



HAL
open science

Games of Incomplete Information: A Framework Based on Belief Functions

Hélène Fargier, Érik Martin-Dorel, Pierre Pomeret-Coquot

► **To cite this version:**

Hélène Fargier, Érik Martin-Dorel, Pierre Pomeret-Coquot. Games of Incomplete Information: A Framework Based on Belief Functions. 16th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU 2021), Sep 2021, Prague, Czech Republic. pp.328-341, 10.1007/978-3-030-86772-0_24 . hal-03382397

HAL Id: hal-03382397

<https://hal.science/hal-03382397v1>

Submitted on 15 Apr 2022

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 destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Games of Incomplete Information: a Framework Based on Belief Functions^{*}

Hélène Fargier¹, Érik Martin-Dorel², and Pierre Pomeret-Coquot²

¹CNRS, IRIT; ²University Paul Sabatier, IRIT; University of Toulouse
{helene.fargier,erik.martin-dorel,pierre.pomeret}@irit.fr

Abstract. This paper proposes a model for incomplete games where the knowledge of the players is represented by a Dempster-Shafer belief function. Beyond an extension of the classical definitions, it shows such a game can be transformed into an equivalent hypergraphical complete game (without uncertainty), thus generalizing Howson and Rosenthal's theorem to the framework of belief functions and to any number of players. The complexity of this transformation is finally studied and shown to be polynomial in the degree of k -additivity of the mass function.

Keywords: Game theory · Incomplete games · Belief functions

The final publication is available at Springer via:

https://dx.doi.org/10.1007/978-3-030-86772-0_24

© 2021 Springer Nature Switzerland AG

1 Introduction

Game theory [21,25] proposes a powerful framework to capture decision problems involving several agents. In non-cooperative games of complete information, the players do not coordinate their actions but each of them knows everything about the game: the players, their available actions and all their utilities. This assumption of complete knowledge cannot always be satisfied. In the real world indeed, players are not so well informed and have only limited knowledge about the game. This is why Bayesian games of incomplete information have been proposed [16]. Nevertheless, the Bayesian hypothesis is strong, and requires a good knowledge of the environment. For instance, in case of ignorance, the Bayesian way is to suppose equiprobability, but this can lead to a model that does not fit with the agents' behavior (e.g. see Ellsberg's paradox [9]).

In the present paper, we propose a new kind of game of incomplete information, which we call *credal game*. Agents have a partial knowledge, represented by a Dempster-Shafer belief function [5,30], and cardinal utilities, but do not necessarily make the equiprobability assumption. The underlying decision rule is the

^{*} Pierre Pomeret-Coquot and Helene Fargier have benefited from the AI Interdisciplinary Institute ANITI. ANITI is funded by the French "Investing for the Future – PIA3" program under the Grant agreement n°ANR-19-PI3A-0004.

Choquet integral based on the *Bel* measure [4], in order to capture the agents' aversion for ambiguity [13,12,22]. We then follow the line defined by Howson and Rosenthal [17] who have shown that any 2-player Bayesian game can be transformed into a complete knowledge *polymatrix game* [33]. In this paper, we show that such a transformation is possible for credal games, and for any number of agents, producing a *hypergraphical game* [26]. An important consequence of this result is that the algorithmics developed for hypergraphical games [32,3] can be reused for the search of Nash equilibria in credal games.

2 Background and motivations

To illustrate and motivate our work, we will use the following example inspired by the *murder of Mr. Jones* [31], where the suspects are *Peter, Paul and Mary*.

Example 1 (Peter, Quentin and Rose). Two agents, named Agent 1 and Agent 2, are independently looking for a business association, with either Peter (*P*), Quentin (*Q*), or Rose (*R*). The point is that a crime has been committed, for which these three people are suspected. Several testimonies, not very reliable, allowed to estimate that there is 50% of chance that the culprit is a man (*P* or *Q*), and 50% of chance that it is a woman (*R*).

As to the interest of the associations, making the deal with an innocent leads to a payoff of \$6k (to be shared between the people making the deal), while associating with a guilty person produces no payoff (\$0k). Moreover, Agent 1 is investigating about *P* and will know whether he is guilty or not. Similarly, Agent 2 will know whether *R* is guilty before making the decision.

The Bayesian approach is not relevant here. Indeed, if Agent 1 learns that *P* is innocent, the probability of guilt should become 1/2 for *Q* and 1/2 for *R*. However, in a purely Bayesian view, equiprobability would be applied and the prior probability of guilt would be 1/4 for *P* and 1/4 for *Q*. Then, after conditioning, Agent 1 would get a probability of 1/3 for *Q* and 2/3 for *R*.

2.1 Dempster-Shafer's theory of evidence

Let us first look at the epistemic aspect of the problem. The prior knowledge is simply that $P(\{P, Q\}) = P(\{M\}) = \frac{1}{2}$, and nothing more. The kind of knowledge at work here is well captured in Dempster-Shafer's theory of evidence, that does not restrict probability assignments to elements of the frame of discernment:

Definition 1 (Mass function). *A mass function for a frame of discernment Ω (or "bpa" for basic probability assignment) is a function $m : 2^\Omega \rightarrow [0, 1]$ such that $m(\emptyset) = 0$ and $\sum_{A \subseteq \Omega} m(A) = 1$.*

A set with a nonzero mass is called a *focal element* and the set of focal elements is denoted \mathcal{S}_m . Two dual measures on 2^Ω derive from m :

$$Bel(A) = \sum_{B \in \mathcal{S}_m, B \subseteq A} m(B) \quad \text{and} \quad Pl(A) = \sum_{B \in \mathcal{S}_m, B \cap A \neq \emptyset} m(B).$$

$Bel(A)$ (resp. $Pl(A)$) estimates to what extent A is implied by (resp. is compatible with) the knowledge captured by m .

Probabilities are belief functions where the focal elements are then singletons – where m is "1-additive". k -additivity is more generally defined as follows:

Definition 2 (k -additivity). *A mass function is k -additive if all its focal elements are at most of size k , i.e., $\forall B \in \mathcal{S}_m, |B| \leq k$.*

Even if the agents share the same prior knowledge, as in our example, they may acquire different pieces of information and thus have a different posterior knowledge. In the example, Agent 1 eventually learns whether P is guilty or not while Agent 2 will acquire information about R . Each agent thus revises his/her knowledge on the basis on the information C she learns (e.g. R for Agent 2). States where C is false are considered as impossible by the agent, so she modifies the initial belief function in such a way that $Pl(\bar{C}) = 0$: the conditioning at work here is Dempster's rule [5] (see also [7,1] for more details about the conditioning).

Definition 3 (Dempster conditioning). *For any nonempty $A, C \subseteq \Omega$,*

$$m_{|C}(A) := K_C \cdot \sum_{B \in \mathcal{S}_m, C \cap B = A} m(B),$$

where $K_C = 1 / \sum_{B \in \mathcal{S}_m, B \cap C \neq \emptyset} m(B)$ is a normalization factor.

2.2 Decision making with belief functions

Let us now consider belief functions in a (single-agent) decision making context. Following Savage's modelling of decision making under uncertainty [28], a decision (or "action") is a function $a : \Omega \rightarrow X$ where Ω is the set of possible states, as previously, and X is the set of possible outcomes. The preferences of an agent are represented by an utility function $u : X \rightarrow \mathbb{R}$. When the knowledge about Ω is captured by a belief function, the discrete Choquet integral [4] based on the Bel measure is classically advocated because of its ability to capture the agents' aversion for ambiguity [13,12,22]:

Definition 4 (Discrete Choquet integral).

