



HAL
open science

Possibilistic Games with Incomplete Information (IJCAI 2019)

Nahla Ben Amor, H el ene Fargier, R egis Sabbadin, Meriem Trabelsi

► **To cite this version:**

Nahla Ben Amor, H el ene Fargier, R egis Sabbadin, Meriem Trabelsi. Possibilistic Games with Incomplete Information (IJCAI 2019). Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI 2019), Aug 2019, Macao, China. pp.1544-1550, 10.24963/ijcai.2019/214 . hal-02879715v2

HAL Id: hal-02879715

<https://hal.science/hal-02879715v2>

Submitted on 28 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destin ee au d ep ot et  a la diffusion de documents scientifiques de niveau recherche, publi es ou non,  emanant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv es.

Possibilistic Games with Incomplete Information

Nahla Ben Amor¹, H el ene Fargier², R egis Sabbadin³ and Meriem Trabelsi^{1,2}

¹ISG-Tunis, Universit e de Tunis

²IRIT-CNRS, Universit e de Toulouse

³INRA-MIAT, Universit e de Toulouse

nahla.benamor@gmx.fr, helene.fargier@irit.fr, regis.sabbadin@inra.fr, merieme.trabelsi@gmail.com

Abstract

Bayesian games offer a suitable framework for games where the utility degrees are additive. This approach does nevertheless not apply to ordinal games, where the utility degrees do not capture more than a ranking, nor to situations of a decision under qualitative uncertainty. This paper proposes a representation framework for ordinal games under possibilistic incomplete information and extends the fundamental notion of Nash equilibrium (NE) to this framework. We show that deciding whether a NE exists is a difficult problem (NP-hard) and propose a Mixed Integer Linear Programming encoding. Experiments on variants of the GAMUT problems confirm the feasibility of this approach.

1 Introduction

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

The preferences of the agents among the outcomes are captured by utility functions. The term “payoff” is often used to designate their utility - this suits to problems where the satisfaction can be expressed on a cardinal scale, typically a monetary scale. But there are instances where the assumption of cardinality can be questioned (see, e.g., [Ouenniche *et al.*, 2016; Cruz and Simaan, 2000]). Moreover, several fundamental notions of game theory, e.g., *Nash Equilibrium* and *Dominance*, do not require such an assumption - they are basically *ordinal* notions. A cardinal notion is called in two cases at least: (i) when the game is repeated (and outcomes are “collected” and assumed to be additive), and (ii) when the outcomes depend on a probabilistic event (e.g., in the prisoner dilemma, if the verdict does not only depend on the confession of the prisoners, but also on the result of the trial). To capture such incomplete information situations, *Bayesian games* have been proposed by Harsanyi [1967]. This approach does nevertheless not apply to ordinal games, where the utility degrees capture no more than a ranking, nor to situations of a decision under qualitative uncertainty.

Following the seminal work of [De Clercq *et al.*, 2018] on possibilistic Boolean games, we propose to use possibility theory to model qualitative uncertainty in *ordinal* games. Unlike [De Clercq *et al.*, 2018], we do not develop a complex language. We stay at the semantic level and work out the idea of possibilistic games, extending the notion of Nash Equilibrium [Nash, 1950] to such games and showing how to transform them into classical games.

Section 2 presents the notions on which the paper relies. Section 3 and 4 propose a possibilistic model for ordinal games with incomplete information and generalizes the notion of a Nash Equilibrium (NE) to this framework. Section 5 states the complexity of the problem of the existence of a NE and proposes a Mixed Integer Linear Programming encoding of the problem. Experiments are reported in Section 6. The proofs are omitted for the sake of brevity; they can be found at [Ben Amor *et al.*, 2019].

2 Background

2.1 Normal Form Games

A *strategic game* or *normal form game* is classically defined as follows [Morgenstern and Von Neumann, 1944]:

Definition 1 (Normal Form Game) A normal form game is a triple $G = \langle N, A, \mu \rangle$, where:

- $N = \{1, \dots, n\}$ is a finite set of players;
- $A = \times_{i \in N} A_i$, 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 $\mu_i : A \mapsto \Delta$.

A joint (or “pure”) strategy is a vector $a \in A$ which specifies an action for each player and $\mu_i(a)$ is the utility of agent i for the joint strategy a ; it takes values in a finite totally ordered scale Δ (the higher $\mu_i(a)$, the more satisfied agent i). Typically $\Delta = [0, 1]$, but any ordered scale may be used.

A joint strategy is a Nash Equilibrium (NE) if no player can improve her utility by unilaterally changing her strategy:

Definition 2 (Nash Equilibrium) A strategy profile $a \in A$ is a NE, iff $\forall i \in N, \forall a'_i \in A_i: \mu_i(a_i.a_{-i}) \geq \mu_i(a'_i.a_{-i})$, where $\forall a \in A, a_i$ is the action of i in $a, a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_n) \in A_{-i} = \times_{j \neq i} A_j$ is its restriction to all the players but i and “.” denotes the concatenation (e.g., $\forall (a'_i, a_{-i}), a'_i.a_{-i} = (a_1, \dots, a_{i-1}, a'_i, a_{i+1}, \dots, a_n)$).

2.2 Possibility Theory

The basic building block in possibility theory [Dubois and Prade, 1988] is the notion of *possibility distribution*. A possibility distribution π is a mapping from a set of states S to an ordered scale Δ (here to $\Delta = [0, 1]$). $\pi(s) = 1$ means that 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: there is at least one totally possible state (a s such that $\pi(s) = 1$). From π , one can compute the possibility $\Pi(E)$ and the necessity $N(E)$ of any event $E \subseteq S$: $\Pi(E) = \sup_{s \in E} \pi(s)$ evaluates to what extent E is *consistent* with the knowledge represented by π while $N(E) = 1 - \Pi(\bar{E}) = 1 - \sup_{s \notin E} \pi(s)$ corresponds to the extent to which $\neg E$ is inconsistent and thus evaluates at which level E is certainly implied by the knowledge.

