



HAL
open science

Games of Influence

Umberto Grandi, Emiliano Lorini, Arianna Novaro, Laurent Perrussel

► **To cite this version:**

Umberto Grandi, Emiliano Lorini, Arianna Novaro, Laurent Perrussel. Games of Influence. Journal of Logic and Computation, 2021, exaa088, pp.1-29. <10.1093/logcom/exaa088>. <hal-03117720>

HAL Id: hal-03117720

<https://hal.science/hal-03117720v1>

Submitted on 21 Jan 2021

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

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



HAL Authorization

Games of Influence

Umberto Grandi¹, Emiliano Lorini², Arianna Novaro¹, Laurent Perrussel¹

¹IRIT, University of Toulouse

²CNRS, IRIT, Toulouse

{grandi, lorini, novaro, perrussel}@irit.fr

Abstract

In this paper we present two models for reasoning about strategic actions in opinion diffusion. In both models a set of agents are endowed with goals expressed compactly in a suitably defined language of Linear Temporal Logic, and are connected in an influence network which defines the underlying opinion diffusion process. The agents can act by exerting their influence or retain from it: in the first case we assume an initial state of incomplete information about the agents' opinions, while in the second we assume that the agents have complete information. We investigate the interplay between network topologies and the existence of game-theoretic solution concepts, and we give bounds for the computational complexity of strategic reasoning in both settings.

Keywords: Opinion diffusion; Nash equilibria; strategic reasoning; social influence

1 Introduction

We want to model a community in which agents can influence other agents. Namely, if agent i has influencing power over another agent j , then she can exert it by overtly expressing her opinions or publicly acting in a certain way, thereby influencing agent j 's opinions and behaviors.¹ We take the notion of influencing power as primitive, without decomposing it at a finer-grained level. Certainly, in the social reality, there are multiple factors that determine whether an agent has influencing power over another. One component is *trust*, which positively correlates with influencing power: the higher the trust agent j has in agent i 's judgment, the higher the power of agent i to influence agent j 's opinions by disclosing hers. Another component is that of *social tie* (Granovetter, 1973;

¹See Lorini et al. (2013) for a logical formalization of different varieties of social power, including influencing power.

Attanasi et al., 2016): agent i can influence agent j in case j has a certain level of social proximity with i .² The existence of a social tie between two individuals has been shown to play a fundamental role in the diffusion of opinions (Granovetter, 1973) and innovations (Rogers, 1983), imitation (Ellwood, 1901), behavioral contagion (Dimant, 2019) and social learning (Bandura, 1986). For instance, according to social learning theory, human agents are more willing to learn from socially similar agents with whom they can identify more easily (e.g., pop stars, sports personalities, business magnates, etc.) and who are socially closer to them (e.g., parents or friends). As emphasized in innovation diffusion theory, “an obvious principle of human communication is that the transfer of ideas occurs most frequently between two individuals who are alike, similar, or homophilous. Homophily is the degree to which pairs of individuals who interact are similar in certain attributes, such as beliefs, education, social status, and the like” (Rogers, 1983, p. 18).

Our model works under three general assumptions. First, the influencing power relationships between agents are dichotomous: i.e., we do not specify an agent i 's degree of power over another agent j , but we only say whether agent i has influencing power over j . Secondly, the graph describing the agents' influencing power relationships is static and does not change through interaction. Thirdly, an agent i 's opinions are affected by another agent j 's opinions only if j exerts her influence over i , despite the fact that i may know j 's opinions. The latter corresponds to the idea that influence requires an act of explicit communication by the influencer. Therefore, we exclude from our model all those forms of implicit influence that are pervasive in the human society whereby the influencee acquires knowledge about the influencer's opinions through inference and is latently affected by them.

The strategic aspect of influence is a central aspect of the model. Agents in the community have individual goals that they try to achieve by using their influencing power over others, as well as the structure of the influence graph. For example, suppose agent i has influencing power over agent j and agent j has influencing power over agent z . Then, i may decide to overtly express her opinion in order to induce j to adopt it, or she may stay tacit in order to prevent j from changing his current opinion. Furthermore, i may decide to affect j 's opinions in order to indirectly affect z 's opinions, given the influence relationship between j and z .

