

Definition 4 (Possibilistic Mixed equilibrium)

Let $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$ be an ordinal game. $\pi^* = (\pi_1^*, \dots, \pi_n^*)$ is a *possibilistic mixed equilibrium* for G if and only if, for any *possibilistic mixed strategy* π , we have:

$$\mu_i^{PES}(\pi^*) \geq \mu_i^{PES}(\pi_i, \pi_{-i}^*), \forall i \in N.$$

In this definition, the double nature of the joint possibility distribution π^* appears. Indeed, while the π_i^* are obviously interpreted as likelihood rankings in Equation (1), they are defined as solutions to equations involving preference rankings only in the above definition. In this respect, as aggregations of preference rankings, they should be interpreted as preference rankings themselves.

Note that, *deterministic* equilibria can be interpreted in the same way in both possibilistic game theory and probabilistic game theory settings. However, in the mixed-equilibrium case, while usual mixed Nash equilibria are interpreted as random plays that give good outcomes "on average", in the possibilistic setting these should be interpreted as "commitments" in a negotiation process (see Section 2.6), that is the fixed point of a negotiation where players indicate which options they may consider in actual play and which they are more prone to play than others. As usual in possibility theory, this is a "one-shot game" interpretation, where games are not repeated and possible outcomes cannot be "averaged"...

Let us now look at some properties of possibilistic mixed equilibria and their computation. We are going to give a constructive proof of the existence of a PME for any ordinal game in the next section. But before this, let us define a *least-specific possibilistic Nash equilibrium*. First, recall the usual definition of *specificity relation* of possibility distributions.

Definition 5 (Specificity relation)

Let π_i and π'_i be two (normalized) possibility distributions over A_i . We say that π'_i is more specific than π_i (written $\pi'_i \preceq \pi_i$), if the following holds: $\pi'_i(a_i) \leq \pi_i(a_i), \forall a_i \in A_i$.

\prec is the corresponding strict order. By extension, for mixed strategies $\pi = (\pi_i)_{i \in N}$ and $\pi' = (\pi'_i)_{i \in N}$, we say that π' is more specific than π when the distributions π'_i are all more specific than the distributions $\pi_i, \forall i \in N$. That is:

$$\pi' \preceq \pi \Leftrightarrow \pi'_i(a_i) \leq \pi_i(a_i), \forall a_i \in A_i, \forall i \in N.$$

Assume that a possibilistic joint strategy π is given. We can show that the *pessimistic* utility of a player never decreases when the joint strategy of all players is made more specific, i.e. if players precise their intentions.

Proposition 1 (Improvement through specificity) Let $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$ be an ordinal game. Let π be an arbitrary possibilistic mixed strategy and let $\pi' \preceq \pi$ be a more specific possibilistic mixed strategy. Then, we can show that:

$$\mu_i^{PES}(\pi) \leq \mu_i^{PES}(\pi').$$

Proof: Notice that

$$\mu_i^{PES}(\pi) = \min_a \max \left(\max_{j \in N} n(\pi_j(a_j)), \mu_i(a) \right). \text{ So,}$$

$$\begin{aligned} \pi' \preceq \pi &\Leftrightarrow \pi'_j(a_j) \leq \pi_j(a_j), \forall j \in N, a_j \in A_j, \\ &\Leftrightarrow n(\pi_j(a_j)) \leq n(\pi'_j(a_j)), \forall j \in N, a_j \in A_j, \\ &\Rightarrow \max_{j \in N} n(\pi_j(a_j)) \leq \max_{j \in N} n(\pi'_j(a_j)), \forall a \in A, \\ &\Rightarrow \min_a \max \left(\max_{j \in N} n(\pi_j(a_j)), \mu_i(a) \right) \\ &\leq \min_a \max \left(\max_{j \in N} n(\pi'_j(a_j)), \mu_i(a) \right). \end{aligned}$$

So, $\pi' \preceq \pi \Rightarrow \mu_i^{PES}(\pi) \leq \mu_i^{PES}(\pi')$. \square

Updating a player's mixed strategy for a more specific one cannot decrease his own utility if all other players stick to their own strategy. However, changing one's strategy for a more specific one may give incentive to other players to change their own, as the following example shows.

