

Adaptation Procedure in Misinformation Games

Konstantinos Varsos^{1*}, Merkouris Papamichail^{1,2},
Giorgos Flouris¹, Marina Bitsaki²

¹Institute of Computer Science, Foundation for Research and Technology -
Hellas, N.Plastira 100, Heraklion, GR70013, Crete, Greece.

²Computer Science Department, University of Greece, Voutes Campus,
Heraklion, GR70013, Crete, Greece.

*Corresponding author(s). E-mail(s): varsosk@ics.forth.gr;
Contributing authors: mercoyris@ics.forth.gr; fgeo@ics.forth.gr;
ecbitsaki@gmail.com;

Abstract

We study interactions between agents in multi-agent systems, in which the agents are misinformed with regards to the game that they play, essentially having a subjective and incorrect understanding of the setting, without being aware of it. For that, we introduce a new game-theoretic concept, called misinformation games, that provides the necessary toolkit to study this situation. Subsequently, we enhance this framework by developing a time-discrete procedure (called the Adaptation Procedure) that captures iterative interactions in the above context. During the Adaptation Procedure, the agents update their information and reassess their behaviour in each step. We demonstrate our ideas through an implementation, which is used to study the efficiency and characteristics of the Adaptation Procedure.

Keywords: Misinformation Games, Adaptation Procedure, natural misinformed equilibrium, stable misinformed equilibrium

1 Introduction

The importance of multi-agent systems has been heavily acknowledged by the research community and industry. A multi-agent system is a system composed of multiple

interacting autonomous, self-interested, and intelligent agents¹, and their environment. In such settings, agents need to be incentivized to choose a plan of action, and game theory [1] provides a suitable framework for analyzing these interactions.

A usual assumption of game theory is that the game specifications (i.e., the rules of interaction, which include the number of players, their strategies and the expected payoff for each strategic choice) are common knowledge among the players. In other words, players have correct (although not necessarily complete) information regarding the game.

However, in realistic situations, a reasoning agent is often faced with erroneous, misleading and unverifiable information regarding the state of the world or the possible outcomes of her actions (see games with misperception [2–4]). This *misinformation* may affect the interaction’s outcomes.

Misinformation in an interaction can occur due to various reasons. For example, the designer may choose to deceive (some of) the agents in order to obtain improved social outcomes or for other reasons (this falls under the general area of mechanism design, see [5]). Or an agent may communicate deceptive information in order to lead the other agents to suboptimal choices and obtain improved outcomes (e.g., fake financial reports misleading investors [6]). Furthermore, when the agents operate in a remote and/or hostile environment, noise and other random effects may distort communication, causing agents to receive a game specification that is different from the intended one (e.g., in the case of autonomous vehicles operating in Mars [7]). Another scenario that could lead to misinformation refers to cases where the environment changes without the agents’ knowledge, causing them to have outdated information regarding the rules of interaction (e.g., an accident causing a major and unexpected disruption in the flow of different roads in a city). Last but not least, endogenous reasons (e.g., limited awareness, bounded computational capacity, cognitive restrictions, biases [8] etc.) may cause agents to misinterpret the situation and assign their own (mistaken) payoffs to different actions.

Thereupon, some aspects of the situation, or the modeling and reasoning regarding the situation, leads players to incorporate a possibly incorrect viewpoint of the real aspects. Thus, they may miss crucial specifications, and interact relying on a restricted and incorrect perception of the situation. The key characteristic of the described scenario is that the players do not question the rules of interaction given to them; this differentiates this scenario from standard settings of games with other forms of uncertainty (such as Bayesian games [9, 10], uncertainty theory [11] etc.), in which players are well-aware of the fact that the information given to them is incomplete, uncertain or flawed in various ways, and this knowledge is incorporated in their reasoning (see also Figure 1). In other words, in misinformation games *agents don’t know that they don’t know*, as opposed to incomplete or imperfect games where the *agents know that they don’t know*.

The study of settings where any participant has (possibly) wrong knowledge with regards to the real situation, has been recently considered in defining the concept of *misinformation games* [5, 12]; this work falls under the hood of games with misperception [2–4]. In misinformation games, each player has a subjective view of the abstract

¹Note that we use the terms “agent” and “player” interchangeably throughout this paper.

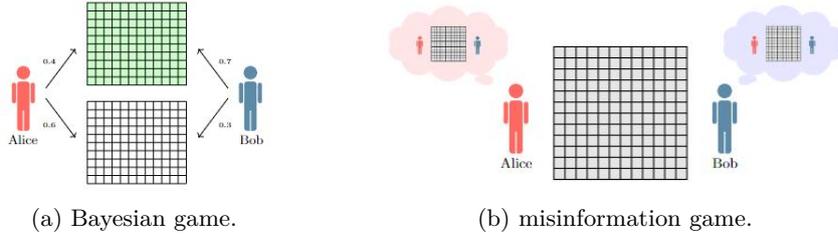


Fig. 1: Example of games without common and correct information.

game’s specifications, that may not coincide with the specifications of the real interaction, modelling the fact that agents may operate under an erroneous specification. For that, the proposed method agglomerates both the real situation (*actual game*), and the subjective (misinformed) views of the players. A key characteristic of misinformation games is that the equilibrium (i.e., the set of strategic choices where no player wants to deviate from) is determined by the subjective views of the players, rather than the actual game [12]. This is called the *natural misinformed equilibrium* in [12]. On the contrary, the actual payoffs received by the players is determined by the actual game, and may be different than the expected ones.

Given the discrepancy between the actual and the perceived payoffs, a natural question is how the players will react upon their realisation that the received payoffs are different than expected. However, the approach presented in [5, 12] focuses on one-shot interactions, and thus misses this crucial and interesting part of the problem. To address this limitation, this work *enhances misinformation games with an iterative time-discrete methodology, called the Adaptation Procedure, which models the evolution of the strategic behaviour of rational players in a misinformation game, as they obtain new information and update their (erroneous) game specifications*. The first steps towards this direction were held in [13]; here we extend this approach.

More specifically, at each time point, players choose a strategic action (using the methodology of [5, 12]), receive the corresponding payoffs, and engulf them into their game specifications. The procedure is then repeated. Importantly, the information received in each time point may lead players to a different choice in the next time point (because they now operate under a different payoff matrix), so the procedure is iterative and stabilises when the players have no incentives to deviate from their current choices, based on what they know so far (which may or may not coincide with the actual game). When the procedure stabilizes we reach a refinement of the natural misinformed equilibrium, that we call the *stable misinformed equilibrium*.

The structure of this paper follows. In Section 2 we present some preliminaries, including the concept of misinformation games. Then, in Section 3 we introduce the Adaptation Procedure process, and in Section 4 we present an algorithmic analysis and tools to compute the Adaptation Procedure and provide some experimental evaluation. In Section 5 we review the literature related to this paper and conclude in Section 6. Finally, we provide Appendix A where we present all the proofs of our results.

This paper is an extended and revised version, combining and extending previous works of ours [5, 12, 13]. More specifically, the main new contributions of the paper, compared to the previous ones, are the following:

- we elaborate on and describe in detail the inflation process described superficially in [12] (see Propositions 1-6, and Algorithms 1, 2, 3 in Subsection 2.2, as well as Algorithms 4, 5 in Subsection 2.3);
- we provide new theoretical and computational results regarding the adaptation procedure (see Subsection 3.4);
- we provide novel implementations for the Adaptation Procedure (including a parallel one), and new experimental results (see Section 4).

2 Normal-form misinformation games

In this section we describe misinformation games; we first present a quick review of normal-form games.

2.1 Normal-form games

In this paper, we consider *normal-form games* G . A normal-form game includes a set of players N . Each player $i \in N$ has a set of pure strategies that is denoted by S_i , and the Cartesian product of all S_i s is the set of joint decisions of the players, which is denoted by S . Finally, each player i has a payoff matrix $P_i \in \mathbb{R}^{|S_i| \times \dots \times |S_{|N|}|}$, assigning a payoff to that player depending on the joint strategy decisions of all players. The payoff matrix of the game is $P = (P_1, \dots, P_{|N|})$. We follow the usual notation and we denote G as a tuple $\langle N, S, P \rangle$.

If player i randomly selects a pure strategy, then she plays a mixed strategy $\sigma_i = (\sigma_{i,1}, \dots, \sigma_{i,|S_i|})$ which is a discrete probability distribution over S_i . Let Σ_i be the set of all possible mixed strategies σ_i . A strategy profile $\sigma = (\sigma_1, \dots, \sigma_{|N|})$ is an $|N|$ -tuple in $\Sigma = \Sigma_1 \times \dots \times \Sigma_{|N|}$. We denote by σ_{-i} the joint decision of all other players except for player i in σ . The payoff function of player i is defined as: $h_i : \Sigma \rightarrow \mathbb{R}$, such that:

$$h_i(\sigma) = \sum_{k \in S_1} \dots \sum_{j \in S_{|N|}} P_i(k, \dots, j) \cdot \sigma_{1,k} \cdot \dots \cdot \sigma_{|N|,j}, \quad (1)$$

where $P_i(k, \dots, j)$ is the payoff of player i in the pure strategy profile (k, \dots, j) . Further, $h_i(\sigma_i, \sigma_{-i})$ represents the expected payoff of player i as a function of σ . The social welfare function is derived from h as $SW(\sigma) := \sum_i h_i(\sigma)$.

In the following, we consider the concept of Nash equilibrium, that is a strategy profile in which no player possesses an improving deviation. Formally, a strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_{|N|}^*)$ is a Nash equilibrium, iff, for any i and for any $\hat{\sigma}_i \in \Sigma_i$, $h_i(\sigma_i^*, \sigma_{-i}^*) \geq h_i(\hat{\sigma}_i, \sigma_{-i}^*)$. Further, we denote by $NE(G)$ the set of Nash equilibria of G (or simply NE , when G is clear from the context).

A useful notion in equilibrium computation is the support of a strategy σ_i , denoted by $\text{supp}(\sigma_i)$, which is the set of pure strategies that are played with positive probability under σ_i . Formally, $\text{supp}(\sigma_i) := \{j \in [|S_i|] : \sigma_{i,j} > 0\}$.

2.2 Inflation process for normal form games

We consider the process of *inflation*, which can be used to “enlarge” a normal form game, adding new players and strategies, in ways that do not alter the game’s important properties. In particular, the *inflated version* of a game G , say G' , is a “bigger” game (in terms of number of players and available strategies), such that the common players will have the same strategic behaviour and will receive the same payoffs in both games.

Interestingly, this process has been used in the past in the literature in order to add “dummy” players or strategies, but, to the best of our knowledge, without any appropriate formalisation. Here we introduce the formal machinery necessary to show that this process indeed does not alter the game’s properties.

The process of inflation will become relevant later, when we define “canonical misinformation games” (see Definition 4). In particular, the process of inflation will allow us to restrict our attention to a specific class of misinformation games, which significantly simplifies our formalization and definitions.

Algorithmically, to construct an inflated version of a game, we add strategies and players as appropriate, while respecting the above properties. To achieve this, when we add a strategy, we make it dominated (and thus irrelevant from a strategic perspective), and when we add a player, this player is dummy, in the sense that her actions do not affect the strategic choices of other players.

Before defining inflation, we start with an auxiliary concept, namely *compatible strategy profiles*. Given two games G and G' , and two strategy profiles σ, σ' of G, G' respectively, we say that σ, σ' are compatible if they agree (i.e., contain the same strategic choice) on all players and strategies that are common in the two games. Formally:

Definition 1. Consider two normal-form games $G = \langle N, S, P \rangle, G' = \langle N', S', P' \rangle$, and let S_i (respectively S'_i) be the set of strategies available for player $i \in N$ in G (respectively $i \in N'$ in G'). Consider also a strategy profile $\sigma = (\sigma_{ij})$ of G and a strategy profile $\sigma' = (\sigma'_{ij})$ of G' . Then, we say that σ, σ' are compatible, denoted by $\sigma \bowtie \sigma'$, if and only if $\sigma_{ij} = \sigma'_{ij}$ for all $i \in N \cap N', j \in S_i \cap S'_i$.

It is easy to see that \bowtie is an equivalence relation. Now we can define the notion of inflation:

Definition 2. Consider two normal-form games $G = \langle N, S, P \rangle, G' = \langle N', S', P' \rangle$, and let S_i (respectively S'_i) be the set of strategies available for player $i \in N$ in G (respectively $i \in N'$ in G'). Then, we say that G' is an inflated version of G , denoted by $G \blacktriangleleft G'$, if and only if:

- $N \subseteq N'$
- For all $i \in N, S_i \subseteq S'_i$
- $P_i(s) = P'_i(s')$ whenever $i \in N$ and $s = (s_j)_{j \in N} \in S, s' = (s'_j)_{j \in N'} \in S'$ such that $s_j = s'_j$ for all $j \in N$.
- If $\sigma \in NE(G)$, then there exists some $\sigma' \in NE(G')$ such that $\sigma \bowtie \sigma'$.
- If $\sigma' \in NE(G')$, then there exists some $\sigma \in NE(G)$ such that $\sigma \bowtie \sigma'$.

Some analysis on Definition 2 and its consequences is in order. The first two bullets of the definition are straightforward: in order for G' to be an inflated version of G , it

must have at least the same players, and at least the same strategies for the (common) players.

The third bullet guarantees that the two games provide the same payoffs “whenever possible”. In other words, P is, in a sense, a submatrix of P' . The effect of this is that players that are common in G and G' , while playing common (pure) strategies, will receive the same payoffs in the two games. An immediate consequence of this requirement is that the payoffs will also be the same when common players play any common strategy (pure or mixed):

Proposition 1. *Consider two normal-form games $G = \langle N, S, P \rangle$, $G' = \langle N', S', P' \rangle$, such that $G \blacktriangleleft G'$, and let S_i (respectively S'_i) be the set of strategies available for player $i \in N$ in G (respectively $i \in N'$ in G'). Take also some $i \in N$, and σ, σ' strategy profiles in G, G' respectively, such that $\sigma \bowtie \sigma'$. Then, $h_i(\sigma) = h'_i(\sigma')$, where h_i, h'_i are the payoff functions of player i in G, G' respectively.*

A consequence of Proposition 1 is that, in G' , the actions of players that exist only in G' do not affect the payoffs of players that are common in G and G' . Indeed, for any two strategy profiles $\hat{\sigma}, \tilde{\sigma}$ in G' , if these profiles agree on the actions of the common players, then there exists some strategy profile σ in G such that $\sigma \bowtie \hat{\sigma}$ and $\sigma \bowtie \tilde{\sigma}$, and thus, by Proposition 1, $h'_i(\hat{\sigma}) = h_i(\sigma) = h'_i(\tilde{\sigma})$.

As a result, P' can be seen as consisting of multiple “copies” of P along the extra dimensions imposed by the extra players. Note that this bullet does not impose any immediate requirements for the payoffs of players that exist only in G' , or for the payoffs of common players when some (common) player employs a strategy that does not exist in G .

A further immediate consequence of the third bullet is that, if two games have the same number of players, and the same number of strategies per player, then each one is an inflated version of the other if and only if they are the same game.

Proposition 2. *Consider two normal-form games $G = \langle N, S, P \rangle$, $G' = \langle N', S', P' \rangle$, such that $N = N'$ and $S = S'$. Then, $G \blacktriangleleft G'$ if and only if $G = G'$.*

The last two bullets on Definition 2 ensure that the strategic behaviour of players in both games is the “same”. In particular, the fourth bullet guarantees that any Nash equilibrium σ of G will have a “counterpart” in G' , i.e., there will be a Nash equilibrium in G' , say σ' , in which all common players employ the same strategies as in σ (as usual, we don’t care about the actions of the non-common players). Analogously, the fifth bullet guarantees that any Nash equilibrium σ' in G' will have a “counterpart” in G , i.e., a Nash equilibrium, say σ in which the common players employ the same strategies as in σ' . Note that this requirement disallows the existence of Nash equilibria in G' for which any common player employs a non-common strategy with a probability higher than 0, because, in such a strategy profile, the probabilities assigned to the common strategies do not sum to 1, and thus there cannot exist a compatible strategy profile in G . Note also the use of the compatibility relation between strategy profiles to formalise the notion of “counterpart”.

Clearly, the relation \blacktriangleleft is a partial order, as shown in the proposition below:

Proposition 3. *The relation \blacktriangleleft is a partial order, i.e., for any normal-form games G_1, G_2, G_3 the following hold:*

1. $G_1 \blacktriangleleft G_1$ (reflexivity)

2. If $G_1 \triangleleft G_2$ and $G_2 \triangleleft G_1$ then $G_1 = G_2$ (antisymmetry)
3. If $G_1 \triangleleft G_2$ and $G_2 \triangleleft G_3$ then $G_1 \triangleleft G_3$ (transitivity)

The next question we will consider is how we create an inflated version of a game. For this, we developed the `InflateGame` algorithm (Algorithm 1), along with two supplementary routines, `AddPlayer` (Algorithm 2) and `AddStrategy` (Algorithm 3). In Algorithm 1, the input is a finite normal-form game G , a set of players N' and a strategy space S' , and the goal is to enlarge (i.e., inflate) G so as to have N' players and strategy space S' . In lines 1-3 we add players to the game using Algorithm 2, and in lines 4-10 we embed the strategy space using Algorithm 3.

In supplementary Algorithm 2, we add a new player i in G . Besides the enlargement of the set of players (line 1), we make sure that the strategy space is properly modified in order to support a (single) choice for the new player (line 2). We modify the payoff matrices so as to provide a payoff value for the new player (line 3). More specifically, as each entry of P is an array, we enlarge each such array by one dimension. This is accomplished in line 3 of the algorithm as follows: given an array $a \in \mathbb{R}^n$ we multiply it from the right with the matrix $[I_{n \times n} \ 0_{n \times 1}] \in \mathbb{R}^{n \times n+1}$. This provides a new array with $n + 1$ dimensions. We perform this operation for every entry of P . Observe that the new player will receive a zero payoff for any strategy profile.

Finally, in supplementary Algorithm 3, we add a new strategy j for player i in G . Again, we enhance the set of strategies S_i of player i in order to support the new strategic choice j , resulting to a new set of strategies S'_i (line 2). Afterwards, we modify the payoff matrices so as to provide a payoff value for the new strategy. For that, first we initialize the matrix P' with dimensions $|S_{-i} \times S'_i|$ (line 3). Then, in lines 4-10 we exhaustively run through all the pure strategy profiles s of the strategic space $S_{-i} \times S'_i$. If $s \in S$ then P' has the same values as in P , otherwise P' has values less than the minimum value, m (line 1), of P for all players (line 8), in order to make this strategy dominated and thus irrelevant with respect to the Nash equilibrium, as required by the definition; this is ensured by the operation $m\mathbb{1}_{|N|}$. Observe that this addition affects all axes in the payoff matrices. Also, the payoff values for the pure strategy profiles that do not include j (aka indexes in P) are unchanged, but the remaining entries of the payoff matrices have to be such that they do not affect the strategic behavior of the players.