We present two variants of our model, under complete and incomplete information. In the complete information variant, it is assumed that agents' opinions are common knowledge. Nonetheless, since tacit opinions have no influence force, an agent can only exert her power over the others by explicitly supporting her opinions. The following example illustrates this variant of the model.

²According to Granovetter's theory, the strength of a social tie between two individuals is determined by multiple parameters including the amount of time, the emotional intensity, the intimacy (mutual confiding), and the reciprocal services characterizing the social relationship.




Figure 1: An influence network in which Ann, Cathy, and Bob have a shared influence on Jesse (in case of unanimous choice), and Ann has a direct influence on Bob.

Jesse has three kids: Ann, Bob and Cathy. Ann is the oldest, Bob is the youngest and Cathy is the one in the middle. Ann is strongly admired by her brother Bob, thereby having an influencing power over him (i.e., Bob is usually influenced by what her sister says). The family has to decide whether to go to the cinema or to make an excursion to the countryside. Jesse asks the kids what they would like to do. If they unanimously agree on the same option, then this option will be chosen. Otherwise, if an agreement is not reached, the default option will be to go to the cinema, which is what Jesse prefers. Each kid can decide either to abstain or to participate in the vote by declaring her/his preference.³ It is common knowledge in the family that Jesse and Bob prefer the cinema, while Ann and Cathy prefer the excursion. Nonetheless, according to the rule of the game, if a kid does not explicitly say what she/he prefers, her/his preference will not count. In this situation, Ann should explicitly declare her preference to be sure that her goal of making an excursion will be achieved. Otherwise, Jesse’s preference of going to the cinema will possibly prevail. In other words, choosing to explicitly declare her preference is for Ann a *winning strategy* relative to her goal of making an excursion. In fact, by participating in the vote, she will shape Bob’s preference such that an unanimous agreement to make an excursion will necessarily be reached.⁴

In the incomplete information variant of the model, an agent’s tacit opinions are epistemically inaccessible to the others: i.e., if an agent does not reveal her opinions, the others will be uncertain about them. In the incomplete information variant of the previous example, Ann would be uncertain about Cathy’s preference before Cathy has publicly revealed it. Consequently, Ann does not know whether, by choosing to explicitly declare her preference, a unanimous agreement to make an excursion will necessarily be reached. In particular, there is no *uniform strategy* for Ann such that Ann knows that, by choosing it, she will achieve her goal—where an agent’s uniform strategy is a strategy that assigns the same choices in states that are epistemically indistinguishable to the agent (van Benthem, 2001; Schobbens, 2004; Jamroga and Ågotnes, 2016). However,

³For the sake of simplicity, we assume that an expressed preference is a special kind of expressed opinion about the goodness (or desirability) of a given option.

⁴We suppose that, although Bob knows Ann’s tacit preference, he will not be affected by it unless Ann makes it explicit by participating to the vote.

Ann knows that the choice of explicitly declaring her opinions (weakly) dominates the choice of abstaining from participating in the vote: i.e., she knows that for any possible choice of Bob and Cathy, if she achieves her goal by choosing abstention, she will also achieve it by choosing to participate in the vote. The action of influencing the other agents about one’s own opinions for the first time has thus the double role of eliminating the uncertainty and influencing others about them.

Our contribution. The definition of our two variants of strategic opinion diffusion under influence, with complete and incomplete information, is presented in Section 2. The diffusion of opinions is unanimous: an agent changes her opinion if all her influencers who are using their influencing power are unanimous about it. This forms the basic bricks of a state-transition model in which agents are able to control the accessibility of their opinions towards the satisfaction of their goals, expressed in a logical language based on Linear Temporal Logic. In Section 3 we give the full definition of influence games, and we introduce game-theoretic concepts adapted to our information model (perfect-recall and memory-less *uniform* strategies), as well as solution concepts that take into consideration the epistemic aspect of an agent’s goal satisfaction (*subjective* winning and weak-dominant strategies, and *subjective* Nash equilibrium). We show some simple initial results, that stress the importance of the complete information model, which we study in Section 4. We then study the computational complexity of finding game-theoretic solutions for games of influence, both under complete and incomplete information, in Section 5. Our findings show that the problem of recognising if a given strategy profile is a (subjective) Nash equilibrium reduces to the satisfaction of Linear Temporal Logic formulas, and thus can be solved in polynomial space. All other tasks we consider are significantly harder, and we give upper bounds above exponential time. Section 6 concludes the paper.