Example 2 (IIME and specificity) Let us consider the following ordinal game with two players, two actions and a simple binary scale $\{0, 1\}$: Since utilities are binary, possibility

		N	
		X	Y
E	X	(0,1)	(1,0)
	Y	(1,0)	(0,1)

distributions correspond to subsets of $\{X, Y\}$. Now, we can check that $\pi = \{\pi_E = \{X, Y\}, \pi_N = \{X, Y\}\}$ has utility zero to both players. If π_E is made more specific (for example $\pi'_E = \{X\}$) and π_N is left unchanged, this does not modify the utility to player E. By symmetry, no player can unilaterally improve his utility. So, π is a IIME. Note that $\{\pi'_E, \pi_N\}$ is no more a IIME: If player N sees that player E has switched to $\pi'_E = \{X\}$, he will be better off if he changes his own strategy to $\pi'_N = \{X\}$ and gets utility 1. In turn, this will give an incentive to E to switch to $\pi'_E = \{Y\}$, etc. Indeed, the only IIME of this game is π .

It will be useful, when interpreting and computing mixed equilibria, to focus on those which are the least-specific, i.e. put the less restrictions on players' rankings of alternatives:

Definition 6 (Least-specific IIME)

Let $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$ be an ordinal game. $\pi^* = (\pi_1^*, \dots, \pi_n^*)$ is a least-specific possibilistic mixed equilibrium for G if and only if: (i) π^* is a possibilistic mixed equilibrium for G and (ii) There exists no π' s.t. $\pi^* \prec \pi'$ and π' is a IIME.

2.6 A Negotiation Interpretation of Least-Specific Possibilistic Mixed Equilibria

We will give a constructive proof of the existence of a possibilistic mixed equilibrium for any ordinal game in the following section. However, in order to introduce the proof, it is useful to recall the double nature of possibilistic mixed strategies. These are first *preference rankings* of alternatives expressed by each player. On the other hand, preference rankings of other players can be interpreted by a player as likelihood rankings of their available strategies.

So, when talking about the improvement and the specificity of a IIME one is faced with a double interpretation, again:

- Interpreting other players' strategies as likelihood rankings, a given player shows a kind of *aversion for uncertainty*: It is more comfortable to see other players commit to a given play. This is reflected in Proposition 1 where it is shown that more specific mixed strategies lead to a better utility.
- On the other hand, as far as player i is concerned, his mixed strategy π_i can be viewed as a *commitment* to play: The more specific π_i the more constraints i puts

on himself. Thus, one can expect players to prefer to put the *lightest* constraints on themselves, i.e. to keep π_i as less specific as possible.

According to these thoughts, one can see the search for a possibilistic mixed equilibrium as a form of *negotiation* where players trade, in turn, commitments on their own strategies, so as to improve their position in the joint equilibrium, until no player has interest to trade anymore. We will see that this iterative negotiation procedure leads to a (least-specific) Nash equilibrium. Furthermore, we will show that a IIME can be computed in time polynomial in the size of the utility tables μ_i , unlike "quantitative" mixed Nash equilibria.

3 A Polynomial Time Algorithm for Building Possibilistic Mixed Equilibria

The proof of the main result of this paper goes through three steps. First, we define the problem of deciding, for a given ordinal game, joint mixed possibilistic strategy and player, whether this player can improve his utility by making (unilaterally) his mixed strategy more specific. Then we propose a polynomial-time procedure which answers this question and, if the answer is positive, returns an improved strategy. Finally, we use this procedure as a building block for a polynomial-time algorithm which computes a IIME for any arbitrary ordinal game.

3.1 Possibilistic Mixed Strategy Improvement

Assume ordinal game $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$, and mixed strategy $\pi = (\pi_1, \dots, \pi_n)$ are given. Let us consider the point of view of player i . We want to check whether π_i can be improved by making it *minimally* more specific.