Algorithm 1 `InflateGame`(G, N', S')

Require:

- $G = \langle N, S, P \rangle$
- N' such that $N \subseteq N'$, and $|N'| < +\infty$
- S' such that, $\forall i \in N', S_i \subseteq S'_i$ and $|S'_i| < +\infty$

- 1: **for** $i \in N' \setminus N$ **do**
 - 2: Set $G \leftarrow \text{AddPlayer}(G, i)$
 - 3: **end for**
 - 4: **for** $i \in N'$ **do**
 - 5: **for** $j \in S'_i \setminus S_i$ **do**
 - 6: Set $G \leftarrow \text{AddStrategy}(G, j, i)$
 - 7: **end for**
 - 8: **end for**
 - 9: **return** $\langle N', S', P' \rangle$
-

Algorithm 2 AddPlayer(G, i)

Require:

A game $G = \langle N, S, P \rangle$
A player $i \notin N$

- 1: $N' \leftarrow N \cup \{i\}$
 - 2: $S' \leftarrow S \times \{1\}$
 - 3: $P'_j \leftarrow P_j [I_{|P_j| \times |P_j|} \ 0_{|P_j| \times 1}]$,
 $\forall j$ entry of P
 - 4: **return** $\langle N', S', P' \rangle$
-

Algorithm 3 AddStrategy(G, j, i)

Require:

A game $G = \langle N, S, P \rangle$
A player $i \in N$, and a strategy $j \notin S_i$

- 1: $m \leftarrow \min_{s \in S, k \in N} \{P_k(s)\} - 1$
 - 2: $S'_i \leftarrow S_i \cup \{j\}$
 - 3: $P' \leftarrow 0_{S_{-i} \times S'_i}$
 - 4: **for** $s \in S_{-i} \times S'_i$ **do**
 - 5: **if** $s \in S$ **then**
 - 6: $P'(s) \leftarrow P(s)$
 - 7: **else**
 - 8: $P'(s) \leftarrow m \mathbb{1}_{|N|}$
 - 9: **end if**
 - 10: **end for**
 - 11: **return** $\langle N, S', P' \rangle$
-

We can show that these algorithms correctly produce an inflated version of a game. Given the transitivity of \blacktriangleleft , the results below hold also for repetitive applications of the above algorithms.

Proposition 4. *Let G be a normal form game, and let G' be the normal form game that is the output of AddPlayer(G, i) (Algorithm 2). Then $G \blacktriangleleft G'$.*

Proposition 5. *Let G be a normal form game, and let G' be the normal form game that is the output of AddStrategy(G, j, i) (Algorithm 3). Then $G \blacktriangleleft G'$.*

Proposition 6. *Let G be a normal form game, and let G' be the normal form game that is the output of InflateGame(G, N', S') (Algorithm 1). Then $G \blacktriangleleft G'$.*

2.3 Misinformation games

Misinformation captures the concept that different players may have a specific, subjective, and thus different view of the game that they play.

Definition 3. *A misinformation normal-form game (or simply misinformation game) is a tuple $mG = \langle G^0, G^1, \dots, G^{|N|} \rangle$, where all G^i are normal-form games and G^0 contains $|N|$ players.*

G^0 is called the *actual game* and represents the game that is actually being played, whereas G^i (for $i \in \{1, \dots, |N|\}$) represents the subjective view of player i , called the *game of player i* . We make no assumptions as to the relation among G^0 and G^i , and allow all types of misinformation to occur. An interesting special class of misinformation games is the following:

Definition 4. *A misinformation game $mG = \langle G^0, G^1, \dots, G^{|N|} \rangle$ is called canonical iff:*

- For any i , G^0, G^i differ only in their payoffs.
- In any G^i , all players have an equal number of pure strategies.

The first bullet of Definition 4 is usually true in practice. In order for a player to miss a strategy or player, or to have additional, imaginary, strategies or players, the communication failures should concern entire sections of a payoff matrix, and thus be

quite substantial. The second bullet is more commonly violated, because the actual game, G^0 , might violate it.

In any case, canonical misinformation games are simpler to grasp, and also significantly simplify the formalisation. We can restrict our focus to canonical misinformation games without loss of generality, because we can transform any non-canonical misinformation game into a canonical one by repetitive applications of the inflation process for normal form games that was described in Subsection 2.2. As explained there, the process of inflation does not alter the strategic properties of the normal form games (which in turn determine the strategic properties of the misinformation game).

To be more precise, we use the **Inflation process** algorithm (Algorithm 5) to transform a non-canonical mG into a canonical mG without altering its strategic properties. In a nutshell, Algorithm 5 first computes the full list of players and strategies that each game should have in order to ensure the properties of Definition 4 (lines 1-2). Then, it utilizes the **AddGame** algorithm (Algorithm 4) to plug in a new game in the misinformation game, in order to ensure that any player, either real or fictional, has her own subjective view of the interaction (lines 3-5). In **AddGame**, we create such a game for any player that is not in N^0 (and thus does not already exist in mG), ensuring that it has the correct players and strategies (N_U, S_U) . Afterwards, Algorithm 5 inflates each of the normal-form games in the non-canonical misinformation game, to ensure that they, too, have the correct players and strategies (lines 6-8). The final resulting mG is returned in line 9.

Algorithm 4 $\text{AddGame}(mG, N', S')$

Require:

- A misinformation game mG
- A set of players N' such that $|N'| < +\infty$
- A set of strategies S'

- 1: Set $G' \leftarrow \langle \emptyset, \emptyset, 0 \rangle$
 - 2: **InflateGame** (G', N', S')
 - 3: **return** $\langle G^0, G^1, \dots, G^N, G' \rangle$
-

Algorithm 5 Inflation process(mG)

Require:

$$mG = \langle G^0, G^1, \dots, G^{|N|} \rangle, \text{ with } |N| < +\infty$$

- 1: Set $N_U \leftarrow \bigcup_{j=0, \dots, |N|} N^j$
 - 2: Set $S_U \leftarrow \times_{j \in N_U} \left(\bigcup_{k \in N_U \cup \{0\}} S_j^k \right)$
 - 3: **for** $i \in N_U \setminus N^0$ **do**
 - 4: $mG \leftarrow \text{AddGame}(mG, N_U, S_U)$
 - 5: **end for**
 - 6: **for** $i \in N_U \cup \{0\}$ **do**
 - 7: Set $G^i \leftarrow \text{InflateGame}(G^i, N_U, S_U)$
 - 8: **end for**
 - 9: **return** mG
-

Since the inflation process does not alter the strategic properties of a game, and any non-canonical misinformation game can be turned into a canonical one using Algorithm 5, we will establish strategy profiles and equilibrium concepts in Subsection 2.4 below, focusing on canonical misinformation games only. As we will see, introducing these concepts for the non-canonical case, brings about all sorts of unrelated technical difficulties that are distracting and complicate the formalism, without contributing to the main intuition.

2.4 Strategy profiles in misinformation games

The definition of misinformed strategies and strategy profiles is straightforward, once noticing that they refer to each player's own game:

Definition 5. A misinformed strategy, $m\sigma_i$ is a strategy of player i in game G^i . We denote the set of all possible misinformed strategies of player i as Σ_i^i . A misinformed strategy profile of mG is an $|N|$ -tuple of misinformed strategies $m\sigma = (m\sigma_1, \dots, m\sigma_{|N|})$, where $m\sigma_i \in \Sigma_i^i$.

As usual, we denote by $m\sigma_{-i}$ the $(|N| - 1)$ -tuple strategy profile of all other players except for player i in a misinformed strategy $m\sigma$. The payoff function h_i of player i under a given profile $m\sigma$ is determined by the payoff matrix of G^0 , and is defined as $h_i : \Sigma_1^1 \times \dots \times \Sigma_{|N|}^{|N|} \rightarrow \mathbb{R}$, such that:

$$h_i(m\sigma_i, m\sigma_{-i}) = \sum_{k \in S_1^1} \dots \sum_{j \in S_{|N|}^{|N|}} P_i^0(k, \dots, j) \cdot m\sigma_{1,k} \cdot \dots \cdot m\sigma_{|N|,j},$$

where $P_i^0(k, \dots, j)$ is the payoff of player i in the pure strategy profile (k, \dots, j) under the actual game G^0 . Also, S_i^j denotes the set of pure strategies of player i in game G^j .

Observe that, although each player's strategic decisions are driven by the information in her own game (G^i), the received payoffs are that of the actual game G^0 , that may differ than G^i . Further, the payoff function would be ill-defined in case of non-canonical misinformation games.

Next, we define the solution concept of a misinformation game, where each player chooses a Nash strategy, neglecting what other players know or play:

Definition 6. A *misinformed strategy*, $m\sigma_i$, of player i , is a misinformed equilibrium strategy, iff, it is a Nash equilibrium strategy for game G^i . A *misinformed strategy profile* $m\sigma$ is called a *natural misinformed equilibrium* iff it consists of misinformed equilibrium strategies.

In the following, we denote by $NME(mG)$ (or simply NME, when mG is obvious from the context) the set of natural misinformed equilibria of mG , and by nme the elements of $NME(mG)$.

Example 1 (Running Example). Consider the canonical misinformation game $mG = \langle G^0, G^1, G^2 \rangle$, where $G^i = \langle \{1, 2\}, S = \{s_1, s_2\} \times \{s_1, s_2\}, P^i \rangle$, with $i \in \{0, 1, 2\}$ and payoff matrices as depicted in Table 1.

Table 1: Payoff matrices in the misinformation game of Example 1.

	s_1	s_2
s_1	(6,6)	(2,7)
s_2	(7,2)	(1,1)

(a) P^0

	s_1	s_2
s_1	(2,2)	(0,3)
s_2	(3,0)	(1,1)

(b) P^1

	s_1	s_2
s_1	(-1,1)	(2,-2)
s_2	(1,-1)	(0,0)

(c) P^2

Player's 1 equilibrium strategy is s_2 in G^1 (i.e., player 1 will play (0,1)), while player's 2 mixed equilibrium strategy in G^2 is (1/2, 1/2). Thus, we have $NME = \{(s_2, s_1), (s_2, s_1)\}$ with $nmes$ ((0, 1), (1, 0)) and ((0, 1), (0, 1)).

To understand the computational properties of calculating an nme , we observe that any natural misinformed equilibrium consists of the agglomeration of the different Nash equilibrium strategy profiles, one for each G^i . The computation of a Nash equilibrium for each G^i is **PPAD**-complete [14], and computing the nme amounts to repeating this computation for each G^i . Thus, the computation of the natural misinformed equilibrium is also **PPAD**-complete.

As a side note, let us consider how the definitions of this subsection would apply (or not) for the case of non-canonical misinformation games. In a non-canonical misinformation game, a player i may have a subjective view G^i that diverts from G^0 in every aspect, i.e., number of players, number of pure strategies, payoffs. We observe that her strategic behavior will still be guided by the contents of her subjective game, G^i , so Definition 5 applies normally. Consequently, Definition 6 for the nme is also well-defined. Contrarily, we observe that the formula for computing the payoff values h_i for some misinformed strategy profile $m\sigma$ is no longer well-defined, because there may exist misinformed strategy profiles (either pure or mixed) that are unknown to G^0 , and thus correspond to no entry in the payoff matrix P^0 . Moreover, from the perspective of individual players, it is unclear what a player should do when being informed that another player used a strategy that is unknown to her (or that an unknown player participated in the game), and how the received payoff should be interpreted (since it corresponds to no entry in her own payoff matrix). Addressing this interesting case will be considered in future work.

3 Adaptation Procedure

In previous work (and in our analysis on misinformation games above), only one-shot interactions were considered; this is a major limitation, as we expect players to realise (partly) the shortcomings of their knowledge following the reception of the (unexpected) payoffs and to revise accordingly. To address this shortcoming, we introduce the Adaptation Procedure that enables us to model iterative interactions. Abusing notation, we write $mG^{(t)}$ instead of mG in order to plug in the step of interaction (this change in the notation is transfused in every concept).

We start with an informal description of the Adaptation Procedure, and continue by presenting the respective formalisation and showing some properties.

3.1 Informal description of the Adaptation Procedure

The Adaptation Procedure is initialized with a misinformation game, say $mG^{(0)}$. As explained in Section 2.4, this will cause each player to employ one of the equilibrium strategies in her own game, resulting to a natural misinformed equilibrium. The payoff received from the joint strategic choices will be provided by the actual game, G^0 , and this may be different from what each player knows (and expects) from her subjective game. Notably, we assume that the payoffs received by each player for their strategic choices are publicly announced, thus are common knowledge. As a result, players will update their payoff matrices by replacing the erroneous payoffs with the correct ones just received, leading to a new misinformation game.

Interestingly, the above process is not, in general, linear. When the misinformation game has more than one natural misinformed equilibria, and/or when there exist mixed strategic choices in them, each of these choices will be considered in a separate branch of the process, because the players may choose any of those strategies to play. Any such branch leads to a different new misinformation game, and, thus, our process should take into account all those potential scenarios. To do this, $mG^{(0)}$ will spawn several new misinformation games, one for each element of the support of the natural misinformed equilibria, eventually forming a tree (more precisely, a directed graph, as there may be self-loops or mergers, as we will show later). Note that this spawning should not be interpreted as leading to multiple parallel processes; instead, it represents different potential ways that the process can unfold.

The process continues recursively for each branch, creating new misinformation games. When no new misinformation games are spawned (i.e., when the players learn nothing new from the environment and no longer change their payoff matrices), the Adaptation Procedure *terminates*.

Observe that the Adaptation Procedure produces new games (and thus new *nmes*) in each time step. In general an *nme* is not a stable equilibrium concept when the Adaptation Procedure is considered, because, after the players learn something new about the reality as explained above, they might choose to change their choices. Nevertheless, the *nmes* of the games appearing in the leaves of the recursive tree at the time when the Adaptation Procedure terminates are stable (nothing new is learnt). Given that, we define a new equilibrium concept, that is, a strategy profile that is an *nme* of one of the leaves, and we call it *stable misinformed equilibrium (sme)*; players

do not have incentives to deviate from an *sme*, even in the presence of the payoff information received by the environment (game).

Also note that this process would be significantly more complex for non-canonical misinformation games, where the agents could potentially have erroneous information for every aspect of the game, including the available strategies or players. In such scenarios, an agent would be unable to introspect the decisions they made when they learn about a new strategy or player in the game, and would have to expand their game to include the newly learned strategies or players. That would significantly complicate the description of the Adaptation Process, both at the intuitive level above and in the formalisation that follows.

3.2 Formal definition

Consider a multidimensional matrix A , and a vector \vec{v} . We denote by $A_{\vec{v}}$ the element of A in position \vec{v} . For example, $A_{(1,2)}$ is the top right element of the 2×2 matrix A . More general, for a $n_1 \times n_2 \times \dots \times n_m$ matrix, a *position* is an m -dimensional vector $\vec{v} \in [n_1] \times [n_2] \times \dots \times [n_m]$, where $[n_i] = \{1, 2, \dots, n_i\}$. Furthermore, the set of all position vectors in a multidimensional payoff matrix is the set of joint strategies S . In what follows, we will use the term position vectors when referring in a specific point in the payoff matrix. Thereupon, we define the operation of replacement of element $A_{\vec{v}}$ with b as follows:

Definition 7. Consider set F , matrix $A \in F^{n_1 \times n_2 \times \dots \times n_m}$, vector \vec{v} indicating a position in A and some $b \in F$. We denote by $A \oplus_{\vec{v}} b$ the matrix $B \in F^{n_1 \times n_2 \times \dots \times n_m}$, such that $B_{\vec{v}} = b$ and $B_{\vec{u}} = A_{\vec{u}}$ for all $\vec{u} \neq \vec{v}$.

In practice, the operator $\oplus_{\vec{v}}$ will be used to replace a vector of payoffs (that appears in the position \vec{v}) with some other vector. So the set F in our case will consist of all appropriately-sized vectors of real numbers, and b will be one such vector. This replacement reflects the effect of the Adaptation Procedure in the subjective payoff matrices of players, as it leads to the updating of an erroneous payoff vector with the actual one that is communicated to them by the game (and appears in the actual payoff matrix). These ideas are formalised below:

Definition 8. Consider a canonical misinformation game $mG = \langle G^0, G^1, \dots, G^{|N|} \rangle$, where $G^i = \langle N, S, P^i \rangle$ (for $0 \leq i \leq |N|$), and some vector \vec{v} . We define the \vec{v} -update of mG , denoted by $mG_{\vec{v}}$, to be the misinformation game $\langle G^0, G^1, \dots, G^{|N|} \rangle$, where $G^{i'} = \langle N, S, P^i \oplus_{\vec{v}} P_{\vec{v}}^0 \rangle$, for $1 \leq i \leq |N|$.

Definition 8 tells us how to perform the update process that the Adaptation Procedure requires, in particular by replacing the payoffs of the newly learnt position in each of the subjective payoff matrices (P^i) by the actual payoff given by P^0 (i.e., by $P_{\vec{v}}^0$). Interestingly, the \vec{v} -update of mG is idempotent, $(mG_{\vec{u}_1})_{\vec{u}_1} = mG_{\vec{u}_1}$, and commutative, $(mG_{\vec{u}_1})_{\vec{u}_2} = (mG_{\vec{u}_2})_{\vec{u}_1}$.

Abusing notation, for a set of positions $X = \{\vec{u}_1, \dots, \vec{u}_k\}$, we denote by mG_X the game $mG_X = (\dots (mG_{\vec{u}_1})_{\vec{u}_2} \dots)_{\vec{u}_k}$. Given the properties above, the notation mG_X is well-defined.

The position where the update takes place (denoted by \vec{v} in Definition 8) is determined by the strategic choices of players, and can be “extracted” from a strategy profile that is an *nme*, using the following definition:

Definition 9. Consider a strategy profile $\sigma = (\sigma_1, \dots, \sigma_N)$ with $\sigma_i \in [0, 1]^{|S_i|}$ and $S_i = \{s_{i1}, \dots, s_{i|S_i|}\}$. The characteristic strategy set of vectors of σ is $\chi(\sigma) = \chi(\text{supp}(\sigma_1)) \times \dots \times \chi(\text{supp}(\sigma_N))$, with $\chi(\text{supp}(\sigma_j)) = \{i | s_{ji} \in \text{supp}(\sigma_j)\}$.

In practice, χ identifies the indices of the strategies that each player plays with non-zero probability, and then returns a set containing all possible vectors that can be formed with these indices (where each point in the vector corresponds to a player). This will be further clarified with the following example:

Example 2. Assume a 4×3 bimatrix game. Then, the characteristic strategy set of vectors of $\sigma = ((1/2, 0, 1/3, 1/6), (0, 0, 1))$ is $\chi(\sigma) = \{(1, 3), (3, 3), (4, 3)\}$. \square

As explained above, the Adaptation Procedure occurs in discrete time steps $t \in \mathbb{N}_0 = \mathbb{N} \cup \{0\}$. It starts from $t = 0$ where player i has the view $G^{i,(0)}$, $\forall i \in [N]$, and in each time step t we implement the update operation described in Definition 8 for the payoff vector(s) that correspond to the strategic choices of the players. In other words, we update the strategy values of the strategy profiles of the players. The following example illustrates this procedure using the above notions, and is also visualised in Figure 3:

Example 1 (continued). We write mG as $mG^{(t)} = \langle G^0, G^{1,(t)}, G^{2,(t)} \rangle$. Using the support of the nme and the characteristic strategy vector we take $\chi\{((0, 1), (1/2, 1/2))\} = \{(2, 1), (2, 2)\}$. So, we have the strategy value vector $(2, 1)$ for the strategy profile $((0, 1), (1, 0))$, and the strategy value vector $(2, 2)$ for the strategy profile $((0, 1), (0, 1))$.

