Solving possibilistic games with incomplete information

Nahla Ben Amor, Hélène Fargier, Régis Sabbadin, Meriem Trabelsi

International Journal of Approximate Reasoning 143 (2022) 139–158

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

Solving possibilistic games with incomplete information

Nahla Ben Amor [a], Hélène Fargier [b], Régis Sabbadin [c], Meriem Trabelsi [a,b,*]

[a] LARODEC, ISG-Tunis, Université de Tunis, Tunisia
[b] IRIT, Université de Toulouse, France
[c] INRAE-MIAT, Université de Toulouse, France

ARTICLE INFO

Article history:
Received 20 August 2021
Received in revised form 3 December 2021
Accepted 17 January 2022
Available online 25 January 2022

Keywords:
Game theory
Ordinal games
Games with incomplete information
Possibility theory
Nash equilibrium

ABSTRACT

Bayesian games offer a suitable framework for games where the utility degrees are additive in essence. This approach does nevertheless not apply to ordinal games, where the utility degrees do not capture more than a ranking, nor to situations of decision under qualitative uncertainty. The present paper proposes a representation framework for ordinal games under possibilistic incomplete information and extends the fundamental notions of pure and mixed Nash equilibrium to this framework. We show that deciding whether a pure Nash equilibrium exists is a difficult task (NP-hard) and propose a Mixed Integer Linear Programming (MILP) encoding of the problem; as to the problem of computing a possibilistic mixed equilibrium, we show that it can be solved in polynomial time. An experimental study based on the GAMUT game generator confirms the feasibility of the approach.

1. Introduction

Game theory [1,2] proposes a simple but powerful framework to capture decision problems involving several agents: in a game with complete information, each agent (or "player") chooses an action among a set of possible actions, and the final outcome depends on the actions chosen by all the players.

The preferences of the players among the outcomes are captured by utility functions. The term "payoff" is often used to designate their utility levels- this terminology targets problems where the satisfaction can be expressed on a cardinal, additive, scale, typically a monetary scale. But there are situations where the assumption of cardinality can be questioned, hence the development of ordinal games (see, e.g., [3–6]). Such ordinal approaches cohere with many fundamental notions of game theory which are basically *ordinal* notions: the notions of pure *Nash equilibrium*, *secure strategy*, *dominance*, for instance, do not require the utilities to be cardinal payoffs.

A cardinal notion is invoked in two cases at least: (i) when the game is repeated (and outcomes are "collected" and their utilities are additive and can compensate each other), and (ii) when the outcomes depend on a probabilistic event (for example, in the prisoner's dilemma if the verdict does not only depend on the confession of the prisoners, but also on the uncertain result of the trial). To capture such situations, *Bayesian games* have been proposed by Harsanyi [7]. In these games, the players' knowledge about the game is assumed to be probabilistic. This approach does not apply to ordinal games, where the utility degrees capture no more than a ranking, nor to situations of decision under qualitative uncertainty. Following the

* Corresponding author at: LARODEC, ISG-Tunis, Université de Tunis, Tunisia.
E-mail addresses: nahla.benamor@gmx.fr (N. Ben Amor), fargier@irit.fr (H. Fargier), Regis.Sabbadin@toulouse.inra.fr (R. Sabbadin), merieme.trabelsi@gmail.com (M. Trabelsi).

https://doi.org/10.1016/j.ijar.2022.01.009
0888-613X/

seminal work of [8] on possibilistic Boolean games, we propose to use possibility theory [9] to model qualitative uncertainty in ordinal games since it offers a natural and flexible model to represent and handle uncertainty information, especially qualitative uncertainty, and total ignorance. Unlike [8], we do not develop a complex representation language. We stay at the semantic level and work out the idea of possibilistic games, extending the notion of pure Nash equilibrium, (possibilistic) mixed Nash equilibrium to such games.

The paper[1] is organized as follows: Section 2 presents the basic notions on which the paper relies, namely possibility theory on the one hand, and ordinal game theory on the other hand. Section 3 proposes a possibilistic model for ordinal games with incomplete information, including the definition of pure and mixed equilibria. Section 4 is devoted to the search of such strategies - we show that deciding the existence of a pure equilibrium is a NP-complete problem and propose a Mixed Integer Linear Programming (MILP) formulation of this problem; as to mixed equilibrium strategies, we show that their computation is tractable and provide a polytime algorithm. Finally, experiments are reported in Section 5.

2. Background

2.1. Possibility theory

The basic building block in possibility theory [9] is the notion of *possibility distribution*. A possibility distribution π is a mapping from a set of states S to an ordered scale Δ (in the remaining, we consider that Δ is a subset of $[0, 1]$, but any ordered scale with lowest and greatest elements and order-reversing function may be used). $\pi(s) = 1$ means that state s is totally plausible, $\pi(s) = 0$ means that s is impossible and $\pi(s) > \pi(s')$ means that s is more plausible than s'. π is assumed to be normalized, i.e. there is at least one totally possible state ($\exists s, \pi(s) = 1$).

From π, one can compute the possibility $\Pi(E)$ and the necessity $N(E)$ of any event $E \subseteq S$: $\Pi(E) = \max_{s \in E} \pi(s)$ evaluates to what extent E is *consistent* with the knowledge represented by π while $N(E) = 1 - \Pi(\bar{E}) = 1 - \max_{s \notin E} \pi(s)$ corresponds to the extent to which $\bar{E}$ is inconsistent and thus evaluates at which level E is implied by the knowledge.

Considering qualitative (possibilistic) problems of decision under uncertainty, where each decision is evaluated by a utility function $\mu : S \mapsto \Delta$, [11,12] have proposed two dual utilities measures:

$$U^{pes}(\pi) = \min_{s \in S} \max(1 - \pi(s), \mu(s)). \tag{1}$$

$$U^{opt}(\pi) = \max_{s \in S} \min(\pi(s), \mu(s)). \tag{2}$$

The so-called pessimistic qualitative utility, U^{pes}, generalizes the Wald criterion and estimates to what extent it is certain (i.e., necessary according to measure N) that μ reaches a good utility while the optimistic qualitative utility, U^{opt}, estimates to what extent it is possible that μ reaches a good utility. U^{opt} is rather unnatural (too adventurous), while U^{pes} conveniently models the behavior of *uncertainty adverse* decision-makers (see [11,12] for more details). As many other models, this model makes a commensurability assumption between the utility levels and the levels of likelihood. This assumption is common to all the models which consider that the agent's preference relation is complete and transitive (this is the case in many models, be they qualitative or quantitative: expected utility [13,1], multi-prior non expected utility [14], Sugeno integrals [15], etc.).

In the following, we shall also use the notion of conditional possibility measure proposed by [16,17] in order to stay in a purely ordinal framework.[2] Formally, for any events E and $F \subseteq S$ the possibility of E given F is defined by:

$$\Pi(E|F) = \begin{cases} 1 & \text{if } \Pi(E \cap F) = \Pi(F) \\ \Pi(E \cap F) & \text{otherwise.} \end{cases} \tag{3}$$

This means that the event E is totally possible, knowing that F occurred, if at least one of the most plausible states making F true also makes E true. Otherwise, $\Pi(E|F)$ is the possibility degree of the most plausible state making both E and F true.

2.2. Normal form games, ordinal games and possibilistic mixed equilibria

A *standard normal form game* (SNF) is classically defined as follows [1]:

Definition 1 *(Standard normal form game).* A *standard normal form game* is a triple $G = \langle N, A, \mu \rangle$, where:

- $N = \{1, \ldots, n\}$ is a finite set of players;

[1] This paper is an extended version of [10]. It presents additional notions e.g. possibilistic mixed equilibrium and provide equilibria computation algorithms. It also includes an experimental study.

[2] For a discussion about the alternative definitions of conditional possibilities, see [18].

Table 1
The landlord-tenant SNF game.

		Landlord C	Landlord $\overline{C}$
Tenant	P	1 , 1	0 , 0.25
	$\overline{P}$	0.25 , 0	0 , 0.75

- $A = \times_{i \in N} A_i$, where A_i is a finite set of actions available to player $i \in N$;
- $\mu = \{(\mu_i)_{i \in N}\}$ is a set of utility functions μ_i on A.

In cardinal games, the utility function of a player associates a payoff, that takes its values in a subset of $\mathbb{R}$, to every joint action $a \in A$ - this suits problems where the satisfaction can be expressed on an additive scale, typically a monetary scale. But, as mentioned above, there are instances where the assumption of additivity can be questioned, hence the development of ordinal games (see, e.g., [3–6]). In such games, the utility functions range in a totally ordered scale Δ (the higher $\mu_i(a)$, more satisfied player i) - no other assumption is made than Δ being totally ordered.

Example *(The landlord and tenant game).* A landlord L can lend a house to a tenant T. She can Confirm (C) or not ($\overline{C}$), i.e., $A_L = \{C, \overline{C}\}$ and the tenant can Pay the landlord (P) or not ($\overline{P}$), i.e., $A_T = \{P, \overline{P}\}$. In this example we consider the utility scale $\Delta = \{0, 0.25, 0.5, 0.75, 1\}$ (such that 0 means: "completely dissatisfied", 1: "completely satisfied", 0.5: "neutral", 0.25: "dissatisfied" and 0.75: "satisfied".). For instance, the tenant is completely satisfied, i.e., she gets a utility equal to 1, if she pays the rent and the landlord confirms. The utilities of both players are given in Table 1.

A joint (pure) strategy is a vector $a = (a_1, \ldots, a_i, \ldots, a_n) \in A$. In the following for any action $a \in A$, we will denote $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n) \in A_{-i} = \times_{j \neq i} A_j$ its restriction to all the players but i and denote "." the concatenation (e.g., $\forall (a'_i, a_{-i})$, $a'_i.a_{-i} = (a_1, \ldots, a_{i-1}, a'_i, a_{i+1}, \ldots, a_n)$).

When a player i knows the strategy a_{-i} of the other players, the principle of rationality leads her to choose an action that maximizes her utility. Such actions are called the "*Best responses*" of i to a_{-i}. Formally:

Definition 2 *(Best response).* For any game $G = \langle N, A, \mu \rangle$, an action $a_i \in A_i$ is a best response of player i to a_{-i}, iff $\forall a'_i \in A_i$:

$$\mu_i(a_i.a_{-i}) \geq \mu_i(a'_i.a_{-i}). \tag{4}$$

Note that the best response of player i to a_{-i} is not necessarily unique.

Generally each player responds in the best way when knowing the other players' actions. A joint strategy where each player responds best to the actions of the other players is a *pure Nash equilibrium* [19]. In other terms, a joint strategy a^* is a pure Nash equilibrium (PNE) iff no player can improve her utility by unilaterally modifying her strategy:

Definition 3 *(Pure Nash equilibrium).* For any game $G = \langle N, A, \mu \rangle$, the strategy profile $a^* \in A$ is a pure Nash equilibrium (PNE), iff $\forall i \in N$, $\forall a'_i \in A_i$:

$$\mu_i(a^*_i.a^*_{-i}) \geq \mu_i(a'_i.a^*_{-i}). \tag{5}$$

It is easily checked that the above definition is equivalent to writing that, in a pure Nash equilibrium, all players choose their best response to other players' best responses.

Example. In the landlord-tenant game (Table 1), the best response of the landlord when the tenant pays the rent (plays P) is C, indeed: $\mu_L(P.C) = 1 \geq \mu_L(P.\overline{C}) = 0.5$. On the other hand, the best response of the tenant to the landlord who confirms is P, since $\mu_T(P.C) = 1 > \mu_T(\overline{P}.C) = 0.25$

So, in situation (P, C) the tenant has no incentive to deviate from P and the landlord has no interest to move from C to $\overline{C}$: (P, C) is a pure Nash equilibrium.

Notice that the notion of pure Nash equilibrium is purely ordinal - nothing more sophisticated than an ordinal ranking of the outcomes is needed. The notion of (possibilistic) mixed strategy proposed in [20,21,6] remains in the ordinal framework. A possibilistic mixed strategy for a player i is indeed a *normalized distribution* $\upsilon_i : A_i \mapsto \Delta$, that is a ranking on her set of actions. This ranking has a dual interpretation in terms of preference on the one hand and likelihood on the other hand. Indeed, for player i distribution υ_i models the ranking in terms of preference or commitment. Under this interpretation, $\upsilon_i(a_i) = 1$ means that action a_i is fully satisfactory/conceivable to player i, while $\upsilon_i(a_i) = 0$ means that it is absolutely not an option for her.

However, the other players may interpret the preference ranking υ_i as a likelihood. Assuming that player i is rational, $\upsilon_i(a_i) = 1$ is interpreted, by other players, as "action a_i is a completely plausible play of player i" while $\upsilon_i(a_i) = 0$ means

that a_i is an impossible play of player i. This dual preference/likelihood interpretation is natural in game theory: according to the other players, the most preferred alternatives of player i should be the most likely to be played by player i.

The joint non-correlated strategy $(\upsilon_1, \ldots, \upsilon_n)$ defines a possibility distribution over the action profiles. Since the joint action $a = (a_1, \ldots, a_n)$ is played iff each player i plays action a_i, the possibility that a is played is computed in a conjunctive way, i.e., as the *minimum* of the $\upsilon_i(a_i)$:

$$\forall a = (a_1, \ldots, a_n) \in A, \upsilon(a) = \min_{i \in N} \upsilon_i(a_i). \tag{6}$$

By abuse of notation υ designates both the above joint possibility distribution and the list $(\upsilon_1, \ldots, \upsilon_n)$ of individual players' possibility distributions. Notice that υ is normalized since all the υ_i are assumed to be normalized: there exists $a^* = (a_1^*, \ldots, a_n^*)$ such that $\upsilon(a^*) = 1$.