Let $\Lambda(a) = \{\lambda_1 \leq \dots \leq \lambda_{|\Lambda(a)|}\}$ be the set of utility values reached by an action a , labelled by increasing order, and $E_{\lambda_i}(a) = \{\omega \mid u(a(\omega)) \geq \lambda_i\}$ denote the set of worlds for which the utility of action a is at least λ_i . The discrete Choquet utility value (or CEU) of a is :

$$CEU(a) = \lambda_1 + \sum_{i=2}^{|\Lambda(a)|} (\lambda_i - \lambda_{i-1}) \times Bel(E_{\lambda_i}(a)).$$

The CEU has a simple expression based on the mass function:

$$CEU(a) = \sum_{B \in \mathcal{S}_m} m(B) \times \min_{\omega \in B} u(a(\omega)).$$

2.3 Game theory

A simultaneous game of complete information models a situation where several agents make a decision (the term “action” is rather used in game theory) without coordination with the other agents – the final utility of each agent depending on the actions chosen by all agents.

Definition 5 (Complete game). *A simultaneous game of complete information (also called complete game) is a tuple $G = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ where:*

- $N = \{1, \dots, n\}$ is a finite set of agents (or “players”),
- A_i is the set of actions of Agent i ; $A := \prod_{i \in N} A_i$ contains all the possible combinations of actions or “profiles”,
- $u_i : A \rightarrow \mathbb{R}$ is the utility function of Agent i .

A mixed strategy for player i is a probability distribution on A_i . The strategy is said to be pure when only one action receives a non-zero probability.

A pure (resp. mixed) joint strategy (or strategy profile) is a vector $p = (p_1, \dots, p_n)$ which specifies a pure (resp. mixed) strategy for each player.

A mixed strategy is classically interpreted as distributions the players use to randomly choose among available actions in order to avoid being predictable – this is especially useful in repeated games. An alternative view is to consider that each p_i represents the knowledge that the other agents have about i ’s decision.

In the following, we will use the following notations: for any vector $v = (v_1, \dots, v_n)$ in some product domain $V = \prod_{i \in N} V_i$ and for any $e \subseteq N$, v_e is the restriction of v to e and $V_e = \prod_{i \in e} V_i$. By abuse of notation, we write v_i for $v_{\{i\}}$. For any i , $-i$ denotes the set $N \setminus \{i\}$, i.e. $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n) \in V_{-i} = \prod_{j \neq i} V_j$. Thus, v_{-i} is the restriction of v to all players but i . Finally, “.” denotes the concatenation, e.g., $v'_i.v_{-i} = (v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n)$. Hence $a = a_i.a_{-i}$ belongs to A and given two profiles $a, a' \in A$, $a'_i.a_{-i}$ denotes the profile a where a_i is replaced with a'_i .

Because the strategies can be randomized, the global utility for a player of a joint mixed strategy p is defined as the expected utility of u_i according to the probability distribution it induces over A (Obviously, when the strategy is pure, EU_i is equal to the utility value given by u_i):

Definition 6 (Utility of a strategy). *Given a joint strategy p in a complete game $(N, (A_i)_{i \in N}, (u_i)_{i \in N})$, the expected utility of player i is defined by:*

$$EU_i(p) = \sum_{a \in A} \left(\prod_{i \in N} p_i(a_i) \right) \times u_i(a).$$

Among the profiles of interest, Nash equilibria are emphasized, i.e., profiles in which no player can increase his/her utility by changing his/her own strategy.

Definition 7 (Nash equilibrium [25]). *A strategy profile p is a Nash equilibrium iff for all $i \in N$, there exists no p'_i such that $u_i(p'_i.p_{-i}) > u_i(p)$.*

A pure strategy profile p is a Nash equilibrium iff for all $i \in N$, there exists no pure strategy p'_i such that $u_i(p'_i.p_{-i}) > u_i(p)$.

When the utility functions are described explicitly, G is said to be in standard normal form (SNF). SNF representations become spatially costly when the number of players increases ($O(n\alpha^n)$ for a game with n players and α actions per player). More succinct forms have been proposed, that suit cases where utility functions can be decomposed as a sum of smaller utility functions – namely hypergraphical games [26] and polymatrix games [33].

Definition 8 (Hypergraphical game). *A hypergraphical game is a tuple $G = (N, E, (A_i)_{i \in N}, (u_i^e)_{e \in E, i \in e})$ where N is a set of players, $E = \{e_1, \dots, e_m\}$ is a multiset of subsets of N ((N, E) is an hypergraph) and for each $e \in E$, $(e, (A_i)_{i \in e}, (u_i^e)_{i \in e})$ is a classical standard normal form game. The global utility of Agent i is the sum of i 's local utilities: $u_i(a) = \sum_{e \in E} u_i^e(a_e)$. Polymatrix games are hypergraphical games with 2-player local games.*

This framework assumes that each player knows everything about the game: the players, the actions available to each player, all their utilities for each combination of actions, etc. The assumption of complete knowledge cannot always be satisfied. In the real world indeed, players have only a limited knowledge about the outcomes of their strategies – the final outcomes may depend on an ill-known event (in our example, the payoff for making the deal with one of P , Q , or R depends on whether they are guilty or innocent).

Harsanyi [16] proposed games of incomplete information as a way to capture such situations (see also [24], for more details). A game of incomplete information can be first understood as a set of possible classical games (of complete information) – one for each possible world $\omega \in \Omega$. Players don't know exactly which world is the real one, but may have some knowledge about it. Just before playing, each player i will receive some information $\tau_i(\omega^*)$ about the real world ω^* . τ_i maps any world to an element θ_i of a set Θ_i called the set of “types” of Agent i . After having observed $\tau_i(\omega^*)$, Agent i knows more about the real game, but several games may still be plausible. The player then conditions his/her knowledge on $\tau_i(\omega^*)$ and decides which action to play. Notice that the different agents may receive different pieces of information and thus have a different posterior knowledge. The question is then, for each player, to determine a strategy (either an action, or a probabilistic strategy) for each of his/her possible types.

Harsanyi has shown that such games can be described on the space of types $\Theta = \Theta_1 \times \dots \times \Theta_n$ (the underlying worlds are omitted). The idea of Harsanyi when defining types is that this concept can encapsulate every piece of information agents may have access to. It includes the agent-observable world status, but also their beliefs on other agents and their introspective mental states.

3 Credal Games

Bayesian games are games of incomplete information where prior knowledge is captured by a probability measure. To capture problems where the Bayesian assumption is not obeyed (as in our motivating example), we propose in this paper the framework of credal games:

Definition 9 (Credal game). A credal game G is defined as a tuple composed of $(N, (A_i)_{i \in N}, (\Theta_i)_{i \in N}, (u_i)_{i \in N}, m)$ where:

- $N = \{1, \dots, n\}$ is a finite set of players,
- A_i is the set of actions of player i ; $A = \prod_{i \in N} A_i$ denotes the set of all action profiles,
- Θ_i is the set of types of player i ; let $\Theta = \prod_{i \in N} \theta_i$,
- $m : 2^\Theta \rightarrow [0, 1]$ is the mass function describing the common prior knowledge,
- $u_i : A \times \Theta \rightarrow \mathbb{R}$ is the utility function of Agent i .

G is said to be in standard normal form iff the utility functions u_i and the mass function m are given in extenso.

Following Harsanyi’s approach of incomplete games, we consider the “*ex interim*” setting where each player plans a strategy for each of the types he/she can receive, ideally a strategy which is a best response to that of the other players. We thus adopt the definition of strategy proposed by Harsanyi’s in the general context of incomplete games:

Definition 10 (Pure and mixed strategies [16]). A pure (resp. mixed) strategy for player i in a credal game is a function ρ_i which maps each “type” $\theta_i \in \Theta_i$ to an action of (resp. a probability over) A_i .