To be more precise, we are looking for $\pi'_i \prec \pi_i$, such that:

$$\mu_i^{PES}(\pi'_i, \pi_{-i}) > \mu_i^{PES}(\pi).$$

Furthermore, we want π'_i to be a least-specific such distribution (there may exist more than one), i.e.

$$\mu_i^{PES}(\pi''_i, \pi_{-i}) \leq \mu_i^{PES}(\pi), \forall \pi'_i \prec \pi''_i \preceq \pi_i.$$

First, let us notice that

$$\begin{aligned} \mu_i^{PES}(\pi) &= \min_{a \in A} \max \left(\max_{j \in N} n(\pi_j(a_j)), \mu_i(a) \right), \\ &= \min_{a_i \in A_i} \max \left(n(\pi_i(a_i)), \mu_i^{PES}(a_i, \pi_{-i}) \right) \end{aligned} \quad (2)$$

where $\mu_i^{PES}(a_i, \pi_{-i})$ is the utility player i gets when playing a_i , when the other players follow the joint strategy π_{-i} :

$$\mu_i^{PES}(a_i, \pi_{-i}) = \min_{a_{-i} \in A_{-i}} \max \left(\max_{j \in N \setminus i} n(\pi_j(a_j)), \mu_i(a) \right).$$

Note that the computation of $\mu_i^{PES}(\cdot, \pi_{-i})$, $A_i \rightarrow L$ takes polynomial time in the game expression. Now, for a given triple (G, π, i) , let $D_i \subseteq A_i$ define the set of *dominated actions* of player i , given π :

$$D_i = \{a_i \in A_i \text{ s.t. } \mu_i^{PES}(a_i, \pi_{-i}) \leq \mu_i^{PES}(\pi)\}.$$

That is, when switching from the distribution π_i to the *fixed* action a_i , player i does not improve his utility.

Intuitively, in order to improve his situation in the game, player i will make new commitments by decreasing, when possible, the plausibility of his dominated actions, $a_i \in D_i$, leaving the possibility of non-dominated actions unchanged.

Note that the following two facts hold:

$$\mu_i^{PES}(\pi) \geq \mu_i^{PES}(a_i, \pi_{-i}), \forall a_i \in D_i, \quad (3)$$

$$\mu_i^{PES}(\pi) = \min_{a_i \in D_i} \max(n(\pi_i(a_i)), \mu_i^{PES}(a_i, \pi_{-i})) \quad (4)$$

Fact (3) holds by definition of D_i , while fact (4) holds since,

$$\max(n(\pi_i(a_i)), \mu_i^{PES}(a_i, \pi_{-i})) > \mu_i^{PES}(\pi), \forall a_i \in A_i \setminus D_i.$$

Thus, if we want to transform π into a more specific distribution π' , by changing only $\pi_i(a_i)$ for $a_i \in D_i$, in order to get an increased utility $\mu_i^{PES}(\pi') > \mu_i^{PES}(\pi)$, we have to ensure that:

$$n(\pi'_i(a_i)) > \mu_i^{PES}(\pi) \Leftrightarrow \pi'_i(a_i) < n(\mu_i^{PES}(\pi)), \forall a_i \in D_i.$$

This leads to the definition of $\pi' = \text{IMPROVE}(G, \pi, i)$:

- $\pi'_{-i} = \pi_{-i}$: Only player i changes his commitment.
- If (i) $\Pi(A_i \setminus D_i) < 1_L$ or (ii) $\mu_i^{PES}(\pi) = 1_L$, then π cannot be improved and $\pi' = \pi$. Indeed, in the first case, we cannot decrease the possibility of all a_i in D_i while keeping a normalized distribution π_i . In the second case, there is no hope to improve $\mu_i^{PES}(\pi)$.
- If $\Pi(A_i \setminus D_i) = 1_L$ and $\mu_i^{PES}(\pi) < 1_L$, then

$$\pi'_i(a_i) = \pi_i(a_i), \forall a_i \in A_i \setminus D_i \text{ and}$$

$$\pi'_i(a_i) = \min\left(\pi_i(a_i), \underline{n(\mu_i^{PES}(\pi))}\right), \forall a_i \in D_i.$$

Given the above definition of the IMPROVE function, one can prove the following results:

Proposition 2 (Improvement function)

Let $G = \langle N, (A_i)_{i \in N}, (\mu_i^{PES})_{i \in N} \rangle$ be an ordinal game. Let also π be a joint possibilistic mixed strategy and $i \in N$ be a player. Then, the following facts hold:

1. $\text{IMPROVE}(G, \pi, i)$ is normalized.
2. $\text{IMPROVE}(G, \pi, i) \preceq \pi$.
3. If $\text{IMPROVE}(G, \pi, i) \prec \pi$, $\mu_i^{PES}(\text{IMPROVE}(G, \pi, i)) > \mu_i^{PES}(\pi)$. Else, $\text{IMPROVE}(G, \pi, i) = \pi$.
4. For any joint possibilistic mixed strategy π'' such that $\text{IMPROVE}(G, \pi, i) \prec \pi'' \preceq \pi$, we have:

$$\mu_i^{PES}(\pi'') = \mu_i^{PES}(\pi) < \mu_i^{PES}(\text{IMPROVE}(G, \pi, i)).$$

5. If $\pi' = \text{IMPROVE}(G, \pi, i)$, then, $\forall j \in N, \forall a_j \in A_j$
 $\mu_j^{PES}(a_j, \pi'_{-j}) \geq \mu_j^{PES}(a_j, \pi_{-j})$ and $\mu_j^{PES}(\pi') \geq \mu_j^{PES}(\pi)$.

Proof:

1. $\text{IMPROVE}(G, \pi, i)$ only changes (potentially) π_i . And when it does, we have taken the caution that the corresponding π'_i remains normalized.
2. The only potential change in $\pi' = \text{IMPROVE}(G, \pi, i)$ occurs when $\pi'_i(a_i) \leftarrow \min\left(\pi_i(a_i), \underline{n(\mu_i^{PES}(\pi))}\right)$. And, obviously, $\pi'_i(a_i) \leq \pi_i(a_i)$. So, $\pi'_i \preceq \pi_i$ and $\pi' \preceq \pi$.

3. Note that if $\text{IMPROVE}(G, \pi, i) = \pi' \prec \pi$, then

- (i) $\pi'_{-i} = \pi_{-i}$,
- (ii) $\forall a_i \in A_i \setminus D_i, \pi'_i(a_i) = \pi_i(a_i)$ and
- (iii) $\forall a_i \in D_i, \pi'_i(a_i) \leq \underline{n(\mu_i^{PES}(\pi))}$ and

$$\begin{aligned} \mu_i^{PES}(\pi') &= \mu_i^{PES}(\pi'_i, \pi_{-i}), \\ &= \min_{a_i \in A_i} \max(n(\pi'_i(a_i)), \mu_i^{PES}(a_i, \pi_{-i})), \\ &\geq \min_{a_i \in D_i} \max(n(\pi'_i(a_i)), \mu_i^{PES}(a_i, \pi_{-i})), \\ &\geq \min_{a_i \in D_i} \max\left(\underline{n(\mu_i^{PES}(\pi))}, \mu_i^{PES}(a_i, \pi_{-i})\right), \\ &\geq \underline{n(\mu_i^{PES}(\pi))} > \mu_i^{PES}(\pi). \end{aligned}$$

The ‘‘else’’ part is trivial since $(\pi' \preceq \pi \text{ and } \pi' \not\prec \pi) \Rightarrow \pi' = \pi$.

4. We proved in 3. that $\forall a_i \in D_i, \pi'_i(a_i) = \underline{n(\mu_i^{PES}(\pi))}$. Then, since $\pi'_{-i} = \pi_{-i}$ and $\forall a_i \in A_i \setminus D_i, \pi'_i(a_i) = \pi_i(a_i)$, we get: $\pi' \prec \pi'' \Rightarrow \exists a_i \in D_i$,

$$\begin{aligned} \pi''_i(a_i) &> \pi'_i(a_i), \\ \Leftrightarrow n(\pi''_i(a_i)) &\leq \mu_i^{PES}(\pi), \\ \Leftrightarrow \max(n(\pi''_i(a_i)), \mu_i^{PES}(a_i, \pi_{-i})) &\leq \max(\mu_i^{PES}(\pi), \mu_i^{PES}(\pi_{-i}(a_i))) = \mu_i^{PES}(\pi). \end{aligned}$$

Thus, $\mu_i^{PES}(\pi'') \leq \mu_i^{PES}(\pi)$. The other direction $\mu_i^{PES}(\pi'') \geq \mu_i^{PES}(\pi)$, comes from the specificity relation $\pi'' \preceq \pi$.