Considering qualitative (possibilistic) problems of decision under uncertainty, where each consequence is evaluated by a utility function $\mu : S \mapsto \Delta$, [Dubois and Prade, 1995; Dubois *et al.*, 2001] have proposed qualitative utilities, called pessimistic and optimistic utility degrees, as counterparts to von Neumann and Morgenstern's [1944] expected utility:

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

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

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. U^{opt} estimates to what extent it is possible that μ reaches a good utility. In the framework of competitive game theory, U^{opt} is rather unnatural (too adventurous), while U^{pes} conveniently models the cautious decision makers. As many other models this models makes a commensurability assumption between the utility levels and the levels of likelihood. This assumption is common to all the models which considers that the agent's preference relation is complete and transitive (this is the case in many models, be they qualitative or quantitative: expected utility [Savage, 1954; Morgenstern and Von Neumann, 1944], multi-prior non expected utility [Gilboa and Schmeidler, 1989], Sugeno integrals [Dubois *et al.*, 1998], etc.).

In the following, we shall finally use the notion of conditional possibility measure proposed by [Hisdal, 1978] in order to stay in the pure ordinal framework (for a discussion about the alternative definitions of conditional possibilities, see [Walley and De Cooman, 1999]). For any events E and $F \subseteq S$ the possibility of E given F is:

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

3 Possibilistic Games with Incomplete Information

The classical framework of games under perfect information assumes that each player knows everything about the game: the players, the actions available to every player, all their utilities, etc. What is not known by a player is the decision 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, a landlord may ignore if a potential tenant is honest and has a moral incentive to pay his loan or not, 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 appointment).

Bayesian games [Harsanyi, 1967] offer a suitable framework for cardinal games under incomplete knowledge; they 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 a decision under qualitative uncertainty. We propose in the following a model for (ordinal) games under possibilistic information.

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 an agent i depends not only on the actions of the other players but also on the real world. Originally, none of the players know which is the real world, but all of them share a common knowledge about the game. Just before playing, each agent i will receive from *Nature* some information $\tau_i(s)$ about the real world - τ_i maps any $s \in S$ to an element θ_i of a set Θ_i called the set of "types" of i . After having observed $\tau_i(s)$, agent i knows more about the real game, but several games may still be plausible. The agent, thus, conditions her knowledge on $\tau_i(s)$. The use of a space type is classic in game theory - it is the basis of Bayesian games. The set of possible types of i , Θ_i , can be considered as the local state space for the player i and $\Theta = \Theta_1 \times \dots \times \Theta_n$ as the effective global state space. 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 his introspective mental states (see [Brandenburger, 1993] [J. Aumann and Brandenburger, 1995] [Battigalli and Bonanno, 1999], [Brandenburger, 2008], [Dekel and Siniscalchi, 2015] for the links between belief states and types, and more generally for further developments about epistemic game theory.

In the present paper we follow Harsanyi's approach, based on types, extending it to qualitative games with possibilistic uncertainty. Thus, the action of an agent i depends only on the information $\theta_i \in \Theta_i$ she receives. A joint strategy $\sigma = (\sigma_1, \dots, \sigma_n)$ is thus a tuple of functions σ_i that map each possible information (each "type" θ_i in Θ_i) to action in A_i ¹.

Definition 3 A strategy is a vector $\sigma = (\sigma_1, \dots, \sigma_n)$ of functions $\sigma_i : \Theta_i \rightarrow A_i$.

$\sigma_i(\theta_i)$ specifies the action that player i will execute when receiving the private information θ_i . Given a strategy σ and a configuration of the players types $\theta \in \Theta = \Theta_1 \times \dots \times \Theta_n$, $\sigma(\theta) = (\sigma_1(\theta_1), \dots, \sigma_n(\theta_n))$ denotes the joint action (the element of A) prescribed by strategy σ when θ occurs. In the

¹In this paper, we consider *pure* strategies only: in our qualitative context, the notion of a mixed (Bayesian) strategy, which relies on additive probabilities, is not pertinent.

following, Σ_i denotes the set of all functions from Θ_i to A_i and $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ the set of all joint strategies.

In the Bayesian framework, the game under incomplete information is most of the time described on the space of types (underlying worlds are omitted). In the present work, we proceed in the same way and define a possibilistic game with incomplete information as follows:

Definition 4 (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 $\{1, \dots, n\}$;
- $A = \times_{i \in N} A_i$ 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 gathers all the configurations of types, i.e. vectors $\theta = (\theta_1, \dots, \theta_n)$. θ_i denotes the type of i in θ and $\theta_{-i} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_n)$. Likewise $\Theta_{-i} = \times_{j \neq i} \Theta_j$;
- $\pi: \Theta \rightarrow \Delta$ is a joint possibility distribution over the combinations of types;
- $\mu = \{(\mu_i)_{i \in N}\}$ where $\mu_i: A \times \Theta \rightarrow \Delta$ 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: i does not know θ , but only θ_i (θ_i is the private knowledge of agent i). $\pi(\cdot | \theta_i)$ represents the knowledge that agent i has when learning θ_i . On the other hand, utility $\mu_i(a, \theta)$ (utility of joint action a for player i when learning θ) will be obtained once all players have played their action and revealed their type (thus μ_i depends on the whole θ and not only on θ_i).