Related work. There exist several contributions on the logical representation of the dynamics of opinion. Seligman et al. (2013) propose a Facebook logic, mixing propositional dynamic logic and dynamic epistemic logic. The aim is to represent how messages circulate over a social network and the associated impact: how agents gain information and how the network changes. The main difference with our work is that the logic does not focus on how agents aggregate incoming information and the strategic aspect of diffusion is also not taken into consideration. Ghosh and Velázquez-Quesada (2015) show how consensus may be reached among a set of agents influencing each other; an important difference with Facebook logic is the opinion type, i.e., preferences, and the aggregation procedure, which is represented as preference change methods. However, no strategic aspect is considered in the influence process. Our work is also close to that of Christoff and Hansen (2015) and Christoff et al. (2016) which adopts a similar approach with the above mentioned contribution: embedding in logic the dynamics of opinion. The authors

propose a sound and complete axiomatisation for reasoning about diffusion and its impact represented by model transformation, making a distinction between private and public opinions. We build on their model adding a strategic component with the introduction of the individual goals. We note in passing that earlier work by Seligman et al. (2011) and Christoff and Hansen (2013) already made use of the unanimity aggregation rule in the definition of their influence models.

Our work draws inspiration from psychological studies on group pressure, informal social communication and influence from majority groups to minorities (Asch, 1956; Festinger, 1950; Moscovici and Personnaz, 1980). Problems close to opinion diffusion are those of information cascades and knowledge diffusion, which have been given formal treatment in a logical setting (see, e.g., Ruan and Thielscher (2011); Baltag et al. (2013)).

Our model of opinion diffusion is inspired by a recent stream of papers that aims at conceiving diffusion models for *structured* information, being it belief sets (Schwind et al., 2015, 2016), preferences (Brill et al., 2016), or multiple binary issues (Grandi et al., 2015; Botan et al., 2019). In these papers, the diffusion rule is based on aggregation procedures imported from the theories of belief merging, voting, or judgment aggregation. We refer to the discussion in the cited works for the relation with classical models of diffusion, such as the De Groot or Lehrer-Wagner model (de Groot, 1974; Lehrer and Wagner, 1981) and threshold models (Granovetter, 1978).

Influence games can be considered a variation of the iterated version of the well-known Boolean games (Harrenstein et al., 2001; Bonzon et al., 2006; Grant et al., 2011; Wooldridge et al., 2013; Ågotnes et al., 2013; Gutierrez et al., 2015)—which have recently been extended with a social network structure where agents choose actions depending on the actions of their neighbors (Seligman and Thompson, 2015).

This paper builds on our previous work in which the assumption was that of complete information (Grandi et al., 2017). In complementary work, we studied similar models at a higher perspective, that of concurrent game structures in which the control of Boolean variables is shared among sets of agents, showing that the computational complexity of strategic reasoning is left unchanged—with respect to structures with exclusive propositional control (Belardinelli et al., 2017). Such structures are used to interpret formulas of Alternating-time Temporal Logic (Alur et al., 2002), a logic that can be used to reason about strategic aspects of iterated games. Similarly, Strategy Logic (Mogavero et al., 2010) is a formalism allowing to express the existence of solution concepts, such as Nash equilibria, in iterated games. We will use both formalisms for our complexity results.

2 Opinion diffusion

In this section we take inspiration from the framework of propositional opinion diffusion by Grandi et al. (2015), to model agents who have control over the influence they can

exert on other agents. We then adapt the diffusion process to this more complex setting.