Following [6] the players evaluate the merit of a strategy by its pessimistic utility:

$$U_i^{pes}(\upsilon) = \min_{a \in A} \max(1 - \upsilon(a), \mu_i(a)). \tag{7}$$

A possibilistic mixed strategy is a mixed equilibrium if no player has an incentive to deviate from her υ_i - hence the following definition:

Definition 4 *(Possibilistic mixed equilibrium).* $\upsilon^* = (\upsilon_1^*, \ldots, \upsilon_n^*)$ is a possibilistic mixed equilibrium (ΠME) in $G = \langle N, A, \mu \rangle$ iff, $\forall i \in N$, $\forall \upsilon_i'$ on A_i:

$$U_i^{pes}(\upsilon_i^*.\upsilon_{-i}^*) \geq U_i^{pes}(\upsilon_i'.\upsilon_{-i}^*). \tag{8}$$

[6] have shown that an ordinal game always admits at least one mixed equilibrium - and generally admits a full range of more or less specific mixed equilibria. They put the emphasis on the least specific of them, claiming that, as soon as a player plays one of her best response (and there may be several sharing the same level of pessimistic utility), the responses which lead to the least commitment are preferable for her. Formally:

Definition 5 *(Specificity relation).* Given two possibility distributions υ and υ' over some set A, υ is more specific than υ', denoted by $\upsilon \prec \upsilon'$ iff (i) $\forall a \in A, \upsilon(a) \leq \upsilon'(a)$ and (ii) $\exists a \in A, \upsilon(a) < \upsilon'(a)$.

Definition 6 *(Least-specific possibilistic mixed equilibrium).* $\upsilon^* = (\upsilon_1^*, \ldots, \upsilon_n^*)$ is a *least-specific* possibilistic mixed equilibrium for $G = \langle N, A, \mu \rangle$ iff υ^* is a ΠME for G and there exists no ΠME υ, s.t. $\upsilon^* \prec \upsilon$.

Example. Let $\upsilon = (\upsilon_T.\upsilon_L)$ be the possibilistic mixed strategy defined by:

$$\upsilon_T(P) = 1, \quad \upsilon_T(\overline{P}) = 1, \quad \upsilon_L(C) = 0.75, \quad \upsilon_L(\overline{C}) = 1$$

The utilities of both players for υ are $\mu_T^{pes}(\upsilon) = 0$ and $\mu_L^{pes}(\upsilon) = 0.25$.

υ is a ΠME since:

- $\forall \upsilon' \neq \upsilon$: $U_T^{pes}(\upsilon_T'.\upsilon_L) \leq U_T^{pes}(\upsilon_T.\upsilon_L) = 0$;
- $\forall \upsilon' \neq \upsilon$: $U_L^{pes}(\upsilon_L'.\upsilon_T) \leq U_L^{pes}(\upsilon_L.\upsilon_T) = 0.25$.

It can be shown that υ is a least-specific ΠME. Indeed, one can check that $U_T^{pes}(\upsilon) \leq \max(1 - \upsilon_T(\bar{P}), 1 - \upsilon_L(C), \mu_T(\bar{P}, C))$. So, since $1 - \upsilon_P(\bar{P}) = 0 < \mu_T(\bar{P}, C) = 0.25 \leq 1 - \upsilon_L(C) = 0.25$, $U_T^{pes}(\upsilon) \leq 1 - \upsilon_L(C)$. Then, if we increase $\upsilon_L(C)$ so as to get a less specific mixed strategy (for player L), we decrease the utility of L. On the other hand, the mixed strategy of player T is already the least-specific possible. υ is thus a least-specific ΠME.

3. Possibilistic games with incomplete information

The classical framework of games under complete information assumes that every player knows everything about the game: the players, the actions available to all players, all their utilities, etc. What is not known by a player are the decisions of the other ones, but the game itself is assumed to be perfectly known by everyone (hence the name "perfect information").

This assumption cannot always be satisfied. In the real world indeed, players are not so well informed, and the knowledge about the game is often incomplete. For instance, if we consider the classical *battle of sexes* example, a person who has doubts about the faithfulness of his/her partner has some uncertainty about his/her priorities (being together or missing the date).

Bayesian games [7] offer a suitable framework for cardinal games under incomplete knowledge. In these games, the players' knowledge about the game is assumed to be probabilistic. This probabilistic knowledge is common to all the players,

but each player may receive some private information, her personal view of the game: Bayes' rule of conditioning is used to derive the knowledge of each player, hence the denomination "Bayesian" game.

Bayesian games assume that the utility degrees are additive in essence and that the knowledge of the players can be quantified in a probabilistic way. This kind of approach does not apply to ordinal games, where the utility degrees do not capture more than a ranking, nor to situations of decision under qualitative uncertainty. We propose in the following a model for (ordinal) games under possibilistic information.

3.1. Possibilistic games with incomplete information: the framework

Following Harsanyi's seminal work [7], a game with incomplete information can be first understood as a set S of states of nature, each state corresponding to a classical game. The utility μ_i of a player i does not depend only on the actions of the other players but also on the game played, i.e., the actual state. Originally, none of the players knows which is the real state, but all of them share a common knowledge about it. Just before playing, each player i will receive some information $\tau_i(s)$ about the real state, i.e., τ_i maps any $s \in S$ to an element $\theta_i \in \Theta_i$ called the set of "types" of player i. After having observed $\tau_i(s)$, player i knows more about the game played, but several games may still remain plausible. The player thus conditions her knowledge on $\tau_i(s)$.

The set of possible types of player i, Θ_i, can be considered as a local state space and $\Theta = \Theta_1 \times \cdots \times \Theta_n$ as the global state space.[3] The idea of Harsanyi when defining types was that a player's local state can encapsulate all the information to which the player has access: it contains not only the status of the external world that the player has observed but also her introspective mental state.[4] Harsanyi shows that any game under incomplete information can be described on the space of types (underlying worlds are omitted).

Now, in such a game, the action of a player i only depends on the information $\theta_i \in \Theta_i$ she receives. A joint pure strategy $\sigma = (\sigma_1, \ldots, \sigma_n)$ is thus a vector of functions σ_i that map possible information (each type $\theta_i \in \Theta_i$) to an action in A_i: $\sigma_i(\theta_i)$ specifies the action that player i will execute when receiving the private information θ_i.

Definition 7 *(Pure strategy).* A pure strategy is a vector $\sigma = (\sigma_1, \ldots, \sigma_n)$ of functions $\sigma_i : \Theta_i \to A_i$.

Given a pure strategy σ and a configuration of the players' types $\theta \in \Theta$, $\sigma(\theta) = (\sigma_1(\theta_1), \ldots, \sigma_n(\theta_n))$ denotes the joint action (the element of A) prescribed by strategy σ when θ occurs. In the following, Σ_i denotes the set of all functions from Θ_i to A_i and $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ the set of all joint strategies.

We follow Harsanyi's approach based on types and define a possibilistic game with incomplete information as follows:

Definition 8 *(Π-game).* A possibilistic game with incomplete information (Π-game) G is a tuple $\langle N, A, \Theta, \pi, \mu \rangle$ where:

- N is a finite set of n players $\{1, ..., n\}$;
- $A = \times_{i \in N} A_i$ where A_i is the set of actions of player i;
- $\Theta = \times_{i \in N} \Theta_i$, where Θ_i is the set of types of player i;
- π: $\Theta \to \Delta$ is a joint normalized possibility distribution over Θ;
- $\mu = \{(\mu_i)_{i \in N}\}$ where $\mu_i : A \times \Theta \to \Delta$ is the utility function of player i.

Possibility distribution π captures the common knowledge of the players. The information that the players have about the real world corresponds to a $\theta \in \Theta$ but is not common: player i does not know θ, but only θ_i (θ_i is the private knowledge of player i). $\pi(.|\theta_i)$ captures the knowledge that player i has when learning θ_i. On the other hand, utility $\mu_i(a, \theta)$ (utility of the joint action a for player i when learning θ) will be obtained once all players have learned their types and played their actions - that is why μ_i depends on the whole θ and not only on θ_i.

A Π-game G can be equivalently defined as a set of $|\Theta|$ normal form games with the same set of players N and the same set of actions A. More precisely, for each $\theta \in \Theta$, there is a normal form game $G^\theta = \langle N, A, \{\mu_i^\theta\}_{i \in N} \rangle$ where $\forall i \in N$:

$$\mu_i^\theta(.) = \mu_i(., \theta) \text{ and } \pi(G^\theta) = \max_{\theta' \in \Theta, \text{ s.t., } G^{\theta'} = G^\theta} \pi(\theta'). \tag{9}$$

Example *(The landlord and tenant Π-game).* We consider a Π-game variant version of the landlord and tenant game introduced in Section 2. In this extension of the complete information game, the tenant can be honest (H) or dishonest (D), i.e., $\Theta_T = \{H, D\}$. The landlord does not know the type of the tenant but makes investigations and can conclude the tenant is honest (IH) or dishonest (ID), i.e., $\Theta_L = \{IH, ID\}$. Generally, the tenant is honest and the landlord is informed about this.

[3] We use the same notations for the projection and concatenation of vectors as the ones used in the previous sections: for any $\theta = (\theta_1, \ldots, \theta_i, \ldots, \theta_n) \in \Theta$, θ_i is the type of i in θ and $\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n)$; $\forall\, \theta, \theta' \in \Theta$, $\theta_i'.\theta_{-i}$ belongs to Θ). Likewise $\Theta_{-i} = \times_{j \neq i} \Theta_j$.

[4] See [22–26] for the links between belief states and types, and more generally for further developments about epistemic game theory. This kind of interpretation also complies with the run-based semantics of epistemic logic [27].

Table 2
A landlord-tenant Π-game with 4 types combinations.

		landlord			
		IH		ID	
tenant		C	$\overline{C}$	C	$\overline{C}$
H	P	1 , 1	0 , 0.5	1 , 1	0 , 0.5
	$\overline{P}$	0.25 , 0	0 , 0.5	0.25 , 0.25	0 , 0.5
		$\pi(H, IH) = 1$		$\pi(H, ID) = 0.25$	
D	P	0.25 , 1	0.25 , 0.5	0.25 , 1	0.25 , 0.5
	$\overline{P}$	1 , 0	0.25 , 0.5	1 , 0.25	0.25 , 0.5
		$\pi(D, IH) = 0.5$		$\pi(D, ID) = 0.75$	

But the tenant may be dishonest and the landlord well informed. There is a low possibility that the landlord is informed that the tenant is honest while it is not the case. Finally, it cannot be excluded that the landlord is informed that the tenant is dishonest while she is honest. This knowledge is captured by the following possibility distribution π:

$$\pi(H, IH) = 1, \quad \pi(D, ID) = 0.75, \quad \pi(D, IH) = 0.5, \quad \pi(H, ID) = 0.25.$$

The projections of π on Θ_T and Θ_L lead to:

$$\pi(H) = 1, \quad \pi(D) = 0.75, \quad \pi(IH) = 1, \quad \pi(ID) = 0.75.$$

As to the utility degrees, scale Δ is used with the same meaning as previously. When the landlord does not confirm to rent, the utility of the tenant is equal to 0 if honest and equal to 0.25 if dishonest. When the landlord confirms, the tenant's utility of $\overline{P}$ (resp. P) is equal to 1 (resp. 0.25) when she is dishonest and equal to 0.25 (resp. 1) when she is honest. The utility of the landlord is neutral (0.5) when not confirming. When confirming, the landlord's utility is high (1) when the tenant pays, low when she does not pay (0 if the landlord believes that the tenant is honest, 0.25 otherwise).

There are four combinations of types, and thus four possible games (see Table 2) the possibility degrees of which are:

$$\pi(G^{H,IH}) = \pi(H, IH) = 1, \qquad \pi(G^{H,ID}) = \pi(H, ID) = 0.25,$$
$$\pi(G^{D,IH}) = \pi(D, IH) = 0.5, \qquad \pi(G^{D,ID}) = \pi(D, ID) = 0.75.$$

Clearly, Π-games properly generalize classical games (with complete information). Indeed:

Proposition 1. *Any classical normal form game $G = \langle N, A, \mu \rangle$ is a Π-game with $|\Theta_i| = 1, \forall i \in N$.*

Π-games and Boolean games with incomplete information Π-games can be related to the framework proposed by [8] as a semantics for possibilistic Boolean games where the prior knowledge of each player i can be captured by a distinct possibility distribution π_i (the knowledge is not common in [8]) and the players do not receive any private information (there is no type). In [8] the knowledge of each player i is captured by a possibility distribution π_i. Then the authors assume, mainly for simplicity and because of the ordinal nature of the π_i, that the fusion of the π_i's leads to a unique distribution. $\pi = \min_{i \in N} \pi_i$ over the joint types of players. The possibility and the necessity of a given profile of actions being a PNE is then computed. The approach of [8] considers the problem from the external point of view of an observer who proceeds to the fusion of these distributions and deduces a unique possibility distribution over the types of the players. The likelihood of an equilibrium is computed *ex-ante*, on the basis of this prior possibility distribution.