Notice that, as one player is randomized, $\chi(nme)$ has more than one elements, and the Adaptation Procedure branches result to two new misinformation games, say $mG^{(1a)}$, $mG^{(1b)}$. The first one becomes, $mG^{(1a)} = \langle G^0, G^{1,(1a)}, G^{2,(1a)} \rangle$ with payoff matrices as shown in Table 2 (note how the bottom-left payoff has been updated).

Table 2: Payoff matrices in the misinformation game of the Example 1 using the nme $((0, 1), (1, 0))$ for the update.

	s_1	s_2
s_1	(6,6)	(2,7)
s_2	(7,2)	(1, 1)

(a) P^0

	s_1	s_2
s_1	(2,2)	(0,3)
s_2	(7, 2)	(1,1)

(b) $P^{1,(1a)}$

	s_1	s_2
s_1	(-1,1)	(2,-2)
s_2	(7, 2)	(0,0)

(c) $P^{1,(1a)}$

Similarly, the second one becomes, $mG^{(1b)} = \langle G^0, G^{1,(1b)}, G^{2,(1b)} \rangle$ with the payoff matrices shown in Table 3 (note that the bottom-right payoff has been updated).

The procedure shown in Example 1 is formalised as follows, taking into account the fact that the process may branch when $\chi(\sigma)$ is not a singleton set:

Definition 10. For a set M of misinformation games, we set:

$$AD(M) = \{mG_{\vec{u}} \mid mG \in M, \vec{u} \in \chi(\sigma), \sigma \in NME(mG)\}$$

Table 3: Payoff matrices in the misinformation game of the Example 1 using the nme $((0, 1), (0, 1))$ for the update.

	s_1	s_2
s_1	(6,6)	(2,7)
s_2	(7,2)	(1,1)

(a) P^0

	s_1	s_2
s_1	(2,2)	(0,3)
s_2	(3,0)	(1, 1)

(b) $P^{1,(1b)}$

	s_1	s_2
s_1	(-1,1)	(2,-2)
s_2	(1,-1)	(1, 1)

(c) $P^{2,(1b)}$

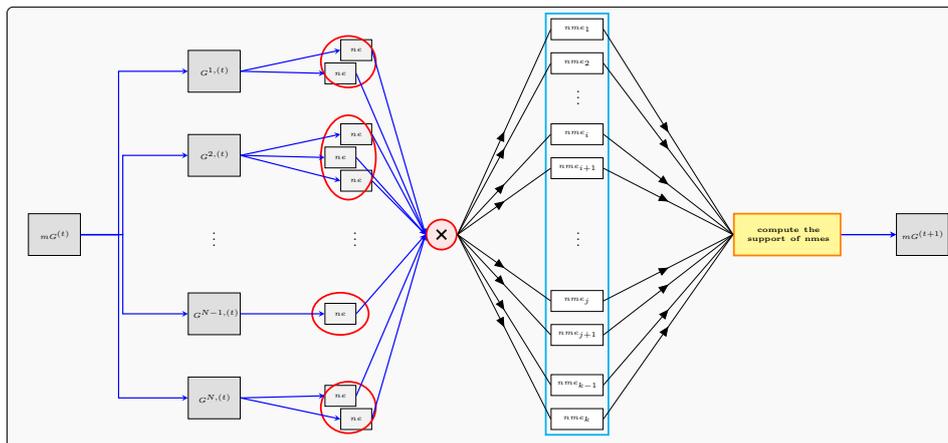


Fig. 2: Schematic representation of the functionality of Adaptation Procedure from time step t to time step $t + 1$, that is $\mathcal{AD}^{t+1}(mG)$.

We define the Adaptation Procedure as the following iterative process:

$$\begin{cases} \mathcal{AD}^{(0)}(M) = M \\ \mathcal{AD}^{(t+1)}(M) = \mathcal{AD}^{(t)}(\mathcal{AD}(M)) \end{cases} \quad (2)$$

for $t \in \mathbb{N}_0$.

The functionality of the Adaptation Procedure between two consecutive time steps t and $t + 1$, as provided by Definition 10, is depicted in Figure 2.

Note that the Adaptation Procedure is defined over a set of misinformation games. Although our intent is basically to apply it over a single misinformation game, the branching process, along with the recursive nature of the definition, forces us to consider the more general case right from the start. Note also that we will often abuse notation and write $\mathcal{AD}(mG)$ (or $\mathcal{AD}^{(t)}(mG)$) instead of $\mathcal{AD}(\{mG\})$ (or $\mathcal{AD}^{(t)}(\{mG\})$).

The following example shows the second step of the Adaptation Procedure of Example 1. Interestingly, $mG^{(1)}$ includes a hybrid natural misinformed equilibrium (i.e., one in which one player plays a pure strategy and the other a mixed one), thus illustrating the branching process mentioned above.

Example 1 (continued). At $t = 1$, the Adaptation Procedure is branched. We demonstrate each branch separately.

The first branch results from updating profile (s_2, s_1) in $mG^{(0)}$. In the new game, player 1 has equilibrium strategy s_2 in $G^{1,(1a)}$ and player 2 has equilibrium strategy s_1 in $G^{2,(1a)}$. Thus, the Adaptation Procedure provides a new misinformation game, say $mG^{(2a)}$. Observe that the payoff matrices of $mG^{(1a)}$ are already updated with the correct value with respect to the bottom-left element, therefore $mG^{(2a)} = mG^{(1a)}$.

The second branch results from updating profile (s_2, s_2) in $mG^{(0)}$. In the new game, we have one pure Nash equilibrium in $G^{1,(1b)}$ for player 1, that is $((0, 1), (0, 1))$ and a mixed Nash equilibrium in $G^{2,(1b)}$ with strategy profile $((2/5, 3/5), (1/3, 2/3))$ for player 2. Thus, the nme is $((0, 1), (1/3, 2/3))$. Applying the characteristic strategy vector $\chi((0, 1), (1/3, 2/3))$ yields $\{(2, 1), (2, 2)\}$, leading to further branching into new misinformation games $mG^{(2b)}$ and $mG^{(2c)}$.

Let us consider the element $(2, 2)$ of $\chi((0, 1), (1/3, 2/3))$ (which leads to $mG^{(2c)}$). We note that the payoff matrices of $mG^{(1b)}$ are already updated, specifically the bottom-left element, making $mG^{(2c)} = mG^{(1b)}$.

Similarly, for the element $(2, 1)$ of $\chi((0, 1), (1/3, 2/3))$, we update the bottom-right element of $P^{1,(1b)}$ and $P^{2,(1b)}$, so $mG^{(1b)}$ leads to $mG^{(2b)} = \langle G^0, G^{1,(2b)}, G^{2,(2b)} \rangle$ with payoff matrices as shown in Table 4.

Table 4: Payoff matrices in the misinformation game of the Example 1 at time $t = 1$, using the nme $((0, 1), (1/3, 2/3))$.

	s_1	s_2
s_1	(2,2)	(0,3)
s_2	(7, 2)	(1, 1)

(a) $P^{1,(2b)}$

	s_1	s_2
s_1	(-1,1)	(2,-2)
s_2	(7, 2)	(1, 1)

(b) $P^{2,(2b)}$

In summary, we deduce that $AD^{(2)}(\{mG^{(0)}\}) = \{mG^{(2a)}, mG^{(1b)}, mG^{(2c)}\}$.

3.3 Stabilisation of the Adaptation Procedure

The following definition determines the final step of the procedure; this corresponds to the time point where any further iterations do not provide new information to the players:

Definition 11. We say that the Adaptation Procedure terminates (or stops) at step t , iff t is the smallest non-negative integer for which

$$AD^{(t+1)}(M) = AD^{(t)}(M)$$

for $t \in \mathbb{N}_0$. We call t the length of the Adaptation Procedure and is denoted by $\mathfrak{L}_{AD}(M)$.

In essence, the Adaptation Procedure concludes when all spawned misinformation games appear already in the input. We will demonstrate (see Theorem 11) that termination is a well-defined concept for all finite misinformation games, as they invariably reach a termination point.

To simplify notation, let $AD^*(M)$ represent the set of all misinformation games created by M through successive applications of the $AD(\cdot)$ function, i.e., $AD^*(M) = \bigcup_{t=0}^{\infty} AD^{(t)}(M)$. Let $AD^\infty(M)$ be the set of misinformation games produced by $AD(\cdot)$ post-termination, i.e., $AD^\infty(M) = AD^{(t)}(M)$ for $t = \mathfrak{L}_{AD}(M)$. $AD^\infty(M)$ will be called the *Stable Set*.

Definition 12. Consider a misinformation game mG . Then, σ is a stable misinformed equilibrium (or *sme* for short) of mG , iff there exists some $\widehat{mG} \in AD^\infty(\{mG\})$ such that $\sigma \in NME(\widehat{mG})$ and, for all $\vec{v} \in \chi(\sigma)$, $\widehat{mG}_{\vec{v}} = \widehat{mG}$.

We denote by $SME(mG)$ the *smes* of mG . The careful reader might wonder why we didn't define *smes* as the *nmes* of the stable set. The extra condition (for all $\vec{v} \in \chi(\sigma)$, $\widehat{mG}_{\vec{v}} = \widehat{mG}$) ensures that the *nme* satisfies some condition of "stability" in the sense that the players already know the actual payoffs of the positions being played. Note that this is not necessarily true for all *nmes*, even when restricting ourselves to the *nmes* of mGs that appear in the stable set: an *nme* could lead to learning new positions, creating different mGs , which happen to be in the stable set. Such an *nme* is not an *sme*, according to Definition 12; instead, an *nme* is an *sme* if and only if the players learn nothing because of it. The following example illustrates this scenario.

Example 1 (continued). For $t = 2$, let us first consider $mG^{(2a)}$. As mentioned, $mG^{(2a)} = mG^{(1a)}$, leading to $mG^{(3a)} = mG^{(1a)}$.

Analogously for $mG^{(2c)}$, we observe that $mG^{(2c)} = mG^{(1b)}$ so, again, from the analysis in the previous steps of Example 1, it branches into $mG^{(3a)} = mG^{(1b)}$, $mG^{(3b)} = mG^{(2b)}$.

Finally, for $mG^{(2b)}$, we have that both players' subjective games, $G^{1,(2b)}$ and $G^{2,(2b)}$, have a pure Nash equilibrium with strategy profile $((0, 1), (1, 0))$. Thus, the *nme* of $mG^{(2b)}$ is $\sigma = ((0, 1), (1, 0))$, for which $\chi(\sigma) = \{(2, 1)\}$. Observe that this position in $\chi(\sigma)$ is known to the players, i.e., it holds that $P_{\vec{v}}^{i,(2b)} = P_{\vec{v}}^0$, for $i \in \{1, 2\}$, $\vec{v} = \{(2, 1)\}$. Thus, $mG^{(3c)} = mG^{(2b)}$.

Combining the above, we observe that $AD^{(3)}(\{mG^{(0)}\}) = \{mG^{(1)}, mG^{(2b)}\} = AD^{(2)}(\{mG^{(0)}\})$, so the Adaptation Procedure terminates at step 2, i.e., $\mathfrak{L}_{AD}(\{mG^{(0)}\}) = 2$.

Now let us identify the *smes* of $mG^{(0)}$. As explained above, $NME(mG^{(1a)}) = \{((0, 1), (1, 0))\}$, and $NME(mG^{(2a)}) = \{((0, 1), (1, 0))\}$. Thus, profile $\sigma = ((0, 1), (1, 0))$ is an *sme*. However, the sole *nme* of $mG^{(1b)}$, $((0, 1), (1/3, 2/3))$, does not meet the criteria for an *sme* since $mG_{\vec{v}} \notin AD^\infty(mG)$ for $\vec{v} = (2, 1)$, and thus $mG^{(1b)}$ fails the second condition of Definition 12.

In Figure 4, we demonstrate the relationships between the three equilibrium concepts, namely NE, *nme*, and *sme*. The top layer represents each player's subjective view, where they choose according to NE in their perceived game. In the second layer (the misinformation game at $t = 0$), *nmes* are formed by combining the supports of NE strategies (see Definition 6). Players then update their information based on these NMEs and corresponding payoffs, iterating through the Adaptation Procedure (the next two layers in the figure), according to Definitions 8 and 10. Finally, at termination (i.e., the bottom layer in the figure), the surviving *nmes* become *smes*.



Fig. 3: Schematic representation of Example 1.

3.4 Adaptation Procedure: Properties

3.4.1 General properties

We will start by providing two useful properties of the Adaptation Procedure. The first one shows that the Adaptation Procedure is “local”, i.e., the different misinformation games in the input of $AD(\cdot)$ do not interact with each other:

Proposition 7. *For any set of misinformation games M , $AD(M) = \bigcup_{mG \in M} AD(\{mG\})$.*

The next two results essentially show that an mG that appears in any step of the Adaptation Procedure cannot reappear in any subsequent step, unless it stays there permanently:

Proposition 8. *Take some finite sequence mG_1, \dots, mG_n such that $mG_{i+1} \in AD(\{mG_i\})$, for $i \in [n-1]$ and $mG_1 \in AD(\{mG_n\})$. Then, $mG_i = mG_j$ for all i, j .*

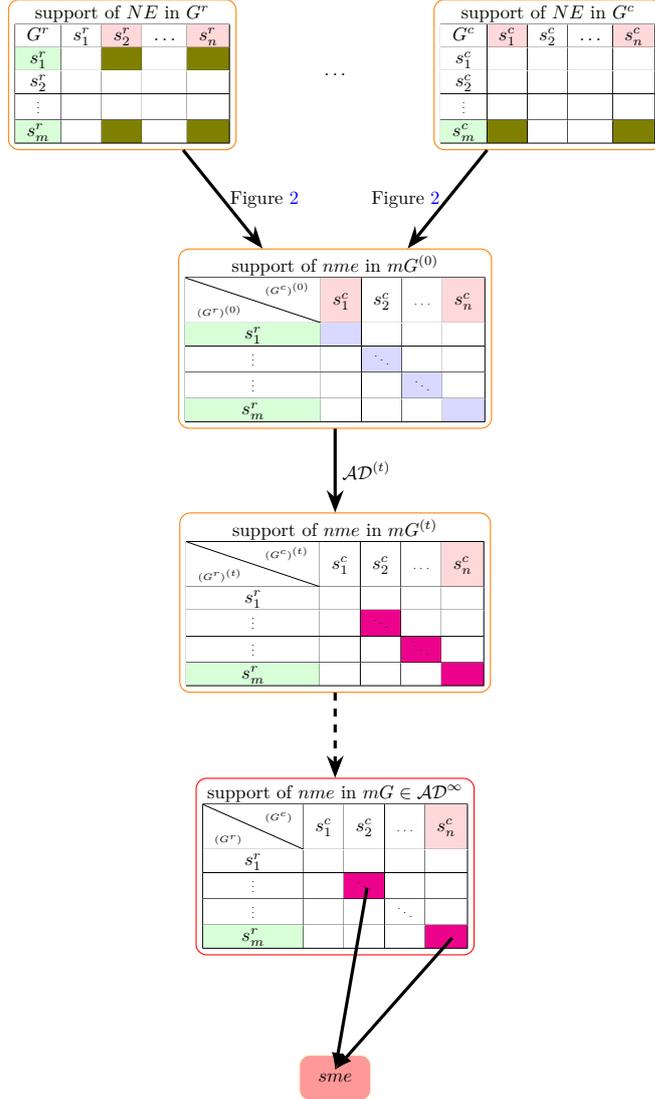


Fig. 4: Schematic summarization for the relationship among NE, *nme* and *sme*. An *nme* is constructed by the supports, of the Nash equilibria, of each player’s subjective game. Thus, the *nme* constitutes the joined decision the players reach operating individually. An *sme* is an *nme* that is stable in the sense that the players do not have an incentive move away from, based on their subjective information.

Proposition 9. For any two misinformation games mG , mG' and $t \geq 0$, if $mG' \in AD^{(t)}(mG)$ and $mG' \in AD(mG')$ then $mG' \in AD^\infty(mG)$.

From Definition 12, not all misinformation games in the Stable Set contribute an *sme*. In other words, some of the games in $AD^\infty(mG)$ are irrelevant when it comes

to computing the *smes*. Moreover, the termination criterion outlined in Definition 11 indicates a “global” situation, where the sets of misinformation games are identical across two successive steps. The set-theoretic comparison of two sets of misinformation games would be computationally expensive, so a more “local” termination criterion is preferable. Therefore, an interesting question is whether it is safe to discard a misinformation game from future iterations. This problem is addressed jointly with the notion of the terminal set.

Definition 13. Let $mG^{(0)}$ be a misinformation game. We define the terminal set of the Adaptation Procedure on $mG^{(0)}$ as follows:

$$\mathcal{T} = \{mG \in \mathcal{AD}^*(mG^{(0)}) \mid mG \in \mathcal{AD}(mG)\}.$$

Example 1 (continued). The terminal set in Example 1 is $\mathcal{T} = \{mG^{(1a)}, mG^{(2b)}\}$.

The next proposition shows that the games in the terminal set are also in the Stable Set; in addition, they are the only members of the Stable Set that matter when it comes to computing *smes*:

Proposition 10. Let \mathcal{T} be the terminal set of the Adaptation Procedure on $mG^{(0)}$. Then:

- $\mathcal{T} \subseteq \mathcal{AD}^\infty(mG^{(0)})$
- For any $\sigma \in \text{SME}(mG^{(0)})$, there exists $mG \in \mathcal{T}$ such that $\sigma \in \text{NME}(mG)$.

Note that the opposite of the second bullet of Proposition 10 does not hold: a misinformation game in the terminal set may contain *NMEs* that are not themselves *SMEs*, because they “lead” to another misinformation game, thereby violating the condition “for all $\vec{v} \in \chi(\sigma)$, $\widehat{mG}_{\vec{v}} = \widehat{mG}$ ” of Definition 12.

Consider the implications of Proposition 10 on the computational aspects of the Adaptation Procedure. When a game mG is not a member of the terminal set \mathcal{T} , our interest shifts to its outcomes, specifically $\mathcal{AD}(mG)$, rather than the game itself, since it will not yield an *sme*. In contrast, if mG is part of the terminal set, we retain it without further processing, as it will perpetually recur (as well as its outcomes), rendering any additional processing superfluous. On the other hand, the games in $\mathcal{AD}(mG) \setminus \{mG\}$ may be relevant, if any of them (or any of their descendants) is in the terminal set.

Finally, Figure 5 shows the relationship between the Stable Set and the Terminal Set (\mathcal{T}) in the Adaptation Procedure for misinformation games. The rectangles represent sets of misinformation games at different stages, from the initial set $\mathcal{AD}^{(0)}(mG^{(0)})$ to the Stable Set $\mathcal{AD}^\infty(mG^{(0)})$ after the procedure terminates. The smaller red rectangle within the Stable Set is the \mathcal{T} , containing the misinformation games critical for computing *smes*. Only *mGs* in the \mathcal{T} recur and are directly relevant to *sme* computation, while not all *mGs* in the Stable Set contribute. *mGs* outside the \mathcal{T} may lead to further steps, but those within \mathcal{T} recur without changes, eliminating the need for further processing.