A Π -game 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 game $G^\theta = \langle N, A, \{(\mu_i^\theta)_{i \in N}\} \rangle$ where:

$$\mu_i^\theta(\cdot) = \mu_i(\cdot, \theta) \text{ and } \pi(G^\theta) = \max_{\theta' \text{ s.t. } G^{\theta'} = G^\theta} \pi(\theta') \quad (4)$$

Example 1 (A landlord and a tenant) A landlord L can rent a house to a tenant T . She can Accept (A) or Refuse (R), i.e., $A_L = \{A, R\}$ and the tenant can Pay the landlord (P) or not (\bar{P}), i.e., $A_T = \{P, \bar{P}\}$. 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 she receives information indicating that the tenant is honest (IH) or dishonest (ID), i.e., $\Theta_L = \{IH, ID\}$. Generally, the tenant is honest and the landlord is informed that the tenant is honest. 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 is captured by the following common π : $\pi(H, IH) = 1$ $\pi(D, ID) = 0.75$ $\pi(D, IH) = 0.5$ $\pi(H, ID) = 0.25$. The projections of π on Θ_T and Θ_L yield: $\pi(H) = 1$, $\pi(D) = 0.75$, $\pi(IH) = 1$ $\pi(ID) = 0.75$. For each strategy, each player has a utility in $\{0, 0.25, 0.5, 0.75, 1\}$ (from 0 meaning “completely dissatisfied” to 1 meaning “completely satisfied”, 0.5 meaning “neutral”, etc.) When the landlord refuses to rent, the utility of the tenant is equal to 0 if honest and equal to 0.25 if dishonest. When the landlord accepts,

		landlord				
		IH		ID		
t e n a n t	H	A	R	A	R	
		P	1,1	0,0.5	P	1,1
	\bar{P}	0.25,0	0,0.5	\bar{P}	0.25,0.25	0,0.5
			$\pi(H, IH) = 1$		$\pi(H, ID) = 0.25$	
D	A	R	A	R		
	P	0.25,1	0.25,0.5	P	0.25,1	0.25,0.5
\bar{P}	1,0	0.25,0.5	\bar{P}	1,0.25	0.25,0.5	
		$\pi(D, IH) = 0.5$		$\pi(D, ID) = 0.75$		

Table 1. A Π -game with 4 types combinations

the tenant’s utility of \bar{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 refusing. When accepting, 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 1) the possibility degrees of which are:

$$\begin{aligned} \pi(G^{H,IH}) &= \pi(H, IH) = 1 & \pi(G^{H,ID}) &= \pi(H, ID) = 0.25 \\ \pi(G^{D,IH}) &= \pi(D, IH) = 0.5 & \pi(G^{D,ID}) &= \pi(D, ID) = 0.75 \end{aligned}$$

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

Proposition 1 Any classical normal form game is a Π -game with $|\Theta_i| = 1, \forall i \in N$.

Π -games can be related to the framework proposed by [De Clercq *et al.*, 2018] as a semantics for possibilistic Boolean games. In possibilistic Boolean games, the knowledge of each player can be captured by a distinct possibility distribution π_i (the knowledge is not common) and the players do not receive any private information (there is no type). [De Clercq *et al.*, 2018] propose to compute the possibility and the necessity of a given profile of actions being a NE in the usual sense (Definition 2): each agent computes these indices according to her own knowledge. The authors then consider the problem from the external point of view of an observer who proceeds to a fusion of these distributions and deduces a unique $\pi = \min_{i \in N} \pi_i$ over types of players.

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

$$\begin{aligned} \Pi(a \text{ is a NE}) &= \max_{\theta, a \text{ is NE for } G^\theta} \pi(\theta) \\ N(a \text{ is a NE}) &= 1 - \max_{\theta, a \text{ is not a NE for } G^\theta} \pi(\theta) \end{aligned}$$

Contrarily to the framework of [De Clercq *et al.*, 2018], ours can handle the knowledge that each agent i has when receiving some private information θ_i . Then, the posterior necessity (resp. possibility) that a is a NE can be different from one agent to another:

$$\begin{aligned} \Pi_i(a \text{ is a NE} | \theta_i) &= \min_{\theta_{-i}, a \text{ is a NE for } G^{\theta_i, \theta_{-i}}} \pi(\theta_{-i} | \theta_i) \\ N_i(a \text{ is a NE} | \theta_i) &= 1 - \max_{\theta_{-i}, a \text{ is not a NE for } G^{\theta_i, \theta_{-i}}} \pi(\theta_{-i} | \theta_i) \end{aligned}$$

4 Nash Equilibrium

One problem with the concepts proposed in [De Clercq *et al.*, 2018] above is that they are based on the definition of strategies as profiles of actions, as usual in a classical normal form game. This definition does not take uncertainty about the types of the players into account: it supposes that the full θ is eventually known by all the players, and thus that the utility of each agent is known by each other before the actions are played. In other words, the definition of the above necessity and possibility indices only suits games which are presently incomplete but will eventually become completely known.

In the present work, an agent i only knows $\pi(\cdot|\theta_i)$ when receiving θ_i and deciding what to play. As in any incomplete information game, strategy σ specifies an action for each θ_i . σ is a NE when i has no incentive to change her choice of action for θ_i , given $\pi(\cdot|\theta_i)$ and σ . Using possibilistic qualitative decision theory, the utility of $\sigma(\theta_i)$ is evaluated, using Equation (1) (we suppose that the agent is cautious).