Adapting these notions to Π-games, where the prior knowledge is common, we can compute, for any profile a of actions:

$$\Pi(a \text{ is a PNE}) = \max_{\theta, a \text{ is PNE for } G^\theta} \pi(\theta). \tag{10}$$
$$N(a \text{ is a PNE}) = 1 - \max_{\theta, a \text{ is not a PNE for } G^\theta} \pi(\theta). \tag{11}$$

The concepts proposed in [8] are based on the definition of strategies as profiles of actions - each agent selects one and only one action. This assumes that no more information about θ can be acquired and thus that each player computes her *ex-ante* utility.

The notion of type present in our framework allows to handle the knowledge that each player i has *when receiving some private information* θ_i, i.e., adopt an *ex-interim* point of view.[5] Then, the possibility (resp. necessity) that a is a NE can be different between players - in our *ex-interim* framework, equations (10) and (11) are replaced by:

[5] For a discussion of the merits of the *ex-ante*, *ex-interim* and *ex-post* approaches of incomplete information games, see [7,28].

$$\Pi_i(a \text{ is a PNE}|\theta_i) = \min_{\theta_{-i}, a \text{ is a PNE for } G^{\theta_i.\theta_{-i}}} \pi(\theta_{-i}|\theta_i).$$
$$N_i(a \text{ is a PNE}|\theta_i) = 1 - \max_{\theta_{-i}, a \text{ is not a PNE for } G^{\theta_i.\theta_{-i}}} \pi(\theta_{-i}|\theta_i).$$

3.2. Nash equilibria in possibilistic games

When making her decision on the basis of the information she receives (θ_i), the knowledge of player i becomes $\pi(.|\theta_i)$ - player i then defines her strategy for θ_i, according to this posterior knowledge. Let us first consider pure strategies - the case of mixed possibilistic strategies will be developed later.

Pure Nash equilibrium in Π*-games* As in any incomplete information game, a pure strategy σ specifies an action for each player i for each θ_i. Using possibilistic qualitative decision theory, the utility of $\sigma(\theta_i)$ is evaluated by its possibilistic pessimistic utility (we assume that the player is cautious).

Definition 9 *(Utility of an action, Utility of a pure strategy).* The utility of player $i \in N$ of type $\theta_i \in \Theta_i$ for an action $a_i \in A_i$ in the context of σ_{-i} is:

$$U_i^{pes}(a_i, \sigma_{-i}, \theta_i) = \min_{\theta_{-i} \in \Theta_{-i}} \max\Big(1 - \pi(\theta_{-i}|\theta_i), \mu_i\big(a_i.\sigma_{-i}(\theta_{-i}), \theta_i.\theta_{-i}\big)\Big). \tag{12}$$

The utility of the pure strategy $\sigma \in \Sigma$ to player $i \in N$ of type $\theta_i \in \Theta_i$ is:

$$U_i^{pes}(\sigma, \theta_i) = U_i^{pes}\big(\sigma_i(\theta_i), \sigma_{-i}, \theta_i\big). \tag{13}$$

Note that $U_i^{pes}(\sigma_i(\theta_i), \sigma_{-i}, \theta_i)$ is independent of the choices of player i when her type is different from θ_i.

A *best response* for player i of type θ_i is computed knowing the provisional strategies of the other players, i.e., knowing σ_{-i}. This is the action a_i which maximizes $U_i^{pes}(a_i, \sigma_{-i}, \theta_i)$.

Definition 10 *(Best response in a* Π*-game).* The best response of player $i \in N$ of type $\theta_i \in \Theta_i$ to σ_{-i} in $G = \langle N, A, \Theta, \pi, \mu\rangle$ is:

$$BR_i(\sigma_{-i}, \theta_i) = \operatorname*{argmax}_{a_i \in A_i} U_i^{pes}(a_i, \sigma_{-i}, \theta_i). \tag{14}$$

σ is a PNE when no player i has interest to deviate from $\sigma_i(\theta_i)$ for some θ_i, given $\pi(.|\theta_i)$ and σ_{-i}, i.e., every player plays her best response to the other players' best strategies.

Definition 11 *(Pure Nash Equilibrium in a* Π*-game).* σ is a pure Nash equilibrium (PNE) in $G = \langle N, A, \Theta, \pi, \mu\rangle$ iff: $\forall i \in N, \forall \theta_i \in \Theta_i, \forall a_i' \in A_i$:

$$U_i^{pes}(\sigma_i(\theta_i), \sigma_{-i}, \theta_i) \geq U_i^{pes}(a_i', \sigma_{-i}, \theta_i). \tag{15}$$

This definition is consistent with the case of complete information (Definition 3), which is recovered when $|\Theta| = 1$ (only one possible type per player). As a consequence, a pure Nash equilibrium may not exist for a Π-game (a complete information game does not always admit a PNE).

Example. Let us first compute the conditional possibility distributions:

$$\pi(H|IH) = \pi(IH|H) = 1, \qquad \pi(D|IH) = \pi(IH|D) = 0.5,$$
$$\pi(H|ID) = \pi(ID|H) = 0.25, \qquad \pi(D|ID) = \pi(ID|D) = 1.$$

For instance, when the landlord learns that the tenant is honest, i.e., when the landlord is of type IH, the most plausible type for the tenant is H.

Consider the joint strategy where the tenant pays if she is honest, and does not pay otherwise, and the landlord confirms iff learning that the tenant is honest:

$$\sigma_T^*(H) = P, \quad \sigma_T^*(D) = \overline{P}, \quad \sigma_L^*(IH) = C, \quad \sigma_L^*(ID) = \overline{C}.$$

We have, $U_T^{pes}(P, \sigma_L^*, H) = 0.75$ and $U_T^{pes}(\overline{P}, \sigma_L^*, H) = 0.25$. Hence, the best response to σ_L^* for the honest tenant is P.

It can be checked that strategy σ^* is a PNE. Indeed, as we have seen, the best response of the honest tenant to strategy σ_L^* is to pay. Moreover, the dishonest tenant is indifferent between P and $\overline{P}$ ($U_T^{pes}(P, \sigma_L^*, D) = U_T^{pes}(\overline{P}, \sigma_L^*, D) = 0.25$); the informed honest ($IH$) landlord is indifferent between C and $\overline{C}$ ($U_L^{pes}(C, \sigma_T^*, IH) = U_L^{pes}(\overline{C}, \sigma_T^*, IH)) = 0.5$); and finally, the informed dishonest (ID) landlord prefers $\overline{C}$ to C (since $U_L^{pes}(C, \sigma_T^*, ID) = 0.25$ is lower than $U_L^{pes}(\overline{C}, \sigma_T^*, IH)) = 0.5$). So, whatever her type, no player has any incentive to change her strategy: σ^* is a Nash equilibrium.

Mixed Nash equilibrium in Π-games Let us now extend the notion of possibilistic mixed joint strategy to Π-games: in such a game, a mixed strategy will be a tuple $\upsilon = (\upsilon_1, ..., \upsilon_n)$ where $\upsilon_i : \Theta_i \mapsto \Psi_i$ maps each type of player i to a mixed possibilistic strategy in Ψ_i (Ψ_i is the set of normalized possibility distributions over actions in A_i). Formally:

Definition 12 *(Possibilistic mixed strategy in a Π-game).* A possibilistic mixed joint strategy is a vector $\upsilon = (\upsilon_1, \ldots, \upsilon_n)$ of functions $\upsilon_i : \Theta_i \to \Psi_i$ where $\Psi_i = \{\pi : A_i \to \Delta\}$ is the set of the normalized possibility distributions over A_i.
For the sake of readability, let $\upsilon_i(a_i|\theta_i) = \upsilon_i(\theta_i)(a_i)$ denote the possibility that player i plays action a_i when her type is θ_i.

Assuming that player i is rational and of type θ_i, the common knowledge and rationality assumptions mean that $\upsilon_i(a_i|\theta_i) = 1$ may be interpreted by other players as a_i being a completely plausible play of player i, while $\upsilon_i(a_i|\theta_i) = 0$ means that a_i is an impossible play. Therefore, the ranking over strategies induced by $\upsilon_i(\cdot|\theta_i)$ may be interpreted as a plausibility ranking by other players.

Example. We shall for instance consider the joint mixed strategy $(\upsilon_T.\upsilon_L)$ defined by:

$$\upsilon_T(P|H) = 1, \quad \upsilon_T(\overline{P}|H) = 0.5, \quad \upsilon_T(P|D) = 0.25, \quad \upsilon_T(\overline{P}|D) = 1,$$
$$\upsilon_L(C|IH) = 1, \quad \upsilon_L(\overline{C}|IH) = 0.75, \quad \upsilon_L(C|ID) = 0.25, \quad \upsilon_L(\overline{C}|ID) = 1.$$

In short, the honest tenant prefers to pay ($\upsilon_T(.|H)$), the dishonest tenant prefers not to pay ($\upsilon_T(.|D)$), the informed honest landlord generally confirms ($\upsilon_L(.|IH)$) and the informed dishonest landlord prefers not no ($\upsilon_L(.|ID)$)

According to [6], when the configuration of types is θ, the joint possibility distribution over the profiles of actions is defined as the minimum of the individual players' possibility distributions over individual actions. Formally:

$$\upsilon(a|\theta) = \min_{i \in N} \upsilon_i(a_i|\theta_i). \tag{16}$$

Let $\pi_\upsilon(a, \theta_{-i}|\theta_i)$ be the possibility distribution on $A \times \Theta_{-i}$ defined by υ, given the type θ_i and prior knowledge $\pi : \Theta \to \Delta$. $\pi_\upsilon(a, \theta_{-i}|\theta_i)$ is defined as the minimum over $\pi(\theta_{-i}|\theta_i)$ and the joint possibility distribution over a[6]:

$$\pi_\upsilon(a, \theta_{-i}|\theta_i) = \min(\pi(\theta_{-i}|\theta_i), \upsilon(a|\theta_i.\theta_{-i})). \tag{17}$$

Let us now study the evaluation of mixed strategies. In order to stay in accordance with the previous assumptions of cautiousness and ordinality, strategies will be evaluated using the pessimistic possibilistic utility.

Definition 13 *(Utility of a mixed strategy).* The *pessimistic* utility for player $i \in N$ of type $\theta_i \in \Theta_i$ of the mixed strategy υ in $G = \langle N, A, \Theta, \pi, \mu \rangle$ is:

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i} \in \Theta_{-i}} \min_{a \in A} \max\big(1 - \pi_\upsilon(a, \theta_{-i}|\theta_i), \mu_i(a, \theta)\big). \tag{18}$$

It can be checked that:

Proposition 2.

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i} \in \Theta_{-i}} \max\Big(1 - \pi(\theta_{-i}|\theta_i), \min_{a \in A} \max\big(1 - \upsilon(a|\theta_i.\theta_{-i}), \mu_i(a, \theta)\big)\Big). \tag{19}$$

Thus, $U_i^{pes}(\upsilon, \theta_i)$ can be written as:

Proposition 3. *The pessimistic utility for player $i \in N$ of type $\theta_i \in \Theta_i$ of the mixed strategy υ in $G = \langle N, A, \Theta, \pi, \mu \rangle$ is:*

$$U_i^{pes}(\upsilon, \theta_i) = \min_{a_i \in A_i} \max\big(1 - \upsilon_i(a_i|\theta_i), U_i^{pes}(a_i, \upsilon_{-i}, \theta_i)\big) \tag{20}$$

where: $U_i^{pes}(a_i, \upsilon_{-i}, \theta_i) =$

$$\min_{a_{-i} \in A_{-i} \theta_{-i} \in \Theta_{-i},} \min \max\Big(1 - \pi(\theta_{-i}|\theta_i), 1 - \upsilon_{-i}(a_{-i}|\theta_{-i}), \mu_i(a_i.a_{-i}, \theta_i.\theta_{-i})\Big). \tag{21}$$

[6] By the possibilistic chaining rule.

A possibilistic mixed Nash equilibrium in a Π-game is a mixed strategy $\upsilon^* = (\upsilon_1^*, ..., \upsilon_n^*)$ where no player i of type θ_i can improve her pessimistic utility by changing her mixed strategy υ_i. Formally:

Definition 14 *(Possibilistic mixed Nash equilibrium in a Π-game).* υ^* is a possibilistic mixed Nash equilibrium (Π-MNE) in a Π-game $G = \langle N, A, \Theta, \pi, \mu \rangle$ iff:

$$\forall i \in N, \forall \theta_i \in \Theta_i, \forall \upsilon_i', U_i^{pes}(\upsilon^*, \theta_i) \geq U_i^{pes}(\upsilon_i'.\upsilon_{-i}^*, \theta_i). \tag{22}$$

Example. Consider the following joint mixed strategy $\upsilon^* = (\upsilon_T^*.\upsilon_L^*)$:

$$\upsilon_T^*(P|H) = 1, \quad \upsilon_T^*(\overline{P}|H) = 1, \upsilon_T^*(P|D) = 1, \quad \upsilon_T^*(\overline{P}|D) = 1,$$
$$\upsilon_L^*(C|IH) = 0.5, \quad \upsilon_L^*(\overline{C}|IH) = 1, \quad \upsilon_L^*(C|ID) = 0.5, \quad \upsilon_L^*(\overline{C}|ID) = 1.$$

The pessimistic utilities of both players are:

$$U_T^{pes}(\upsilon^*, H) = \min\Bigg(\max\Big(1 - \upsilon_T^*(P|H), U_T^{pes}(P, \upsilon_L^*, H)\Big), \max\Big(1 - \upsilon_T^*(\overline{P}|H), U_T^{pes}(\overline{P}, \upsilon_L^*, H)\Big)\Bigg).$$

Since $U_T^{pes}(P, \upsilon_L^*, H)\Big) = 0$, then: $U_T^{pes}(\upsilon^*, H) = 0$.