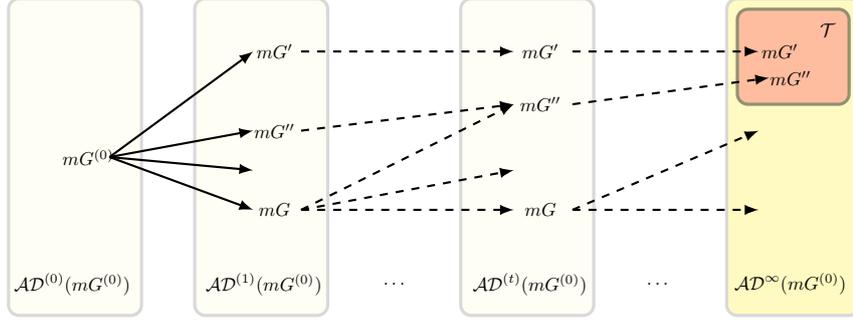


Fig. 5: Relationship between the Stable and Terminal sets. As we show in Proposition 11 the Adaptation Procedure reaches a fixed point after a finite number of steps. We call this fixed point the *Stable Set* of the Adaption Procedure $\mathcal{AD}^\infty(mG^{(0)})$. The Stable Set $\mathcal{AD}^\infty(mG^{(0)})$ will be repeated ad infinitum. The Terminal Set is the set of the misinformation games the generate themselves. In Theorem 14 we will see that in a sense the Terminal Set *generates* the Stable Set, thus being responsible for the Adaptation Procedure reaching a fixed point.

3.4.2 Termination, and existence of smes

In this subsection we show some results related to the termination of the Adaptation Procedure.

We start by showing that the Adaptation Procedure will always terminate when mG is finite. In particular, when mG is finite, then at any given point in the Adaptation Procedure a finite number of positions can be learnt, and the positions that remain to be learnt are also finite. Therefore the length of the procedure is bounded, as well as the total number of mG s that can be generated at any step, so the set generated by the Adaptation Procedure is finite and the Adaptation Procedure terminates. This result is formally phrased in the next propositions.

Proposition 11. *Consider a finite $|N|$ -player canonical misinformation game mG on $|S|$ strategies. For the length $\mathfrak{L}_{AD}(mG)$ of the Adaptation Procedure on mG , we have,*

$$\mathfrak{L}_{AD}(mG) \leq |S| \quad (3)$$

Proposition 12. *Let mG be a $|N|$ -player canonical misinformation game on $|S|$ strategy profiles. Then for the number $|\mathcal{AD}^*(\{mG\})|$ of all the misinformation games produced during the Adaptation Procedure on mG , it holds that,*

$$|\mathcal{AD}^*(\{mG\})| = O\left(\min\left[|S|^{\mathfrak{L}_{AD}(mG)+1}, 2^{|S|}\right]\right)^2. \quad (4)$$

²When $\mathfrak{L}_{AD}(mG) \ll |S|$, then $|S|^{\mathfrak{L}_{AD}(mG)+1} < 2^{|S|}$. On the other hand, when $|S| \approx \mathfrak{L}_{AD}(mG)$ or $|S| \gg \mathfrak{L}_{AD}(mG)$, we have $2^{|S|} < |S|^{\mathfrak{L}_{AD}(mG)+1}$.

Intuitively, Proposition 12 proves that the total number of misinformation games produced during the Adaptation Procedure on a given initial mG depends on two factors: the number of strategy profiles and the length of the Adaptation Procedure.

Corollary 13. *Assume a $|N|$ -player canonical misinformation game mG , on $|S|$ strategies. Then the stable set $AD^\infty(mG)$ is also finite.*

Unfortunately, a similar result cannot be shown for infinite games. The following counter-example proves this fact:

Example 3. *Consider $G^0 = \langle N, S, P^0 \rangle$ such that $N = \{r, c\}$, $S_r = S_c = \{1, 2, \dots\}$, $S = S_r \times S_c$, and the payoff for a position $(x, y) \in S_r \times S_c$ is computed as follows:*

$$P_{(x,y)}^0 = \left(\frac{1}{x}, \frac{1}{y} \right)$$

Observe that the only (pure) Nash equilibrium for G^0 is in position $(1, 1)$, where the payoff is $(1, 1)$.

Now consider the canonical misinformation game $mG = \langle G^0, G^r, G^c \rangle$, where $G^r = \langle N, S, P^r \rangle$ such that $P_{(x,y)}^r = (\frac{1}{x} + 1, \frac{1}{y})$, and $G^c = G^0$. By the definition of mG we note that player c knows the correct payoffs and will always play strategy 1 (she will never learn anything by the Adaptation Procedure). On the other hand, player r knows the correct payoffs as far as player c is concerned, but her own payoffs are distorted, and she believes that the value is 1 point more than the actual one. The key observation is that, for player r , any of her subjective payoffs is better than any of her actual ones. Therefore, when she learns any position, this position becomes highly unattractive and cannot be selected again.

Formally, note that $NME(mG) = \{\sigma_0\}$, where $\chi(\sigma_0) = \{(1, 1)\}$. Thus, $AD^{(1)}(\{mG\}) = \{mG^{(1)}\}$, where $mG^{(1)} = mG_{(1,1)}$. It is easy to see that $NME(mG^{(1)}) = \{\sigma_1\}$, where $\chi(\sigma_1) = \{2, 1\}$. Continuing this process, we observe that $AD^{(i)}(\{mG\}) = \{mG^{(i)}\}$, where $mG^{(i)}$ is such that positions $(1, 1), (2, 1), \dots, (i, 1)$ have been learnt. But the only (pure) nme of $mG^{(i)}$ corresponds to position $(i + 1, 1)$. As a result, the Adaptation Procedure will continuously lead to the learning of new positions (and, thus, to new misinformation games), which shows that the Adaptation Procedure will not terminate. \square

Due to this negative result, all subsequent analysis focuses on finite misinformation games.

The following theorem characterises a Stable Set as the closure of the terminal set with respect to the adaptation operator $AD(\cdot)$.

Theorem 14 (Stable Set Characterisation). *Let $AD^\infty(mG)$ be the Stable Set and \mathcal{T} the terminal set of the Adaptation Procedure on $mG^{(0)}$. Then,*

$$AD^\infty(mG) = AD^*(\mathcal{T}).$$

Intuitively the above theorem shows that a misinformation game mG belongs to the Stable Set of the Adaptation Procedure if it is either in the terminal set itself or has an ancestor that belongs to the terminal set. Moreover, Theorem 14 shows that the terminal set contains all the information of the Adaptation Procedure on a

misinformation game mG^0 . Not only all the *smes* of the Adaptation Procedure are included in the elements of \mathcal{T} (see Proposition 10), but also we can *reconstruct* the Stable Set $\mathcal{AD}^\infty(mG^0)$, from the terminal set \mathcal{T} . As we shall see in Section 4, keeping track of the terminal set suffices to compute the Adaptation Procedure.

We will now show that all finite misinformation games have an *sme*. We start with a result that describes a condition sufficient for the existence of an *sme*. In particular, Proposition 15 states that if the players in some misinformation game mG' learn nothing new from the respective *nmes* (i.e., if $\mathcal{AD}(mG') = \{mG'\}$), then the *nmes* of mG' are also *smes*. Formally:

Proposition 15. *If $mG' \in \mathcal{AD}^*(\{mG\})$ and $\mathcal{AD}(\{mG'\}) = \{mG'\}$, then $NME(mG') \subseteq SME(mG)$.*

Proposition 16. *If mG is finite, then $SME(mG) \neq \emptyset$.*

An immediate observation from Proposition 16 is the following:

Corollary 17. *If mG is finite, then $\mathcal{T} \neq \emptyset$.*

Proposition 18 below shows that the upper bound on $\mathfrak{L}_{\mathcal{AD}}(mG)$ provided by equation 3 is tight.

Proposition 18. *Let N be a set of players and S a set of strategies. Then there is a $|N|$ -player canonical misinformation game mG , on $|S|$ strategies, such that,*

$$\mathfrak{L}_{\mathcal{AD}}(mG) = |S|.$$

4 Computing the Adaptation Procedure

In this section, we develop an algorithm to test the properties of the Adaptation Procedure in practice, and investigate its behaviour across a spectrum of randomly generated misinformation games. We also run experiments and discuss about the results.

4.1 The adaptation graph

We encode the computation of the Adaptation Procedure as a graph search problem. We will see that the Adaptation Procedure operator induces a structure on the set of all misinformation games $\mathcal{AD}^*(mG)$ derived from the Adaptation Procedure on mG . We call this structure *adaptation graph*.

Definition 14. *Let mG be a misinformation game, and the set of all misinformation games $\mathcal{AD}^*(mG)$ derived from the Adaptation Procedure on mG . We denote by $\Gamma = (\mathcal{AD}^*(mG), E)$ the directed graph where,*

$$E = \left\{ (mG^1, mG^2) \mid \text{there is } \vec{v} \in \chi(NME(mG^1)), \text{ s.t. } mG^2 = (mG^1)_{\vec{v}} \right\}. \quad (5)$$

Further, we denote by $\Gamma' = (\mathcal{AD}^*(mG), E')$ the loopless version of Γ where,

$$E' = E \setminus \left\{ (mG^1, mG^2) \in E \mid mG^1 = mG^2 \right\}. \quad (6)$$

Some notes on Definition 14 are in order. First, by Corollary 13, Γ and Γ' are finite. Moreover, comparing Definition 14 with Definition 10, we note that the edge (mG^1, mG^2) belongs to the adaptation graph Γ if and only if $mG^2 \in \mathcal{AD}(mG^1)$. Further, the definition of terminal set \mathcal{T} (Definition 13) implies that each misinformation game mG that belongs to \mathcal{T} has a self-loop in Γ . Moreover, as the following proposition shows, Γ' is a *directed acyclic graph (DAG)*, i.e., Γ is “almost” a DAG.

Proposition 19. *Let $\Gamma = (\mathcal{AD}^*(mG), E)$ be the adaptation graph of the Adaptation Procedure on a misinformation game mG and Γ' its loopless version. Then Γ' is a directed acyclic graph.*

Focusing on Γ' , we can explore more effectively the adaptation graph Γ . In the following analysis we will use the concept of the degree of a vertex, where $d_{\Gamma}^-(v)$ is the number of the incoming edges and $d_{\Gamma}^+(v)$ is the number of the outgoing edges in graph Γ .

Proposition 20. *Let $\Gamma = (\mathcal{AD}^*(mG), E)$ be the adaptation graph, of the Adaptation Procedure, on a misinformation game mG and Γ' its loopless version. Then, mG is the only source node of Γ' . That is, mG is the only node in Γ' such that*

$$d_{\Gamma'}^-(mG) = 0. \quad (7)$$

Moreover, we are able to characterize the *sink* nodes of Γ' as members of the terminal set \mathcal{T} , and we argue that *every nme* of such a node is an *sme*.

Proposition 21. *Let $\Gamma = (\mathcal{AD}^*(mG), E)$ be the adaptation graph of the Adaptation Procedure on a misinformation game mG and Γ' its loopless version. Let K be the set of sink nodes in Γ' that is*

$$K = \{mG' \in \mathcal{AD}^*(mG) \mid d_{\Gamma'}^+(mG') = 0\}. \quad (8)$$

Then, $K \subseteq \mathcal{T}$. Moreover, every nme of some $mG' \in K$ is an sme of the Adaptation Procedure on mG .

Note that the reverse direction of Proposition 21 is not true. That is, in general, $K \subsetneq \mathcal{T}$, and there are *smes* which are not *nmes* of a sink node. Continuing our exploration of the structure of Γ' , we remark that the *longest path* in Γ' is $\mathfrak{L}_{\mathcal{AD}}(mG)$.

Proposition 22. *Let $\Gamma = (\mathcal{AD}^*(mG), E)$ be the adaptation graph of the Adaptation Procedure on a misinformation game mG and Γ' its loopless version. Then, the length of the longest path in Γ' is $\mathfrak{L}_{\mathcal{AD}}(mG)$.*

Let us summarize what the above results say about the form of the adaptation graph Γ . First, the adaptation graph is “almost” a DAG; when we omit the self-loops from Γ , the resulting graph Γ' is a DAG. The graph Γ' begins with a single source node, namely the root of the Adaptation Procedure. Additionally, Γ' extends for $\mathfrak{L}_{\mathcal{AD}}(mG)$ length units, concluding to a set of sink nodes, each of which belongs to the terminal set \mathcal{T} . Moreover, every *nme* of a sink node is an *sme*. The ideas highlighted here will be exploited in the next subsection to develop algorithms for computing either one, or all *smes*; as we shall see, there is an intriguing asymmetry between these two problems (computing a single *sme* or all *smes*), that we explore next.

4.2 Set-based representation: efficiently detecting self-loops

In the previous subsection we captured the graph structure induced on the space of misinformation games by the Adaptation Procedure. We named this structure the *adaptation graph* Γ . We also showed that we can characterise the terminal set of the Adaptation Procedure as the nodes in Γ that contain self-loops. Moreover, we observed that we can omit these self-loops during graph exploration, resulting in the graph Γ' . Since the self-loop detection constitutes one of the main operations that will be performed at each step of the graph exploration, we want it to be as efficient as possible. In this direction, we define the *set-based representation* $\tilde{\Gamma}$ of the adaptation graph Γ , showing that this reduces the computational cost of self-loop detection. Lastly, we will explore $\tilde{\Gamma}$, the same way we did with Γ , omitting the self-loops and thus resulting in the graph $\tilde{\Gamma}'$.

Assume a misinformation game mG' in the adaptation graph Γ . In order to detect whether mG contains a self-loop, we need to determine whether there exists some $\sigma \in NME(mG)$ and some $\vec{v} \in \chi(\sigma)$ such that $mG = mG_{\vec{v}}$. To determine whether $mG = mG_{\vec{v}}$, we need to verify that in both misinformation games, the players have equal subjective and objective payoffs. The computational complexity of this operation is given by the following result.

Proposition 23. *Let mG^1, mG^2 be two canonical $|N|$ -player misinformation games with $|S|$ joint strategies. Then the computational complexity to determine if $mG^1 = mG^2$ is*

$$\Theta(|N|^2 \cdot |S|). \quad (9)$$

To reduce the computational cost of self-loop detection, we will adopt a representation of a misinformation game that allows us to omit equality checks, when possible. In this direction, we will consider the *set-based representation* $\tilde{\Gamma}$ of the adaptation graph Γ . Recall that S is the set of all position vectors in a misinformation game mG . Consider a mapping $\theta: 2^S \rightarrow \mathcal{AD}^*(mG)$ where,

$$\theta[X] = mG_X. \quad (10)$$

The mapping $\theta[\cdot]$ is surjective by construction but not injective, as different sets of position vectors may result in the same misinformation game, e.g., when the subjective payoffs of the players coincide with the objective payoffs at this position vector. Using $\theta[\cdot]$ we can define the *set-based representation* $\tilde{\Gamma}$.

Definition 15. *Let $\Gamma = (\mathcal{AD}^*(mG), E)$ be the adaptation graph, of the Adaptation Procedure, on a misinformation game mG . We denote with $\tilde{\Gamma} = (2^S, \tilde{E})$ the set-based representation of Γ , where*

$$\tilde{E} = \{(X, Y) \mid Y = X \cup \{\vec{v}\}, X \neq Y \text{ and the edge } (\theta[X], \theta[Y]) \text{ exists in } \Gamma\}.$$

Further, we denote by $\tilde{\Gamma}' = (2^S, \tilde{E}')$ the loopless version of $\tilde{\Gamma}$ where,

$$\tilde{E}' = \tilde{E} \setminus \{(X, Y) \in \tilde{E} \mid \theta[X] = \theta[Y]\}.$$

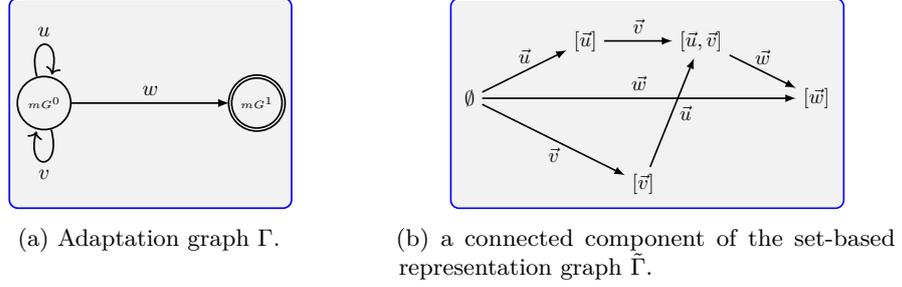


Fig. 6: Correspondence between the adaptation graph Γ and $\tilde{\Gamma}$.

Note that, with respect to Definition 15, the vertex set of $\tilde{\Gamma}$ contains all the position vectors sets, i.e. the whole powerset 2^S . From Proposition 12, in the worst case, all the position vector sets will be connected to the root of the Adaptation Procedure. On the other hand, in general, in $\tilde{\Gamma}$ there will be isolated nodes, i.e. nodes with zero in and out degree. Subsequently, we implicitly discard the isolated nodes, and only consider the connected component of $\tilde{\Gamma}$ to which the root of the Adaptation Procedure belongs. The existence of isolated nodes, does not affect the subsequent worst case analysis of our algorithms.

In Figure 6 we depict an adaptation graph Γ and its corresponding set-based representation $\tilde{\Gamma}$. In $\tilde{\Gamma}$, the misinformation games are represented as position vector sets, which are implemented as sorted lists. Additionally, we keep a dictionary that implements $\theta[\cdot]$, by mapping position vector sets to misinformation games. Before performing an equivalence check between two misinformation games, we check whether their set-based representations are equal. Note that two position vector sets X, Y that are connected with an edge in $\tilde{\Gamma}$, differ by *at most* one element \vec{v} . Given the position vector \vec{v} , we can perform a set equivalence check between X, Y , by only checking whether \vec{v} belongs to X . This can be done in sub-linear time, with a simple binary search (given that position vector sets are implemented as sorted lists).

Proposition 24. *Let $\Gamma = (\mathcal{AD}^*(mG), E)$ be the adaptation graph of the Adaptation Procedure on a canonical misinformation game mG , and $\tilde{\Gamma} = (2^S, \tilde{E})$ be the set-based representation of Γ . Assume two nodes X, Y of $\tilde{\Gamma}$, and a position vector $\vec{v} \in \chi(NME(\theta[X]))$, such that $\theta[Y] = \theta[X]_{\vec{v}}$. The set equivalence $X = Y$ can be performed in*

$$O(\log |S|) \tag{11}$$

time, where $|S|$ is the number of joint strategies.