Definition 5 (Utility of an action Utility of joint strategy)

The utility of a_i for player i of type θ_i in the context of σ_{-i} is

$$U_i^{pes}(a_i, \sigma_{-i}, \theta_i) = \min_{\theta_{-i} \in \Theta_{-i}} \max(1 - \pi(\theta_{-i}|\theta_i), \mu_i(a_i, \sigma(\theta_{-i}), \theta_i, \theta_{-i})) \quad (5)$$

The utility of strategy σ to agent i of type θ_i is:

$$U_i^{pes}(\sigma, \theta_i) = U_i^{pes}(\sigma_i(\theta_i), \sigma_{-i}, \theta_i) \quad (6)$$

Of course $U_i^{pes}(\sigma_i(\theta_i), \sigma_{-i}, \theta_i)$ 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 $U_i^{pes}(a_i, \sigma_{-i}, \theta_i)$ (again, when considering a joint strategy σ , the uncertainty of i only bears on the types of the other players. What they plan to do depending on their type is known, prescribed by σ).

Definition 6 (Best Response in a Π -game)

$$BR_i(\sigma_{-i}, \theta_i) = \operatorname{argmax}_{a_i \in A_i} U_i^{pes}(a_i, \sigma_{-i}, \theta_i) \quad (7)$$

A (*Nash*) *Equilibrium* (NE) is a joint strategy from which no player i will deviate unilaterally knowing σ_{-i} . In the possibilistic context, we thus set the following definition:

Definition 7 (Nash Equilibrium in a Π -game) σ is a Π -Nash equilibrium (Π -NE) iff $\forall i \in N, \forall \theta_i \in \Theta_i, \forall a'_i \in A_i$,

$$U_i^{pes}(a_i, \sigma_{-i}, \theta_i) \geq U_i^{pes}(a'_i, \sigma_{-i}, \theta_i)$$

This definition generalizes Definition 2, which is recovered when $|\Theta| = 1$ (only one possible type per player). As a consequence, a pure equilibrium may not exist for a Π -game (a classical normal form game does not always admit a NE).

Example 2 (cont' Example 1) Let us first compute the conditional possibility distribution (Equation (3)):

$$\begin{aligned} \pi(H|IH) &= \pi(IH|H) = 1 & \pi(D|IH) &= \pi(IH|D) = 0.5 \\ \pi(H|ID) &= \pi(ID|H) = 0.25 & \pi(D|ID) &= \pi(ID|D) = 1 \end{aligned}$$

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 ($\pi(H|IH) = 1$) but it can not be excluded that is not the case ($\pi(D|IH) = 0.5$).

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

$$\begin{aligned} \sigma_T^*(H) &= P, \quad \sigma_T^*(D) = \bar{P} \quad \sigma_L^*(IH) = A, \quad \sigma_L^*(ID) = R. \\ \text{The pessimistic utility of the honest tenant is: } &U_T^{pes}(\sigma^*, H) \\ &= \min \max \left\{ 1 - \pi(IH|H), \mu_T((P,A), (H,IH)) \right\}, \\ &\quad \max \left\{ 1 - \pi(ID|H), \mu_T((P,R), (H,ID)) \right\} \\ &= \min \max \left\{ 1 - 1, 1 \right\}, \max \left\{ 1 - 0.25, 0 \right\} = 0.75. \end{aligned}$$

In the same way, we get: $U_T^{pes}(\sigma^*, D) = 0.25$; $U_L^{pes}(\sigma^*, IH) = 0.5$; $U_L^{pes}(\sigma^*, ID) = 0.5$.

The best response to σ_L^* for the tenant of type H is P ; indeed: $U_T^{pes}(P, \sigma_L^*, H) = 0.75$, $U_T^{pes}(\bar{P}, \sigma_L^*, H) = 0.25$.

The two responses \bar{P} and P are equally interesting for tenant of type D : $U_T^{pes}(P, \sigma_L^*, D) = 0.25$, and $U_T^{pes}(\bar{P}, \sigma_L^*, D) = 0.25$.

It can also be checked that strategy σ^* is a Π -NE; the pessimistic utility of the H tenant playing P (resp. \bar{P}) is equal to 0.75 (resp. 0.25). So, in this case, the tenant prefers P to \bar{P} . The dishonest tenant is indifferent between P and \bar{P} ($U_T(P, \sigma_L^*, D) = U_T(\bar{P}, \sigma_L^*, D) = 0.25$). Also, an IH landlord prefers A to R ($U_L(A, \sigma_T^*, IH) \geq U_L(R, \sigma_T^*, IH)$). Unlike the IH landlord, the ID landlord prefers R to A (utility 0.5 instead of 0.25).

We finally show that any incomplete game can be transformed into an equivalent normal form game with *complete information*, the pure equilibria of which are in bijection with the pure equilibria of the Π -game. This representation result is a qualitative counterpart of Harsanyi's about the representation of Bayesian games by normal form games under complete information. The idea is to consider as many players as the number of pairs (i, θ_i) , each (i, θ_i) having A_i as a set of available actions. A joint strategy in this (classical) normal form game then corresponds to a strategy σ of the Π -game.

Definition 8 (Transformed Game) Given a Π -game $G = \langle N, A, \Theta, \mu \rangle$ $\tilde{G} = \langle \tilde{N}, \tilde{A}, \tilde{\mu} \rangle$ is the normal form game where:

- $\tilde{N} = \{(i, \theta_i), i \in N, \theta_i \in \Theta_i\}$;
- $\tilde{A}_{(i, \theta_i)} = A_i$ for all $(i, \theta_i) \in \tilde{N}$;
- $\tilde{\mu}_{(i, \theta_i)}(a) = U_i^{pes}(\sigma^a, \theta_i) \forall a, i, \theta_i$, where σ^a is the strategy of the original game defined by: $\sigma_i^a(\theta_i) = a_{(i, \theta_i)}$.