Similarly, $U_T^{pes}(\upsilon^*, D) = 0.25$, $U_L^{pes}(\upsilon^*, IH) = 0.5$, $U_L^{pes}(\upsilon^*, ID) = 0.5$.

It can be checked that υ^* is a Π-MNE of the Π-game:

- $\forall \upsilon_T'(.|H) \neq \upsilon_T^*(.|H)$, $U_T^{pes}(\upsilon^*, H) = U_T^{pes}(\upsilon_T'(.|H).\upsilon_L^*, H) = 0$. Thus, the honest tenant has no incentive to deviate to another strategy;
- $\forall \upsilon_T'(.|D) \neq \upsilon_T^*(.|D)$, $U_T^{pes}(\upsilon^*, D) \geq U_T^{pes}(\upsilon_T'(.|D).\upsilon_L^*, D)$. Thus, the dishonest tenant has no incentive to deviate to another mixed strategy;
- $\forall \upsilon_L'(.|IH) \neq \upsilon_L^*(.|IH)$, $U_L^{pes}(\upsilon^*, IH) \geq U_L^{pes}(\upsilon_T^*.\upsilon_L'(.|IH), IH)$. Thus, the informed honest landlord has no incentive to deviate;
- $\forall \upsilon_L'(.|ID) \neq \upsilon_L^*(.|ID)$, $U_L^{pes}(\upsilon^*, ID) \geq U_L^{pes}(\upsilon_T^*.\upsilon_L'(.|ID), ID)$. Thus, the informed dishonest landlord has no incentive to deviate.

3.3. Transforming a Π-game into an equivalent ordinal normal form game

In this Section, we show that any incomplete information game can be transformed into an equivalent normal form game with *complete information*, such that its strategies and equilibria are in bijection with the ones of the original Π-game. This result is a qualitative counterpart of Harsanyi's result about the transformation of Bayesian games into normal form games under complete information [7]. We consider as many players as the number of pairs (i, θ_i) in the transformed game, each player (i, θ_i) having set of available actions A_i. A pure strategy in the Π-game associates an action to each type θ_i of each player i. In the "Complete information game representation" of the Π-game, a pure strategy associates an action to each player (i, θ_i):

Definition 15 *(Complete standard normal form representation of a Π-game).* The complete standard normal form representation (C-SNF-representation) of Π-game $G = \langle N, A, \Theta, \pi, \mu \rangle$ is the SNF game $\tilde{G} = \langle \tilde{N}, \tilde{A}, \tilde{\mu} \rangle$, where:

- $\tilde{N} = \big\{(i, \theta_i), i \in N, \theta_i \in \Theta_i\big\}$;
- $\tilde{A}_{(i,\theta_i)} = A_i$, $\forall (i, \theta_i) \in \tilde{N}$;
- $\tilde{\mu}_{(i,\theta_i)}(\tilde{a}) = U_i^{pes}(\sigma, \theta_i)$, $\forall \tilde{a} \in \tilde{A}, \forall (i, \theta_i) \in \tilde{N}$, where $\sigma_i(\theta_i) = \tilde{a}_{(i,\theta_i)}, \forall (i, \theta_i) \in \tilde{N}$.

Definition 16 *(Complete standard normal form representation of a pure strategy).* Let σ be a pure strategy in $G = \langle N, A, \Theta, \pi, \mu \rangle$. The C-SNF-representation of σ in $\tilde{G}$ is the profile of actions $\tilde{a}$ defined by:

$$\forall (i, \theta_i) \in \tilde{N}, \tilde{a}_{(i,\theta_i)} = \sigma_i(\theta_i). \tag{23}$$

It is easy to see that $\tilde{a}$ is a profile of actions of $\tilde{G}$ and that the strategies of G and $\tilde{G}$ are in bijection. It follows that:

Proposition 4. *The pure strategy σ is a PNE for $G = \langle N, A, \Theta, \pi, \mu \rangle$ iff $\tilde{a}$ is a PNE in its C-SNF representation $\tilde{G} = \langle \tilde{N}, \tilde{A}, \tilde{\mu} \rangle$.*

Example. The C-SNF-representation of the landlord and tenant game, $\tilde{G}$, has 4 players: $\tilde{N} = \{(T, H), (T, D), (L, IH), (L, ID)\}$, with $\tilde{A}_{(T,H)} = \tilde{A}_{(T,D)} = \{P, \overline{P}\}$, $\tilde{A}_{(L,IH)} = \tilde{A}_{(L,ID)} = \{C, \overline{C}\}$.

The joint action $\tilde{a}^* = (P.\overline{P}.C.\overline{C})$ in $\tilde{G}$ is the C-SNF representation of the joint pure strategy σ^* of G defined by $\sigma^*_T(H) = P$, $\sigma^*_T(D) = \overline{P}$, $\sigma^*_L(IH) = C$, $\sigma^*_L(ID) = \overline{C}$.

The utilities of $\tilde{a}^*$ in $\tilde{G}$ are:

$$\tilde{\mu}_{(T,H)}(\tilde{a}^*) = U^{pes}_T(\sigma^*, H) = 0.75, \quad \tilde{\mu}_{(T,D)}(\tilde{a}^*) = U^{pes}_T(\sigma^*, D) = 0.25,$$
$$\tilde{\mu}_{(L,IH)}(\tilde{a}^*) = U^{pes}_L(\sigma^*, IH) = 0.5, \quad \tilde{\mu}_{(L,ID)}(\tilde{a}^*) = U^{pes}_L(\sigma^*, ID) = 0.5.$$

Finally, every mixed strategy in the Π-game G has an equivalent mixed strategy in its C-SNF-representation $\tilde{G}$ and the mixed equilibria of the two games are in bijection.

Definition 17 *(Complete normal form representation of a mixed strategy).* Let υ be a mixed strategy in $G = \langle N, A, \Theta, \pi, \mu\rangle$. The C-SNF-representation of υ is the mixed strategy $\tilde{\upsilon}$ such that:

$$\forall (i, \theta_i) \in \tilde{N}, \tilde{\upsilon}_{(i,\theta_i)}(\tilde{a}_{(i,\theta_i)}) = \upsilon_i(a_i|\theta_i). \tag{24}$$

It is easy to show that: $\tilde{\upsilon}(\tilde{a}) = \upsilon(a|\theta)$.

As a consequence of Definition 17, we get:

Proposition 5. *The mixed strategy υ is a ΠMNE for $G = \langle N, A, \Theta, \pi, \mu\rangle$ iff $\tilde{\upsilon}$ is a ΠME in $\tilde{G} = \langle \tilde{N}, \tilde{A}, \tilde{\mu}\rangle$.*

Regarding the complexity of the transformation, consider that the original game is extensively represented by tables (the game is in standard normal form). The transformed game contains $\tilde{n} = \sum_{i=1,n} |\Theta_i|$ utility functions of size $\prod_{(i,\theta_i)\in\tilde{N}} |\tilde{A}_{(i,\theta_i)}| = \prod_{i=1,n} \prod_{\theta_i\in\Theta_i} |A_i| = \prod_{i=1,n} |A_i|^{|\Theta_i|}$.

For simplification purpose, let $|A_i| = d$ and $|\Theta_i| = t$, $\forall i \in N$ (i.e., the number of actions and the number of types are the same for all players). A game containing n utility functions of size $(d \cdot t)^n$ is transformed into a game containing $\tilde{n} = n \cdot t$ utility functions of size $d^{\tilde{n}} = (d^n)^t$, since the number of players $\tilde{n}$ in the transformed game is equal to $(n \cdot t)$. This transformation is exponential in time and space.

So, except when the number of types is very small, the transformation does not provide a convenient way to solve the game. Proposition 5 is, as in the Bayesian case, more a representation result than a solving tool.

3.4. Complexity results

In the previous Section (Section 3.3), we have shown that a Π-game can be transformed into an equivalent normal form game with complete information. Since it is not guaranteed to get a PNE in a normal form game, it is also not guaranteed that a Π-game admits a PNE.

In the following, we study the complexity of the problem of deciding whether a Π-game admits a PNE.

Definition 18 *(**PNE** problem).* The **PNE** problem consists in determining whether any given Π-game $G = \langle N, A, \Theta, \pi, \mu\rangle$ G admits a pure Nash equilibrium.

Deciding the existence of a pure Nash equilibrium in a Π-game is a difficult problem as stated by the following proposition:

Proposition 6. ***PNE*** *is NP-Hard, even in symmetric[7] 2-player games where π corresponds to total ignorance, i.e., $\forall\theta_1 \in \Theta_1$ and $\forall\theta_2 \in \Theta_2, \pi(\theta_1, \theta_2) = 1$.*

Finally, in [6], authors have shown that an ordinal SNF game admits at least one Π-MNE. Since any Π-game can be transformed into an equivalent SNF game with complete information (this is the result of the previous Section), it follows that a Π-game always admits at least one Π-MNE. We show in Section 4.2 that such an equilibrium can be computed in polynomial time.

[7] A game is symmetric if all players have the same set of actions, and the utilities of playing a given action depends only on the actions being played, not on who plays them.

4. Computing equilibria in Π-games

4.1. Finding a pure Nash equilibrium: a MILP formulation

Taking advantage of the efficiency of modern solvers, we propose a Mixed Integer Linear Programming (MILP) formulation of the problem of finding, if it exists, a PNE in a Π-game (we follow in this the idea explored by [29] for solving Bayesian games).

- the main decision variables are Boolean variables encoding the strategy searched for: each σ_{i,a_i,θ_i} is a Boolean variable indicating whether action a_i is prescribed for type θ_i of player i: $\forall i \in N, \forall a_i \in A_i, \forall \theta_i \in \Theta_i : \sigma_{i,a_i,\theta_i} \in \{0,1\}$;
- the utilities are encoded by continuous variables: U_{i,a_i,θ_i} is the utility (according to σ_{-i}) of player i if action a_i is chosen for type θ_i (i.e., if $\sigma_i(\theta_i) = a_i$): $\forall i \in N, \forall a_i \in A_i, \forall \theta_i \in \Theta_i : U_{i,a_i,\theta_i} \in [0,1]$.
- We will also use Boolean variables to constrain the U_{i,a_i,θ_i} to be equal to the $\min_{\theta_{-i} \in \Theta_{-i}}$ of $\max\big(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta)\big)$ (and not only lower than the *min*): $\forall i \in N, \forall a_i \in A_i, \forall \theta \in \Theta : M_{i,a_i,\theta} \in \{0,1\}$.

The MILP formulation we propose is the following:

$$\forall i \in N, \forall \theta_i \in \Theta_i, \sum_{a_i \in A_i} \sigma_{i,a_i,\theta_i} = 1. \tag{25}$$

$$\forall i \in N, \forall a_i, a_i' \in A_i, \text{s.t}, a_i \neq a_i', \forall \theta_i \in \Theta_i, U_{i,a_i,\theta_i} - U_{i,a_i',\theta_i} \geq \sigma_{i,a_i,\theta_i} - 1. \tag{26}$$

$\forall i \in N, \forall a \in A, \forall \theta \in \Theta$,

$$U_{i,a_i,\theta_i} \leq \max\big(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta)\big) + \sum_{j \in N,\ j \neq i} (1-\sigma_{j,a_j,\theta_j}). \tag{27}$$

$\forall i \in N, \forall a \in A, \forall \theta \in \Theta$,

$$U_{i,a_i,\theta_i} + M_{i,a_i,\theta} + \sum_{j \in N,\ j \neq i} (1-\sigma_{j,a_j,\theta_j}) \geq \max(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta)). \tag{28}$$

$$\forall i \in N, \forall a_i \in A_i, \forall \theta_i \in \Theta_i, \sum_{\theta_{-i} \in \Theta_{-i}} (1 - M_{i,a_i,\theta_i,\theta_{-i}}) = 1. \tag{29}$$

- Constraints (25) ensure that the strategy σ searched for specifies exactly one action per type, for each player i.
- Constraints (26) require that σ is a PNE: when $\sigma_{i,a_i,\theta_i} = 1$, this constraint requires $U_{i,a_i,\theta_i} \geq U_{i,a_i',\theta_i}$, i.e., that player i has no incentive to deviate from a_i. When action a_i is not chosen for θ_i, ($\sigma_{i,a_i,\theta_i} = 0$) the constraint is always satisfied ($U_{i,a_i,\theta_i} - U_{i,a_i',\theta_i}$ is always greater than -1).
- Constraints (27) implement Definition (9). They ensure that the utility of player i playing $\sigma(\theta_i) = a_i$ is lower than all, i.e., the minimum over the θ_{-i} of the $\max(1-\pi(\theta_{-i}|\theta_i), \mu_i(a_i.\sigma_{-i}(\theta_{-i}),\theta))$. Indeed, for any profile of action a that does not correspond to what is prescribed by σ, $\Sigma_{j \neq i}(1-\sigma_{j,a_j,\theta_j}) \geq 1$, the constraint is always satisfied ($U_{i,a_i,\theta_i} \leq 1$). If a_{-i} is chosen for θ_{-i}, then $\sigma_{j,a_j,\theta_j} = 1$, $\forall j \neq i$ and $\Sigma_{j \neq i}(1-\sigma_{j,a_j,\theta_j}) = 0$: the constraint becomes $U_{i,a_i,\theta_i} \leq \max(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta))$.
- Constraints (28) and (29) ensure that U_{i,a_i,θ_i} is equal to the *min* over θ_{-i} of $\max(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta))$: If a_{-i} does not correspond to σ_{-i}, $\Sigma_{j \neq i}(1-\sigma_{j,a_j,\theta_j})$ is at least equal to 1 and the constraints (28) are always satisfied. Otherwise, (a_{-i} correspond to σ_{-i}) the *sum* is equal to 0 and does not annihilate the constraint. The *min* is reached if $U_{i,a_i,\theta_i} = \max(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta))$. Whenever $M_{i,a_i,\theta} = 1$, Equation (28) holds, and Equation (29) ensures that (28) is an equality for one θ_{-i} (minimizing $\max(1-\pi(\theta_{-i}|\theta_i), \mu_i(a,\theta))$).