2.1 Opinion and influence functions

Let $\mathcal{I} = \{p_1, \dots, p_m\}$ be a finite set of propositions, or *issues*, and $\mathcal{N} = \{1, \dots, n\}$ a finite set of individuals, or *agents*. Agents have opinions on the issues in \mathcal{I} in the form of a propositional evaluation (equivalently, a binary vector):

Definition 1. *The opinion of agent i is a function $B_i : \mathcal{I} \rightarrow \{1, 0\}$ where $B_i(p) = 1$ and $B_i(p) = 0$ express, respectively, the agent's opinion that p is true and that p is false.*

Let $\mathbf{B} = (B_1, \dots, B_n)$ denote the profile of opinions of agents in \mathcal{N} . Propositional evaluations can be used to represent ballots in multiple referenda, expressions of preference over alternatives, or judgments over correlated issues (Christian et al., 2007; Grandi and Endriss, 2011). Depending on the application at hand, an integrity constraint may be introduced to model correlations among the issues: for the sake of simplicity we do not introduce here any such constraint (see the recent work by Grandi et al. (2015); Schwind et al. (2015); Botan et al. (2019) on opinion diffusion with constraints).

Moreover, for each issue an agent has the option to use her influence power or not:

Definition 2. *We call influence function of agent i a map $I_i : \mathcal{I} \rightarrow \{1, 0\}$ where $I_i(p) = 1$ expresses that agent i is using her influence on p , and $I_i(p) = 0$ that she is not.*

We denote by $\mathbf{I} = (I_1, \dots, I_n)$ the profile composed of the agents' influence functions. By combining the opinions with the influence functions of an agent, we can build her expressed opinion as a three-valued function on the issues:

Definition 3. *Let B_i be the opinion of agent i and I_i her influence function. The expressed opinion of i is a function $P_i : \mathcal{I} \rightarrow \{1, 0, ?\}$ such that*

$$P_i(p) = \begin{cases} B_i(p) & \text{if } I_i(p) = 1 \\ ? & \text{if } I_i(p) = 0 \end{cases}$$

Again, $\mathbf{P} = (P_1, \dots, P_n)$ is the profile of expressed opinions of all the agents in \mathcal{N} . We denote by \mathbf{P}_C the restriction of profile \mathbf{P} to the individuals in $C \subseteq \mathcal{N}$.

2.2 Unanimous opinion diffusion

We define the process of influence among agents in our model starting from the setting of propositional opinion diffusion proposed by Grandi et al. (2015), that we adapt to take into account the presence of the influence function. Firstly, the agents are linked by an *influence network*, modeled as a directed irreflexive graph:

Definition 4. We call an influence network a directed irreflexive graph $E \subseteq \mathcal{N} \times \mathcal{N}$, where $(i, j) \in E$ reads as “agent j is influenced by agent i ”.

We also refer to E as the *influence graph* and to the individuals in \mathcal{N} as the nodes of the graph. Let $\text{Inf}(i) = \{k \in \mathcal{N} \mid (k, i) \in E\}$ be the set of *influencers* of agent i in the network E . As the graph is irreflexive, $i \notin \text{Inf}(i)$ for any $i \in \mathcal{N}$, i.e., agents do not influence themselves.

Given a profile of expressed opinions and an influence network, we model the process of opinion diffusion by means of an aggregation function, which shapes the opinion of an agent from the expressed opinions of other agents.

Definition 5. An aggregation procedure for agent i is a class of functions

$$F_{i,C} : \{0, 1\}^{\mathcal{I}} \times \{0, 1, ?\}^{\mathcal{I} \times C} \longrightarrow \{0, 1\}^{\mathcal{I}} \text{ for all } C \subseteq \mathcal{N} \setminus \{i\}$$

that maps agent i 's opinion and the expressed opinions of a set of agents C to agent i 's opinion.