Example 3 (Cont' Example 1) The transformed game \tilde{G} has 4 players: $\tilde{N} = \{(T, H), (T, D), (L, IH), (L, ID)\}$, with $\tilde{A}_{(T, H)} = \tilde{A}_{(T, D)} = \{P, \bar{P}\}$, $\tilde{A}_{(L, IH)} = \tilde{A}_{(L, ID)} = \{A, R\}$. The joint action $a = (P, \bar{P}, A, R)$ of \tilde{G} corresponds to σ^* in G . The utilities of a in \tilde{G} are:

$$\begin{aligned} \tilde{\mu}_{(T, H)}(a) &= U_T^{pes}(\sigma^*, H) = 0.75 & \tilde{\mu}_{(T, D)}(a) &= U_T^{pes}(\sigma^*, D) = 0.25 \\ \tilde{\mu}_{(L, IH)}(a) &= U_L^{pes}(\sigma^*, IH) = 0.5 & \tilde{\mu}_{(L, ID)}(a) &= U_L^{pes}(\sigma^*, ID) = 0.5 \end{aligned}$$

Let us check that a is a NE in \tilde{G} ; the utility of player (T, H) playing P (resp. \bar{P}) is equal to 0.75 (resp. 0.25): she prefers to play P than \bar{P} . For player (T, D) , $\tilde{\mu}_{(T, D)}(a) = U_T^{pes}(\sigma, D) = 0.25$ and $\tilde{\mu}_{(T, D)}(P, a_{-i}) = U_T^{pes}(P, \sigma_{-i}, D) = 0.25$. So, she is indifferent between P and \bar{P} , etc.

The following proposition holds.

Proposition 2 a is a pure NE in $\tilde{G} = \langle \tilde{N}, \tilde{A}, \tilde{\mu} \rangle$ iff σ_a is a Π -NE in $G = \langle N, A, \Theta, \mu \rangle$.

As to the complexity of the transformation, consider that the games are extensively represented, by tables. The transformed game contains $\sum_{i=1,n} |\Theta_i|$ utility functions of size $\prod_{i=1,n} |A_i|^{|\Theta_i|}$. If we write for simplification purpose that $|A_i| = a$ and $|\Theta_i| = d \forall i \in N$ (same number of types and the same number of actions for all players), this means that a game containing n utility functions of a size $(a \cdot d)^n$ is transformed into a game containing $n \cdot d$ utility functions of size $(a^n)^d$. Except when the number of types is very small, the transformation does not provide a convenient way to solve the game. Proposition 2 is, as in the Bayesian case, more a representation result than the premise of a solving tool.

5 Solving the Possibilistic NE Problem

Finding a Π -NE in a Π -game is a difficult problem. Indeed:

Definition 9 (Π -NE problem) *Given a Π -game $G = \langle N, A, \Theta, \pi, \mu \rangle$, the Π -NE problem consists in determining whether there exists a Π -Nash equilibrium.*

Proposition 3 Π -NE is NP-Hard, even in symmetric 2-player games where π corresponds to total ignorance, i.e., $\forall \theta_1 \in \Theta_1$ and $\forall \theta_2 \in \Theta_2, \pi(\theta_1, \theta_2) = 1$.

Taking advantage of the efficiency of modern solvers, we propose a Mixed Integer Linear Programming (MILP) formulation for finding Π -NE in Π -game (we follow in this line opened by [Ceppi *et al.*, 2009] 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 : $\forall i = 1, n, \forall a_i \in A_i, \forall \theta_i \in \Theta_i : \sigma_{i,a_i,\theta_i} \in \{0, 1\}$. The utilities are encoded by continuous variables: U_{i,a_i,θ_i} is the utility (according to σ_{-i}) of agent i if action a_i is chosen for type θ_i (i.e., if $\sigma_i(\theta_i) = a_i$):

$$\forall i = 1, n, \forall a_i \in A_i, \forall \theta_i \in \Theta_i : U_{i,a_i,\theta_i} \in [0, 1]$$

We will also use the following to constrain the U_{i,a_i,θ_i} to be equal to the $\min_{\theta_{-i} \in \Theta_{-i}}$ of $\max(1 - \pi(\theta_{-i}|\theta_i), \mu_i(a, \theta))$ (and not only lower than the \min):

$$\forall i = 1, n, \forall a_i \in A_i, \forall \theta \in \Theta : M_{i,a_i,\theta} \in \{0, 1\}$$

Let $MAX_{i,a,\theta} = \max(1 - \pi(\theta_{-i}|\theta_i), \mu_i(a, \theta))$, the constraints are the following:

$$\forall i = 1, n, \forall \theta_i \in \Theta_i, \sum_{a_i \in A_i} \sigma_{i,a_i,\theta_i} = 1 \quad (8)$$

$$\forall i = 1, n, \forall a_i, a'_i \in A_i \text{ s.t } a_i \neq a'_i, \forall \theta_i \in \Theta_i, U_{i,a_i,\theta_i} - U_{i,a'_i,\theta_i} \geq \sigma_{i,a_i,\theta_i} - 1 \quad (9)$$

$$\forall i = 1, n, \forall a \in A, \forall \theta \in \Theta, U_{i,a_i,\theta_i} \leq MAX_{i,a,\theta} + \sum_{j=1,n, j \neq i} (1 - \sigma_{j,a_j,\theta_j}) \quad (10)$$

$$U_{i,a_i,\theta_i} + M_{i,a_i,\theta} + \sum_{j=1,n, j \neq i} (1 - \sigma_{j,a_j,\theta_j}) \geq MAX_{i,a,\theta} \quad (11)$$

$$\forall i = 1, n, \forall a_i \in A_i, \forall \theta_i \in \Theta_i, \sum_{\theta_{-i} \in \Theta_{-i}} (1 - M_{i,a_i,\theta_i,\theta_{-i}}) = 1 \quad (12)$$

Constraints (8) ensure that the strategy σ searched for specifies exactly one action per type, for each player i .