5. First, note that if $\pi' = \text{IMPROVE}(G, \pi, i)$, π' only differs from π in its component π'_i and $\pi'_i \preceq \pi_i$. Then, note that, $\forall j \in N \setminus i, a_j \in A_j$:

$$\begin{aligned} \mu_j^{PES}(a_j, \pi'_{-j}) &= \min_{a_{-j} \in A_{-j}} \max\left(n(\pi'_i(a_i)), \right. \\ &\quad \left. \max_{k \in N \setminus \{j, i\}} n(\pi_k(a_k)), \mu_k(a)\right), \\ &\geq \min_{a_{-j} \in A_{-j}} \max\left(n(\pi_i(a_i)), \right. \\ &\quad \left. \max_{k \in N \setminus \{j, i\}} n(\pi_k(a_k)), \mu_k(a)\right), \\ &\geq \mu_j^{PES}(a_j, \pi_{-j}). \end{aligned}$$

In addition, since $\mu_i^{PES}(a_i, \pi'_{-i})$ does not need π'_i to be computed, we have $\mu_i^{PES}(a_i, \pi'_{-i}) = \mu_i^{PES}(a_i, \pi_{-i})$. Now, since, $\forall j \in N$,

$$\mu_j^{PES}(\pi') = \min_{a_j \in A_j} \max(n(\pi'_j(a_j)), \mu_j^{PES}(a_j, \pi'_{-j})),$$

and we have both $\pi' \preceq \pi$ (implying $n(\pi'_j(a_j)) \geq n(\pi_j(a_j)), \forall j, a_j$) and $\mu_j^{PES}(a_j, \pi'_{-j}) \geq \mu_j^{PES}(a_j, \pi_{-j})$, then: $\mu_j^{PES}(\pi') \geq \mu_j^{PES}(\pi), \forall j \in N$. \square

Properties 1, 2 and 3 in Proposition 2 show that $\text{IMPROVE}(G, \pi, i)$ is at least as specific as π and strictly improves the utility of player i or leaves π unchanged. Property

5 shows that $\text{IMPROVE}(G, \pi, i)$ does not decrease the utilities of other players than i . Property 5 will be especially useful for proving Proposition 3. Finally, property 4 shows that $\text{IMPROVE}(G, \pi, i)$ is the least specific improvement of π to player i , when it changes π . Note that the complexity of IMPROVE is dominated by that of the computation of the set of tables $\{\mu_i^{PES}(a_i, \pi_{-i})\}_{i \in N, a_i \in A_i}$. These can be computed in time polynomial in $|G|$. Therefore, $\text{IMPROVE}(G, \pi, i)$ can itself be performed in polynomial time.

3.2 An Algorithm for the Construction of Possibilistic Nash Equilibria

We exploit procedure IMPROVE in an algorithm which computes a possibilistic mixed equilibrium of an ordinal game.

Data: $G = \langle N, (A_i)_{i \in N}, (\mu_i)_{i \in N} \rangle$
Result: $\pi^* = (\pi_1^*, \dots, \pi_n^*)$, a IIME
 $\pi^0 \leftarrow (\pi_1^0, \dots, \pi_n^0)$; $/* \pi_i^0(a_i) = 1_L, \forall i, a_i * /$
 $t \leftarrow 0$;
repeat
 $\pi_{loc} \leftarrow \pi^t$;
 for $i = 1 \dots n$ **do** $\pi_{loc} \leftarrow \text{IMPROVE}(G, \pi_{loc}, i)$;
 $\pi^{t+1} \leftarrow \pi_{loc}$;
 $t \leftarrow t + 1$;
until $\pi^t = \pi^{t-1}$;
return $\pi^* = \pi^t$;

Algorithm 1: Algorithm which computes a possibilistic mixed equilibrium.

We have to show that Algorithm 1, which performs iterates of the IMPROVE procedure, converges. In addition, we will show that it converges towards a possibilistic mixed equilibrium of the ordinal game G and that convergence occurs within time polynomial in the expression of G .

Proposition 3 (Convergence)

Algorithm 1 converges in a finite number of steps and convergence occurs in time polynomial in the size of G .

Proof: First, note that the outer loop of the algorithm requires n calls to the function IMPROVE per iteration. Then, note that by Prop. 2, result 3, $\pi' = \text{IMPROVE}(G, \pi, i)$ is either more specific than π , or equal to π and that π' can only differ from π in its component π'_i . Finally, note that the number of possible *strict* improvements of a player i 's mixed strategy $\pi_i : A_i \rightarrow L$ is upper bounded by $|L| \times |A_i|$: For each improvement, one of the $|A_i|$ coordinates of the mixed strategy is decreased, and each coordinate belongs to L .