A pure (resp. mixed) joint strategy is a vector $p = (p_1, \dots, p_n)$ which specifies a pure (resp. mixed) strategy for each player.

$\rho(\theta) = (\rho_1(\theta_1), \dots, \rho_n(\theta_n))$ denotes the profile which will be played if the configuration of types is θ .

Let us first consider pure strategies. In the *ex interim* approach of incomplete games, Agent i may consider the strategy ρ_{-i} planned by the other players for each of their types. When receiving her type θ_i , he/she revises his/her knowledge – in a credal game, his/her posterior knowledge over the joint type configuration is $m_{|\theta_i}$. According to the definition of the Choquet expected utility, the utility of a pure strategy profile for Agent i of type θ_i , shall thus be defined as:

Definition 11 (Choquet Expected Utility of a pure strategy profile).

The utility of a pure strategy profile $\rho = (\rho_1, \dots, \rho_n)$, for Agent i of type θ_i , is: $CEU_{(i, \theta_i)}(\rho) = \sum_{B \in \mathcal{S}_{m_{|\theta_i}}} m_{|\theta_i}(B) \times \min_{\theta' \in B} u_i(\rho(\theta'), \theta')$.

Let us now consider mixed strategies. For each configuration θ , $\rho(\theta)$ defines a probability distribution over A : the probability that a is played when agents have types θ is equal to $\prod_{i \in N} \rho_i(\theta_i)(a_i)$. If we now consider all the θ ’s, we get a bpa m^ρ over $A \times \Theta$. To any focal element $B \in \mathcal{S}_m$ and any pure strategy profile $\sigma : \Theta \rightarrow A$ corresponds an element $\sigma.B := \{(\sigma(\theta), \theta) \mid \theta \in B\}$ of mass $m^\rho(\sigma.B) = m(B) \times \prod_{\theta \in B} \prod_{i \in N} \rho_i(\theta_i)(\sigma_i(\theta_i))$ – $\sigma.B$ is focal iff B is focal and σ is possible according to ρ . Finally, Agent i receiving type θ_i conditions his/her knowledge which becomes $m_{|\theta_i}^\rho$. Hence the following definition of the utility of a mixed strategy profile:

Definition 12 (Choquet Expected Utility of a mixed strategy profile).

The utility of a mixed strategy profile $\rho = (\rho_1, \dots, \rho_n)$, for player i of type θ_i , is:

$$CEU_{(i,\theta_i)}(\rho) = \sum_{\sigma.B \in \mathcal{S}_{m_{|\theta_i}^\rho}} \left(m_{|\theta_i}(B) \times \prod_{\theta' \in B} \prod_{j \in N} \rho_j(\theta'_j)(\sigma_j(\theta'_j)) \right) \times \min_{\theta' \in B} u_i(\sigma(\theta'), \theta')$$

Obviously, Definition 12 leads to Definition 11 when ρ is a pure strategy profile. Now, recall that a strategy profile is a Nash equilibrium if no player can improve unilaterally his/her utility. This concept straightforwardly extends to credal games:

Definition 13 (Nash equilibrium). A mixed (resp. pure) strategy profile ρ is a Nash equilibrium for CEU iff, whatever i , there exists no mixed (resp. pure) strategy ρ'_i such that for any θ_i , $CEU_{(i,\theta_i)}(\rho'_i, \rho_{-i}) > CEU_{(i,\theta_i)}(\rho)$.

Example Our running example (see Example 1) is captured by the credal game $G = (N, (A_i)_{i \in N}(\Theta_i)_{i \in N}, (u_i)_{i \in N}, m)$ where:

- $N = \{1, 2\}$;
- $A_1 = \{P_1, Q_1, R_1\}$, $A_2 = \{P_2, Q_2, R_2\}$ (each agent chooses an associate).
- $\Theta_1 = \{P, \bar{P}\}$, $\Theta_2 = \{R, \bar{R}\}$ (Agent 1 investigates on Peter, Agent 2 investigates on Rose).
- $m : 2^\Theta \rightarrow [0, 1]$ has two focal elements: $m(\{\bar{P}, R\}) = 1/2$ (the murderer is a woman, thus necessarily Rose – in this case Agent 1 will learn \bar{P} and Agent 2 will learn R) and $m(\{(P, \bar{R}), (\bar{P}, R)\}) = 1/2$ (the murderer is a man: Agent 2 necessarily learns \bar{R} but Agent 1 can learn either \bar{P} – which happens when Quentin is the murderer – or P – Peter is the murderer).
- Making a deal with a murderer has a utility value of 0, making a deal with an innocent leads to a utility of $\frac{6}{2} = 3$, unless the other agent approaches the same associate, in which case each agent receives $\frac{6}{3} = 2$. The utility functions are summarized below (Table 1). Null values (in gray) are given for the case where $\theta = (P, R)$ (both R and P are guilty) which is not a possible world.

Let ρ be the pure strategy where Agent 1 makes the deal with R when learning that P is guilty and with P otherwise, and Agent 2 joins Q when learning that R is guilty and R otherwise: $\rho_1(P) = R_1$, $\rho_1(\bar{P}) = P_1$, $\rho_2(R) = Q_2$ and $\rho_2(\bar{R}) = R_2$.

- Consider Agent 1 receiving type P : the conditioned bpa, $m_{|P}$, has only one focal element $\{(P, \bar{R})\}$, $K_P = 1/\frac{1}{2}$ and $m_{|P}(\{(P, \bar{R})\}) = 1$. In short, Agent 1 knows that P is guilty and R is not. In the only possible configuration, (P, \bar{R}) , ρ prescribes $\rho_1(P) = R_1$ for Agent 1 and $\rho_2(\bar{R}) = R_2$ for Agent 2. Then $CEU_{(1,P)}(\rho) = m_{|P}(\{(P, \bar{R})\}) \times u_1((R_1, R_2), (P, \bar{R})) = 1 \times 2 = 2$.
- Consider now Agent 1 receiving \bar{P} : his/her revised knowledge, $m_{|\bar{P}}$, has two focal elements, $\{\bar{P}, R\}$ and $\{\bar{P}, \bar{R}\}$ (each with probability $\frac{1}{2}$, thus $K_{\bar{P}} = 1$). The strategy prescribes $\rho(\bar{P}) = P_1$ for Agent 1, who doesn't know whether Agent 2 learns R (and plays $\rho(R) = Q_2$) or \bar{R} (and plays $\rho(\bar{R}) = R_2$). Hence $CEU_{(1,\bar{P})}(\rho) = \frac{1}{2} \times u_1((P_1, R_2), (\bar{P}, \bar{R})) + \frac{1}{2} \times u_1((P_1, Q_2), (\bar{P}, R)) = 3$.