Constraints (9) require that σ is a Π -NE: 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 actions 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 (10) implement Definition (5). 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 $MAX_{i,a_i,\sigma(\theta_{-i}),\theta}$. Indeed, for any profile of action a that does not correspond to what is prescribed by σ , $\sum_{j \neq i} (1 - \sigma_{j,a_j,\theta_j}) \geq 1$ and 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 $\sum_{j \neq i} (1 - \sigma_{j,a_j,\theta_j}) = 0$: the constraint becomes $U_{i,a_i,\theta_i} \leq MAX_{i,a,\theta}$.

Constraints (11) and (12) ensure that U_{i,a_i,θ_i} is equal to the \min of the $MAX_{i,a,\theta}$: when a_{-i} does not correspond to σ_{-i} , $\sum_{j \neq i} (1 - \sigma_{j,a_j,\theta_j})$ is at least equal to 1 and the constraint is always satisfied. When a_{-i} does correspond to σ_{-i} , the \sum is equal to 0 and does not annihilate the constraint. The \min is reached if $U_{i,a_i,\theta_i} = MAX_{i,a,\theta}$. $M_{i,a_i,\theta_i,\theta_{-i}}$ is another way to annihilate the constraint, and reaching the \min means not annihilating the constraint for one θ_{-i} - hence Equation (12).

The above problem is linear (the \max operator which appears in constraints (10) and (11) deals with constants only). The above formulation does not lead to a combinatorial explosion of the space required. Recall that the size of the original problem is $size = n \cdot |\Theta| \cdot |A| + |\Theta|$. Let us denote a (resp. d) the number of actions (resp. types) of each player. The MILP formulation contains:

- $n \cdot d \cdot a$ continuous variables U_{p,a_p,θ_p} , $n \cdot d \cdot a$ Boolean variables σ_{p,a_p,θ_p} and $n \cdot a \cdot |\Theta|$ Boolean variables $M_{p,a_p,\theta}$;
- $n \cdot d$ constraints (8), each involving a variables;
- $n \cdot d \cdot a \cdot (a-1)$ constraints (9), each involving 3 variables;
- $n \cdot |\Theta| \cdot |A|$ constraints (10) each involving n variables;
- $n \cdot |\Theta| \cdot |A|$ constraints (11) each involving $n+1$ variables;
- $n \cdot a \cdot d$ constraints (12) each involving $\frac{|\Theta|}{d}$ variables.

The size of the MILP encoding is thus in $O(n^2 \cdot |\Theta| \cdot |A|)$ (polynomial in the size of the original $size$).

6 Experimental Study

To conduct our experimental study, we developed a novel Π -game generator based on GAMUT [Nudelman *et al.*, 2004], a suite of classical normal form games (with complete information) generators (following the approach of [Ceppi *et al.*, 2009] for the generation of Bayesian games). More precisely, to generate a Π -game version of a GAMUT problem (e.g., the *Covariant game*), we need as inputs: the number of degrees in Δ , the number n of players, the class of game and if necessary the number of actions $|A_i|$ and of types $|\Theta_i|$ for each player i . 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). The utilities $\mu_i(a, \theta)$ is simply the

Game class	Θ_i	A_i	2	3	4	5	6	7	8	9	10	
Covariant Game	2	2	.00	.02	.05	.10	.30	.70	1.44	5.72	6.94	TG
		3	.30	13.34	11.72	17.08	13.3	19.54	24.74	26.22	25.62	MILP
	3	2	.01	.24	1.76	13.70	42.40	69.96	38.76	103.72	160.66	TG
		3	.02	14.48	10.26	10.92	14.82	17.08	19.16	20.42	23.6	MILP
Dispersion Game	2	2	.00	.05	.10	.13	.31	.74	5.62	1.36	7.02	TG
		3	.30	.50	.60	1.00	5.38	1.44	6.78	2.34	7.72	MILP
	3	2	.16	.32	1.76	15.76	39.12	69.94	29.64	88.46	126.38	TG
		3	.48	.64	1.2	1.68	2.4	11.28	7.12	11.84	13.04	MILP
Travelers Dilemma Game	2	2	.20	6.9	135.28	600.74	478.22	818.52	1.1k	1.5k	1.8k	TG
		3	3.24	1.28	4.02	4.84	10.08	10.82	10.52	13.2	20.26	MILP
	3	2	.01	.11	.13	.17	.20	.60	.96	1.52	6.94	TG
		3	.23	.32	.41	.90	.56	1.44	7.38	2.94	10.74	MILP
Travelers Dilemma Game	3	2	.05	.32	1.84	10.22	42.38	70.1	30.38	60.18	173.82	TG
		3	.32	.64	.88	8.48	11.96	11.68	21.08	22.36	33.52	MILP
	4	2	.09	3.84	50.62	546.46	1.5k	1.5k	1.6k	2.8k	3.4k	TG
		3	.03	2.36	5.34	9.92	14.82	15.14	19.42	21.32	23.00	MILP