We drop C from the subscript when clear from context. Many aggregation procedures have been considered in the literature on judgment aggregation (Grandi and Endriss, 2011; Lang and Slavkovik, 2013), and they can be adapted to our setting. Notable examples are quota rules, where agents change their opinion if the number of people disagreeing with them is higher than a given quota, such as the majority rule—see the class of threshold models studied in the literature on opinion diffusion (Granovetter, 1978; Schelling, 1978). Unanimity is also an instance of a quota rule, and we assume it to be the aggregation function defining influence for the agents for the remainder of this paper:

Definition 6. The unanimous issue-by-issue aggregation procedure is defined as follows:

$$F_{i,C}^U(B_i, \mathbf{P}_C)(p) = \begin{cases} x & \text{if there is } k \in C \text{ such that } P_k(p) = x \in \{0, 1\} \text{ and} \\ & P_j(p) \in \{x, ?\} \text{ for all } j \in C \\ B_i(p) & \text{otherwise} \end{cases}$$

That is, an individual will change her opinion about issue p if and only if all agents in C (usually among her influencers) that are using their influence on their opinion are unanimous (i.e., some of the agents in C may not be using their influence on p , but those who do are all accepting it or rejecting it).

2.3 Strategic actions and state transitions

In our model of strategic diffusion, the central notion of *state* consists of the profiles of opinions and influence functions of the agents. The formal definition is as follows:

Definition 7. A state is a tuple $S = (\mathbf{B}, \mathbf{I})$ where \mathbf{B} is a profile of opinions and \mathbf{I} is a profile of influence functions. The set of all states is denoted by \mathcal{S} .

Different states can correspond to the same profile of expressed opinions \mathbf{P} : namely, if some agents are not using their influence on an issue (i.e., its value in \mathbf{P} is ‘?’), then there will be multiple corresponding states depending on their opinions in \mathbf{B} on the issue.

In our model, agents can use specific actions of type $\text{exert}(J)$ and $\text{quiet}(J)$ for $J \subseteq \mathcal{I}$ to push (or not) their opinions on the issues in J to the other agents. We thus allow for simultaneous influence on multiple propositions.

The set of individual actions for the agents is then defined as:

$$\mathcal{A} = \{(\text{exert}(J), \text{quiet}(J')) \mid J, J' \subseteq \mathcal{I} \text{ and } J \cap J' = \emptyset\}.$$

Each joint action, or action profile, $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A}^n$, consisting of an individual action for each agent, induces a deterministic transition function between states:

Definition 8. The transition function $\text{succ} : \mathcal{S} \times \mathcal{A}^n \rightarrow \mathcal{S}$ associates to each state $S = (\mathbf{B}, \mathbf{I})$ and joint action \mathbf{a} a new state $S' = (\mathbf{B}', \mathbf{I}')$ as follows, for all $i \in \mathcal{N}$ and $p \in \mathcal{I}$. For $a_i = (\text{exert}(J), \text{quiet}(J')) \in \mathcal{A}$:

$$\begin{aligned} \bullet \quad I'_i(p) &= \begin{cases} 1 & \text{if } p \in J \\ 0 & \text{if } p \in J' \\ I_i(p) & \text{otherwise} \end{cases} \\ \bullet \quad B'_i &= F_i^U(B_i, \mathbf{P}'_{\text{Inf}(i)}) \end{aligned}$$

Where \mathbf{P}' is the expressed profile obtained from opinion profile \mathbf{B} and influence profile \mathbf{V}' .

Observe that, according to Definition 8, for all issues $p \in (\mathcal{I} \setminus (J \cup J'))$ their value for the influence function does not change between S and S' . By a slight abuse of notation we denote with $\mathbf{a}(S)$ the state $\text{succ}(S, \mathbf{a})$ obtained from S and \mathbf{a} by applying the transition function succ . We also use the following abbreviations: $\text{skip} = (\text{exert}(\emptyset), \text{quiet}(\emptyset))$ for doing nothing, $\text{exert}(J) = (\text{exert}(J), \text{quiet}(\emptyset))$, $\text{quiet}(J) = (\text{exert}(\emptyset), \text{quiet}(J))$, and we drop curly parentheses in $\text{exert}(\{p\})$ and $\text{quiet}(\{p\})$. The influence process thus occurs after the actions have changed the influence status of the agents’ opinions: i.e., first, actions affect the influence of opinions, and then each agent modifies her opinion on the basis of those opinions of her influencers that are now pushed by them.