		$\theta_2 = \bar{\mathbf{R}}$			$\theta_2 = \mathbf{R}$		
		P_2	Q_2	R_2	P_2	Q_2	R_2
$\theta_1 = \mathbf{P}$	P_1	(0, 0)	(0, 3)	(0, 3)	P_1	(0, 0)	(0, 0)
	Q_1	(3, 0)	(2, 2)	(3, 3)	Q_1	(0, 0)	(0, 0)
	R_1	(3, 0)	(3, 3)	(2, 2)	R_1	(0, 0)	(0, 0)
$\theta_1 = \bar{\mathbf{P}}$	P_1	(2, 2)	(3, 0)	(3, 3)	P_1	(2, 2)	(3, 3)
	Q_1	(0, 3)	(0, 0)	(0, 3)	Q_1	(3, 3)	(2, 2)
	R_1	(3, 3)	(3, 0)	(2, 2)	R_1	(0, 3)	(0, 3)

Table 1. Example 1: Utility matrices for each configuration of the types

- Similarly, the bpa of Agent 2 receiving \mathbf{R} , $m_{|\mathbf{R}}$, has only one focal element, $\{(\bar{\mathbf{P}}, \mathbf{R})\}$ (thus $K_{\mathbf{R}} = 1/\frac{1}{2}$) in which ρ prescribes P_1 for Agent 1 and Q_2 for Agent 2. Then $CEU_{(2,\mathbf{R})}(\rho) = 1 \times u_2((P_1, Q_2), (\bar{\mathbf{P}}, \mathbf{R})) = 1 \times 3 = 3$.
- Finally, the bpa of Agent 2 receiving $\bar{\mathbf{R}}$, $m_{|\bar{\mathbf{R}}}$, has one focal element, $\{(\bar{\mathbf{P}}, \bar{\mathbf{R}}), (\mathbf{P}, \bar{\mathbf{R}})\}$ and $K_{\bar{\mathbf{R}}} = 1/\frac{1}{2}$. Agent 2 does not know whether Agent 1 receives $\bar{\mathbf{P}}$ or \mathbf{P} . Since ρ prescribes Agent 1 to play P_1 in the first case, R_1 in the second one, and prescribes Agent 2 to play R_2 in both cases, $CEU_{(2,\bar{\mathbf{R}})}(\rho) = 1 \times \min [u_2((P_1, R_2), (\bar{\mathbf{P}}, \bar{\mathbf{R}})), u_2((R_1, R_2), (\mathbf{P}, \bar{\mathbf{R}}))] = 1 \times \min(3, 2) = 2$.

In this strategy, Agent 1 does not give the best possible response to Agent 2's strategy: when learning that P is guilty, he/she plays R_1 while knowing that in this case Agent 2 learns $\bar{\mathbf{R}}$ and thus plays R_2 . Let Agent 1 modify his/her strategy and play Q_1 when learning $\bar{\mathbf{P}}$ – hence the strategy ρ' :

$$\rho'_1(\mathbf{P}) = Q_1, \rho'_1(\bar{\mathbf{P}}) = P_1, \rho'_2(\mathbf{R}) = Q_2, \rho'_2(\bar{\mathbf{R}}) = R_2$$

- $CEU_{(1,\mathbf{P})}(\rho') = K_{\mathbf{P}} \times u_1((Q_1, R_1), (\mathbf{P}, \bar{\mathbf{R}})) = 1 \times 3 = 3$,
- $CEU_{(1,\bar{\mathbf{P}})}(\rho') = K_{\bar{\mathbf{P}}} \times u_1((P_1, R_2), (\bar{\mathbf{P}}, \bar{\mathbf{R}})) + K_{\bar{\mathbf{P}}} \times u_1((P_1, Q_2), (\bar{\mathbf{P}}, \mathbf{R})) = 3$,
- $CEU_{(2,\mathbf{R})}(\rho') = K_{\mathbf{R}} \times u_2((P_1, R_2), (\bar{\mathbf{P}}, \mathbf{R})) = 1 \times 3 = 3$,
- $CEU_{(2,\bar{\mathbf{R}})}(\rho') = K_{\bar{\mathbf{R}}} \times \min (u_2((P_1, R_2), (\bar{\mathbf{P}}, \bar{\mathbf{R}})), u_2((Q_1, R_2), (\mathbf{P}, \bar{\mathbf{R}}))) = 3$.

It can be checked in ρ' each player has his/her maximal possible utility (§3k)
– no player has an incentive to change: ρ' is a pure Nash equilibrium.

4 From credal games to complete games

One of the most prominent results about Bayesian games is Howson's and Rosenthal's theorem [17]: any 2-player Bayesian game can be transformed into a (complete information) polymatrix game equivalent to the original one. This result is important from the computational point of view since it provides 2-player Bayesian games with practical resolution tools: to solve a 2-player Bayesian games, it is enough to use this theorem and to solve the resulting polymatrix game by one of algorithms proposed for such games [32,3]. In the sequel, we generalize this theorem to credal games and extend it to any number of players.

4.1 The direct transform

A first idea is to define from a credal game G , an hypergraphical game \tilde{G} , the vertices (players) of which are pairs (i, θ_i) with action set A_i – to each pure strategy ρ of G corresponds a unique pure strategy $\tilde{\rho}$ of \tilde{G} and conversely – we call $\tilde{\rho}$ the Selten transform of ρ :¹²

Definition 14 (Selten transform of a pure strategy). *For any pure strategy ρ of G , the Selten transform of ρ is the vector $\tilde{\rho}$ defined by $\tilde{\rho}_{(i, \theta_i)} = \rho_i(\theta_i)$.*

The local games correspond to the focal elements of m . Roughly, (i, θ_i) plays in the local game corresponding to the focal element B if the type θ_i is plausible for B – technically, if there exists $\theta' \in B$ such that $\theta'_i = \theta_i$. In this local game, (i, θ_i) obtains a local utility $K_{|\theta_i} \cdot m(B) \cdot \min_{\theta' \in B, \theta'_i = \theta_i} u_i(\rho(\theta'), \theta')$.

Given a profile of actions $\tilde{\rho}$, and a player (i, θ_i) , the hypergraphical game sums these local utilities over all the focal elements for which θ_i is plausible. Hence the global utility for (i, θ_i) is equal to the *CEU* of the joint ρ .

One may note that two pairs (i, θ_i) and (i, θ'_i) may play in the same local game – this happens when θ and θ' belong to the same focal set. In this case, the utility of (i, θ'_i) does not depend on the action played by (i, θ_i) and conversely.

For any focal element B of m , let $Players(B) := \{(i, \theta_i) \mid \theta \in B, i \in N\}$ – $Players(B)$ denotes the future players involved in the local game corresponding to B . Let \tilde{E} be the multiset $\tilde{E} := [Players(B) \mid B \in S_m]$. The elements e of \tilde{E} and the focal elements in S_m are in bijection and we denote B_e the focal element of m which leads to e . These notations allow us to propose a first, direct generalization to credal games of the Howson's and Rosenthal's transform:

Definition 15 (Direct transform of a credal game). *The direct transform of a credal game $G = (N, (A_i, \Theta_i, u_i)_{i \in N}, m)$ is the hypergraphical game $\tilde{G} = (\tilde{N}, \tilde{E}, (\tilde{A}_{(i, \theta_i)})_{(i, \theta_i) \in \tilde{N}}, (\tilde{u}_{(i, \theta_i)}^e)_{e \in \tilde{E}, (i, \theta_i) \in e})$ where:*

- $\tilde{N} = \{(i, \theta_i) \mid i \in N, \theta_i \in \Theta_i\}$,
- For each $i \in N$, $\theta_i \in \Theta_i$, $\tilde{A}_{(i, \theta_i)} = A_i$,
- $\tilde{A} = \prod_{i \in N, \theta_i \in \Theta_i} \tilde{A}_{(i, \theta_i)}$ denotes the set of all the pure strategy profiles in \tilde{G} .
- $\tilde{E} = [Players(B) \mid B \in S_m]$,
- For each $e \in \tilde{E}$, $(i, \theta_i) \in e$ and $\tilde{\rho} \in \tilde{A}$,
- $\tilde{u}_{(i, \theta_i)}^e(\tilde{\rho}_e) = K_{|\theta_i} \cdot m(B_e) \cdot \min_{\theta' \in B_e, \theta'_i = \theta_i} u_i(\rho(\theta'), \theta')$.

It is easy to show that the *CEU* value of a pure strategy ρ in G and the global utility of $\tilde{\rho}$ in \tilde{G} are equal, whatever is the couple (i, θ_i) considered.