 Table 2. Avg. memory usage of MILP & $T\tilde{G}$ (MB), $n = 2$

Game class	Θ_i	A_i	2	3	4	5	6	7	8	9	10	
Covariant Game	2	2	.5	.4	.8	1.9	4.3	10.3	35.7	61.3	144.1	TG
		3	12.8	850.2	633.3	1.6k	2k	2.2k	2.4k	1.9k	2.1k	MILP
	3	2	.3	.34	43	261.8	947.2	2.4k	5.9k	14.7k	22.6k	TG
		3	21.8	1.1k	1.5k	2.1k	1.8k	1.3k	1.8k	1.7k	1.6k	MILP
Dispersion Game	2	2	1.9	139.9	2.4k	3.9k	85.6k	318.5k	-	-	-	TG
		3	23.9	853.6	1.3k	885.8	1.7k	1.3k	1.6k	867.8	1.1k	MILP
	3	2	.4	.3	.7	1.8	4.6	8.8	28.3	64.8	116.8	TG
		3	11.8	44.5	53.5	81.8	78.3	80.3	115.5	96	110.2	MILP
Travelers Dilemma Game	2	2	.3	.0	35.8	220.9	813	2.3k	5.7k	14.3k	27.4k	TG
		3	62.5	65.1	46.6	63.1	50.1	60.1	63.5	97.6	109.2	MILP
	3	2	1.9	116.1	1.3k	9.7k	41.5k	113.2k	295.6k	478.1k	-	TG
		3	87.5	58.7	34.3	35.3	54.3	65.0	88.0	98.9	145.6	MILP
Travelers Dilemma Game	2	2	.2	.2	.6	1.5	3.7	8.4	18.9	40.5	76	TG
		3	21.2	59.3	78.4	89.1	76.6	95.8	103.2	93.9	182.8	MILP
	3	2	.3	2.7	20.3	104.2	365.4	932.9	2.3k	6.3k	13.1k	TG
		3	51.0	71.8	59.5	643.4	578.8	217.8	689.2	1k	1.6k	MILP
Travelers Dilemma Game	4	2	1.5	47.4	535.7	3.9k	18.9k	77.4k	271.8k	-	-	TG
		3	18.8	26.6	257.4	525.9	794.4	647.3	1.3k	632.8	670.5	MILP

 Table 3. Avg. execution time of MILP & $T\tilde{G}$ (ms), $n = 2$

utility of a for player i in game G^θ . In our experiments, we vary the number of players from 2 to 6, the number of types from 2 to 8 and the number of actions from 2 to 15. For each combination of parameters, we have generated 50 instances and measured the time necessary to get a Π -NE (or a negative result). We present in the following results of 3 game classes: *Covariant games*, *Dispersion games* and *Travelers Dilemma game*. All experiments were conducted on an Intel Xeon E5540 processor and 64GB RAM workstation. We used CPLEX [CPLEX, 2009] as a MILP solver. We also implemented the transformation of the Π -game as a normal form game ($T\tilde{G}$) in Java 8. This method, which is exponential in time and space, cannot be considered as a solving method, and this is supported by the experimental results. The implementation of the $T\tilde{G}$ and MILP solver are available online [Ben Amor *et al.*, 2019]. In our evaluation, we bounded the execution time to 10 minutes as in [Sandholm *et al.*, 2005; Porter *et al.*, 2008] experiments.

Table 2 and Table 3 present, respectively, the average memory (in MB) required to decide whether the problem admits a Π -NE or not and the average of execution time (in milliseconds) needed to find one Π -NE (best results are in bold). We vary the $|A_i|$ from 2 to 10 and $|\Theta_i|$ from 2 to 4 for player i while fixing $n = 2$. “-” mentions that the execution time exceeds 10 minutes. Table 2 confirms that $T\tilde{G}$ is a very naive (exponential) way and that cannot scale up contrarily to MILP which requires less memory, e.g., $T\tilde{G}$ requires more than 3,4MB when MILP needs just 23MB to solve *Travelers Dilemma* game with 2 players, 10 actions and 4 types per


 Figure 1. Avg. execution time of MILP (s), $n = 2$

	Θ_i	n	2	3	4	5	6	7
Covariant Game	2	2	13	50	244	720	4k	49.5k
		3	22	587	1k	8k	26.5k	-
	3	2	24	1.5k	3.6k	50.1k	-	-
		3	33	2.3k	30.2k	213.6k	-	-
Dispersion Game	2	2	12	49	2.5k	6.8k	82.9k	33.7k
		3	63	350	9.5k	32.2k	14.8k	-
	3	2	47	788	14.8k	42.9k	-	-
		3	48	480	19.5k	107.1k	-	-
Travelers Dilemma Game	2	2	21	73	35	1k	4.4k	35.5k
		3	19	224	648	5.4k	155.5k	-
	3	2	19	139	1.5k	40.1k	-	-
		3	51	173	2.3k	185.6k	-	-

 Table 4. Avg. execution time of MILP (ms), $|A_i| = 2$ $i = 1..7$

player. The non-scalability of $T\tilde{G}$ is also observable in Table 3 showing that this latter cannot be used in practice. We also deepened our study of MILP by varying its complexity parameters. More precisely, we first vary the number of actions from 2 to 15 and the number of types from 2 to 8 with two players ($n = 2$) (see Figure 1). Then, we vary n from 2 to 7 and the number of types from 2 to 5 while fixing the number of actions to 2 (see Table 4). Globally MILP results confirm the feasibility of the qualitative approach of incomplete game advocated by this paper. Results show that MILP has almost the same behavior with different games and that the number of players highly impacts the execution time which is consistent with the theoretical complexity since adding a player directly increases $|\Theta|$ and $|A|$.

7 Conclusion

The main contributions of this paper are threefold: first, it proposes possibilistic games as new representation framework for (ordinal) games under possibilistic incomplete information and extends the standard notion of Nash Equilibrium to such games. Second, it proposes a transformation from possibilistic games to normal form games with complete information and shows the equivalence between pure Nash Equilibria in both frameworks. Third, it proves that finding a pure Nash Equilibrium in a Π -game is a NP-hard problem and proposes a Mixed-integer linear constraints system model to find a possibilistic Nash Equilibrium if any. The experimental study we led shows that, except when the number of types and actions are very small, the MILP approach outperforms the transformation to normal form game approach. As future work, we first plan to define possibilistic mixed NE in possibilistic games and design polynomial time solution algorithms to find them, in the line of [Ben Amor *et al.*, 2017]. Another interesting line of research is to explore concisely expressed possibilistic games, in the line of Bayesian action-graph games [Jiang and Leyton-Brown, 2010], for example.