The above formulation is linear (the max operator which appears in constraints (27) and (28) deals with constants only). Furthermore, it does not lead to a combinatorial explosion of the required space. Recall that the size of the original problem is $n \cdot |\Theta| \cdot |A| + |\Theta|$. Let us denote d (resp. t) the number of actions (resp. types) of each player. The MILP formulation contains: $O(n \cdot t \cdot d)$ continuous variables U_{i,a_i,θ_i}; $O(n \cdot t \cdot d)$ Boolean variables σ_{i,a_i,θ_i}; $O(n \cdot d \cdot t^n)$ Boolean variables $M_{i,a_i,\theta}$; $O(n \cdot t)$ constraints (25), each involving $O(a)$ variables; $O(n \cdot t \cdot d \cdot (d-1))$ constraints (26), each involving 3 variables; $O(n \cdot t^n \cdot d^n)$ constraints (27) each involving $O(n)$ variables; $O(n \cdot t^n \cdot d^n)$ constraints (28) each involving $O(n+1)$ variables; $O(n \cdot d \cdot t)$ constraints (29) each involving $O(\frac{t^n}{d})$ variables. The size of the MILP encoding is thus in $O(n^2 \cdot |\Theta| \cdot |A|)$ (polynomial in the size of the original size of the problem (i.e., $n \cdot |\Theta| \cdot |A| + |\Theta|$)).

4.2. Building a possibilistic mixed Nash equilibrium

We have seen in Section 3.3 that any Π-game G can be transformed into an equivalent C-SNF-representation $\tilde{G}$ (Definition 15) and that a mixed strategy π in $\tilde{G}$ corresponds to a unique mixed strategy υ in G (Definition 17). So, it is

Algorithm 1 IMPROVE.

Data: $G = \langle N, A, \Theta, \pi, \mu \rangle$, υ, i, θ_i
Result: $\upsilon' = (\upsilon_1, \ldots, \upsilon_{i-1}, \upsilon'_i, \upsilon_{i+1}, \ldots, \upsilon_n)$
$\tilde{D}_{(i,\theta_i)} \leftarrow \emptyset$
forall $a_i \in A_i$ **do**
  if $U_i^{pes}(a_i, \upsilon_{-i}, \theta_i) \leq U_i^{pes}(\upsilon, \theta_i)$ **then** $\tilde{D}_{(i,\theta_i)} \leftarrow \tilde{D}_{(i,\theta_i)} \cup \{a_i\}$
if $\forall a_i \in A_i \setminus \tilde{D}_{(i,\theta_i)}$, $\upsilon_i(a_i|\theta_i) < 1$ **then** $\upsilon' \leftarrow \upsilon$
else
  forall $a_i \in A_i$ **do**
    if $a_i \in \tilde{D}_{i,\theta_i}$ **then** $\upsilon'_i(a_i|\theta_i) \leftarrow \max\{\alpha \in \Delta; \alpha < 1 - U_i^{pes}(\upsilon, \theta_i)\}$
    else $\upsilon'_i(a_i|\theta_i) \leftarrow \upsilon_i(a_i|\theta_i)$
return υ'

Algorithm 2 FINDING A LEAST-SPECIFIC Π-MNE IN Π-GAME.

Data: $G = \langle N, A, \Theta, \pi, \mu \rangle$
Result: $\upsilon^* = (\upsilon_1^*, \ldots, \upsilon_n^*)$, a Π-MNE
$\upsilon^0 \leftarrow (\upsilon_1^0, \ldots, \upsilon_n^0)$ /* $\upsilon_i^0(\theta_i, a_i) = 1, \forall i \in N, \forall a_i \in A_i, \forall \theta_i \in \Theta_i$ */
$t \leftarrow 0$
repeat
  $\upsilon^{loc} \leftarrow \upsilon^t$
  forall $i \in N$ **do**
    forall $\theta_i \in \Theta_i$ **do** $\upsilon^{loc} \leftarrow$ IMPROVE $(G, \upsilon^{loc}, i, \theta_i)$
  $\upsilon^{t+1} \leftarrow \upsilon^{loc}$
  $t \leftarrow t + 1$
until $\upsilon^t = \upsilon^{t-1}$
$\upsilon^* \leftarrow \upsilon^t$
return υ^*

theoretically possible to apply the algorithm proposed by [6]. This algorithm is guaranteed to return a ΠMNE of $\tilde{G}$, π^*, in time polynomial in the expression of $\tilde{G}$. And from π^*, we immediately get an equilibrium strategy for G as expressed above. But, as noted in Section 3.3, the size of $\tilde{\mu}$ is exponential in that of the initial Π-game. Thus, if we actually construct $\tilde{G}$, the algorithm of [6] will converge to a ΠMNE, but in time exponential in the size of G.

We explore here a more direct (and polynomial time) approach. The algorithm we present (Algorithms 1 and 2) is a generalization of the one proposed in [6]. In this iterative process, at each step, each player i of type θ_i tries to maximize her pessimistic utility by moving her mixed strategy to a more specific one. We can indeed show that when a player moves to a more specific mixed strategy her pessimistic utility cannot decrease.

Proposition 7. *Let υ and υ' be two mixed strategies such that $\upsilon' \preceq \upsilon$. Then: $\forall i \in N, \forall \theta_i \in \Theta_i, U_i^{pes}(\upsilon, \theta_i) \leq U_i^{pes}(\upsilon', \theta_i)$.*

Algorithm 2 mimics a negotiation process. At each stage (each t), given the joint mixed strategy, every player i tries to improve her pessimistic utility by changing her mixed strategy υ_i^{loc}. Thanks to Proposition 7 changing to a more specific mixed strategy never decreases the utility of any of the players. When no player has any more incentive to make her strategy more specific, the result of the negotiation leads to a *least specific* mixed equilibrium.

We first show that Algorithm 2, which performs iterations over the IMPROVE procedure, converges, and convergence occurs in time polynomial in the size of the Π-game G.

Proposition 8 *(Convergence). Let $G = \langle N, A, \Theta, \pi, \mu \rangle$ be a Π-game. Algorithm 2 converges in a finite number of steps*

Proposition 9 *(Complexity of Algorithm 1).* IMPROVE *runs in time polynomial in the size of the input Π-game. The whole complexity of the* IMPROVE *function is: $O(d^{n+1} \times t^{n-1})$*

Algorithm 2 calls $N \times t$ times Algorithm 1. Thus, thanks to Proposition 9, we get:

Corollary 1 *(Complexity of Algorithm 2). The whole complexity of algorithm is 2: $O(n \times t^n \times |\Delta| \times d \times d^n \times n) = O(n^2 \times t^n \times |\Delta| \times d^n)$.*

We finally show that the algorithm converges towards a possibilistic mixed equilibrium of the Π-game G.

Proposition 10 *(Soundness). Algorithm 2 converges towards a least-specific possibilistic mixed equilibrium of $G = \langle N, A, \Theta, \pi, \mu \rangle$.*

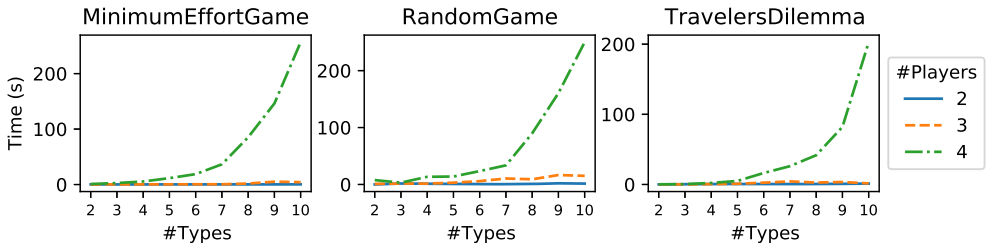

Fig. 1. Avg. execution time, MILP-PNE, $|A_i| = 2$, $|\Theta_i| \in [2, 10]$, $n \in [2, 4]$.

5. Experimental study

The goal of this experimental study is to evaluate the efficiency and feasibility of the algorithms described in the previous Section. To this end, we adapted the GAMUT game generator [30] to generate Π-games (GAMUT produces exclusively normal form games with complete information). In the present Section, we detail this Π-game generator, the experimental protocol and our experimental results.

5.1. A Π-game generator

Our generator is based on the idea that every Π-game can be equivalently defined as a set of $|\Theta|$ normal form games with the same set of players N and actions A, following the approach of [29] for the generation of Bayesian games. The idea is to generate, using GAMUT, for each combination of types $\theta \in \Theta$ a normal form game G^θ and a possibility distribution over the combinations of types. For the sake of simplicity, we assume that the number of actions and types are equal for all players, i.e., $\forall i, j \in N, |A_i| = |A_j| = d$ and $\forall i, j \in N, |\Theta_i| = |\Theta_j| = t$.

To generate a Π-game, $G = \langle N, A, \Theta, \pi, \mu \rangle$, we need as inputs: *(1)* the class and the name of the game, *(2)* the number of players, *(3)* the number of degrees in Δ, *(4)* the number of types per player and *(5)* the number of actions per player. Then, we ask GAMUT to generate $|\Theta|$ normal form games of the class given in input, the range of utility of which is Δ and we generate a normalized distribution $\pi : \Theta \mapsto \Delta$ (a randomly selected θ receives degree 1; the degrees of the other elements of Θ are selected in Δ following a uniform distribution). Finally, the utility $\mu_i(a, \theta)$ is simply the utility of the joint action a for player i in the normal form game $G^\theta = \langle N, A, \{\{\mu_i^\theta\}_{i \in N}\} \rangle$.

5.2. Experimental protocol

In our experiments, we set $\Delta = \{0, 0.25, 0.5, 0.75, 1\}$ and let the number of players vary from 2 to 10, the number of types from 2 to 10 and the number of actions from 2 to 10. For each combination of parameters, 100 instances are generated and the times necessary to get *(i)* a pure equilibrium (or a negative result) and *(ii)* a possibilistic mixed equilibrium, are measured.

We present results for 3 classes of games: *Minimum Effort game*, *Random game* and *Travelers Dilemma game*. In our evaluation, we bounded the execution time for a single game to 10 minutes, as in [31,32].

All experiments were conducted on an Intel Xeon E5540 processor and 64 GB RAM workstation. We used CPLEX [33] as a MILP solver and Java 8 as a programming language. The implementation of the transformation of the Π-game as a normal form game and the MILP solver are available online.[8]

5.3. Experimental results: pure Nash equilibria

Globally the experiments based on the MILP formulation of Section 4.1 confirm the feasibility of the qualitative approach for incomplete information games. In the experiment reported on Figs. 1 and 2, we let the number of types vary from 2 to 10 and n vary from 2 to 6 - setting the number of actions to 2. Unsurprisingly, the CPU time increases with the number of players and the number of types. This is explained by the fact that when the number of types and players increases, $|\Theta|$ and $|A|$ increase directly (as in any incomplete information games). Thus, the size of the MILP increases.

However, we note a non-monotonic behavior for the *Minimum Effort game*: the average execution time for 6 players is lower than 5 players (in the case of $|\Theta_i| = 5, \ \forall i \in N$). This can be explained by the fact that among the 100 games randomly generated at each point, some may have no PNE. Indeed, from Figs. 3 and 4 showing the average execution time for games that admit at least a PNE and for those that do not admit a PNE, respectively, we can see that the existence of a PNE in a game highly impacts the time needed to solve it. In particular, for the *Minimum Effort game*, with 5 players and 5 types per player, the average execution time to solve 100 games is around 300*s* (see Fig. 2), this value is obtained from 63

[8] https://www.irit.fr/~Helene.Fargier/PossibilisticGames.html.