Proposition 1. ³ *Let G be a credal game and \tilde{G} its direct transform. For any pure strategy ρ of G , it holds that $CEU_{(i, \theta_i)}(\rho) = \tilde{u}_{(i, \theta)}(\tilde{\rho})$.*

¹ Named after Selten, who proposed a similar definition for Bayesian games [16].

² We could use the notation ρ for both, but the pure strategy profiles of the credal game are vectors of functions $\rho_i : \Theta_i \mapsto A_i$ while the pure strategy profiles of \tilde{G} are vectors in $\prod_{i \in N, \theta_i \in \Theta_i} A_i$. So, we keep the two notations $\tilde{\rho}$ and ρ .

³ The proofs can be found in the appendix of the present document [11], page 15

Let us extend the Selten transform to mixed strategy profiles ρ of G : each $\tilde{\rho}_{(i,\theta_i)} = \rho_i(\theta_i)$ is then a probability distribution over A_i , and $\tilde{\rho}$ is then a vector of such distributions.

Corollary 1. *Let G be a credal game and \tilde{G} its direct transform. For any mixed strategy profile ρ of G , it holds that $CEU_{(i,\theta_i)}(\rho) = \tilde{u}_{(i,\theta)}(\tilde{\rho})$.*

When m is a probability distribution and G is a 2-player game, we get at most $|\Theta|$ local games, each involving two players (i, θ_i) and (j, θ_j) : \tilde{G} is a polymatrix game, and Howson and Rosenthal's Theorem is recovered. More generally, we get:

Theorem 1 (Generalized Howson-Rosenthal Theorem). *For any credal game G , there exists an hypergraphical game \tilde{G} such that ρ is a pure (resp. mixed) Nash equilibrium of G iff $\tilde{\rho}$ is a pure (resp. mixed) Nash equilibrium of \tilde{G} .*

Example 2. Let us define the direct transform of the credal game G corresponding to our running example. The set of players is: $\tilde{N} = \{(1, P), (1, \bar{P}), (2, R), (2, \bar{R})\}$. The set of actions are $\tilde{A}_{(i,\theta_i)} = \{P_i, Q_i, R_i\}$.

Because m has two focal elements $B_1 = \{(\bar{P}, R)\}$ and $B_2 = \{(P, \bar{R}), (\bar{P}, \bar{R})\}$ each with probability $\frac{1}{2}$, \tilde{G} involves two local games. The set of players involved are respectively $e_1 = \{(1, \bar{P}), (2, R)\}$ and $e_2 = \{(1, \bar{P}), (1, P), (2, \bar{R})\}$. \tilde{G} 's hypergraph is drawn on Figure 1.




Fig. 1. G 's direct transform. Gray circles denote vertices (players; one shade per agent), white boxes denote hyperedges (local games; linked to the players involved).

Player $(2, \bar{R})$ plays only in e_2 , we have for instance:

$$\tilde{u}_{(2,\bar{R})}^{e_2}(R_1, P_1, R_2) = K_{\bar{R}} \cdot m(B_2) \times \min \left[u_2((R_1, R_2), (P, \bar{R})), u_2((P_1, R_2), (\bar{P}, \bar{R})) \right] = 2.$$

For player $(1, \bar{P})$, which plays in both local games, we have for instance:

$$\begin{aligned} \tilde{u}_{(1,\bar{P})}^{e_1}(P_1, P_2) &= K_{\bar{P}} \cdot m(B_1) \times u_1((P_1, P_2), (\bar{P}, R)) = 0.5 \times 2 = 1 \\ \tilde{u}_{(1,\bar{P})}^{e_2}(-, P_1, Q_2) &= K_{\bar{P}} \cdot m(B_2) \times u_1((P_1, Q_2), (\bar{P}, \bar{R})) = 0.5 \times 3 = 1.5. \end{aligned}$$

The Selten transform of the Nash equilibrium ρ' described in the previous section is $\tilde{\rho}'((1, \bar{P})) = P_1$, $\tilde{\rho}'((1, P)) = Q_1$, $\tilde{\rho}'((2, \bar{R})) = R_1$, $\tilde{\rho}'((2, R)) = Q_2$. It is easy to check that: $\tilde{u}_{(1,\bar{P})}(\rho') = \tilde{u}_{(1,\bar{P})}^{e_1}((P_1, Q_2)) + \tilde{u}_{(1,\bar{P})}^{e_2}((Q_1, P_1, R_2)) = CEU_{(1,\bar{P})}(\rho')$. Notice that in the sum, one part of the utility of $(1, \bar{P})$ comes from the local game e_1 (i.e., from B_1) and the other part comes from e_2 (i.e., from B_2).

As to the complexity of the transform, let α (resp. β) be the maximum number of actions (resp. types) per player in G and k the degree of additivity of m . G contains n utility tables of size $(\alpha\beta)^n$ and the size of the description of m is bounded by $k \cdot n \cdot |\mathcal{S}_m|$. So, $Size(G)$ is in $O(n(\alpha\beta)^n + kn \cdot |\mathcal{S}_m|)$.

\tilde{G} contains $|\mathcal{S}_m|$ local games. Each of them involves at most kn players (i, θ_i) – the size of their SNF representation is thus at most $kn\alpha^{kn}$ – hence a spatial cost for the representation of \tilde{G} in $O(|\mathcal{S}_m| \cdot kn\alpha^{kn})$. Notice now that since m is k -additive, $|\mathcal{S}_m| < \beta^{kn}$. So, $Size(\tilde{G})$ is bounded by $kn(\alpha\beta)^{kn}$. In short, we get:

Proposition 2. *The direct transform of a credal game G has a temporal complexity in $O(|\mathcal{S}_m| \cdot kn^2\alpha^{kn}\beta)$, also bounded by $k^2 \cdot Size(G)^{k+1}$ and a spatial complexity in $O(|\mathcal{S}_m| \cdot kn\alpha^{kn})$, also bounded by $k \cdot Size(G)^k$.*

So, the degree of additivity of the bpa is the main factor of complexity. Hopefully, low degrees of additivity can be assumed – it has indeed been shown [19,15] that such low values (typically, $k \leq 3$) allow the description of many cases of interest. In such situations, the transform is quadratic or, at worst, cubic.

4.2 The conditioned transform

Now, when for each focal element B and each θ_i , only a few types are compatible for the other players, one shall use a more sophisticated transform. In the following, we propose to condition each focal element B for which θ_i is plausible – we thus get a *subset* of B for which a local game is created. This transform leads to smaller local games as soon as this subset involves less players than B .

Formally, for any $B \in \mathcal{S}_m$, $i \in N$ and $\theta_i \in \Theta_i$, let $B|_{\theta_i} = \{\theta' \in B \mid \theta'_i = \theta_i\}$. We thus define the conditioned transform of G as follows:

Definition 16 (The conditioned transform of a credal game). *The conditioned transform of a credal game $(N, (A_i)_{i \in N}, (\Theta_i)_{i \in N}, (u_i)_{i \in N}, m)$ is the hypergraphical game $\tilde{G} = (\tilde{N}, \tilde{E}, (\tilde{A}_{(i, \theta_i)})_{(i, \theta_i) \in \tilde{N}}, (\tilde{u}_{(i, \theta_i)}^e)_{e \in \tilde{E}, (i, \theta_i) \in e})$ where:*

- $\tilde{N} = \{(i, \theta_i) \mid i \in N, \theta_i \in \Theta_i\}$.
- For each $i \in N, \theta_i \in \Theta_i$, $\tilde{A}_{(i, \theta_i)} = A_i$.
- $\tilde{A} = \prod_{i \in N, \theta_i \in \Theta_i} \tilde{A}_{(i, \theta_i)}$ denotes the set of all the pure strategy profiles in \tilde{G} .
- $\tilde{E} = [Players(B|_{\theta_i}) \mid B \in \mathcal{S}_m, i \in N, \theta_i \in \Theta_i]$.
- For each $e \in \tilde{E}$, $(i, \theta_i) \in e$, $\tilde{\rho}_e \in \tilde{A}_e$, $\tilde{u}_{(i, \theta_i)}^e(\tilde{\rho}_e) = m|_{\theta_i}(B_e) \cdot \min_{\theta' \in B_e} u_i(\rho_e(\theta), \theta)$.