As we see in the sequel, the position vector \vec{v} , required by Proposition 24, is always known, as it is the position vector used to apply the update operation to the parent to obtain the child. Our previous observations regarding Γ and Γ' apply also for $\tilde{\Gamma}$ and $\tilde{\Gamma}'$. In Table 5 we summarize the characteristics of the graphs $\Gamma, \Gamma', \tilde{\Gamma}$, and $\tilde{\Gamma}'$ (visualised in Figure 7).

Before concluding this section, it is important to note that the use of the set equivalence checks does not improve the worst-time analysis. If a set equivalence check fails, a subsequent equality check between the misinformation games is necessary. In

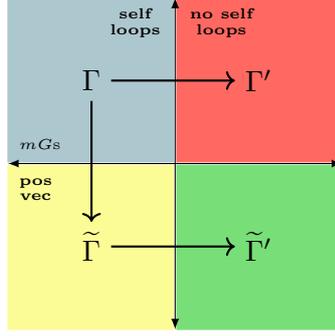


Fig. 7: Correspondence between all variations of adaptation graph.

Table 5: Taxonomy of the adaptation graphs.

Graph	Edges
$\Gamma = (\mathcal{AD}^*(mG), E)$	$E = \left\{ (mG^1, mG^2) \mid \begin{array}{l} \text{there is } \vec{v} \in \chi(\text{NME}(mG^1)), \\ \text{s.t. } mG^2 = (mG^1)_{\vec{v}} \end{array} \right\}$
$\Gamma' = (\mathcal{AD}^*(mG), E')$	$E' = E \setminus \{(mG^1, mG^2) \in E \mid mG^1 = mG^2\}$
$\tilde{\Gamma} = (2^F, \tilde{E})$	$\tilde{E} = \left\{ (X, Y) \mid \begin{array}{l} Y = X \cup \{\vec{v}\}, X \neq Y \\ \text{and the edge } (\theta[X], \theta[Y]) \text{ exists in } \Gamma \end{array} \right\}$
$\tilde{\Gamma}' = (2^F, \tilde{E}')$	$\tilde{E}' = \tilde{E} \setminus \{(X, Y) \in \tilde{E} \mid \theta[X] = \theta[Y]\}$

the worst-case, each set equivalence check would be followed by a misinformation game equality check. However, in practice, we have observed considerable runtime improvements. This practical benefit justifies our decision to base our subsequent algorithms on this approach.

Additionally, recall that we cannot explore the Adaptation Procedure solely based on the set-based representation graph $\tilde{\Gamma}$. Since we need the underlying misinformation games both for computing the *nmes* and for performing equality checks, we need to store the $\theta[\cdot]$ mapping. In the algorithms presented below, $\theta[\cdot]$ is dynamically mapped to a dictionary. The dictionary is expanded whenever a new misinformation game is discovered. For the details involved in implementing $\theta[\cdot]$ see [13]. Moreover, in the sections below, we insist on the notation we used here: given a set of misinformation games \mathcal{X} , the set of position vectors that correspond to a set of strategy profiles \mathcal{S} in \mathcal{X} is $\tilde{\mathcal{X}} := \{\chi(\sigma) \mid \text{for } \sigma \in \mathcal{S} \text{ in } \mathcal{X}\}$. For simplicity we will refer to $\tilde{\mathcal{X}}$ as $\chi(\mathcal{S}(\mathcal{X}))$. Furthermore, in the following we abuse the notation $\theta[\mathcal{X}]$, where \mathcal{X} is a set of position vector sets. Naturally, $\theta[\mathcal{X}] = \{\theta[X] \mid X \in \mathcal{X}\}$.

4.3 Computing ALL-SMEs

In this subsection, we present and analyse an algorithm for computing all the stable misinformed equilibria of a misinformation game. We state the problem formally below.

ALL-SMEs	
Input:	mG , a N -player canonical misinformation game on $ S $ joint strategies.
Output:	SME , the set of all stable misinformed equilibria of mG .

We organize our algorithm in three routines. The main routine is `AdaptationProcedure` (Algorithm 6), which uses two subroutines, `TraverseAdaptationGraph(mG)` (Algorithm 7) and `ComputeSME(\mathcal{T})` (Algorithm 8).

Main routine: Adaptation Procedure

Algorithm 6 implements the Adaptation Procedure in two phases. In the first phase (line 1), we use Algorithm 7 to traverse the adaptation graph using the set-based representation described previously. That is we traverse $\tilde{\Gamma}$. The output of that routine is the terminal set (see Definition 13) and the mapping $\theta[\cdot]$. Subsequently, the second phase (line 2) scans the terminal set to identify the stable misinformed equilibria by calling Algorithm 8. In order to do that, the second phase utilizes the mapping $\theta[\cdot]$ computed in the first phase.

Algorithm 6 `AdaptationProcedure`

Require: A root misinformation game mG .

- 1: $\tilde{\mathcal{T}}, \theta[\cdot] \leftarrow \text{TraverseAdaptationGraph}(mG)$
 - 2: $SME \leftarrow \text{ComputeSME}(\tilde{\mathcal{T}}, \theta[\cdot])$
 - 3: **return** SME
-

Phase 1: Traverse Adaptation Graph

Algorithm 7 (`TraverseAdaptationGraph(mG)`) essentially employs a breadth-first traversal of the set-based representation $\tilde{\Gamma}$ of the adaptation graph to compute the terminal set. To do so, we maintain two sets of sets of position vectors, namely $\tilde{\mathcal{Q}}$ (a queue of the nodes that remain to be processed) and $\tilde{\mathcal{V}}$ (the set of visited nodes). Moreover, we incrementally compute the $\theta[\cdot]$ mapping as we traverse the graph.

In the beginning, the set-based representation of the terminal set $\tilde{\mathcal{T}}$ is empty and the sets $\tilde{\mathcal{Q}}, \tilde{\mathcal{V}}$ start with the value $\{\emptyset\}$, as we initially have to process the root misinformation game mG . We initialize the $\theta[\cdot]$ dictionary with the root misinformation game and compute the root's *nmes*. In each iteration of the **while** loop in lines 4-26, we select (arbitrarily) a set of position vectors $W \in \tilde{\mathcal{Q}}$, and remove W from $\tilde{\mathcal{Q}}$. The set W denotes the node undergoing processing during the current iteration of the **while**

loop. Note that for each new set X , we compute $NME(\theta[X])$ before inserting X to the queue \tilde{Q} .

To process node $W \in \tilde{Q}$, we compute the expansions X of W . An expansion X of W corresponds to $W \cup \{\vec{v}\}$, for some position vector $\vec{v} \in \chi(NME(\theta[W]))$. Note that it is possible to have a position vector $\vec{v} \in \chi(NME(\theta[W]))$, which also belongs to W .

In this case, W has a self-loop in $\tilde{\Gamma}$, thus constituting a terminal node. Lines 7–9 perform this check, and if positive, add W to the terminal set. Additionally, X is another set-based representation for the same misinformation game. Thus $\theta[\cdot]$ is updated to reflect that fact.

For the remaining expansions $X = W \cup \{\vec{v}\}$, with $\vec{v} \notin W$, we firstly check if X is visited, in lines 12–14. If X is visited, we proceed to the next expansion. Otherwise, we append X to the set of visited nodes \tilde{V} , line 15. Recall that, we may have $\theta[W] = \theta[X]$, even if $W \neq X$. For instance, the position vector \vec{v} may correspond to a position of the payoff matrices, already known to the players. Hence, an additional equivalence check between misinformation games should be performed in lines 17–19. If the equivalence check proves negative, then mG_X is a new misinformation game. mG_X is added to the dictionary $\theta[\cdot]$, and its $nmes$ are computed. Lastly, X is inserted into the queue \tilde{Q} .

The algorithm concludes when there are no further unprocessed nodes ($\tilde{Q} = \emptyset$), in which case the complete adaptation graph has been explored, and all misinformation games that must be added to the terminal set are already there. The terminal set is returned in line 27.

Phase 2: Compute SMEs

The role of Algorithm 8 is to determine the *smes* of the Adaptation Procedure. Given the set-based representation of the terminal set $\tilde{\mathcal{T}}$, it traverses the terminal set and checks whether each *nme* in each of the misinformation games in the terminal set satisfies Definition 12; if so, it is an *sme* and it is added to the list of *smes* (maintained in variable *SME*). Proposition 10 guarantees that this procedure will not miss any of the stable misinformed equilibria.

Correctness

Now we discuss the complexity of algorithms 6, 7, 8. Observe that the correctness of algorithms 6, 8 depends solely on the correctness of Algorithm 7. Indeed, assuming that Algorithm 7 correctly computes the terminal set \mathcal{T} , then Algorithm 8 correctly computes the *smes*, since it simply verifies Definition 12 for each element in the terminal set. Moreover, the correctness of Algorithm 6 is implied by Proposition 10. Thus, we only need to validate the correctness of Algorithm 7. In this direction, we need to show that Algorithm 7 terminates when and only when the end criterion of Definition 11 is met, and after considering all the misinformation games in $\mathcal{AD}^*(mG)$.

To simplify the notation of the following proofs we use sets of misinformation games instead of their set-based representations. Consider the sets $\tilde{\mathcal{T}}, \tilde{Q}, \tilde{V}$ of Algorithm 7. In the following we consider the corresponding sets of misinformation games, i.e., $\mathcal{T}, \mathcal{Q}, \mathcal{V}$. Recall that we can transition from a set of sets of position vectors to a set of misinformation games, using the mapping $\theta[\cdot]$, e.g., $\mathcal{Q} = \theta[\tilde{Q}]$.

Algorithm 7 `TraverseAdaptationGraph(mG)`

Require: A root misinformation game mG .

```
1:  $\tilde{\mathcal{T}} \leftarrow \emptyset, \tilde{\mathcal{Q}} \leftarrow \{\emptyset\}, \tilde{\mathcal{V}} \leftarrow \{\emptyset\}$ 
2:  $\theta[\emptyset] \leftarrow mG$ 
3: compute  $NME(\theta[\emptyset])$ 
4: while  $\tilde{\mathcal{Q}} \neq \emptyset$  do
5:   Pick some  $W \in \tilde{\mathcal{Q}}$ 
6:    $\tilde{\mathcal{Q}} \leftarrow \tilde{\mathcal{Q}} \setminus \{W\}$ 
7:   if  $\exists \vec{v} \in \chi(NME(\theta[W])) \cap W$  then
8:      $\tilde{\mathcal{T}} \leftarrow \tilde{\mathcal{T}} \cup \{W\}$  # Self-loop detected
9:   end if
10:  for each position vector  $\vec{v} \in \chi(NME(\theta[W])) \setminus W$  do
11:     $X \leftarrow W \cup \{\vec{v}\}$ 
12:    if  $X \in \tilde{\mathcal{V}}$  then # Check if  $X$  revisited
13:      continue
14:    end if
15:     $\tilde{\mathcal{V}} \leftarrow \tilde{\mathcal{V}} \cup \{X\}$ 
16:    # Compute the update operation resulting in  $mG_X$ .
17:    if  $mG_X = mG_W$  then
18:       $\theta[X] \leftarrow \theta[W]$  # New set-based representation
19:      # for same misinformation games
20:       $\tilde{\mathcal{T}} \leftarrow \tilde{\mathcal{T}} \cup \{W\}$  # Self-loop detected
21:    else
22:       $\theta[X] \leftarrow mG_X$  # New misinformation game
23:      compute  $NME(\theta[X])$ 
24:       $\tilde{\mathcal{Q}} \leftarrow \tilde{\mathcal{Q}} \cup \{X\}$ 
25:    end if
26:  end for
27: end while
28: return  $\tilde{\mathcal{T}}, \theta[\cdot]$ 
```

We begin by stating a recursive definition of the queue \mathcal{Q} , which is respected by Algorithm 7. Observe that only the misinformation games that do not have a self-loop in the adaptation graph Γ are inserted in the queue \mathcal{Q} . Therefore, for the i -th iteration of the **while** loop the following holds.

$$\begin{cases} \mathcal{Q}_0 = \{mG^0\}, \\ \mathcal{Q}_{i+1} = AD(\mathcal{Q}_i) \setminus \mathcal{T} \end{cases} \quad (12)$$

Let $\mathcal{T}_i = \mathcal{Q}_i \cap \mathcal{T}$. Intuitively, \mathcal{T}_i contains the terminal sets computed in the i -th iteration. Using the recursive equations in (12), we can rewrite the recursive equation

Algorithm 8 ComputeSME($\tilde{\mathcal{T}}$)

Require: The terminal set $\tilde{\mathcal{T}}$ and the mapping $\theta[\cdot]$

```
1:  $SME \leftarrow \emptyset$ 
2: for each  $X \in \tilde{\mathcal{T}}$  do
3:   for each  $\sigma \in NME(\theta[X])$  do
4:     if for all position vectors  $\vec{v} \in \chi(\sigma)$ , we have  $\theta[X] = \theta[X \cup \vec{v}]$  then
5:        $SME \leftarrow SME \cup \{\sigma\}$            # Following Definition 12.
6:     end if
7:   end for
8: end for
9: return  $SME$ 
```

in Definition 10 as,

$$\begin{cases} M_0 = \mathcal{Q}_0, \\ M_{i+1} = \mathcal{Q}_{i+1} \cup \mathcal{AD}^{i+1}(\bigcup_{k=1}^{i+1} \mathcal{T}_k). \end{cases} \quad (13)$$

To see that, note that \mathcal{Q}_{i+1} is produced from \mathcal{Q}_i , we apply the adaptation operator on \mathcal{Q}_i , neglecting the terminal nodes in \mathcal{Q}_i . To properly compute M_{i+1} , we need to re-instantiate the omitted terminal nodes and apply the adaptation operator on them. Noteworthy, Algorithm 7 terminates when $\mathcal{Q}_i = \emptyset$. We argue that $\mathcal{Q}_i = \emptyset$ if and only if $M_i = M_{i+1}$, i.e., the end criterion of Definition 11 holds. In order to prove that, we first need the following lemma.

Lemma 25. *If $\mathcal{T}_i = \mathcal{Q}_i \cap \mathcal{T}$, where \mathcal{Q}_i are defined as in equation (12), then it holds $\mathcal{T} = \bigcup_{i=1}^{\infty} \mathcal{T}_i$.*

The above lemma states that every node in the terminal set will be considered by Algorithm 7, since, at some point, will belong to the queue \mathcal{Q} . The following proposition shows that when the queue \mathcal{Q} becomes empty, i.e. the graph exploration terminates, then the end criterion of the Adaptation Procedure is met. Thus proving the correctness of Algorithm 7.

Proposition 26 (Correctness of Algorithm 7). *Consider the equations (12), (13). If $\mathcal{Q}_i = \emptyset$, then M_{i+1} is a Stable Set, and $\bigcup_{k=1}^{i+1} \mathcal{T}_k = \mathcal{T}$.*

Complexity

We now proceed to the time analysis of our method. We begin from Algorithm 7 since it does most of the heavy lifting. For the time complexity of Algorithm 7, it suffices to see that we are essentially doing a breadth-first search in the set-based representation $\tilde{\Gamma}$ of the adaptation graph. Thus, the computational complexity is bounded by the number of vertices. We state this complexity rigorously in the following proposition.

Proposition 27. *Assume a canonical $|N|$ -player misinformation game mG , on $|S|$ joint strategies. Then Algorithm 7 will terminate after*

$$O\left(2^{|S|}|N|^3\mathfrak{t}(\mathbf{NASH})|S|^2\right). \quad (14)$$

The following theorem summarises our analysis of algorithms 6, 7, 8. Note that the total complexity of Algorithm 6 is the sum of the complexities of Algorithms 7, 8. **Theorem 28** (Correctness and Complexity of Algorithm 6). *Algorithm 6 correctly computes all the smes of a canonical $|N|$ -player misinformation game mG^0 , on $|S|$ joint strategies in time*

$$O\left(|N|^2|S|^2\left[|\tilde{\mathcal{T}}| + 2^{|S|}|N|\mathfrak{t}(\mathbf{NASH})\right]\right) \quad (15)$$

In the above theorem, we provide an upper bound for the computation of the *all* the smes. Since the input to the algorithm is a misinformation game, the size of the input is $|N| \cdot |S|$. Therefore, from equation 15 our algorithm is exponential to the size of the input. Moreover, Algorithm 6 performs an exponential number of calls to an oracle computing the Nash equilibria of a $|N|$ -player normal-form game. This computational hardness seems to be unavoidable since the Adaptation Procedure (see Definition 10) leads, in the worst case, to an exponential explosion of the number of misinformation games in $AD^*(mG)$. We will see in the next subsection that computing a single *sme* appears to be a much easier problem.

4.4 Computing 1-SME

In this subsection, we will introduce a simple algorithm for computing a single *sme* of the Adaptation Procedure on a misinformation game mG . First, the computation of a single *sme* problem is articulated as follows

1-SME	
Input:	mG , a canonical $ N $ -player misinformation game on $ S $ joint strategies.
Output:	<i>sme</i> , a (possibly mixed) strategy profile that constitutes a stable misinformed equilibrium.

The rationale of our algorithm is stated in the following corollary.

Corollary 29. *Let $\Gamma = (AD^*(mG), E)$ be the adaptation graph of the Adaptation Procedure on a misinformation game mG . Also, let Γ' be the resulting graph after omitting the self-loops from Γ . Let mG' be the last node on a maximal path in Γ' , starting from mG . Then, all the nmes of mG' are smes.*

From the above corollary, it suffices to find a maximal path in Γ' . Intuitively, we can do that by iteratively adapting a single misinformation game, beginning from the root mG . We present this simple method in Algorithm 9. In the following theorem, we discuss the computational complexity of Algorithm 9.

Algorithm 9 Find-1-SME(mG)

Require: A root misinformation game mG .

```
1: while True do
2:   compute  $NME(mG)$ 
3:   choose some  $\sigma \in NME(mG)$ 
4:   if  $\exists \vec{v} \in \chi(\sigma)$  such that  $mG \neq (mG)_{\vec{v}}$  then
5:      $mG \leftarrow (mG)_{\vec{v}}$ 
6:   else
7:     return  $\sigma$ 
8:   end if
9: end while
```

Theorem 30. Consider the canonical $|N|$ -player misinformation game mG , on $|S|$ joint strategies. Then for the computational complexity of the Algorithm 9 it holds

$$O(|S|^2 \cdot |N|^3 \cdot t(\mathbf{NASH})).$$

Using Theorems 28, 30 we can directly compare the **ALL-SMEs** problem, which involves computing all the *smes* of a misinformation game, and the **1-SME** problem, which involves computing a single *sme*. In both problems, we use an oracle to the **NASH** problem to compute all the Nash equilibria of a normal-form game. Since **NASH** is a **PPAD**-complete problem, both **ALL-SMEs** and **1-SME** are computationally intractable. However, there is a qualitative difference between these problems. In **ALL-SMEs**, we will be making an exponential number of calls to the **NASH** oracle (with respect to the size of the input $|S|$), whereas in **1-SME**, we make only a polynomial number of calls. This difference in the number of **NASH** oracle calls comes from the fact that Algorithm 7 explores the entire adaptation graph, while Algorithm 9 explores only a small portion of it, corresponding to a single path in the adaptation graph. The latter is achieved by *greedily* expanding a single misinformation game. Such a greedy approach cannot be applied to the **ALL-SMEs** problem.