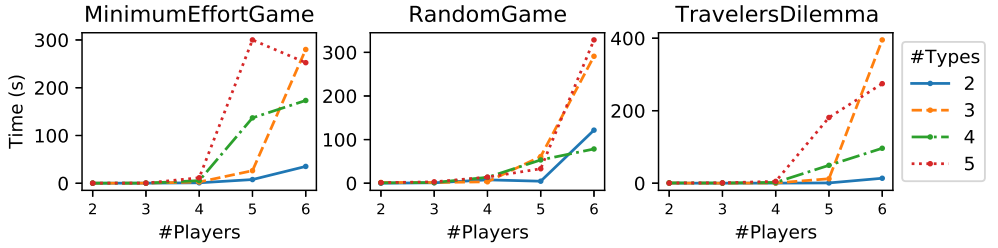

Fig. 2. Avg. execution time (s), MILP-PNE, $|A_i| = 2$, $n \in [2, 6]$, $|\Theta_i| \in [2, 5]$.

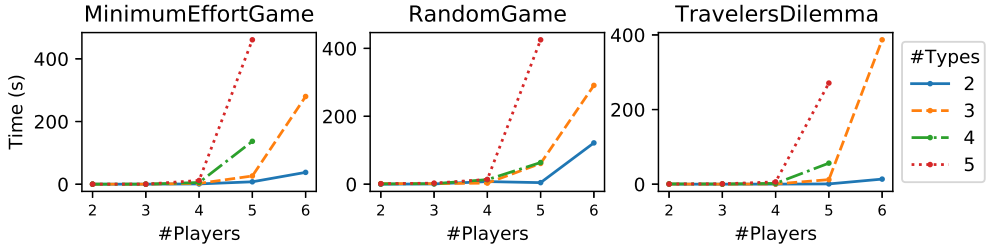

Fig. 3. Avg. execution time (s), MILP-PNE, for the games that admit a PNE, $|A_i| = 2$, $n \in [2, 6]$, $|\Theta_i| \in [2, 5]$. Missing configurations (e.g. 6 players/5 types for the three categories of games) are configurations where no game with a PNE was found.

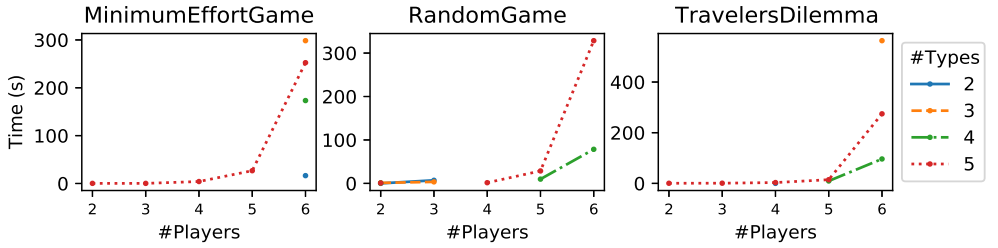

Fig. 4. Avg. execution time (s), MILP-PNE for the games that do not admit a PNE, $|A_i| = 2$, $n \in [2, 6]$, $|\Theta_i| \in [2, 5]$. Missing configurations (e.g. every 2 player games for the three categories of games) are configurations where all games admitted a PNE.

games that admit a PNE with an average time of around 460*s* (see Fig. 3) and 37 games that do not admit a PNE with an average time around 27*s* (see Fig. 4). Moreover, for *Minimum Effort game*, with 6 players and 5 types per player, the decrease of the average execution time to solve 100 games (in Fig. 2) is due to the non-existence of a PNE for all instances. The increasing of the number of players has a twofold effect: the size of the game grows, leading hence to the growth of the MILP formulation - finding a PNE becomes harder when it exists. On the other hand, the MILP formulation becomes more constrained, and for the instances without PNE, the solver detects inconsistent instances faster.

5.4. Experimental results: possibilistic mixed equilibria

In this experimentation of Algorithm 2, we start by setting the number of players equal to 2 and varying the number of actions from 2 to 10 and the number of types from 2 to 10. The execution time needed to find one of the least specific Π-MNE is presented in Figs. 5 and 6, respectively. Then, we set the number of types per player equal to 2 and let the number of players and action vary - see Fig. 7.

The results show that, for the 3 classes of games, Algorithm 2 can return the least specific Π-MNE in a reasonable time (less than 0.4 seconds when the number of players is equal to 2 and the number of actions and types are equal to 10). The number of players in a Π-game is the "main" parameter that influences the average execution time needed to find a Π-MNE. Fig. 7 shows that if the number of actions per player is equal to 4 and the number of types per player is equal to 2, the average execution time needed to find the least specific Π-MNE in a Minimum Effort game with 5 players is around to 2 *s*. When the number of players is equal to 6 (resp., 7) the average execution time is around to 30 *s* (resp., 115 *s*).

The results empirically confirm the complexity analysis of the algorithm - its complexity is polynomial in the size of the game, which is itself exponential in the number of players.

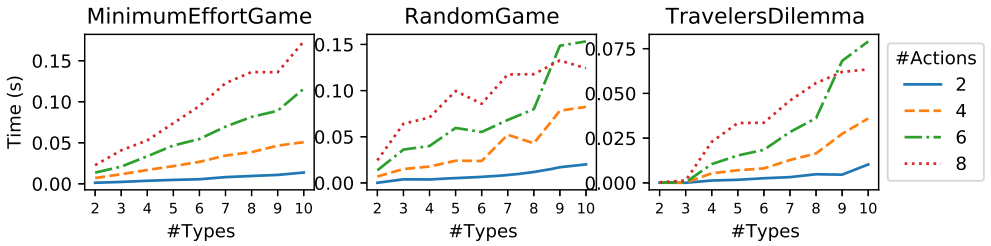

Fig. 5. Avg. execution time (s) to find a Π-MNE, $n = 2$, $|\Theta_i| \in [2, 10]$, $|A_i| \in [2, 8]$.

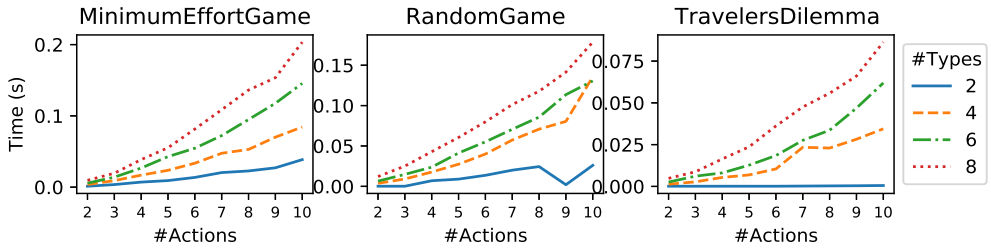

Fig. 6. Avg. execution time (s) to find a Π-MNE, $n = 2$, $|A_i| \in [2, 10]$, $|\Theta_i| \in [2, 8]$.

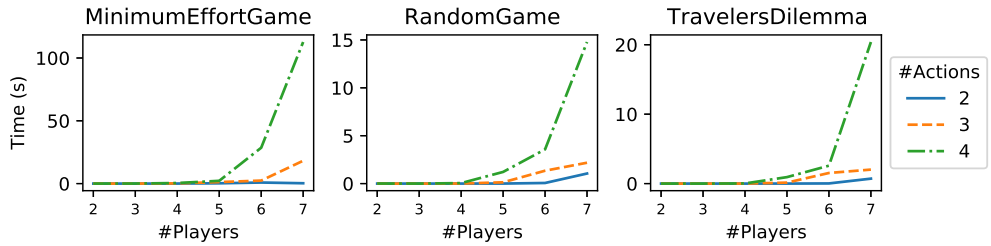

Fig. 7. Avg. execution time (s) to find a Π-MNE, $|\Theta_i| = 2$, $n \in [2, 7]$, $|A_i| \in [2, 4]$.

6. Conclusion

In this paper, possibilistic games have been proposed as a new representation framework for ordinal games under possibilistic incomplete information. The notions of pure and mixed ordinal equilibria have been defined for such games and we have shown that determining whether a pure Nash equilibrium exists in a Π-game is a NP-complete problem. Instead, the problem can be solved in polynomial time when mixed equilibria are considered. The theoretical time needed to reach a mixed equilibrium increases linearly with the size of the input (of the game), which is itself exponential in the number of players ($n \cdot (t \cdot a)^n$). Because the size of the game is highly impacted by the number of players, experiments show that the average execution time increases drastically when the number of players increases. This is observable for all the classes of games we considered in our experiments: *Minimum Effort game*, *Random game* and *Travelers Dilemma game*. This confirms the theoretical results of this paper.

There are several possible followings to this work. The first one would be to define succinct forms of possibilistic games. Indeed, it is even more crucial in incomplete information game than in complete information ones, to offer the possibility of succinct expressions of utility tables, since they depend on joint types in addition to joint actions. The natural extensions would be to consider games where the interactions are local, i.e. where the utility of a player depends only on a subset of players. To explore this path of research, we may extend the *polymatrix* [34], *graphical* [35], *hypergraphical* [36] (complete information) games frameworks to the possibilistic, incomplete information case.

A longer term perspective would be to consider studying sequential possibilistic games. In the probabilistic, quantitative framework, competitive Markov decision processes (MDP) [37] extend both (MDP) [38] and cardinal games, to sequential games under uncertainty. Partially observed stochastic games [39], on their side, extend both partially observed MDP [40] and cardinal games, to sequential incomplete information games. Possibilistic MDP (and POMDP) have been introduced by [41,42], in the possibilistic framework. A fruitful avenue for new research would be to extend possibilistic MDP/POMDP and possibilistic complete/incomplete information games, so as to define possibilistic competitive MDP and to develop a possibilistic version of partially observed stochastic games.

7. Proofs

Proof of Proposition 1. Let $G = \langle N, A, \Theta, \pi, \mu\rangle$ be a Π-game where $\forall i \in N$, $|\Theta_i| = 1$. Then it exists a unique type combination $\theta \in \Theta$, i.e., $|\Theta| = 1 \Rightarrow$ there is one possible game G^θ such that $\pi(G^\theta) = 1$ since π is normalized. Thus G is equal to classical normal form game, i.e., $G = G^\theta = \langle N, A, \{\{\mu_i^\theta\}_{i\in N}\}\rangle$. □

Proof of Proposition 2.

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i},\, a\in A} \max\big(1 - \pi_\upsilon(a, \theta_{-i}|\theta_i), \mu_i(a, \theta)\big).$$

Based on Equation (17):

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i},\, a\in A} \max\big(1 - \min(\pi(\theta_{-i}|\theta_i), \upsilon(a|\theta_i.\theta_{-i})), \mu_i(a, \theta)\big).$$

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i},\, a\in A} \max\big(1 - \pi(\theta_{-i}|\theta_i), 1 - \upsilon(a|\theta_i.\theta_{-i}), \mu_i(a, \theta)\big).$$

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i}} \max\Big(1 - \pi(\theta_{-i}|\theta_i), \min_{a\in A} \max\big(1 - \upsilon(a|\theta_i.\theta_{-i}), \mu_i(a, \theta)\big)\Big). \quad \square$$

Proof of Proposition 3.

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i}} \max\Big(1 - \pi(\theta_{-i}|\theta_i), \min_{a\in A} \max\big(1 - \upsilon(a|\theta), \mu_i(a, \theta)\big)\Big)$$

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i}} \max\Big(1 - \pi(\theta_{-i}|\theta_i),$$

$$\min_{a\in A} \max\big(1 - min(\upsilon_i(a_i|\theta_i), \upsilon_{-i}(a_{-i}|\theta_{-i})), \mu_i(a, \theta)\big)\Big).$$

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i}} \max\Big(1 - \pi(\theta_{-i}|\theta_i), \min_{a\in A}$$

$$\max\big(1 - \upsilon_i(a_i|\theta_i), 1 - \upsilon_{-i}(a_{-i}|\theta_{-i}), \mu_i(a, \theta)\big)\Big).$$

$$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i}, a\in A}$$

$$\max\Big(1 - \pi(\theta_{-i}|\theta_i), 1 - \upsilon_i(a_i|\theta_i), 1 - \upsilon_{-i}(a_{-i}|\theta_{-i}), \mu_i(a, \theta)\Big).$$

$$U_i^{pes}(\upsilon, \theta_i) = \min_{a_i\in A_i} \max(1 - \upsilon_i(a_i|\theta_i), U_i^{pes}(a_i, \upsilon_{-i}, \theta_i)). \quad \square$$

Proof of Proposition 4. $\Rightarrow$ Assume that a is a PNE in $\tilde{G}$. Then,

$$\tilde{\mu}_{(i,\theta_i)}(a'_i.a_{-i}) \le \tilde{\mu}_{(i,\theta_i)}(a), \forall i, \theta_i, a'_i.$$

But, since $\tilde{\mu}_{(i,\theta_i)}(a) =_{def} U_i^{pes}(a^\sigma, \theta_i)$, we get $U_i^{pes}(a'_i, a^\sigma_{-i}, \theta_i) \le U_i^{pes}(a^\sigma, \theta_i)$.
Thus, a^σ is a PNE of G.