The hypergraph of the conditioned transform our running example is drawn on Figure 2. And finally, we can show that:

Proposition 3. *Let G be a credal game and \tilde{G} its conditioned transform. For any pure or mixed strategy ρ of G , it holds that (i) $CEU_{(i, \theta_i)}(\rho) = \tilde{u}_{(i, \theta)}(\tilde{\rho})$ and (ii) ρ is a Nash equilibrium of G iff $\tilde{\rho}$ is a Nash equilibrium of \tilde{G} .*

Proposition 4. *The conditioned transform of a credal game G has a temporal complexity in $O(|\mathcal{S}_m| \cdot kn^2\alpha^{kn}\beta)$ which is bounded by $O(k^2 \cdot Size(G)^{k+1})$ and a spatial complexity in $O(|\bigcup \mathcal{S}_{m|_{\theta_i}}| \cdot kn\alpha^{kn})$ which is bounded by $O(k \cdot Size(G)^k)$.*




Fig. 2. G 's conditioned transform. Gray circles are vertices (players; one color per agent), white boxes are hyperedges (local games; linked to the involved players).

So, the conditioned transform has the same worst case spatial complexity than the direct one. In practice, the size of the transform depends on the structure of the mass function. Typically, if a focal element B involves only one θ_i for Agent i , both transforms contain the same local game $Players(B)$ (as $B = B_{|\theta_i}$), but the conditioned one may produce many more local games. If on the contrary many types are compatible with a focal B for any agent, the local game produced by the direct transform is needlessly large. For instance, from a 2-player credal game where both agents have types a, b, \dots and considering the focal element $\{(a, a), (b, b), \dots\}$, the conditioned transform produces several 2-player games while the direct one has one single local game linked to *all* vertices (i, θ_i) .

5 Conclusion

This article provides two main contributions. On the one hand, we define a model for games of incomplete information based of belief functions. On the other hand, we introduce two transformations that make it possible to build an hypergraphical game (of complete information) equivalent to the initial credal game, generalizing Howson–Rosenthal's theorem. As a result, the algorithmic tools developed for hypergraphical games [3,32] can be used to solve credal games.

This work opens several research directions. First, we shall let the model use other decision rules, e.g. Jaffray's [18,6], which generalize Hurwicz's approach to belief functions, or Gilboa and Schmeider's multiple prior expected utility [29,14]⁴. Beyond belief functions, we aim at extending the model to other Choquet capacities and encompass other decision principles, based e.g. on rank-dependent utility [27], probability intervals [2] or on some neighborhood models such as the PPM model [20]. Finally, we shall also study how belief function based extensions of mixed strategies [23,8] extend to credal games. In this extended framework, the power of representation belief functions will be used not only to capture the uncertainty about the game, but also as a way to describe the agents' knowledge about the others' strategies. Finally, we like to formalize those results with the COQ proof assistant in order to build, with other in-progress results, a modular formal library on incomplete games and decision theory.

⁴ Notice that in the latter approach, the belief function is understood as the lower bound of an imprecise probability – under this interpretation, the conditioning at work must rather be Fagin-Halpern's [10]

References

1. Augustin, T., Schollmeyer, G.: Comment: On Focusing, Soft and Strong Revision of Choquet Capacities and Their Role in Statistics. *Statistical Science* **36**(2), 205–209 (2021)
2. Campos, L.M.D., Huete, J.F., Moral, S.: Probability Intervals: a Tool for Uncertain Reasoning. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.* **2**(2), 167–196 (1994)
3. Chapman, A., Farinelli, A., de Cote, E.M., Rogers, A., Jennings, N.: A Distributed Algorithm for Optimising over Pure Strategy Nash Equilibria. In: *Proc. of the 24th AAAI Conference on Artificial Intelligence*. vol. 24, pp. 749–755 (2010)
4. Choquet, G.: Theory of Capacities. In: *Annales de l’institut Fourier*. vol. 5, pp. 131–295 (1954)
5. Dempster, A.P.: Upper and Lower Probabilities Induced by a Multivalued Mapping. *The Annals of Mathematical Statistics* **38**, 325–339 (1967)
6. Denoeux, T., Shenoy, P.P.: An Interval-Valued Utility Theory for Decision Making with Dempster-Shafer Belief Functions. *International Journal of Approximate Reasoning* **124**, 194–216 (2020)
7. Dubois, D., Denoeux, T.: Conditioning in Dempster-Shafer Theory: Prediction vs. Revision. In: *Proc. of the 2nd International Conference on Belief Functions*, pp. 385–392. Springer (2012)
8. Eichberger, J., Kelsey, D.: Non-Additive Beliefs and Game Theory. Tech. rep., Tilburg University, Center for Economic Research (1994)
9. Ellsberg, D.: Risk, Ambiguity, and the Savage Axioms. *The Quarterly Journal of Economics* pp. 643–669 (1961)
10. Fagin, R., Halpern, J.Y.: A New Approach to Updating Beliefs. In: *Proc. of the 6th conference on Uncertainty in Artificial Intelligence (UAI)*. pp. 347–374 (1990)
11. Fargier, H., Martin-Dorel, E., Pomeret-Coquot, P.: Games of Incomplete Information: a Framework Based on Belief Functions – version with proofs. Available at https://www.irit.fr/~Erik.Martin-Dorel/ecrits/2021_credal-games-ecsqaru_extended.pdf or in HAL.
12. Ghirardato, P., Le Breton, M.: Choquet Rationality. *Journal of Economic Theory* **90**(2), 277–285 (2000)
13. Gilboa, I.: Expected Utility with Purely Subjective Non-Additive Probabilities. *Journal of Mathematical Economics* **16**(1), 65–88 (1987)
14. Gilboa, I., Schmeidler, D.: Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics* **18**(2), 141–153 (1989)
15. Grabisch, M.: Upper Approximation of Non-additive Measures by k -additive Measures—The Case of Belief Functions. In: *Proc. of the 1st International Symposium on Imprecise Probabilities and Their Applications (ISIPTA)*. pp. 158–164 (1999)
16. Harsanyi, J.C.: Games With Incomplete Information Played by “Bayesian” Players, I–III Part I. The Basic Model. *Management Science* **14**(3), 159–182 (1967)
17. Howson Jr, J.T., Rosenthal, R.W.: Bayesian Equilibria of Finite Two-Person Games with Incomplete Information. *Management Science* **21**(3), 313–315 (1974)
18. Jaffray, J.Y.: Linear Utility Theory for Belief Functions. *Operations Research Letters* **8**(2), 107–112 (1989)
19. Miranda, P., Grabisch, M., Gil, P.: Dominance of Capacities by k -Additive Belief Functions. *European Journal of Operational Research* **175**(2), 912–930 (2006)
20. Montes, I., Miranda, E., Destercke, S.: Unifying Neighbourhood and Distortion Models: Part II - New Models and Synthesis. *Int. J. Gen. Syst.* **49**(6), 636–674 (2020)