References

- [Battigalli and Bonanno, 1999] Pierpaolo Battigalli and Giacomo Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics*, 53(2):149–225, 1999.
- [Ben Amor *et al.*, 2017] Nahla Ben Amor, H el ene Fargier, and R egis Sabbadin. Equilibria in ordinal games: A framework based on possibility theory. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 105–111, 2017.
- [Ben Amor *et al.*, 2019] Nahla Ben Amor, H el ene Fargier, R egis Sabbadin, and Meriem Trabelsi. The possibilistic games page. <https://www.irit.fr/~Helene.Fargier/PossibilisticGames.html>, 2019.
- [Brandenburger, 1993] Adam Brandenburger. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1):189 – 198, 1993.
- [Brandenburger, 2008] Adam Brandenburger. *Epistemic Game Theory: An Overview*, pages 1790–1791. Palgrave Macmillan UK, London, 2008.
- [Ceppi *et al.*, 2009] Sofia Ceppi, Nicola Gatti, and Nicola Basilico. Computing bayes-nash equilibria through support enumeration methods in bayesian two-player strategic-form games. In *Proceedings of the International Joint Conferences on Web Intelligence and Intelligent Agent Technologies*, volume 2, pages 541–548. IEEE, 2009.
- [CPLEX, 2009] IBM ILOG CPLEX. V12. 1: User’s manual for CPLEX. *International Business Machines Corporation*, 46(53):157, 2009.
- [Cruz and Simaan, 2000] JB Cruz and Marwan A Simaan. Ordinal games and generalized nash and stackelberg solutions. *Journal of Optimization Theory and Applications*, 107(2):205–222, 2000.
- [De Clercq *et al.*, 2018] Sofie De Clercq, Steven Schockaert, Ann Now e, and Martine De Cock. Modelling incomplete information in boolean games using possibilistic logic. *International Journal of Approximate Reasoning*, 93:1–23, 2018.
- [Dekel and Siniscalchi, 2015] Eddie Dekel and Marciano Siniscalchi. *Epistemic Game Theory*, chapter 12, page Elsevier. 619 - 702, 2015.
- [Dubois and Prade, 1988] Didier Dubois and Henri Prade. *Possibility theory: An approach to computerized processing of uncertainty*. Plenum Press, New York, 1988.
- [Dubois and Prade, 1995] Didier Dubois and Henri Prade. Possibility theory as a basis for qualitative decision theory. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1924–1930. Springer, 1995.
- [Dubois *et al.*, 1998] Didier Dubois, Henri Prade, and R egis Sabbadin. Qualitative decision theory with sugeno integrals. In *Proceedings of Uncertainty in Artificial Intelligence 1998*, pages 121–128, 1998.
- [Dubois *et al.*, 2001] Didier Dubois, Henri Prade, and R egis Sabbadin. Decision-theoretic foundations of qualitative possibility theory. *European Journal of Operational Research*, 128(3):459–478, 2001.
- [Gilboa and Schmeidler, 1989] Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141 – 153, 1989.
- [Harsanyi, 1967] John C Harsanyi. Games with incomplete information played by “Bayesian” players. *Management science*, 14(3):159–182, 1967.
- [Hisdal, 1978] Ellen Hisdal. Conditional possibilities independence and noninteraction. *Fuzzy Sets and Systems*, 1(4):283–297, 1978.
- [J. Aumann and Brandenburger, 1995] Robert J. Aumann and Adam Brandenburger. Epistemic conditions for nash equilibrium. *Econometrica*, 63:1161–80, 02 1995.
- [Jiang and Leyton-Brown, 2010] Albert X Jiang and Kevin Leyton-Brown. Bayesian action-graph games. In *Advances in Neural Information Processing Systems*, pages 991–999, 2010.
- [Morgenstern and Von Neumann, 1944] Oskar Morgenstern and John Von Neumann. *Theory of games and economic behavior*. Princeton university press, 1944.
- [Nash, 1950] John F Nash. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- [Nudelman *et al.*, 2004] Eugene Nudelman, Jennifer Wortman, Yoav Shoham, and Kevin Leyton-Brown. Run the gamut: A comprehensive approach to evaluating game-theoretic algorithms. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 880–887. IEEE Computer Society, 2004.
- [Ouenniche *et al.*, 2016] Jamal Ouenniche, Aristotelis Boukouras, and Mohammad Rajabi. An ordinal game theory approach to the analysis and selection of partners in public–private partnership projects. *Journal of Optimization Theory and Applications*, 169(1):314–343, 2016.
- [Porter *et al.*, 2008] Ryan Porter, Eugene Nudelman, and Yoav Shoham. Simple search methods for finding a nash equilibrium. *Games and Economic Behavior*, 63(2):642–662, 2008.
- [Sandholm *et al.*, 2005] Tuomas Sandholm, Andrew Gilpin, and Vincent Conitzer. Mixed-integer programming methods for finding nash equilibria. In *Proceedings of the Association for the Advancement of Artificial Intelligence 2005*, volume 20, pages 495–501, 2005.
- [Savage, 1954] Leonard Savage. *The Foundations of Statistics*. Wiley, New York, 1954.
- [Walley and De Cooman, 1999] Peter Walley and Gert De Cooman. Coherence of rules for defining conditional possibility. *International Journal of Approximate Reasoning*, 21(1):63–107, 1999.