$\Leftarrow$ Now, let σ be a PNE of G, define $\tilde{G}$ and joint action a: $a_{(i,\theta_i)} = \sigma_i(\theta_i), \forall(i, \theta_i)$. Then, again, $\tilde{\mu}_{(i,\theta_i)}(a) = U_i^{pes}(\sigma, \theta_i)$. And since σ is a PNE in G, we get $U_i^{pes}(a'_i, \sigma_{-i}, \theta_i) \le U_i^{pes}(\sigma, \theta_i), \forall i, \theta_i, a'_i$ and, $\forall i, \theta_i, a'_i$

$$\tilde{\mu}_{(i,\theta_i)}(a'_{(i,\theta_i)}.a_{-(i,\theta_i)}) \le \tilde{\mu}_{(i,\theta_i)}(a).$$

Thus, a is a PNE of $\tilde{G}$. □

Proof of Proposition 5. First, note that the transformation $\upsilon \to \tilde{\upsilon}$ (Definition 17) between the sets of mixed strategies in games G and $\tilde{G}$ is bijective. Thus, in order to prove that mixed equilibria are the same in both games, it is enough to show that

$$U_i^{pes}(\upsilon, \theta_i) = \mu_{i,\theta_i}^{pes}(\tilde{\upsilon}), \forall(i, \theta_i). \tag{30}$$

To do so, first note that:
$U_i^{pes}(\upsilon, \theta_i) = \min_{\theta_{-i}\in\Theta_{-i}, a\in A} max(1 - \pi_\upsilon(a, \theta_{-i}|\theta_i), \mu_i(a, \theta))$, by Definition 13,
$\pi_\upsilon(a, \theta_{-i}|\theta_i) = \min(\pi(\theta_{-i}|\theta_i), \upsilon(a|\theta))$, by Equation (17),
and $\tilde{\upsilon}(\tilde{a}) = \upsilon(a|\theta)$, by Definition 17.

Thus,

$$U_i^{pes}(\upsilon, \theta_i) = \min_{a \in A} max(1 - \tilde{\upsilon}(\tilde{a}), \min_{\theta_{-i} \in \Theta_{-i}} max(1 - \pi(\theta_{-i}|\theta_i), \mu_i(a, \theta))).$$

Note that, by Definition 9,

$$U_i^{pes}(\sigma, \theta_i) = \min_{\theta_{-i} \in \Theta_{-i}} max(1 - \pi(\theta_{-i}|\theta_i), \mu_i(\sigma(\theta), \theta)),$$

where σ is the unique pure strategy in G, defined from any pure strategy a in $\tilde{G}$.

Then,

$$U_i^{pes}(\upsilon, \theta_i) = \min_{a \in A} max(1 - \tilde{\upsilon}(\tilde{a}), U_i^{pes}(\sigma, \theta_i)),$$
$$U_i^{pes}(\upsilon, \theta_i) = \min_{a \in A} max(1 - \tilde{\upsilon}(\tilde{a}), \tilde{\mu}_{(i,\theta_i)}(a)), \text{ by Definition 15},$$
$$U_i^{pes}(\upsilon, \theta_i) = \mu_{i,\theta_i}^{pes}(\tilde{\upsilon}), \text{ by Eq. (7)}.$$

Thus, Proposition 5 holds. □

Proof of Proposition 6. Membership. We prove the membership in NP for the more general case of N unbounded. In this case, the size of the input is exponential in the number of players n. The PNE can be solved by guessing a strategy σ, i.e., guessing an action for each pair player/type, then checking whether σ is a PNE or not. More precisely: For each player i and for each type $\theta_i \in \Theta_i$:

- compute $U_i^{pes}(\sigma, \theta_i)$ and
- for each action $a_i \in A_i$, compute $U_i^{pes}(a_i, \sigma_{-i}, \theta_i)$.

Then check if i has incentive to deviate from $\sigma_i(\theta_i)$, i.e., we should compare $U_i^{pes}(\sigma, \theta_i)$ and $U_i^{pes}(a_i, \sigma_{-i}, \theta_i)$. Under the assumption that π is represented by a table of $|\Theta|$ lines, the complexity of computing $U_i^{pes}(\sigma, \theta_i)$ is in $O(|\Theta_{-i}|)$, from Definition 9. Thus the whole complexity is polynomial $O(n \times |\Theta_i| \times |\Theta_{-i}| \times |A_{max}|) = O(n \times |\Theta| \times |A_{max}|)$ where $|A_{max}| = max(|A_1|, ..., |A_n|)$.

Algorithm 3 details this process:

Algorithm 3 CHECK_EQUILIBRIUM.

```
Data: G = ⟨N, A, Θ, π, μ⟩, σ
Result: IsPNE (Boolean)
IsPNE ← true
forall i in N do
    forall θ_i ∈ Θ_i do
        Compute U_i^pes(σ, θ_i)
        forall a_i ∈ A_i do
            Compute U_i^pes(a_i, σ_{-i}, θ_i)
            if U_i^pes(a_i, σ_{-i}, θ_i) > U_i^pes(σ, θ_i) then IsPNE← false;
return IsPNE
```

σ is a PNE using Algorithm 3 is in $O(n \times |\Theta| \times |A_{max}|)$.

Hardness.[9] The hardness proof uses a reduction from the SET-COVER problem:

Definition 19 (***SET-COVER*** *(SC) problem).* Given a set $S = \{s_1, ..., s_n\}$, subsets $\{S_1, S_2, ..., S_m\}$ of S with $\cup_{1 \le i \le m} S_i = S$ and an integer $K \le m$. We are asked whether there exists a subset of $\{S_1, S_2, ..., S_m\}$ of size K whose union equals S, i.e., $S_{c_1}, ..., S_{c_K}$ such that $\cup_{1 \le i \le K} S_{c_i} = S$.

We reduce an arbitrary $SC = \langle S, \{S_1, S_2, ..., S_m\}, K \rangle$ instance to a PNE instance: Let Π-game $G_{SC} = \langle N, A, \Theta, \pi, \mu \rangle$ where:

- $N = \{1, 2\}$
- $A = A_1 \times A_2$ where $A_1 = A_2 = \{S_1, ..., S_m, s_1, ..., s_n\}$

[9] The proof of hardness is inspired from [43].

- $\Theta = \Theta_1 \times \Theta_2$ where $\Theta_1 = \Theta_2 = \{t_1, ..., t_K\}$ (both players belong to one of K types)
- $\forall \theta_1 \in \Theta_1, \theta_2 \in \Theta_2, \pi(\theta_1, \theta_2) = 1$ (π reflects total ignorance).
- We assume utility functions that do not depend on a specific type $\theta \in \Theta$. They are as follows:

(i) $\mu_1(S_i.S_j) = \mu_2(S_j.S_i) = 0.25\ \forall S_i, S_j$
(ii) $\mu_1(S_i.s_j) = \mu_2(s_j.S_i) = 0.25 + \frac{j}{4n+1}\ \forall S_i, \forall s_j \notin S_i$
(iii) $\mu_1(S_i.s_j) = \mu_2(s_j.S_i) = 0.5\ \forall S_i, \forall s_j \in S_i$
(iv) $\mu_1(s_i.s_j) = \mu_2(s_j.s_i) = 0\ \forall s_i, s_j$
(v) $\mu_1(s_j.S_i) = \mu_2(S_i.s_j) = 0.75\ \forall S_i, \forall s_j \notin S_i$
(vi) $\mu_1(s_j.S_i) = \mu_2(S_j.s_i) = 0\ \forall S_i, \forall s_j \in S_i$

Note that a Set-Cover instance SC can be represented in space $O(mn\log(n))$. The size of the Π-game G_{SC} is the size required to represent π, μ_1 and μ_2. Assuming that π is represented as a table (which is obviously not the most concise way to represent it), $|\pi| = O(K^2) = O(n^2)$ and $|\mu_1| = |\mu_2| = O((m+n)^2)$. The latter size may become $O((m+n)^2K^2)$ if we store $\mu_i(\sigma, \theta)$ even though utilities are independent of θ. Thus, G_{SC} requires space polynomial in that of SC to be represented. And since every $\pi(\theta)$ and $\mu_i(\sigma, \theta)$ require constant time to be computed, the transformation is polynomial (time).

Note that, in Π-game with 2 players and π corresponding to total ignorance, the utility of the 2 players are computed as follows:

$$U_1^{pes}(\sigma, \theta_1) = \min_{\theta_2 \in \Theta_2} \max\left(1 - \pi(\theta_2|\theta_1), \mu_1(\sigma_1(\theta_1).\sigma_2(\theta_2))\right)$$

Since $\pi(\theta_2|\theta_1) = 1$, thus:

$$U_1^{pes}(\sigma, \theta_1) = \min_{\theta_2 \in \Theta_2} \mu_1(\sigma_1(\theta_1).\sigma_2(\theta_2)).$$

In the same way, $U_2^{pes}(\sigma, \theta_2) = \min_{\theta_1 \in \Theta_1} \mu_2(\sigma_1(\theta_1).\sigma_2(\theta_2))$.

Now we show that G admits a PNE $\Leftrightarrow$ SC admits a SET-COVER.

SC admits a SET-COVER $\Rightarrow$ G admits a Π-PNE First suppose there exist $S_{c_1}, ..., S_{c_K}$ such that $\cup_{1 \le i \le K} S_{c_i} = S$. Suppose both players $i = \{1, 2\}$ play $S_{c_{\theta_i}}$ when their type is θ_i, i.e., $\forall \theta_1 \in \Theta_1$, $\sigma_1(\theta_1) = S_{c_{\theta_1}}$ and $\forall \theta_2 \in \Theta_2$, $\sigma_2(\theta_2) = S_{c_{\theta_2}}$. We claim that $\langle \sigma_1.\sigma_2 \rangle$ is a PNE.

Player 1 (resp. 2) supposes that player 2 (resp. 1) employs this strategy. Then, note that for any s_j, there is at least one S_{c_i} such that $s_j \in S_{c_i}$, since SC admits a SET-COVER $\{S_{c_1}, \ldots, S_{c_K}\}$. So, if player 1 of type θ_1, for example, changes her strategy by replacing $S_{c_{\theta_1}}$ with some s_j, this will decrease her utility from 0.25 (i) to 0 (vi), since s_j is covered by some $S_{c_{\theta_2}}$ played by player 2 of type θ_2. Of course, the same holds for the other player, so that no player has interest to deviate from the SET-COVER play. It follows that playing any of the S_j is optimal. So there is a PNE.

G admits a PNE $\Rightarrow$ SC admits a SET-COVER Suppose that G admits a PNE σ^*. We are going to show by contradiction that $\{\sigma_1^*(\theta_1)\}_{\theta_1 \in \Theta_1}$ and $\{\sigma_2^*(\theta_2)\}_{\theta_2 \in \Theta_2}$ form Set covers of SC.

1. Assume player 1 plays some $\sigma_1(\theta_1) \in S$ for some $\theta_1 \in \Theta_1$. We show that we have $\sigma_2(\theta_2) \subseteq S, \forall \theta_2 \in \Theta_2$. Indeed, if any $\sigma_2(\theta_2) \in S$ for some $\theta_2 \in \Theta_2$, then $U_2^{pes}(\sigma, \theta_2) = \mu_2(\sigma_1(\theta_1), \sigma_2(\theta_2)) = 0$ while $U_2^{pes}(\sigma', \theta_2) = 0.25$ if $\sigma_2(\theta_2) \in S$ is replaced with any $\sigma_2'(\theta_2) = S_c$.
2. Now, forget about Player 1 and assume that $\sigma_2(\theta_2) \subseteq S, \forall \theta_2 \in \Theta_2$ and that $S \setminus \cup_{\theta_2} \sigma_2(\theta_2)$ is non-empty. Then, obviously, $\sigma_1(\theta_1) \in S \setminus \cup_{\theta_2} \sigma_2(\theta_2), \forall \theta_1 \in \Theta_1$, since this provides utility 0.75 to player 1 of any type θ_1. However, let s_{j^*} be the state with minimum index for which there exists θ_1 s.t. $\sigma_1(\theta_1) = s_{j^*}$. s_{j^*} is not "covered" by any $\sigma_2(\theta_2)$ (so, $U_2^{pes}(\sigma, \theta_2) = 0.25 + \frac{j^*}{4n+1}$). However, considering (ii), Player 2 of any type θ_2 will be better off trading $\sigma_2(\theta_2)$ for some S_c such that $s_{j^*} \in S_c$, since this will increase the smallest index of uncovered states.
   Thus, we have a contradiction, and $\cup_{\theta_2} \sigma_2(\theta_2) = S$.
3. The final step is the following: In step 1 we proved that, when G admits a PNE, if for some $\theta_1 \in \Theta_1, \sigma_1(\theta_1) \in S$, then $\sigma_2(\theta_2) \subseteq S, \forall \theta_2 \in \Theta_2$. Then, in step 2 we proved that if $\sigma_2(\theta_2) \subseteq S, \forall \theta_2 \in \Theta_2$, then $\cup_{\theta_2} \sigma_2(\theta_2) = S$. However, it may be that $\sigma_1(\theta_1) \subseteq S, \forall \theta_1 \in \Theta_1$. But in this case, symmetrically to step 2, we can show that $\cup_{\theta_1} \sigma_1(\theta_1) = S$. In this case too, we have proved that there exists a set cover. □

Proof of Proposition 7. The Improve function defined in [6] relies on the computation of a set $\tilde{D}_{(i,\theta_i)} \subseteq \tilde{A}_{(i,\theta_i)} = A_i$ of *dominated actions* in $\tilde{G}$: $\tilde{D}_{(i,\theta_i)} = \left\{\tilde{a}_{(i,\theta_i)} \in \tilde{A}_{(i,\theta_i)},\ \text{s.t.,}\ U_{(i,\theta_i)}^{pes}\left(\tilde{a}_{(i,\theta_i)}, \pi_{-i}\right) \le U_{(i,\theta_i)}^{pes}(\pi)\right\}$.