4.5 Implementation

In this section, we review the implementation details of the algorithms outlined in Subsections 4.3 and 4.4. We will focus on three important parts of Algorithms 7, 8. Namely, the implementation of the \mathcal{Q}, \mathcal{V} data structures, the computation of the update operation, and the computation of the Nash equilibria. We close this section by briefly discussing a parallel variant of Algorithms 7, 8.

Before addressing the details mentioned above, we make some remarks on the architecture of our program. The Algorithms 6, 7, 8 are implemented in PYTHON. On the other hand, the update operation and the computation of Nash equilibria form PYTHON subprocesses that communicate with the main program via pipes. This allows for modularity and better maintainability of our code. The subprocess computing the update operation is a logic program, written in the Answer Set Programming language

CLINGO³. Further, we use the GAMBIT [15] package to compute Nash equilibria in $|N|$ -player normal-form games.

The \mathcal{Q} , \mathcal{V} data structures

The implementation of the \mathcal{Q}, \mathcal{V} data structures significantly influences both the behaviour and the efficiency of Algorithm 7. The set \mathcal{Q} essentially denotes a dynamic linear structure of the sets of position vectors to be processed by the algorithm. The implementation of this structure impacts the way the algorithm traverses the graph of misinformation games. Implementing \mathcal{Q} as a *first in-first out* queue will induce a breadth-first search strategy, whereas implementing \mathcal{Q} as a *first in-last out* stack will induce a depth-first search strategy. Both choices are adequate since they offer constant-time insertions and extractions.

Perhaps the most significant concern for the implementation of \mathcal{V} will be to support efficient membership queries (see Algorithm 7, line 15). Therefore, a suitable choice would be a *dictionary*, with $O(\log |\mathcal{V}|)$ membership queries, or a hash table, with *randomized* constant membership queries.

Taking into account the above remarks, we implemented \mathcal{Q} as a simple queue, and \mathcal{V} as a dictionary because it supports efficient worst-case membership queries.

Implementing the update operation

We implemented the update operation of Definition 7 using the Answer Set Programming language CLINGO. As the update operation captures the main logic of the program, it is suitable to encode the axioms of Definition 7 using a Logic Programming language. A notable feature of our algorithm is its ability, via the CLINGO program, to ascertain whether a child misinformation game diverges from its parent. Therefore, the *equality check* of lines 17–19 in Algorithm 7 is realized within the logic program.

In order to encode Definition 7 as a logic program, some technicalities are involved. The payoff function is encoded using the predicate $u(\mathbf{G}, \mathbf{P}, \mathbf{SP}, \mathbf{U})$ of arity four. It signifies the payoff value (\mathbf{U}) awarded to player \mathbf{P} in the context of player \mathbf{G} 's normal-form game when the strategy profile \mathbf{SP} is played. Note that CLINGO does not support lists natively⁴. Therefore, we use symbolic functions to encode finite lists. For example, we encode the strategy profile $(1, 2, 3)$, with $\mathbf{sp}(1, \mathbf{sp}(2, \mathbf{sp}(3, \mathbf{nul})))$, where $\mathbf{sp}/2$ is a symbolic function with two arguments and $\mathbf{nul}/0$ is a constant.

In Listing 1 we present the logic program for applying the update operation of Definition 7. The logic program takes as input a list of payoff predicates $u/4$ and a single predicate $\mathbf{pos}(\mathbf{SP})$ describing a position vector as a finite list. Then, it outputs (proves) a list of updated payoff predicates $v/4$, which follow the syntax of $u/4$. For each strategy profile \mathbf{SP} we check whether the predicate $\mathbf{pos}(\mathbf{SP})$ holds. If it does, we force the new payoff entry $v/4$ of this strategy profile to equal the *actual* game G^0 . If there is an instance where $v \neq u$, then the logic program proves the predicate $\mathbf{changed}/0$. Otherwise, the predicate $\mathbf{unchanged}/0$ holds. This way the Python program can tell whether $mG \neq mG'$, where mG' is a child of mG .

³<https://potassco.org/clingo/>

⁴Unlike other Logic Programming languages, e.g., Prolog.

Listing 1: The Update Operation’s Logic program (Pseudocode).

```

% the new payoff function if pos(SP).
v(G, P, SP, V) :- pos(SP), u(0, P, SP, U), V = U.

% the new payoff function if not pos(SP).
v(G, P, SP, V) :- not pos(SP), u(G, P, SP, U), V = U.

% check if child-mG != parent-mG
changed :- v(G, P, SP, V), u(G, P, SP, U), V != U.
unchanged :- not changed.

```

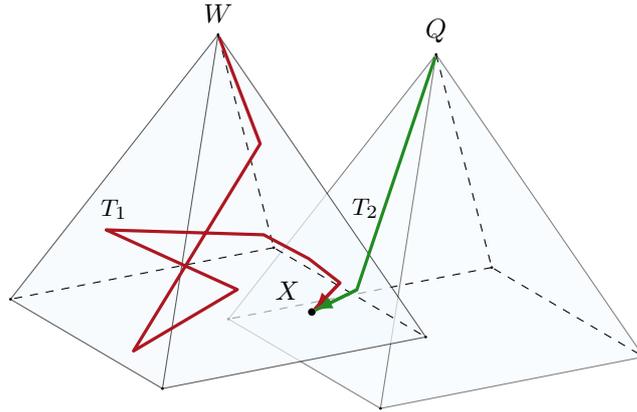


Fig. 8: Two threads T_1, T_2 simultaneously discovering the same node X . T_1 begins from W , while T_2 starts from Q .

Computation of Nash equilibria

To compute the Nash equilibria of a normal-form game we utilized the GAMBIT package. In particular, we used the `gambit-pol` command which implements the Support Enumeration method for computing the Nash equilibria in $|N|$ -player normal-form games [1].

Parallelism

In order to fully utilize the computational resources of a modern CPU, a parallel variant of our ALL-SMEs algorithm was developed. In details, Algorithm 6 does not need any parallelization, whereas the parallelisation of Algorithm 8 is straightforward, since Algorithm 8 simply iterates through the terminal set \mathcal{T} . On the other hand, developing a parallel version of the traversal algorithm, i.e., Algorithm 7, involves some technicalities that need to be addressed (see Algorithm 10); these are described below.

Our main objective is to parallelise the expensive computations of Algorithm 7, i.e., the *computation of the nmes*, which presupposes the computation of Nash equilibria, and the update operation. Recall that these computations are performed

by the two subsystems, namely the package GAMBIT and a logic program in CLINGO, respectively. In this direction, we assume k threads. Each thread will execute a variation of Algorithm 7, thus *independently* traversing the adaptation graph. Each thread will communicate with the others through the *shared* data structures $\tilde{\mathcal{Q}}, \tilde{\mathcal{V}}, \tilde{\mathcal{T}}, \theta[\cdot]$. Observe that for each new misinformation game mG that appears in the adaptation graph we need to compute an *nme*, and compute the result of the update operation. Since each thread independently traverses the adaptation graph, it is possible to rediscover a node that was previously discovered by another thread (see Figure 8). This fact entails two possible inefficiencies of parallel traversal, namely redundant computations and long waiting times. In other words, if a new misinformation game is discovered (almost) simultaneously by two threads, i.e., before the one notifies the other about its discovery, it is possible that the same computations will be performed in both threads, thus nullifying the benefits of parallelization. On the other hand, if one thread waits the other to complete its computations before continuing its traversal, the benefit of parallelization is also nullified.

Our solution to the above problems is described in Algorithm 10. Each of the shared data structures $\tilde{\mathcal{Q}}, \tilde{\mathcal{V}}, \tilde{\mathcal{T}}, \theta[\cdot]$ is protected by a lock. In order to gain access to the data structure the thread needs to successfully acquire the lock. Note that the independent traversal begins after line 4 and ends in line 43, where the threads are joined. We focus our attention in lines 17–23. Observe that the expensive computations we want to be performed in parallel are carried out in lines 25, 36. Assume that two threads T_1, T_2 (almost) simultaneously discover X , with T_1 discovering it slightly earlier. Since T_1 arrives slightly earlier at X , it will achieve to acquire the lock for $\tilde{\mathcal{V}}$. If $X \notin \tilde{\mathcal{V}}$, T_1 will mark X as visited (see line 22) and release the lock. We claim that the expensive computations in lines 25 and 36 will be executed only by T_1 . Additionally, T_2 will only wait for the lines 17–23. These lines only involve a membership query, and (possibly) an insertion to the $\tilde{\mathcal{V}}$ data structure. In order to see that, we follow the execution of T_2 . Since T_2 acquires the lock for $\tilde{\mathcal{V}}$ after T_1 , the membership query in line 18 will be positive. Thus, T_1 will omit all computations in lines 22 to 41, and *continue* to process the next position vector.

4.6 Experiments

In this section, we detail our experimental findings on the Adaptation Procedure in misinformation games, exploring its performance across a spectrum of multiplayer misinformation games. The methodology employed involved Monte Carlo simulations, following a structured approach. In particular, we first considered various configurations regarding the number of players and pure strategies per player; these are shown in the first column of Tables 6 and 7 (under “Setting”). For instance, the notation $3 \times 2 \times 2$ denotes a misinformation game with 3 players, where the first player has 3 pure strategies and the other two have 2 pure strategies, across all games in the misinformation game. Then, we constructed the payoff matrices of the misinformation game by using a random process to generate the payoff matrices of the actual game and of each subjective game. In particular, we populated the payoff matrices with integers drawn randomly from the interval $[-10, 10]$. We repeated this process 100 times and executed the Adaptation Procedure for each generated misinformation game.

Algorithm 10 Parallel-TraversalAdaptationGraph(mG, k)

Require: A root misinformation game mG and the number of threads k .

```
1:  $\tilde{\mathcal{T}} \leftarrow \emptyset, \tilde{\mathcal{Q}} \leftarrow \{\emptyset\}, \tilde{\mathcal{V}} \leftarrow \{\emptyset\}$ 
2:  $\theta[\emptyset] \leftarrow mG$ 
3: compute  $NME(mG)$ 
4: threads-begin( $k$ )
5: while  $\tilde{\mathcal{Q}} \neq \emptyset$  do
6:   lock( $\tilde{\mathcal{Q}}$ )
7:   Pick some  $W \in \tilde{\mathcal{Q}}$ 
8:    $\tilde{\mathcal{Q}} \leftarrow \tilde{\mathcal{Q}} \setminus \{W\}$ 
9:   unlock( $\tilde{\mathcal{Q}}$ )
10:  if  $\exists \vec{v} \in \chi(NME(\theta[W])) \cap W$  then
11:    lock( $\tilde{\mathcal{T}}$ )
12:     $\tilde{\mathcal{T}} \leftarrow \tilde{\mathcal{T}} \cup \{W\}$  # Self-loop detected
13:    unlock( $\tilde{\mathcal{T}}$ )
14:  end if
15:  for each position vector  $\vec{v} \in \chi(NME(\theta[W])) \setminus W$  do
16:     $X \leftarrow W \cup \{\vec{v}\}$ 
17:    lock( $\tilde{\mathcal{V}}$ )
18:    if  $X \in \tilde{\mathcal{V}}$  then # Check if  $X$  unvisited
19:      unlock( $\tilde{\mathcal{V}}$ )
20:      continue
21:    end if
22:     $\tilde{\mathcal{V}} \leftarrow \tilde{\mathcal{V}} \cup \{X\}$ 
23:    unlock( $\tilde{\mathcal{V}}$ )
24:    # Compute the update operation resulting in  $mG_X$ .
25:    if  $mG_X = mG_W$  then
26:      lock( $\theta$ )
27:       $\theta[X] \leftarrow \theta[W]$  # New set-based representation for
28:                               # same misinformation games
29:      unlock( $\theta$ )
30:       $\tilde{\mathcal{T}} \leftarrow \tilde{\mathcal{T}} \cup \{mG_X\}$  # Self-loop detected
31:      unlock( $\tilde{\mathcal{T}}$ )
32:    else
33:      lock( $\theta$ )
34:       $\theta[X] \leftarrow mG_X$  # New misinformation game
35:      unlock( $\theta$ )
36:      compute  $NME(mG_X)$ 
37:      lock( $\tilde{\mathcal{Q}}$ )
38:       $\tilde{\mathcal{Q}} \leftarrow \tilde{\mathcal{Q}} \cup \{X\}$ 
39:      unlock( $\tilde{\mathcal{Q}}$ )
40:    end if
41:  end for
42: end while
43: threads-join()
44: return  $\tilde{\mathcal{T}}, \theta[\cdot]$ 
```

For each randomly generated misinformation game (and subsequent run of the Adaptation Procedure), we collected various data regarding the efficiency and characteristics of the process. We present those results below in two separate tables, one concerning the time efficiency of the implementation (Table 6), and one concerning the characteristics of the adaptation graph (Table 7).

All experiments were conducted on a macOS Ventura version 13.5.2 (22G91) machine, utilizing PYTHON 3.11.4, GAMBIT 15.1.1, and CLINGO 5.4.0. Notably, all experiments were executed using a multi-threaded implementation with eight hardware threads⁵.

Table 6: Efficiency of the Adaptation Procedure for various realizations.

Setting	Total time (sec)	CPU time (sec)
2×2	[0.11, 0.18]	[0.04, 0.1]
3×2	[0.23, 0.29]	[0.17, 0.22]
3×3	[0.17, 0.2]	[0.11, 0.15]
4×3	[1.1, 1.29]	[1.06, 1.28]
4×4	[1424.42, 1476.48]	[163.59, 188.12]
$2 \times 2 \times 2$	[1.83, 2.19]	[1.66, 1.91]
$3 \times 2 \times 2$	[22.76, 30.33]	[18.94, 22.45]
$2 \times 2 \times 2 \times 2$	[3793.66, 4092.59]	[2172.43, 2386.21]

For assessing the efficiency of the process, we used the metrics of Total run time (sec) and CPU run time (sec), as shown in Table 6. CPU time refers to the time consumed only by the PYTHON program, whereas Total time refers to the total elapsed time, including the time consumed by the subprocesses, CLINGO and GAMBIT. For both attributes we provide the interval between the minimum and maximum time as provided by the simulation, in the form $[min, max]$.

We observe that when the misinformation game has two players with three pure strategies each (3×3) the calls to PYTHON constitute the main bottleneck of our pipeline, as shown by the fact that CPU time is close to Total time. On the other hand, in 2×2 and 4×4 settings the elapsed time is dominated by the subprocesses (i.e., CLINGO and GAMBIT), and the CPU time consumes only $\sim 20\%$ and $\sim 12\%$ respectively of the elapsed time on average. An explanation here is that for the 2×2 case the possible misinformation games spawned from the Adaptation Procedure are few. On the contrary, in the 4×4 case, the adaptation graph may have more extensive branches, so CLINGO and GAMBIT take their toll.

Moreover, as we increase the number of players ($2 \times 2 \times 2 \times 2$) we witness that both Total time and CPU time increase rapidly. To comprehend the reason behind this performance degradation, it's crucial to consider the exponential growth of the search space with the increase in available players and/or strategies. Recall that the

⁵The implementation of the Adaptation Procedure is in <https://github.com/merkouris148/adaptation-procedure-misinformation-games>.

computational complexity for computing all Nash equilibria grows exponentially as proved in Theorem 28.

Table 7: Characteristics of the adaptation graph for various realizations.

Setting	#nodes	unique mGs	#leaves	# $smes$	SP
2×2	13	7	6	3	2^4
3×2	50	17	18	3	2^6
3×3	49	16	20	10	2^9
4×3	451	132	205	9	2^{12}
4×4	54122	10653	32779	377	2^{16}
$2 \times 2 \times 2$	832	148	489	373	2^8
$3 \times 2 \times 2$	8226	1287	4655	634	2^{12}
$2 \times 2 \times 2 \times 2$	632302	58617	24346	102338	2^{16}

Regarding the characteristics of the Adaptation Procedure, we analyzed various aspects, as shown in Table 7. Below, we explain the meaning of each of these quantities, and analyse the respective results. Note that the number shown for each quantity in Table 7 is the average over all 100 runs (rounded to the closest integer).

To study the characteristics of the adaptation graph generated by our algorithm, we also implemented the recursive method (2). Using this implementation, we compared the number of nodes (i.e., misinformation games) generated by the implementation (as shown in column “#nodes”) against the number of nodes generated by the implementation of the set-based representation (12) (shown in column “unique mGs ”). From this comparison it is obvious that the number of duplicate computations generated by (2) is huge. For example, in the 2×2 setting, the (2) computed 13 misinformation games, while the (12) one only 7, meaning that 6 (almost half) of the games computed by the (2) were actually duplicates, and thus caused redundant computations. Analogous observations can be made in the other settings (rows of the table).

Further, it is interesting to compare the numbers appearing in columns “#leaves” and “# $smes$ ” in the table. The former (“#leaves”) corresponds to the number of distinct sets of position vectors that lead to an mG that belongs to the terminal set (Definition 13), or, in other words, it denotes the size of the terminal set (i.e., the number of distinct mGs that belong to the terminal set). The latter (“# $smes$ ”) corresponds to the actual number of $smes$ that this terminal set generates (according to Definition 12). Note that all $smes$ are generated by mGs in the terminal set. However, not every mG in the terminal set is guaranteed to yield an sme ; some may lead to none, while others may lead to one or more $smes$. The fact that the numbers in the “#leaves” column are significantly larger than the numbers in the “# $smes$ ” column, signifies that the majority of the mGs in the terminal set produced no sme . Upon reviewing Definitions 12 and 13, we conclude that there are a lot of mGs that “self-replicate” via adaptation (in the sense of Definition 13 on the terminal set), yet they do not achieve “stability” (for their $nmes$ to produce an sme , as described in Definition 12).

Another interesting observation stems by comparing the numbers appearing in Table 7 across different rows. For example, one might anticipate that the settings 4×3 and 4×4 would yield comparable outcomes, but, in fact, the adaptation graph for 4×3 generates only 451 nodes (132 distinct), whereas 4×4 results in 54,122 nodes (10,653 distinct). This counter-intuitive result is explained by looking at the last column of the table (“SP”), which denotes the size of the strategy space of the respective setting, i.e., the total number of sets of position vectors that the setting can generate; this in turn corresponds to the maximum possible number of distinct nodes (misinformation games) that the algorithm implementing the Adaptation Procedure could generate (i.e., the maximum possible value of column “unique mGs ”).

Considering the strategy space (“SP”) column values, the 4×3 scenario indicates that the Adaptation Procedure could generate up to 2^{12} unique misinformed strategy profiles. In contrast, the 4×4 setting could yield 2^{16} distinct misinformed strategy profiles. It’s noteworthy that the “SP” value by itself does not suffice to fully explain the size of the adaptation graph and the related quantities in all settings. It appears that incorporating more players contributes more to this aspect, compared to the addition of strategies, even when such additions (players or strategies) have a comparable effect on the “SP” value.

5 Related work

Here we present the works that are close to the setting where agents with subjective views of the interaction interact in a turn-based procedure. As stated in [12] the misinformation games are close to the Hypergames (HG), the games with (un)awareness (GwU), and the Selten games (see [4, 9, 16–19]). The first two of these classes are under the hood of games with misperception, often abbreviated as MP [2], while the third is the a posteriori lottery model and defined in Harsanyi’s seminal work [9]. However, since the setting in [12] focused only on one-shot interactions, the comparison with MPs (HGs and GwU) was also affected by this limitation.