Therefore, Algorithm 1 converges after at most $n \times |L| \times |A_i|$ calls to the IMPROVE procedure, which itself takes time polynomial in the size of the expression of G . \square

Proposition 4 (Soundness)

If Algorithm 1 has converged towards π^ , then π^* is a possibilistic mixed equilibrium of G .*

Proof: Since π^* has been obtained after convergence of Algorithm 1, it verifies : $\pi^* = \text{IMPROVE}(G, \pi^*, i), \forall i \in N$.

This implies that $\forall i \in N, \forall \pi_i, \mu_i^{PES}(\pi^*) \geq \mu_i^{PES}(\pi_i, \pi_{-i}^*)$, that is, π^* is a Nash equilibrium. \square

Example 3 (Illustration of Algorithm 1) $L = \{0, \dots, 4\}$.
 Let us consider the previous ordinal game: Let us start with

		X	Y
E	X	(2,2)	(1,4)
	Y	(1,3)	(2,0)

$\pi_E^0 = \pi_N^0 = [4; 4]$, that is, $\pi_E^0(X) = \pi_N^0(X) = \pi_E^0(Y) = \pi_N^0(Y) = 4$: *Uncertainty is maximal.* Now,

$$\begin{aligned} \mu_E^{PES}(X, \pi_{-E}^0) &= \min_{a_N \in \{X, Y\}} \max(n(\pi_N^0(a_N)), \mu_E(X, a_N)) \\ &= \min(\mu_E(X, X), \mu_E(X, Y)) = 1. \end{aligned}$$

$$\mu_E^{PES}(Y, \pi_{-E}^0) = \min(\mu_E(Y, X), \mu_E(Y, Y)) = 1.$$

$$\mu_N^{PES}(X, \pi_{-N}^0) = \min(\mu_N(X, X), \mu_N(Y, X)) = 2.$$

$$\mu_N^{PES}(Y, \pi_{-N}^0) = \min(\mu_N(X, Y), \mu_N(Y, Y)) = 0.$$

Furthermore, $\mu_E^{PES}(\pi^0) = 1$, $\mu_N^{PES}(\pi^0) = 0$, $A_N^* = \{Y\}$ and $A_E^* = \{X, Y\}$. Since $A_E \setminus A_E^* = \emptyset$, $\text{IMPROVE}(G, \pi^0, E) = \pi^0$. Since $A_N \setminus A_N^* = \{X\}$, we get $\text{IMPROVE}(G, \pi^0, N) = [\pi_N^0(X); n(\mu_N^{PES}(\pi^0))] = [4; 3]$. Another round of improvement does not give anymore changes. So, $\pi^* = (\pi_E^*, \pi_N^*)$, where $\pi_E^* = [4; 4]$ and $\pi_N^* = [4; 3]$, forms a possibilistic mixed equilibrium of the ordinal game.

4 Concluding Remarks and Future Work

In this paper, we have proposed first steps towards the definition of possibilistic mixed equilibria in ordinal game theory. Our main contribution has been to show that possibilistic mixed Nash equilibria exist for all possibilistic games and can be computed in polynomial time. This contribution offers several perspectives:

Concisely expressed games. Possibility theory can be embedded nicely in logic, to provide enhanced reasoning tools. [De Clercq *et al.*, 2014] have studied the possibility and necessity of the existence of ONE in games expressed in possibilistic logic. Considering mixed possibilistic equilibria in possibilistic boolean games would be natural. Of course, it is likely that possibilistic mixed equilibrium computation becomes NP-hard.

Competitive Markov Decision Processes. These, also called *stochastic games* [Filar and Vrieze, 1997] extend both Markov Decision Processes [Puterman, 1994] and games. Possibilistic Markov Decision Processes have been defined in [Sabbadin *et al.*, 1998; Sabbadin, 2001]. Possibilistic competitive MDP are a natural extension to possibilistic MDP.

Partial observability. Games with incomplete information [Harsanyi, 1967; Myerson, 1991] are a first way to deal with partial observability. A second way is *Partially Observed Stochastic Games (POSG)* [Sorin, 2002], which extend both Partially Observed MDP [Cassandra *et al.*, 1994] and games. Such an extended framework could also be developed in the possibilistic case, by grounding it on the possibilistic POMDP framework [Sabbadin, 2001].