This definition will be useful in the next proofs. Note first that due to the correspondence of utility functions in Π-games ad their SNF representation, we have:

$$\tilde{D}_{(i,\theta_i)} = \left\{ a_i \in A_i,\ \text{s.t.,}\ U_i^{pes}(a_i, \upsilon_{-i}, \theta_i) \leq U_i^{pes}(\upsilon, \theta_i) \right\}.$$

To prove the current Proposition, note first that if $\tilde{D}_{i,\theta_i} = \emptyset$, then $\upsilon' = \upsilon$. Algorithm 1 Does not modify the current strategy.

Then, for a given i, if $\tilde{D}_{i,\theta_i} \neq \emptyset$, then, by application of Proposition 3, we have, $\forall \upsilon, \theta_i$:

$$U_i^{pes}(\upsilon, \theta_i) = \min_{a_i \in \tilde{D}_{i,\theta_i}} 1 - \upsilon_i(a_i|\theta_i).$$

Now, if $\upsilon_i'(a_i|\theta_i) \leftarrow \max\left\{\alpha \in \Delta; \alpha < 1 - U_i^{pes}(\upsilon, \theta_i)\right\}$, we get

$$1 - \upsilon_i'(a_i|\theta_i) > U_i^{pes}(\upsilon, \theta_i), \forall a_i \in \tilde{D}_{i,\theta_i}.$$

Thus, $U_i^{pes}(\upsilon', \theta_i) = \min_{a_i \in \tilde{D}_{i,\theta_i}} 1 - \upsilon_i'(a_i|\theta_i) > U_i^{pes}(\upsilon, \theta_i)$.

In the case where $\upsilon_i(a_i|\theta_i) < 1 - U_i^{pes}(\upsilon, \theta_i)$, $\upsilon_i'(a_i|\theta_i) \leftarrow \upsilon_i(a_i|\theta_i)$. □

Proof of Proposition 8. Algorithm 2 performs iterated calls to function $\upsilon' \leftarrow \textsc{Improve}(G, \upsilon, i, \theta_i) \Leftrightarrow \pi' \leftarrow \textsc{Improve}\Big(\tilde{G}, \pi, (i, \theta_i)\Big)$.

By analogy with [6] it is easy to check that the number of calls to the Improve function is finite. Furthermore, it is bounded by $n \times t^n \times |\Delta| \times d$, thus is polynomial in the size of G. □

Proof of Proposition 9. Note that if calling $\pi' \leftarrow \textsc{Improve}\left(\tilde{G}, \pi, (i, \theta_i)\right)$ requires to compute $\tilde{G}$, which is of size exponential in that of G, on the other hand, computing $\upsilon' \leftarrow \textsc{Improve}(G, \upsilon, i, \theta_i)$ only requires a polynomial size *input*, $(G, \upsilon, i, \theta_i)$.

It then remains to check that the computation of $\upsilon' \leftarrow \textsc{Improve}(G, \upsilon, i, \theta_i)$ is polynomial in the size of G. To do so, recall that the Improve function defined in [6] relies on the computation of the set $\tilde{D}_{(i,\theta_i)} \subseteq \tilde{A}_{(i,\theta_i)} = A_i$ of *dominated actions* in $\tilde{G}$:
$\tilde{D}_{(i,\theta_i)} = \left\{ \tilde{a}_{(i,\theta_i)} \in \tilde{A}_{(i,\theta_i)},\ \text{s.t.,}\ U_{(i,\theta_i)}^{pes}\left(\tilde{a}_{(i,\theta_i)}, \pi_{-i}\right) \leq U_{(i,\theta_i)}^{pes}(\pi) \right\}$.
But, due to the correspondence of utility functions in both games, we have:
$\tilde{D}_{(i,\theta_i)} = \left\{ a_i \in A_i,\ \text{s.t.,}\ U_i^{pes}(a_i, \upsilon_{-i}, \theta_i) \leq U_i^{pes}(\upsilon, \theta_i) \right\}$.
The first expression takes polynomial space to compute since we do not actually need to store $\tilde{\mu}$ to compute $U_{(i,\theta_i)}^{pes}(\pi)$, for any fixed π. However, it takes exponential time.

On the other hand, $U_i^{pes}(\upsilon, \theta_i) = \min_{a_i \in A_i} \max\left(1 - \upsilon_i(a_i|\theta_i), U_i^{pes}(a_i, \upsilon_{-i}, \theta_i)\right)$ can be computed in time polynomial in the size of G and it is also the case for $U_i^{pes}(a_i, \upsilon_{-i}, \theta_i)$, of course.

Thus, the only step of the Improve function of [6] that could potentially take exponential time in the size of G (computation of $\tilde{D}_{(i,\theta_i)}$) can actually be performed in polynomial time.

Under the assumption that π is represented by a table of $|\Theta|$ lines, the complexity of computing $\{U_i^{pes}(a_i, \upsilon_{-i}, \theta_i)\}$ $\forall i \in N, \forall a_i \in A_i, \forall \theta_i \in \Theta_i$ is in $O(d^{(n-1)} \times t^{(n-1)})$ and the complexity of computing $U_i^{pes}(\upsilon, \theta_i)$ is in $O(d \times d^{(n-1)} \times t^{(n-1)}) = O(d^n \times t^{(n-1)})$.

Thus the whole complexity of Algorithm 1 is polynomial in the size of the Π-game: $O(d^{n+1} \times t^{(n-1)})$. □

Proof of Proposition 10. Since υ^* has been obtained after convergence of Algorithm 2, which can only make mixed strategies more specific at every step, it verifies:
$\upsilon^* = \textsc{Improve}(G, \upsilon, i, \theta_i), \forall i \in N, \forall \theta_i \in \Theta_i$. This implies that $\forall i \in N, \forall \theta_i \in \Theta_i, \forall \upsilon_i(\theta_i)$,
$U_i^{pes}(\upsilon^*, \theta_i) \geq U_i^{pes}(\upsilon_i(\theta_i).\upsilon_{-i}^*, \theta_i)$, that is, υ^* is a possibilistic mixed Nash Equilibrium. Furthermore, the definition of the Improve procedure ensures that, whenever it computes a new strategy of player i, it improves over every least-specific strategy of this player. It results that the fixed-point strategy obtained is a least specific equilibrium. □

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] O. Morgenstern, J. Von Neumann, Theory of Games and Economic Behavior, Princeton University Press, 1944.
[2] J.F. Nash, Non-cooperative games, Ann. Math. (1951) 286–295.
[3] J. Ouenniche, A. Boukouras, M. Rajabi, An ordinal game theory approach to the analysis and selection of partners in public-private partnership projects, J. Optim. Theory Appl. 169 (1) (2016) 314–343.

[4] J. Cruz, M.A. Simaan, Ordinal games and generalized Nash and Stackelberg solutions, J. Optim. Theory Appl. 107 (2) (2000) 205–222.
[5] C. Xu, Computation of noncooperative equilibria in ordinal games, Eur. J. Oper. Res. 122 (1) (2000) 115–122.
[6] N. Ben Amor, H. Fargier, R. Sabbadin, Equilibria in ordinal games: a framework based on possibility theory, in: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, 2017, pp. 105–111.
[7] J.C. Harsanyi, Games with incomplete information played by "Bayesian" players, Manag. Sci. 14 (3) (1967) 159–182.
[8] S. De Clercq, S. Schockaert, A. Nowé, M.D. Cock, Modelling incomplete information in Boolean games using possibilistic logic, Int. J. Approx. Reason. 93 (2018) 1–23.
[9] D. Dubois, H. Prade, Possibility Theory - an Approach to Computerized Processing of Uncertainty, Springer Science & Business Media, 1988.
[10] N. Ben Amor, H. Fargier, R. Sabbadin, M. Trabelsi, Possibilistic games with incomplete information, in: Proceedings of the 28th International Joint Conference on Artificial Intelligence, 2019, pp. 1544–1550.
[11] D. Dubois, H. Prade, Possibility theory as a basis for qualitative decision theory, in: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, 1995, pp. 1924–1932.
[12] D. Dubois, H. Prade, R. Sabbadin, Decision-theoretic foundations of qualitative possibility theory, Eur. J. Oper. Res. 128 (3) (2001) 459–478.
[13] L. Savage, The Foundations of Statistics, Wiley, 1954.
[14] I. Gilboa, D. Schmeidler, Maxmin expected utility with non-unique prior, J. Math. Econ. 18 (2) (1989) 141–153.
[15] D. Dubois, H. Prade, R. Sabbadin, Qualitative decision theory with Sugeno integrals, in: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, 1998, pp. 121–128.
[16] D. Dubois, H. Prade, The logical view of conditioning and its application to possibility and evidence theories, Int. J. Approx. Reason. 4 (1) (1990) 23–46.
[17] D. Dubois, L.F. del Cerro, A. Herzig, H. Prade, An ordinal view of independence with application to plausible reasoning, in: Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence, 1994, pp. 195–203.
[18] P. Walley, G.D. Cooman, Coherence of rules for defining conditional possibility, Int. J. Approx. Reason. 21 (1) (1999) 63–107.
[19] J.F. Nash, Equilibrium points in n-person games, Proc. Natl. Acad. Sci. USA 36 (1) (1950) 48–49.
[20] H. Hosni, E. Marchioni, Some notes on ordinal strategic interaction with possibilistic expectation (extended abstract), in: Proceedings of the 34th Linz Seminar on Fuzzy Set Theory, 2013, pp. 82–84.
[21] H. Hosni, E. Marchioni, Possibilistic randomisation in strategic-form games, Int. J. Approx. Reason. 114 (2019) 204–225.
[22] A. Brandenburger, Hierarchies of beliefs and common knowledge, J. Econ. Theory 59 (1) (1993) 189–198.
[23] R.J. Aumann, A. Brandenburger, Epistemic conditions for Nash equilibrium, Econometrica 63 (1995) 1161–1180.
[24] P. Battigalli, G. Bonanno, Recent results on belief, knowledge and the epistemic foundations of game theory, Res. Econ. 53 (2) (1999) 149–225.
[25] A. Brandenburger, Epistemic game theory: an overview, in: L.E. Blume, S.N. Durlauf (Eds.), The New Palgrave Dictionary of Economics, Macmillan, 2008, pp. 1790–1791.
[26] E. Dekel, M. Siniscalchi, Epistemic game theory, in: Handbook of Game Theory with Economic Applications, vol. 4, Elsevier, 2015, pp. 619–702.
[27] R. Fagin, J.Y. Halpern, Y. Moses, M.Y. Vardi, Common knowledge revisited, in: Proceedings of the 6th Conference on Theoretical Aspects of Rationality and Knowledge, 1996, pp. 283–298.
[28] R.B. Myerson, Comments on "games with incomplete information played by 'Bayesian' players, I-III Harsanyi's games with incoplete information", Manag. Sci. 50 (12-Supplement) (2004) 1818–1824.
[29] S. Ceppi, N. Gatti, N. Basilico, Computing Bayes-Nash equilibria through support enumeration methods in Bayesian two-player strategic-form games, in: Proceedings of the IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology, 2009, pp. 541–548.
[30] E. Nudelman, J. Wortman, Y. Shoham, K. Leyton-Brown, Run the GAMUT: a comprehensive approach to evaluating game-theoretic algorithms, in: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004, pp. 880–887.
[31] T. Sandholm, A. Gilpin, V. Conitzer, Mixed-integer programming methods for finding Nash equilibria, in: Proceedings of the Twentieth AAAI National Conference on Artificial Intelligence, 2005, pp. 495–501.
[32] R. Porter, E. Nudelman, Y. Shoham, Simple search methods for finding a Nash equilibrium, Games Econ. Behav. 63 (2) (2008) 642–662.
[33] I. I. CPLEX, V12. 1: user's manual for CPLEX, Int. Bus. Mach. Corp. 46 (53) (2009) 157.
[34] E.B. Yanovskaya, Equilibrium points in polymatrix games, Lith. Math. J. (1968) 381–384.
[35] M.J. Kearns, M.L. Littman, S.P. Singh, Graphical models for game theory, in: Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, 2001, pp. 253–260.
[36] C.H. Papadimitriou, T. Roughgarden, Computing correlated equilibria in multi-player games, J. ACM 55 (3) (2008) 14.
[37] J. Filar, K. Vrieze, Competitive Markov Decision Processes, Springer, 1997.
[38] M.L. Puterman, Markov Decision Processes, John Wiley and Sons, 1994.
[39] S. Sorin, A First Course on Zero-Sum Repeated Games, vol. 37, Springer Science & Business Media, 2002.
[40] A.R. Cassandra, L.P. Kaelbling, M.L. Littman, Acting optimally in partially observable stochastic domains, in: Proceedings of the Twelfth AAAI Conference on Artificial Intelligence, vol. 94, 1994, pp. 1023–1028.
[41] R. Sabbadin, H. Fargier, J. Lang, Towards qualitative approaches to multi-stage decision making, Int. J. Approx. Reason. 19 (3–4) (1998) 441–471.
[42] R. Sabbadin, Possibilistic Markov decision processes, Eng. Appl. Artif. Intell. 14 (3) (2001) 287–300.
[43] V. Conitzer, T. Sandholm, Complexity results about Nash equilibria, in: Proceedings of the 18th International Joint Conference on Artificial Intelligence, 2002, pp. 765–771.