However, MPs do not address only one-shot interactions but have also been proposed as a novel framework to analyze iterative interactions. Towards this direction, there is a significant stream of works that aims in the analysis of players’ beliefs in MP (e.g., [20–27]), where players update their beliefs using probabilistic formulas and techniques. Moreover, the agents have common knowledge about their views, as they update their knowledge in order to converge as close as possible to the outcome that they should have if they did not experience subjectivity. Contrary to these works, in our work the agents update their beliefs in a one-way manner, meaning that the agents update their subjective views plugging into the feedback that they receive from the environment.

In [20], a learning framework was plugged in with an MP, where the players are completely ignorant regarding opponents’ decisions and any new information is automatically integrated. In [28] the authors describe approaches for MP, where the objective views emerged from partially observable cooperative games, but they assume perfect knowledge of other players’ policies. Here, we drop the constraint of the limited observability and the cooperative setting.

Two other classes of games that are closely related to the misinformation games are the Selten games [9] and the Global games [29]. In the first case, the agents learn only their own payoff matrix. So the agents have uncorrelated observations, but this is common knowledge. In the second case, the agents receive private and fuzzy observations of the game, where the latter is influenced by small random fluctuations. To cope with incompleteness, the agents take into account a whole family of a priori possible games; this is common knowledge. In other words, the agents know that they do not know, as opposed to the misinformation games where the agents do not have common knowledge and take into account only the agglomeration of their subjective views and the information that is broadcasted from the environment.

Further, in [30] a class of Bayesian processes for iterated normal form games is studied, where each player knows her own payoff function, but is uncertain about the opponents' payoffs. Another learning mechanism was introduced in [22] for games with misperception. In the same manner, in [25], a behavioral-based model is provided in order to model the ignorance of each player about adversaries. As opposed to the stream of works that rely on probabilistic techniques, in this work, we take no probabilistic considerations regarding the beliefs of the players. Here we focus on the effect of new information in the misinformed views of the players, i.e., the players “do not know that they do not know”, as opposed to probabilistic approaches, where the players “know that they do not know”. This leads to a conceptual and methodological difference between the Nash-Berk equilibrium concept, developed in [22, 31], and the stable misinformed equilibrium.

Now we steer our discussion to the solution concepts. In HGs, the outcome is not influenced by the actual/exact game. On the other hand, in GwUs the solution concept is influenced by the elimination of beliefs (rationalizability) and does not require correct conception and information, or mutual knowledge. Here, we both agglomerate the actual/exact game and we do not eliminate any belief towards the outcome. In MPs, the direction is towards the updates that restrict the subjectivity of the agents, while in this work we do not enforce such kind of condition.

Another equilibrium concept closely related to the solution concept in this work is that of self-confirming equilibrium [32, 33], which captures the idea of “stable” joint decisions in MP. Although it shares the same intuition with the stable misinformed equilibrium, which is that players' choices stabilize when the choices of their opponents confirm their view, these two concepts have two differences. Firstly, the self-confirming equilibrium is applied in extensive-form games, while the stable misinformed equilibrium is in strategic-form games. Second, the former concept requires that the beliefs of the participants are correct along the path of plays, whereas we do not make such an assumption.

The limited behaviour of players while the values of their payoffs are non-stationary, has drawn significant attention. Authors in [31] provide a characterization of asymptotic behaviour, based on a continuous time model, where the player thinks that there are two possible payoff functions, yet none of them is true. Moreover, in [26] authors establish convergence of beliefs, and actions, in a misspecified model with endogenous actions. In our case, players do not have second thoughts regarding the received information.

In a weaker sense, a posteriori lotteries contain any interaction where agents depart from correct/complete information. In this respect, another concept where players have incomplete knowledge is that of Bayesian games [9]. It has attracted considerable attention ([10, 11, 34, 35], etc.), where a key assumption is that of common priors. Nevertheless, Bayesian games are defined in [9] as a priori lotteries.

Furthermore, a conceptual difference with misinformation games is that, in Bayesian games, although agents are unsure as to their actual payoff, they are well aware of that, and they do their best out of the uncertainty that they have, e.g. Figure 1a. On the contrary, in misinformation games, the agents play according to their subjective game definition, without considering mitigation measures, e.g., Figure 1b. In other words, in Bayesian games players “know that they do not know”, whereas in misinformation games players “do not know that they do not know”. Thus, the scenario depicted in Figure 1b cannot be captured by Bayesian games, as the players cannot distinguish the actual situation from the one provided to them.

6 Discussion and conclusions

6.1 A possible use case on disaster management

One possible use case application where we believe that our work on misinformation games and the Adaptation Procedure could be of use is the case of disaster management.

Disaster management constitutes an important potential application of robots and intelligent systems, such as drones, autonomous vehicles, and mobile robotic systems. In such scenarios, robots are deployed in unpredictable and hazardous environments associated with natural and human-made disasters, most importantly in situations where it is impossible, or too dangerous for humans to intervene.

In such unpredictable and chaotic environments, operators’ knowledge about the situation is both incomplete and erroneous. Furthermore, communication infrastructure is likely to be damaged, and network congestion makes post-deployment communication unreliable and erratic. These factors underscore the critical need for robots to be highly adaptable. They must operate autonomously, relying on initially faulty information, and adapt dynamically based on sensor data and trial-and-error, allowing them to learn the “rules of the game” in real-time.

The game-theoretic aspect becomes even more important when multiple robots, with different objectives, are deployed. They will need to both compete and cooperate to achieve their goals, adding a strategic layer to their behavior.

We believe that the characteristics of the disaster management use case highlight the need for more research in misinformation games and on the adaptation of agents in the face of new information that contradicts their initial game specifications. Such scenarios cannot be addressed by classical game settings, in which the game specification is either fully known to the agents, or, when it is not so (e.g., incomplete games), the agents are fully aware of the limitations of their knowledge. Disaster management scenarios challenge those assumptions, as the lack of knowledge is prevalent and cannot be quantified or fully described through, e.g., probabilities.

Table 8: Basic families of methodologies regarding either one-round or multiple-round interactions with either incomplete or incorrect information.

Class	Game setting	iterative process	outcome
HG [36],[37]	incorrect information, erroneously aware, they don't know that they don't know, no exact/actual game	✓	hyper Nash equilibrium
GwU [17]	lack of conception, not lack of information	✗	notions of rationalizability
MP [22]	players know that they don't know the game setting	✓	variants of equilibria
Selten games [9]	players know the game setting, know their own attribute vector but don't know the exact/actual game, know that they don't know	✗	equilibrium selection
Global games [29]	players know the game setting, don't know the exact/actual game, know that they don't know	✓	equilibrium selection
Bayesian games	players know the game setting, don't know the exact/actual game, know that they don't know	✓	Bayes-Nash equilibrium
Misinformation games [12], [5]	incorrect information, erroneously aware, they don't know that they don't know	✗	natural misinformed equilibrium
Misinformation games plus Adaptation Procedure [13]	incorrect information, erroneously aware, they don't know that they don't know the information they learn becomes commonly known	✓	stable misinformed equilibrium

6.2 Closing remarks

This work explores the pervasive nature of misinformation in multi-player interactions. By employing the framework of *misinformation games*, we gain insights into various real-world dynamics in scenarios where the players' information on the interaction are erroneous, but the players have no clue that their specifications are false.

In this setting, we define an iterative process, the *Adaptation Procedure*, which enables players to refine their strategies through continuous interaction and feedback. This process is rooted in the idea that players will adjust their strategies based on the payoffs that are publicly announced in each iteration, which may be different from what they know, thereby progressively reducing misinformation.

The Adaptation Procedure models the evolution of players' decisions as they assimilate new information, culminating in a stable misinformed equilibrium, which is the strategy profile adopted upon process stabilization. We establish that this

procedure invariably concludes, guaranteeing the existence of such an equilibrium and study (and prove) various related properties.

Lastly, we devise algorithmic tools to analyze this evolution. Our analysis reveals the intractability of the stable misinformed equilibrium (*sme*), which is inherited to the problem both by the mechanism the agents reach a joined decision, i.e. by combining the Nash equilibria of each subjective game, and the exponentially large state space they explore. In addition to the theoretical analysis, we provide some experiments demonstrating the characteristics and computational properties of the Adaptation Procedure under various randomly-generated settings.

Interestingly enough our setting exhibits an intriguing combinatorial structure that seems to be independent of the way the agents move from one state to the other. In other words, it seems to be possible to abstract away the method the players transit from one state to the other, based on the current state, and the subjective information of each player. This introduces a family of adaptation mechanisms for a group of agents, that share the same combinatorial structure. Each member of this family is differentiated from the others by the method the agents use to jointly reach a decision (e.g. the notion of NME discussed here, or a voting scheme). Such a generalized analysis falls outside of the scope of this paper, but provides a fruitful ground for future work.

Moreover, we will revisit our implementation, considering additional optimizations, and conduct more extensive experiments, possibly in the context of a real-world scenario. Connections and applications of misinformation games to mechanism design will also be explored (see also [12]).

Appendix A Proofs of results

A.1 Normal-form misinformation games

Proof of Proposition 1. By equation (1), we have that:

$$h_i(\sigma) = \sum_{k \in S_1} \cdots \sum_{j \in S_{|N|}} P_i(k, \dots, j) \cdot \sigma_{1,k} \cdots \sigma_{|N|,j},$$

Replacing the tuple (k, \dots, j) with $s = (s_i) \in S$, we can rewrite this as follows

$$h_i(\sigma) = \sum_{s \in S} P_i(s) \cdot \sigma_{1,s_1} \cdots \sigma_{|N|,s_{|N|}}$$

Working analogously for h'_i , we get:

$$h'_i(\sigma') = \sum_{s' \in S'} P'_i(s') \cdot \sigma'_{1,s'_1} \cdots \sigma'_{|N|,s'_{|N'|}}$$

Now, let us write s' as (\hat{s}, \tilde{s}) , where \hat{s} represents the strategies in s' for the players in N and \tilde{s} the strategies in s' for the players in $N' \setminus N$. Let us also denote with \tilde{S} the different possible values of the tuple \tilde{s} (note that $\tilde{S} = \times_{i \in N' \setminus N} S'_i$). Under this notation, we have that, for any s' , $P'_i(s') = P_i(\hat{s})$, from the third bullet of Definition 2, and thus, rearranging the sums and using associativity, the latter equation on h'_i can be written as:

$$h'_i(\sigma') = \sum_{\hat{s} \in S} P_i(\hat{s}) \cdot \left(\sigma'_{1,\hat{s}_1} \cdots \sigma'_{|N|,\hat{s}_{|N|}} \right) \cdot \left(\sum_{\tilde{s} \in \tilde{S}} \sigma'_{|N|+1,\tilde{s}_{|N|+1}} \cdots \sigma'_{|N'|,\tilde{s}_{|N'|}} \right)$$

The second big parenthesis sums to 1. Indeed, using associativity:

$$\sum_{\tilde{s} \in \tilde{S}} \sigma'_{|N|+1,\tilde{s}_{|N|+1}} \cdots \sigma'_{|N'|,\tilde{s}_{|N'|}} = \prod_{i \in N' \setminus N} \sum_{j \in S'_i} \sigma'_{ij} = 1 \cdots 1 = 1$$

Since $\sigma \boxtimes \sigma'$, it follows that $\sigma_{ij} = \sigma'_{ij}$ for all $i \in N$, $j \in S_i$, and thus, combining the above results, we get:

$$h'_i(\sigma') = \sum_{\hat{s} \in S} P_i(\hat{s}) \cdot \sigma_{1,\hat{s}_1} \cdots \sigma_{|N|,\hat{s}_{|N|}}$$

From the latter equation, and the respective equation for h_i , it is clear that $h_i(\sigma) = h'_i(\sigma')$, which completes the proof. \square

Proof of Proposition 2. Apparently, if $G = G'$ then $G \blacktriangleleft G'$ (by the definition), so let us focus on the reverse. Assume that $G \blacktriangleleft G'$. Then, since $N = N'$ and $S = S'$ by the hypothesis, take any $i \in N$ and $s = (s_i)_{j \in N} \in S$, $s' = (s'_i)_{j \in N'} \in S'$ such that

$s_j = s'_j$ for all $j \in N$. Then $s = s'$ (by the hypothesis that $N = N'$ and $S = S'$, and $P_i(s) = P'_i(s')$ (by the hypothesis that $G \blacktriangleleft G'$), and thus $P = P'$, i.e., $G = G'$. \square

Proof of Proposition 3. Any game G is trivially an inflated version of itself, so reflexivity holds.

For antisymmetry, we observe that $G_1 \blacktriangleleft G_2$ and $G_2 \blacktriangleleft G_1$ implies that G_1 and G_2 have the same set of players and the same set of strategies per player. Antisymmetry now follows by applying Proposition 2.

For transitivity, let us denote, for $i = 1, 2, 3$, N^i, S^i, P^i the set of players, the set of strategies and the payoff matrix, respectively, of G_i . Let us also denote by S_j^i the available strategies for player j in game G_i . Suppose that $G_1 \blacktriangleleft G_2$ and $G_2 \blacktriangleleft G_3$. We will show that $G_1 \blacktriangleleft G_3$.

Indeed, it is clear that $N^1 \subseteq N^2 \subseteq N^3$, so the first bullet of the definition holds.

Similarly, for any $i \in N^1$, $S_i^1 \subseteq S_i^2$. Also $i \in N^1$ implies $i \in N^2$ and thus $S_i^2 \subseteq S_i^3$.

Thus $S_i^1 \subseteq S_i^3$ for all $i \in N$, so the second bullet holds as well.

For the third, pick any $i \in N^1$, and $s^1 = (s_i^1)_{j \in N^1} \in S^1$, $s^3 = (s_i^3)_{j \in N^3} \in S^3$ such that $s_j^1 = s_j^3$ for all $j \in N^1$. Let us pick some tuple $s^2 = (s_i^2)_{j \in N^2} \in S^2$ such that $s_j^2 = s_j^3$ for all $j \in N^2$. Now, by the fact that $G_1 \blacktriangleleft G_2$ and $G_2 \blacktriangleleft G_3$ and the previous results, we get that $P_i^1(s^1) = P_i^2(s^2)$ and $P_i^2(s^2) = P_i^3(s^3)$, and thus the third bullet holds as well.

For the fourth, suppose that $\sigma^1 \in NE(G_1)$. Then, since $G_1 \blacktriangleleft G_2$, there exists some $\sigma^2 \in NE(G_2)$ such that $\sigma^1 \bowtie \sigma^2$. Moreover, since $G_2 \blacktriangleleft G_3$, there exists some $\sigma^3 \in NE(G_3)$ such that $\sigma^2 \bowtie \sigma^3$. Given that $N^1 \subseteq N^2 \subseteq N^3$ and $S^1 \subseteq S^2 \subseteq S^3$, it is easy to see that $\sigma^1 \bowtie \sigma^2$ and $\sigma^2 \bowtie \sigma^3$ imply that $\sigma^1 \bowtie \sigma^3$. Thus, there exists some $\sigma^3 \in NE(G_3)$ such that $\sigma^1 \bowtie \sigma^3$, which shows the fourth bullet.

For the fifth bullet, we proceed analogously. \square

Proof of Proposition 4. Observe that since the payoffs of the new player are zero (line 3), then this player will be indifferent in any choice of the other players. Thus does not affect the outcome of the G' . So, we have $G \blacktriangleleft G'$. \square

Proof of Proposition 5. In Algorithm 3 we increase the pure strategies of player i by adding the j pure strategy. From lines 6-7 of the algorithm the new entries of the payoff matrices will have values smaller than any other value in the payoff matrix without the j strategy. Thus, this strategy is dominated and does not affect the strategic behavior of the players. So, we have $G \blacktriangleleft G'$. \square

Proof of Proposition 6. We observe that Algorithm 1 consists of repetitive calls to Algorithms 2 and 3. The result now follows from Propositions 4, 5, and the transitivity of the \blacktriangleleft relation (Proposition 3). \square

A.2 Adaptation Procedure

Proof of Proposition 7. By definition:

$$\begin{aligned} AD(M) &= \{mG_{\vec{u}} \mid mG \in M, \vec{u} \in \chi(\sigma), \sigma \text{ is a nme of } mG\} \\ &= \bigcup_{mG \in M} \{mG_{\vec{u}} \mid \vec{u} \in \chi(\sigma), \sigma \text{ is a nme of } mG\} = \bigcup_{mG \in M} AD(\{mG\}) \end{aligned}$$

□

Proof of Proposition 8. Observe that $mG_i \in \mathcal{AD}(\mathcal{AD}(\dots \mathcal{AD}(\{mG_i\}))) = \mathcal{AD}^{(n)}(\{mG_i\}) \forall i \in [n]$. Suppose, for the sake of contradiction, that $mG_i \neq mG_j$ for some i, j , and assume, without loss of generality, that $i < j$. Then, it holds that $mG_j \in \mathcal{AD}^{(j-i)}(\{mG_i\})$, i.e., mG_j has resulted from mG_i by updating some (at least 1 and at most $j - i$) elements of the respective payoff matrices of mG_i . But then, we also have that $mG_i \in \mathcal{AD}^{(n-j+i)}(\{mG_i\})$ (by the periodic pattern above), so again, mG_i has resulted from mG_j by updating some (at least 1 and at most $n - j + i$) elements of the respective payoff matrices of mG_j . But this is an absurdity, because replacements are cumulative and cannot be “undone” by subsequent ones (see Definitions 8 and 10). □

Proof of Proposition 9. Let the misinformation game mG' , for which it holds $mG' \in \mathcal{AD}^{(1)}(\{mG'\})$. It is easy to see, through induction, that $mG' \in \mathcal{AD}^{(\tau)}(\{mG'\})$ for all $\tau \geq 1$. Since, $mG' \in \mathcal{AD}^{(t)}(\{mG^{(0)}\})$, then $mG' \in \mathcal{AD}^\infty(\{mG^{(0)}\})$. □

Proof of Proposition 10. The first bullet is a consequence of the Definition 11. For the second bullet we have that, since $\sigma \in \text{SME}(mG^{(0)})$ applying the Definition 12 there must be a misinformation game mG in the terminal set, such that σ is a *nme* in mG . □

Proof of Proposition 11. Observe that in any given branch of the Adaptation Procedure, a certain position can be updated at most once. Given that the number of positions is $|S|$, we conclude. □

Proof of Proposition 12. We use two ways to upper bound the number of misinformation games $|\mathcal{AD}^*(\{mG\})|$. Firstly, note that $|\mathcal{AD}^*(\{mG\})| \leq 2^{|S|}$, since there cannot be generated more misinformation games, throughout the adaptation procedure, than the powerset of position vectors S .