21. Morgenstern, O., Von Neumann, J.: *Theory of Games and Economic Behavior*. Princeton University Press (1953)
22. Mukerji, S.: Understanding the Nonadditive Probability Decision Model. *Economic Theory* **9**(1), 23–46 (1997)
23. Mukerji, S., Shin, H.S.: Equilibrium Departures from Common Knowledge in Games with Non-Additive Expected Utility. *Advances in Theoretical Economics* **2**(1), 1011 (2002)
24. Myerson, R.B.: *Game Theory*. Harvard university press (2013)
25. Nash, J.: Non-Cooperative Games. *Annals of Mathematics* pp. 286–295 (1951)
26. Papadimitriou, C.H., Roughgarden, T.: Computing Correlated Equilibria in Multi-Player Games. *Journal of the ACM* **55**(3), 1–29 (2008)
27. Quiggin, J.: A Theory of Anticipated Utility. *Journal of economic behavior & organization* **3**(4), 323–343 (1982)
28. Savage, L.J.: *The Foundations of Statistics*. Wiley (1954)
29. Schmeidler, D.: Integral Representation without Additivity. *Proceedings of the American Mathematical Society* **97**(2), 255–261 (1986)
30. Shafer, G.: *A Mathematical Theory of Evidence*. Princeton University Press (1976)
31. Smets, P., Kennes, R.: The Transferable Belief Model. *Artificial Intelligence* **66**(2), 191–234 (1994)
32. Wahbi, M., Brown, K.N.: A Distributed Asynchronous Solver for Nash Equilibria in Hypergraphical Games. In: *Frontiers in Artificial Intelligence and Applications*. vol. 285, pp. 1291–1299 (2016)
33. Yanovskaya, E.: Equilibrium Points in Polymatrix Games. *Lithuanian Mathematical Journal* **8**, 381–384 (1968)

A Proofs

Proof (Spatial complexity of a credal game). The size of a credal game G is the size of its utility functions – given in tables. Each $u_i : A \times \Theta \rightarrow \mathbb{R}$ requires a table of size $|A| \times |\Theta|$. Since $|A|$ is bounded by α^n and $|\Theta|$ by β^n , we have $\text{Size}(G) \in O(n(\alpha\beta)^n)$

Proof (Proposition 1 – Direct transform – Pure strategy profiles utility equality). Let us rewrite Dempster-conditioning's definition into CEU 's one. For any pure strategy profile ρ , we have:

$$CEU_{(i,\theta_i)}(\rho) = \sum_{B' \in \mathcal{S}_{m|\theta_i}} \left(K_{\theta_i} \times \sum_{\substack{B \in \mathcal{S}_m \\ B \cap \{\theta' | \theta'_i = \theta_i\} = B'}} m(B) \right) \times \min_{\theta' \in B'} u_i(\rho_i(\theta'), \theta')$$

Any mass $m(B)$ is transferred to 0 or 1 focal element of $m|\theta_i$, so, since multiplication distributes over addition, one can write:

$$CEU_{(i,\theta_i)}(\rho) = \sum_{\substack{B \in \mathcal{S}_m \\ \exists \theta' \in B, \theta'_i = \theta_i}} K_{\theta_i} \times m(B) \times \min_{\substack{\theta' \in B \\ \theta'_i = \theta_i}} u_i(\rho(\theta'), \theta')$$

Having definitions $\text{Players}(B) = \{(i, \theta_i) \mid \theta \in B, i \in N\}$ and $\tilde{E} = [\text{Players}(B) \mid B \in \mathcal{S}_m]$, one can write:

$$CEU_{(i,\theta_i)} = \sum_{\substack{e \in \tilde{E} \\ (i,\theta_i) \in e}} K_{\theta_i} \times m(B) \times \min_{\substack{\theta' \in B \\ \theta'_i = \theta_i}} u_i(\rho(\theta'), \theta') = \tilde{u}_{(i,\theta_i)}(\tilde{\rho})$$

Proof (Corollary 1 – Direct transform – Mixed strategy profiles utility equality). Let G be a credal game and ρ be one of its mixed strategy.

First, note that given a focal element B of m , there may be several pure strategy profile σ, σ', \dots which correspond to the same focal element $\sigma.B = \sigma'.B$ of m^ρ – they are such as $\forall \theta \in B, \sigma(\theta) = \sigma'(\theta)$ and $\exists \theta \notin B, \sigma(\theta) \neq \sigma'(\theta)$. However, ρ defines a probability over pure strategy profiles, by $Pr^\rho(\sigma) = \prod_{i \in N, \theta_i \in \Theta_i} \rho_i(\theta_i)(\sigma_i(\theta_i))$. Thus, we have:

$$m^\rho(\sigma.B) = m(B) \times \prod_{\theta \in B} \prod_{i \in N} \rho_i(\theta_i)(\sigma_i(\theta_i)) = m(B) \times \sum_{\sigma' : \Theta \rightarrow A} Pr^\rho(\sigma')$$

Then, let rewrite this equation into the definition of ρ 's CEU :

$$\begin{aligned}
CEU_{(i,\theta_i)}(\rho) &= \sum_{\sigma \in \mathcal{S}_{m_{|\theta_i}^\rho}} \left(m_{|\theta_i}(B) \times \sum_{\sigma': \Theta \rightarrow A} Pr^\rho(\sigma') \right) \times \min_{\theta \in B} u_i(\sigma(\theta), \theta) \\
&= \sum_{\sigma': \Theta \rightarrow A} Pr^\rho(\sigma') \times \left(\sum_{B \in \mathcal{S}_{m_{|\theta_i}}} m_{|\theta_i}(B) \times \min_{\theta \in B} u_i(\sigma'(\theta), \theta) \right) \\
&= \sum_{\sigma': \Theta \rightarrow A} Pr^\rho(\sigma') \times CEU_{(i,\theta_i)}(\sigma')
\end{aligned}$$

In the previous equation, we can write σ' instead of σ (inside the CEU 's min). Indeed, $Pr^\rho(\sigma') \neq 0 \iff \forall \theta \in B, \sigma'(\theta) = \sigma(\theta)$. Thus, we have expressed the CEU of ρ as an expectation over pure strategy profiles' CEU values.

Finally, let \tilde{G} be the direct transform of G and $\tilde{\rho}$ the Selten transform of ρ , which is a mixed strategy profile of \tilde{G} . We already show utility equality for pure strategy profiles. So, since the Selten transform is a one-to-one correspondence (both for pure and mixed strategy profiles), we have:

$$\begin{aligned}
EU_{(i,\theta_i)}(\tilde{\rho}) &= \sum_{\tilde{a} \in \tilde{A}} \left(\prod_{(i,\theta_i) \in \tilde{N}} \tilde{\rho}_{(i,\theta_i)}(\tilde{a}) \right) \times \tilde{u}_{(i,\theta_i)}(\tilde{a}) \\
&= \sum_{\sigma': \Theta \rightarrow A} \left(\prod_{(i,\theta_i) \in \tilde{N}} \tilde{\rho}_{(i,\theta_i)}(\sigma'_i(\theta_i)) \right) \times CEU_{(i,\theta_i)}(\sigma') \\
&= \sum_{\sigma': \Theta \rightarrow A} \left(\prod_{i \in N} \prod_{\theta_i \in \Theta_i} \rho_i(\theta_i)(\sigma'_i(\theta_i)) \right) \times CEU_{(i,\theta_i)}(\sigma') \\
&= \sum_{\sigma': \Theta \rightarrow A} Pr^\rho(\sigma') \times CEU_{(i,\theta_i)}(\sigma')
\end{aligned}$$

That is, it means $CEU_{(i,\theta_i)}(\rho) = EU_{(i,\theta_i)}(\tilde{\rho})$.