The notion of *history* describes the temporal aspect of the agents’ opinion dynamic:

Definition 9. Given a set of issues \mathcal{I} , a set of agents \mathcal{N} , and aggregation procedures F_i for $i \in \mathcal{N}$ over a network E , a history H is an infinite sequence of states H_j and action profiles $H = (H_0, \mathbf{a}_0, H_1, \mathbf{a}_1 \dots)$, such that $H_{t+1} = \mathbf{a}_t(H_t)$ for all $t \in \mathbb{N}$.

A partial history $H_{\leq t} = (H_0, \mathbf{a}_0, H_1, \mathbf{a}_1 \dots \mathbf{a}_{t-1}, H_t)$ is a finite sequence of $t \geq 0$ states and actions of history H .

We denote the set of all histories by \mathcal{H} and the set of all partial histories by \mathcal{H}^+ . Observe that Definition 9 restricts the set of all possible histories to those that correspond to a *run* of the influence dynamics of Definition 8. For notational convenience, for any $i \in \mathcal{N}$ and for any $t \in \mathbb{N}$, we denote with $H_{i,t}^B$ agent i 's tacit opinion in state H_t and with $H_{i,t}^I$ agent i 's influence in H_t .

Example 1. Consider the example in Figure 2 for a single issue $p \in \mathcal{I}$, where the agents are $\mathcal{N} = \{A, B, C\}$ such that $\text{Inf}(A) = \{B, C\}$. The initial state of the history is $H_0 = ((0, 1, 1), (1, 1, 0))$, i.e., agent A currently has the tacit opinion that p is false, while agents B and C have the tacit opinion that p is true. Let $\mathbf{a}_0 = (\text{skip}, \text{skip}, \text{exert}(p))$ and $\mathbf{a}_1 = (\text{skip}, \text{quiet}(p), \text{skip})$ be the joint actions of the agents at the first two states. Namely, agent C pushes her opinion on p , and at the next step B keeps quiet on hers.

$$\begin{array}{ccccc} ((0, 1, 1), (1, 1, 0)) & \xrightarrow{\hspace{10em}} & ((1, 1, 1), (1, 1, 1)) & \xrightarrow{\hspace{10em}} & ((1, 1, 1), (1, 0, 1)) \\ H_0 & (\text{skip}, \text{skip}, \text{exert}(p)) & H_1 & (\text{skip}, \text{quiet}(p), \text{skip}) & H_2 \end{array}$$

Figure 2: The first three states of a history, with the influence network $B \rightarrow A \leftarrow C$.

If all individuals are using the unanimous aggregation procedure, then in state H_1 , agent A 's opinion about p has changed to 1, as all her influencers are pushing unanimously for p , while in H_2 no opinion is updated.

Under the assumption of incomplete information the opinions are not common knowledge among the agents (unless they decide to push them). Thus, the influence function naturally defines an indistinguishability relation between states:

Definition 10. Given two states S and S' and an agent $i \in \mathcal{N}$ we say that they are indistinguishable for i , written $S \sim_i S'$ if the following is the case:

1. $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I_j(p) = I'_j(p)$, and
2. $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I_j(p) = 1$ implies $B_j(p) = B'_j(p)$.


Intuitively, by condition 1 an agent considers equivalent two states where the influence status is the same, while condition 2 affirms that in equivalent states all the pushed opinions have to be equivalent. We now prove in the following proposition that the indistinguishability relation defined above is indeed an equivalence relation:

Proposition 1. The relation \sim_i of Definition 10 is an equivalence relation.