For the other lower bound, notice that, in the worst case the Adaptation Procedure will produce a tree with $\mathfrak{L}_{\mathcal{AD}}(mG)$ height. Each node in this tree will have at most $|S|$ children. We have,

$$\begin{aligned} |\mathcal{AD}^*(\{mG\})| &\leq \sum_{i=1}^{\mathfrak{L}_{\mathcal{AD}}(mG)} |S|^i \\ &= \frac{|S| \left(|S|^{\mathfrak{L}_{\mathcal{AD}}(mG)} - 1 \right)}{|S| - 1} \\ &\leq |S|^{\mathfrak{L}_{\mathcal{AD}}(mG)+1} \end{aligned}$$

Note that in the last inequality we assumed $|S| \geq 1$, without loss of generality. Thus, we showed the desideratum. □

Proof of Corollary 13. Since $\mathcal{AD}^\infty(mG) \subseteq \mathcal{AD}^*(mG)$, the desideratum holds from Proposition 12. □

Proof of Theorem 14. Let $\mathcal{S} = \mathcal{AD}^\infty(mG)$ to be the Stable Set of the Adaptation Procedure on mG^0 . It holds $\mathcal{AD}(\mathcal{S}) = \mathcal{S}$. We prove that $\mathcal{AD}^*(\mathcal{T}) \subseteq \mathcal{S}$ and $\mathcal{S} \subseteq \mathcal{AD}^*(\mathcal{T})$. The direction $\mathcal{AD}^*(\mathcal{T}) \subseteq \mathcal{S}$ holds trivially. Indeed, from Proposition 10 we have $\mathcal{T} \subseteq \mathcal{S}$. On the other hand, since \mathcal{S} is the Stable Set, we have $\mathcal{AD}(\mathcal{S}) = \mathcal{S}$. From Proposition 7, we take $\mathcal{AD}(\mathcal{S}) = \mathcal{AD}(\mathcal{S} \setminus \mathcal{T}) \cup \mathcal{AD}(\mathcal{T})$. Thus, $\mathcal{AD}^*(\mathcal{T}) \subseteq \mathcal{AD}(\mathcal{S})$.

For the direction $\mathcal{S} \subseteq \mathcal{AD}^*(\mathcal{T})$, we work as follows. For the sake of contradiction, assume $mG \in \mathcal{S} \setminus \mathcal{AD}^*(\mathcal{T})$. We consider the set C of all ancestors of mG that belong to \mathcal{S} .

$$C = \left\{ mG' \in \mathcal{S} \mid mG \in \mathcal{AD}^k(mG'), k \in \mathbb{N} \cup \{0\} \right\}. \quad (\text{A1})$$

Note that $C \neq \emptyset$, since $mG \in \mathcal{AD}^0(mG)$, hence $mG \in C$. We choose some $mG'' \in C$, for which there is not an $mG' \in C$, such that $mG'' \in \mathcal{AD}(mG')$. Observe that $C \cap \mathcal{T} = \emptyset$, since $mG \notin \mathcal{AD}^*(\mathcal{T})$. Therefore, such an mG'' exists. Intuitively, mG'' is the oldest ancestor⁶ of mG , that belongs to the Stable Set. This oldest ancestor exists, since non of the elements of C generates itself. Moreover, mG'' has no ancestor in $\mathcal{AD}^*(mG^{(0)})$. Otherwise, it would also be an ancestor of mG , thus belonging to C , and mG'' would not be the oldest ancestor. From the above, there is no element in \mathcal{S} to generate mG'' . Hence, $mG'' \in \mathcal{S} \setminus \mathcal{AD}(\mathcal{S})$, which is a contradiction. \square

Proof of Proposition 15. Suppose that $mG' \in \mathcal{AD}^{(t_0)}(\{mG\})$ for some $t_0 \geq 0$. Since $mG' \in \mathcal{AD}(\{mG'\})$, it follows that $mG' \in \mathcal{AD}^{(t)}(\{mG\})$ for all $t \geq t_0$, thus, $mG' \in \mathcal{AD}^\infty(\{mG\})$. Since $\mathcal{AD}(\{mG'\}) = \{mG'\}$, it is clear that for all $\sigma \in \text{NME}(mG')$ and for all $\vec{v} \in \chi(\sigma)$, it holds that $mG'_{\vec{v}} = mG'$. Now the result is direct from Definition 12. \square

Proof of Proposition 16. Set $S = \mathcal{AD}^\infty(\{mG\})$. For any given $mG_1, mG_2 \in S$, we define the relation \mapsto , such that $mG_1 \mapsto mG_2$ iff $mG_1 \neq mG_2$ and $mG_2 \in \mathcal{AD}(\{mG_1\})$. Now let us suppose, for the sake of contradiction, that mG has no *sme*. By Proposition 15, it follows that for any $mG' \in S$ there exists some $mG'' \in S$ such that $mG' \mapsto mG''$ (otherwise $\text{SME}(mG) \neq \emptyset$ by Proposition 15, which contradicts our hypothesis). Since S is finite (see Theorem 11), there must exist a sequence of $mG_1, \dots, mG_n \in S$, such that $mG_i \mapsto mG_{i+1}$ (for $i \in [n-1]$) and $mG_n \mapsto mG_1$. Which is an absurdity by the definition of \mapsto and Proposition 8. \square

Proof of Proposition 18. Assume a N -player misinformation game $mG = \langle G^0, G^1, \dots, G^N \rangle$, where $S^1 = S^2 = \dots = S^N$. Further, let P^0 be the payoff matrix for the actual game G^0 , and P the payoff matrices for the G^i , for all i . For P we use the dovetail principle to assign distinct integers, to each cell, form $\sum_{k=1}^{|N|} |S^k|$ to 1, starting from $P[|S^1|, |S^2|, \dots, |S^{|N|}]$. On the other hand, we take $P^0 = -P$. It is easy to see that initially each player i picks the $|S^i|$ -th pure strategy, so the strategy profile $(|S^1|, |S^2|, \dots, |S^{|N|})$ is a pure Nash equilibrium in dominated strategies, for each player, constituting an *nme*. After a single step of the Adaptation Procedure, the players will be informed that $P^0[i_1, i_2, \dots, i_n] = -P[i_1, i_2, \dots, i_n]$. Using the same

⁶Formally, the $\mathcal{AD}(\cdot)$ operator defines a transitive, anti-symmetric (but not reflexive) relation $\preceq_{\mathcal{AD}}$ on the sets of all misinformation games $\mathcal{AD}^*(mG^{(0)})$ generated from $mG^{(0)}$. From Proposition 12 we have that $\mathcal{AD}^*(mG^{(0)})$ is finite. Thus, there are maximal elements in $\preceq_{\mathcal{AD}}$. Moreover, we have maximal elements in C , since $C \subseteq \mathcal{AD}^*(mG^{(0)})$. The "oldest ancestors", we discuss here, are these maximal elements of $\preceq_{\mathcal{AD}}$, restricted on C .

rationale as before, the next nme chosen by the players ($|S^1|-1, |S^2|-1, \dots, |S^{|N}|-1$), which again will be updated to opposite value. Inductively, we can prove that the Adaptation Procedure will take $|S|$ steps. \square

A.3 Computing the Adaptation Procedure

Proof of Proposition 19. It suffices to show that Γ' does not contain directed circles. This is an immediate consequence of Proposition 8. \square

Proof of Proposition 20. We first show that there can be no other source $mG' \in AD^*(mG) \setminus \{mG\}$. For the sake of contradiction, let the misinformation game $mG' \in AD^*(mG)$, with $mG' \neq mG$, and $d_{\Gamma'}(mG') = 0$. Let $k \in \mathbb{N}$ be the smallest integer such that $mG' \in AD^k(mG)$. Therefore, $mG' \notin AD^{k-1}(mG)$. Nevertheless, from Definition 10, $AD^k(mG) = AD(AD^{k-1}(mG))$. Thus, there is some $mG'' \in AD^{k-1}(mG)$, with $mG' \in AD(mG'')$. We reached a contradiction, since we assumed that mG' is a source.

It remains to show that $d_{\Gamma'}(mG) = 0$. Suppose, for the sake of contradiction, that $d_{\Gamma'}(mG) > 0$, so there is some $mG' \in AD^*(mG)$, $mG' \neq mG$ such that $mG \in AD(mG')$. Repeating the above argument, we conclude that there is some sequence of position vectors X such that $mG = mG_X$, which is impossible by Proposition 8. \square

Proof of Proposition 21. Assume some sink node mG' , i.e., $mG' \in K$. Since mG' is a sink in Γ' , any outgoing edges from mG' in Γ are self loops. This shows that $AD(mG') = \{mG'\}$, and, thus, $mG' \in \mathcal{T}$, which proves the first claim of the proposition.

For the second claim, we first observe that, as shown above, $AD(mG') = \{mG'\}$ for all $mG' \in K$, and thus $K \subseteq AD^\infty(mG)$. Now take some $\sigma \in NME(mG')$ for some $mG' \in K \subseteq AD^\infty(mG)$. Given that $AD(mG') = \{mG'\}$, it is easy to confirm that all the conditions of Definition 12 hold, and thus σ is an *sme*. \square

Proof of Proposition 22. An immediate consequence of Proposition 11. \square

Proof of Proposition 23. Since mG^1, mG^2 are canonical misinformation games on the same number of players $|N|$ and the same number of strategies $|S|$, they can be represented by $|N|+1$ matrices each with $|S|$ cells. Each cell of the matrices will contain a $|N|$ -dimensional vector containing the payoffs for each player. To show equality, we have to compare each component of each of those vectors from these misinformation games, which leads us to the result. \square

Proof of Proposition 24. Since $\theta[Y] = \theta[X]_{\vec{v}}$, in order to determine if $Y = X$, it suffices to show that $\vec{v} \in X$. Since, we keep the position vectors sets as ordered list, the membership query $\vec{v} \in X$ can be done in $O(\log |S|)$ time. \square

Proof of Lemma 25. It suffices to show that for each $mG \in \mathcal{T}$ there exists some $i \in \mathbb{N}$ such that $mG \in \mathcal{Q}_i$. For the sake of contradiction assume that there is no such \mathcal{Q}_i for mG . Let mG' be the most recent ancestor of mG that belongs to some \mathcal{Q}_j , $j \geq 0$. There is such mG' , since mG^0 is an ancestor of mG and $mG^0 \in \mathcal{Q}_0$. Observe that all the *proper* descendants $AD(mG') \setminus \{mG'\}$ of mG' belong to \mathcal{Q}_{j+1} . This is a contradiction from the choice of mG' ; thus, the result follows. \square

Proof of Proposition 26. From equation (12), if $\mathcal{Q}_i = \emptyset$, then $\mathcal{Q}_{i+1} = \emptyset$, and $\mathcal{T}_{i+1} = \emptyset$. Thus, $\cup_{k=1}^{i+1} \mathcal{T}_k = \cup_{k=1}^{\infty} \mathcal{T}_k$. From Lemma 25, $\cup_{k=1}^{i+1} \mathcal{T}_k = \mathcal{T}$. Then, from equation (13), we have $M_{i+1} = \mathcal{AD}^{i+1}(\cup_{k=1}^{i+1} \mathcal{T}_k)$, since $\mathcal{Q}_{i+1} = \emptyset$. Moreover, it holds $\mathcal{AD}^{i+1}(\cup_{k=1}^{i+1} \mathcal{T}_k) = \mathcal{AD}^*(\cup_{k=1}^{i+1} \mathcal{T}_k)$. Observe that after the $(i+1)$ -th step, no additional misinformation games will be generated, otherwise, $\mathcal{Q}_{i+1} \neq \emptyset$. Lastly, $\mathcal{AD}^*(\cup_{k=1}^{i+1} \mathcal{T}_k) = \mathcal{AD}^*(\mathcal{T})$, from the above. Therefore $M_{i+1} = \mathcal{AD}^*(\mathcal{T})$. From Theorem 14, M_{i+1} is a Stable Set, and the desideratum holds. \square

Proof of Proposition 27. As we noted in the previous subsection, for the set-based representation approximation $\tilde{\Gamma} = (2^S, \tilde{E})$ of the adaptation graph. For each of these nodes, we compute its *nmes* in $|N|t(\mathbf{NASH})$ time. These *nmes* may result in at most $|S|$ position vectors. Lastly, in the worst case, for each of these position vectors, we will perform an equality check of $O(|N|^2|S|)$ time steps (see Proposition 23). \square

Proof of Theorem 28. The complexity of Algorithm 7 is established in Proposition 27. We argue about the complexity of Algorithm 8. We note that for each misinformation game mG_X , such that X belongs to the set-based representation $\tilde{\mathcal{T}}$ of the terminal set, and for each position vector resulting from an *nme*, we perform an equivalence check. The number of position vectors is at most $|S|$. \square

Proof of Corollary 29. Since mG' is the last node on a maximal path from mG in Γ' , then it is a sink in Γ' . From Proposition 21 we obtain the desideratum. \square

Proof of Theorem 30. Assume a misinformation game mG . The **while** loop of steps 4–26 in Algorithm 7 will make $\mathfrak{L}_{\mathcal{AD}}(mG)$ steps. On the other hand, $\mathfrak{L}_{\mathcal{AD}}(mG) = \Theta(|S|)$. In each step, we compute the *nmes*, consuming $O(|N| \cdot t(\mathbf{NASH}))$ time. A single *nme* σ may have at most $|N|$ position vectors, i.e. $|\chi(\sigma)| \leq |N|$. For each position vector we perform an equality check, consuming time $O(|N|^2 \cdot |S|)$. \square

References

- [1] Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V.: Algorithmic Game Theory. Cambridge University Press, USA (2007)
- [2] Luce, R.D., Raiffa, H.: Games and Decisions: Introduction and Critical Survey. Wiley, USA (1957)
- [3] Bennett, P.G.: Hypergames: Developing a model of conflict. *Futures* **12**(6), 489–507 (1980)
- [4] Yossi, F.: Games with Unawareness. *The B.E. Journal of Theoretical Economics* **21**(2), 433–488 (2021)
- [5] Varsos, C., Fasoulakis, M., Flouris, G., Bitsaki, M.: Coordination Mechanisms with Misinformation. In: ICAART 2022: Proceedings of the 14th International Conference on Agents and Artificial Intelligence - Volume 1, pp. 237–244 (2022)

- [6] Toms, S.: Financial scandals: a historical overview. *Accounting and Business Research* **49**(5), 477–499 (2019)
- [7] Brown, P.N., Borowski, H.P., Marden, J.R.: Are multiagent systems resilient to communication failures? *ArXiv* (2017)
- [8] Meir, R., Parkes, D.C.: Congestion games with distance-based strict uncertainty. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 986–992 (2015)
- [9] 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)
- [10] Zamir, S.: *Bayesian Games: Games with Incomplete Information*, pp. 426–441. Springer, New York (2009)
- [11] Gao, J.: Uncertain bimatrix game with applications. *Fuzzy Optimization and Decision Making* **12**(1), 65–78 (2013)
- [12] Varsos, C., Flouris, G., Bitsaki, M., Fasoulakis, M.: A study of misinformation games. In: *PRICAI 2021: Trends in Artificial Intelligence - 18th Pacific Rim International Conference on Artificial Intelligence*, vol. 13031, pp. 76–87 (2021)
- [13] Papamichail, M., Varsos, C., Flouris, G.: Implementing the adaptation procedure in misinformation games. *Proceedings of the 12th Hellenic Conference on Artificial Intelligence* (2022)
- [14] Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H.: The complexity of computing a nash equilibrium. *SIAM J. Comput.* **39**(1), 195–259 (2009)
- [15] McKelvey, R.D., Turocy, T.L., McLennan, A.: *Gambit: Software Tools for Game Theory*, Version 16.0.0. (2014)
- [16] Copic, J., Galeotti, A.: Awareness as an Equilibrium Notion: Normal-Form Games (2006)
- [17] Schipper, B.C.: Unawareness - A gentle introduction to both the literature and the special issue. *Mathematical Social Sciences* **70**, 1–9 (2014)
- [18] Sasaki, Y.: Generalized Nash equilibrium with stable belief hierarchies in static games with unawareness. *Annals of Operations Research* **256**(2), 271–284 (2017)
- [19] Schipper, B.C.: Discovery and equilibrium in games with unawareness. *Journal of Economic Theory* **198**, 105365 (2021)
- [20] Arrow, K.J., Green, J.R.: Notes on expectations equilibria in bayesian settings. Technical report, Institute for Mathematical Studies in the Social Sciences (1973). Working Paper No. 33, Stanford University

- [21] Esponda, I., Pouzo, D.: Equilibrium in misspecified markov decision processes. *Theoretical Economics* **16**(2), 717–757 (2021)
- [22] Esponda, I., Pouzo, D.: Berk-nash equilibrium: A framework for modeling agents with misspecified models. *Econometrica* **84**(3), 1093–1130 (2016)
- [23] Esponda, I., Pouzo, D., Yamamoto, Y.: Asymptotic behavior of bayesian learners with misspecified models. *J. Econ. Theory* **195** (2019)
- [24] Spiegel, R.: Bayesian Networks and Boundedly Rational Expectations*. *The Quarterly Journal of Economics* **131**(3), 1243–1290 (2016)
- [25] Ray, D., King-Casas, B., Montague, P.R., Dayan, P.: Bayesian model of behaviour in economic games. In: *Advances in Neural Information Processing Systems 21, NIPS08*, pp. 1345–1352 (2008)
- [26] Heidhues, P., Kőszegi, B., Strack, P.: Convergence in models of misspecified learning. *Theoretical Economics* **16**(1), 73–99 (2021)
- [27] Easley, D., Kiefer, N.M.: Controlling a Stochastic Process with Unknown Parameters. *Econometrica* **56**(5), 1045–1064 (1988)
- [28] Lerer, A., Hu, H., Foerster, J., Brown, N.: Improving policies via search in cooperative partially observable games. In: *AAAI Conference on Artificial Intelligence*, vol. 34, pp. 7187–7194 (2020)
- [29] Carlsson, H., Damme, E.: Global games and equilibrium selection. *Econometrica* **61**(5), 989–1018 (1993)
- [30] Jordan, J.S.: Bayesian learning in normal form games. *Games and Economic Behavior* **3**(1), 60–81 (1991)
- [31] Fudenberg, D., Romanyuk, G., Strack, P.: Active learning with a misspecified prior. *Theoretical Economics* **12**(3), 1155–1189 (2017)
- [32] Fudenberg, D., Levine, D.K.: Self-confirming equilibrium. *Econometrica* **61**(3), 523–545 (1993)
- [33] Dekel, E., Fudenberg, D., Levine, D.K.: Payoff information and self-confirming equilibrium. *Journal of Economic Theory* **89**(2), 165–185 (1999)
- [34] Székely, G.J., Rizzo, M.L.: The uncertainty principle of game theory. *The American Mathematical Monthly* **114**(8), 688–702 (2007)
- [35] Gairing, M., Monien, B., Tiemann, K.: Selfish routing with incomplete information. *Theory of Computing Systems* **42**(1), 91–130 (2008)
- [36] Sasaki, Y., Kijima, K.: Hypergames and bayesian games: A theoretical comparison

of the models of games with incomplete information. *Journal of Systems Science and Complexity* **25**(4), 720–735 (2012)

- [37] Gharesifard, B., Cortés, J.: Stealthy strategies for deception in hypergames with asymmetric information. In: 2011 50th IEEE Conference on Decision and Control and European Control Conference, pp. 5762–5767 (2011)