Acknowledgements

This work was supported by the French "Agence Nationale de la Recherche" grant ANR-13-AGRO-0001-04.

References

- [Barany *et al.*, 1992] Imre Barany, Jon Lee, and Martin Shubik. Classification of two-person ordinal bimatrix games. *International Journal of Game Theory*, 21:267–290, 1992.
- [Cassandra *et al.*, 1994] Anthony R. Cassandra, Leslie Pack Kaelbling, and Michael L. Littman. Acting optimally in partially observable stochastic domains. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI'94)*, pages 1023–1028, 1994.
- [Cruz Jr. and Simaan, 2000] Jose B. Cruz Jr. and Marwan A. Simaan. Ordinal games and generalized nash and stackelberg solutions. *Journal of Optimization Theory and Applications*, 107:205–222, 2000.
- [Daskalakis *et al.*, 2006] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 71–78, 2006.
- [De Clercq *et al.*, 2014] Sofie De Clercq, Steven Schockaert, Martin De Cock, and Ann Nowé. Possibilistic boolean games: Strategic reasoning under incomplete information. In *Proceedings of the 14th European conference on Logics in Artificial Intelligence (JELIA 2014)*, pages 196–209, 2014.
- [De Clercq *et al.*, 2015] Sofie De Clercq, Steven Schockaert, Ann Nowé, and Martin De Cock. Multilateral negotiation in boolean games with incomplete information using generalized possibilistic logic. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI 2015)*, pages 2890–2896, 2015.
- [Dubois and Prade, 1988] Didier Dubois and Henri Prade. *Possibility Theory*. Plenum Press, New York, 1988.
- [Dubois and Prade, 1995] Didier Dubois and Henri Prade. Possibility theory as a basis for qualitative decision theory. In *Proceedings of the 14th international joint conference on Artificial intelligence (IJCAI'95)*, pages 1925–1930, 1995.
- [Dubois *et al.*, 2001] Didier Dubois, Henri Prade, and Régis Sabbadin. Decision-theoretic foundations of qualitative possibility theory. *European Journal of Operational Research*, 128:459–478, 2001.
- [Durieu *et al.*, 2008] Jacques Durieu, Hans Haller, Nicolas Quérou, and Philippe Solal. Ordinal games. *International Game Theory Review*, 10(2):177–194, 2008.
- [Filar and Vrieze, 1997] Jerzy Filar and Koos Vrieze. *Competitive Markov Decision Processes*. Springer, 1997.
- [Harsanyi, 1967] John C. Harsanyi. Games with incomplete information played by "Bayesian" players. *Management Science*, 14:159–182, 1967.
- [Monderer and Shapley, 1996] Dov Monderer and Lloyd S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [Myerson, 1991] Roger B. Myerson. *Game theory – Analysis of conflict*. Harvard University Press, 1991.
- [Nash, 1950] John Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- [Nisan *et al.*, 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [Osborne and Rubinstein, 1994] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT, Cambridge, MA, 1994.
- [Ouenniche *et al.*, 2016] Jamal Ouenniche, Aristotelis Boukouras, and Mohammad Rajabi. An ordinal game theory approach to the analysis and selection of partners in public-private partnership projects. *Journal of Optimization Theory and Applications*, 169:314–343, 2016.
- [Puterman, 1994] Martin L. Puterman. *Markov Decision Processes*. John Wiley and Sons, 1994.
- [Sabbadin *et al.*, 1998] Régis Sabbadin, Hélène Fargier, and Jérôme Lang. Towards qualitative approaches to multi-stage decision making. *International Journal of Approximate Reasoning*, 19:441–471, 1998.
- [Sabbadin, 2001] Régis Sabbadin. Possibilistic Markov decision processes. *Engineering Applications of Artificial Intelligence*, 14:287–300, 2001.
- [Sorin, 2002] Sylvain Sorin. *A first course in zero-sum repeated games*. Springer, 2002.
- [Von Neumann and Morgenstern, 1948] John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. 1948.
- [Xu, 2000] Chunhui Xu. Computation of noncooperative equilibria in ordinal games. *European Journal of Operational Research*, 122:115–122, 2000.