Proof (Theorem 1 – Generalized Howson and Rosenthal's Theorem). Let G be a credal game and \tilde{G} its direct transform.⁵ We already prove the equality of utility values of any (pure or mixed) strategy profile ρ of G and of its Selten transform $\tilde{\rho}$ of \tilde{G} . Thus, for any player i of any type θ_i and for any of her strategies ρ'_i , it holds that:

$$\left(CEU_{(i,\theta_i)}(\rho'_i, \rho_{-i}) > CEU_{(i,\theta_i)}(\rho) \right) \iff \left(\tilde{u}_{(i,\theta_i)}(\tilde{\rho}'_i, \tilde{\rho}_{-i}) > \tilde{u}_{(i,\theta_i)}(\tilde{\rho}) \right)$$

That is, ρ is a Nash equilibrium of G iff $\tilde{\rho}$ is a Nash equilibrium of \tilde{G}

⁵ We may also prove the Generalized Howson and Rosenthal's Theorem using the conditioned transform. The proof is exactly the same.

Proof (Proposition 2 – Direct transform – Spatial complexity). The size of an hypergraphical game is bounded by the size of its local utility functions given in tables. Each local game e has a size $|A_e| \leq \alpha^{|e|}$. Let G be a credal game and \tilde{G} its direct transform. \tilde{G} is an hypergraphical game, local games of which are build for each focal element $B \in \mathcal{S}_m$ – one player (i, θ_i) for each component of vectors $\theta \in B$ – so local games have at most kn players, thus a size bounded by $kn\alpha^{kn}$. Since $|\mathcal{S}_m| \leq \beta^{kn}$, it holds that:

$$Size(\tilde{G}) \leq |\mathcal{S}_m| \times kn\alpha^{kn} \leq kn(\alpha\beta)^{kn}$$

Furthermore $Size(G) \in O(n(\alpha\beta)^n)$, so we have:

$$Size(\tilde{G}) \leq kn^{1-k} \times Size(G)^k \in O(k \times Size(G)^k)$$

Proof (Proposition 2 – Direct transform – Temporal complexity). Let us study this naive algorithm:

```

1  for  $i$  in  $N$ :
2    for  $\theta_i$  in  $\Theta_i$ :
3       $K_{\theta_i} := 1 / \sum_{B \in \mathcal{S}_m, B \cap \theta_i \neq \emptyset} m(B)$ 
4      for  $B$  in  $\mathcal{S}_m$ :
5         $e := Players(B)$ 
6        if  $(i, \theta_i)$  in  $e$ :
7          for  $a_e$  in  $A_e$ :
8             $u[e][[(i, \theta_i)]] [a_e] := K_{\theta_i} * m(B) * \min_{\theta' \in B, \theta'_i = \theta_i} u_i(a_e(\theta'), \theta')$ 

```

- Line 3 computes $K_{|\theta_i}$. The \sum operation walks over \mathcal{S}_m and for each focal element B it checks the i -th component of $\theta \in B$ vectors. It requires $T_K \in O(k \times |\mathcal{S}_m|)$ operations.
- Line 5 computes $Players(B) = \{(i, \theta_i) \mid i \in N, \theta \in B\}$, thus in $T_e = O(kn)$ operations.
- Line 8 computes the local utility of a vector a_e of size $|e| \leq kn$. The min operation walks over B and for each vector $\theta \in B$ it may need to project a_e over its $|e|$ relevant components (the components given by θ). It requires $T_u = O(k^2n)$ operations.

Lines 1, 2, 4 and 7 are loops of length bounded by n , β , $|\mathcal{S}_m|$ and α^{kn} respectively. Thus:

$$T = n\beta \times (T_k + |\mathcal{S}_m| \times (T_e + \alpha^{kn} \times T_u)) \in O(|\mathcal{S}_m| \times k^2n^2\alpha^{kn}\beta)$$

Since $|\mathcal{S}_m| \leq \beta^{kn}$, it holds that $T \in O(k^2n^2\alpha^{kn}\beta^{kn+1}) \subset O(k^2 \times Size(G)^{k+1})$

Proof (Proposition 3 – Conditioned transform – Pure strategy profiles utility equality). Direct from the definition.

Proof (Proposition 3 – Conditioned transform – Mixed strategy profiles utility equality). Let G be a credal game, ρ one of its mixed strategy profiles, \tilde{G} the

conditioned transform of G and $\tilde{\rho}$ the Selten transform of ρ , which is a mixed strategy profile of \tilde{G} . As for the direct transform's proof of utility equality, we can rewrite the expected utility of $\tilde{\rho}$ with $EU_{(i,\theta_i)}(\tilde{\rho}) = \sum_{\sigma':\Theta \rightarrow A} Pr^{\rho}(\sigma') \times CEU_{(i,\theta_i)}(\sigma')$. Similarly, we conclude that $CEU_{(i,\theta_i)}(\rho) = EU_{(i,\theta_i)}(\tilde{\rho})$.

Proof (Proposition 4 – Conditioned transform – Spatial complexity). Let G be a credal game and \tilde{G} its conditioned transform. \tilde{G} 's local games are build for each focal element of conditioned mass functions, that is, for any $B \in \bigcup_{i \in N, \theta_i \in \Theta_i} \mathcal{S}_{m|\theta_i}$. Those focal elements are subsets of the original ones, thus their sizes are also bounded by k . With the same arguments as for the direct transform's proof of spatial complexity, it holds that \tilde{G} 's local games size is bounded by $kn(\alpha)^{kn}$, thus:

$$Size(\tilde{G}) \leq \left| \bigcup \mathcal{S}_{m|\theta_i} \right| \times kn\alpha^{kn} \leq kn(\alpha\beta)^kn \in O(k \times Size(G)^k)$$

Proof (Proposition 4 – Conditioned transform – Temporal complexity). Let us study this naive algorithm:

```

1  for i in N:
2    for  $\theta_i$  in  $\Theta_i$ :
3       $m_{|\theta_i} := \text{condition}(m, \theta_i)$ 
4      for B in  $\mathcal{S}_{m_{|\theta_i}}$ :
5         $e := \text{Players}(B)$ 
6        for  $a_e$  in  $A_e$ :
7           $u[e][i, \theta_i][a_e] := m_{|\theta_i}(B) * \min_{\theta \in B} u_i(a_e(\theta), \theta)$ 

```

- Line 3 computes the conditioned mass function given θ_i . It requires walking over all focal elements of m . For each $B \in \mathcal{S}_m$, it checks the i -th component of every vector $\theta \in B$ (in k operations), then updates $m_{|\theta_i}(B_{|\theta_i})$ if needed (in 1 operation). The normalization factor K_{θ_i} can be computed *on the fly*, thus the full conditioning requires $T_{cond} = k \times |\mathcal{S}_m|$ operations.
- Line 5 computes $\text{Players}(B)$ is $T_e \in O(kn)$ operations – see Proof (A).
- Line 7 computes the local utility of a vector a_e for a given player $(i, \theta_i) \in e$ in $T_u \in O(k^2n)$ operation – also see Proof (A).

Lines 1, 2, 4 and 6 are loops of length bounded by n , β , $|\mathcal{S}_m|$ and α^{kn} respectively. Thus:

$$T = n\beta \left(T_{cond} + |\mathcal{S}_m| \times (T_e + \alpha^{kn} \times T_u) \right) \in O(|\mathcal{S}_m| \times k^2 n^2 \alpha^{kn} \beta)$$

Since $|\mathcal{S}_m| \leq \beta^{kn}$, it holds that $T \in O(k^2 n^2 \alpha^{kn} \beta^{kn+1}) \subset O(k^2 \times Size(G)^{k+1})$.