Proof. We need to prove that \sim_i is reflexive, symmetric and transitive. For reflexivity, consider a state S : clearly $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I_j(p) = 1$ implies $B_j(p) = B_j(p)$ and $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I_j(p) = I_j(p)$. Hence, $S \sim_i S$. For symmetry, assume that $S \sim_i S'$: we want to show that $S' \sim_i S$. From condition 1 we immediately get that $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I'_j(p) = I_j(p)$ and thus from condition 2 we get $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I'_j(p) = 1$ implies $B'_j(p) = B_j(p)$: hence, $S' \sim_i S$. For transitivity, assume that $S \sim_i S'$ and $S' \sim_i S''$: we want to show that $S \sim_i S''$. Since $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I_j(p) = I'_j(p)$ and $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I'_j(p) = I''_j(p)$, we have that $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I_j(p) = I''_j(p)$. Hence, we get that $(\forall j \in \mathcal{N})(\forall p \in \mathcal{I}) : I''_j(p) = 1$ implies $B_j(p) = B'_j(p) = B''_j(p)$. Therefore \sim_i is an equivalence relation. \square

Given agent $i \in \mathcal{N}$ and state S , we let $[S]_i = \{S' \mid S \sim_i S'\}$ be the equivalence class of S for i . The set \mathcal{S}_i contains all the representative elements of the equivalence classes of states from the perspective of agent i .

Example 2. Consider the example of the introduction. Let us denote the four agents with capital letters A, B, C, J , and recall that the decision was about one issue p : going to the cinema, denoted with 1, and going to an excursion, denoted with 0. Recall the influence network of Figure 1 and that $B_B(p) = B_J(p) = 1$ and $B_A(p) = B_C(p) = 0$. If only Jesse and Cathy push their opinion, the situation can be described in the following figure, which represents the expressed opinions of the agents, with a dashed line whenever the expressed opinion is equal to ?:



In this example, the equivalence class of the initial state from Ann's point of view is:

$$[S_0]_A = \{((0, \underline{1}, 0, 1), (0, 0, 1, 1)), ((0, \underline{0}, 0, 1), (0, 0, 1, 1))\}.$$

That is, the only uncertainty for Ann is Bob's opinion, i.e., the second element of the first vector of a state. The equivalence class of Jesse contains instead four states, to take into account her uncertainty about the opinions of both Ann and Bob:

$$[S_0]_J = \{((\underline{0}, \underline{1}, 0, 1), (0, 0, 1, 1)), ((\underline{1}, \underline{1}, 0, 1), (0, 0, 1, 1)), ((\underline{0}, \underline{0}, 0, 1), (0, 0, 1, 1)), ((\underline{1}, \underline{0}, 0, 1), (0, 0, 1, 1))\}.$$

The equivalence relation at the basis of our epistemic model is defined from the

assignment of values to a given set of variables contained in the states, departing from a strand of literature that considers epistemic accessibility relations as the primary objects of actions, and grounding these relations to the notion of states (van der Hoek et al., 2011; Herzig et al., 2015). Observe that the successor states of two indistinguishable states are not necessarily indistinguishable from the perspective of an agent:

Proposition 2. *There exists a model with two states $S_1 \sim_i S_2$ and an action profile \mathbf{a} such that $\mathbf{a}(S_1) \not\sim_i \mathbf{a}(S_2)$.*

Proof. Assume that $\mathcal{N} = \{1, 2, 3\}$, $\mathcal{I} = \{p\}$ and the influence network E has no edges, i.e., the agents do not influence one another. Let $S_1 = ((1, 0, 0), (0, 0, 1))$ and $S_2 = (0, 0, 0), (0, 0, 1)$: clearly, $S_1 \sim_3 S_2$, i.e., the two states are indistinguishable from the perspective of agent 3, as they differ on the opinion of agent 1 on p —who is hiding it. Now let $\mathbf{a} = (\text{exert}(p), \text{skip}, \text{skip})$: we have that $\mathbf{a}(S_1) = ((1, 0, 0), (1, 0, 1))$ and $\mathbf{a}(S_2) = ((0, 0, 0), (1, 0, 1))$, which are not in the same equivalence class for agent 3. \square

The following example shows that joint actions reduce the uncertainty of the agents:

Example 3. *Let us continue the discussion in Example 2. Assume that Ann pushes her opinion at the first step, i.e., the action profile is $\mathbf{a} = (\text{exert}(p), \text{skip}, \text{skip}, \text{skip})$. The successor of any state in Ann’s equivalence class, as well as in Jesse’s equivalence class is $((0, 0, 0, 1), (1, 0, 1, 1))$. After Ann has supported her opinion publicly, both Ann and Jesse know that Bob will have updated his opinion to the one of Ann. Hence, even if Bob still does not use his influence, there is no uncertainty for both Ann and Jesse that his opinion is equal to 0.*

We now define a notion of indistinguishability on partial histories to model the *perfect recall* of an agent:

Definition 11. *Let H and H' be two histories, $i \in \mathcal{N}$ an agent and $t, t' \in \mathbb{N}$. We say that two partial histories $H_{\leq t}$ and $H'_{\leq t'}$ are indistinguishable for agent i if:*

- $t = t'$,
- for all $0 \leq k \leq t$ we have that $H_k \sim_i H'_k$,
- for all $0 \leq k \leq t$ we have that $\mathbf{a}_k = \mathbf{a}'_k$.

In such case we denote $(H_{\leq t}) \sim_i^H (H'_{\leq t'})$.

Intuitively, two histories are indistinguishable for an agent i if they are generated by the same set of actions and generate states that are indistinguishable for agent i . Consider the following example:

Example 4. Let us consider the usual example, and assume that in the initial state H_0 both Bob and Cathy push their opinion. We are in state $H_0 = ((0, 1, 0, 1), (0, 1, 1, 0))$. However, the following states cannot be distinguished by Ann:

$$((0, 1, 0, 1), (0, 1, 1, 0)) \sim_A ((0, 1, 0, 0), (0, 1, 1, 0))$$

Assume $\mathbf{a}_0 = (\text{exert}(p), \text{skip}, \text{skip}, \text{skip})$: Ann pushes her opinion. Ann still cannot distinguish between the two successor states:

$$\begin{aligned} \mathbf{a}_0\left(\left((0, 1, 0, 1), (0, 1, 1, 0)\right)\right) &= ((0, 0, 0, 1), (1, 1, 1, 0)) \\ \mathbf{a}_0\left(\left((0, 1, 0, 0), (0, 1, 1, 0)\right)\right) &= ((0, 0, 0, 0), (1, 1, 1, 0)) \end{aligned}$$

However, whatever action will be played next—we assume all players in \mathbf{a}_1 do nothing—Ann will stop being uncertain about Jesse’s opinion, since her influence has “reached” Jesse. Figure 3 summarizes the indistinguishability for Ann between the two partial histories $H_{\leq 2} \sim_A^H H'_{\leq 2}$, which are finally converging in state $H_2 = ((0, 0, 0, 0), (1, 1, 1, 0))$.




Figure 3: Indistinguishability relation between two partial histories.

Observe that our notion of indistinguishability relation for partial histories is in line with formal semantics for temporal epistemic logic in which agents are assumed to have perfect memory of the past (see, e.g., van Benthem et al. (2009)).

2.4 Individual goals

By using her influence power, an agent influences others towards the satisfaction of her goal. To account for the temporal aspect of our model, we follow the work on iterated Boolean games (Gutierrez et al., 2015) and we define a language $\mathcal{L}_{\text{LTL-}I}$ to express individual goals on influence using Linear Temporal Logic LTL.

Let therefore $\mathcal{L}_{\text{LTL-}I}$ be defined as follows:

$$\varphi ::= \text{op}(i, p) \mid \text{inf}(i, p) \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \bigcirc\varphi \mid \varphi_1 \mathcal{U} \varphi_2$$

where i ranges over \mathcal{N} and p ranges over \mathcal{I} . We read $\text{op}(i, p)$ as “agent i ’s opinion is that

p is true”, while $\neg\text{op}(i, p)$ reads “agent i ’s opinion is that p is not true” (as agents have binary opinions). Moreover, $\text{inf}(i, p)$ reads “agent i is using her influence on p ”.

The reading of $\bigcirc\varphi$ is “ φ is going to be true at the next state” and of $\varphi_1\mathcal{U}\varphi_2$ is “ φ_1 will be true until φ_2 is true”: they are the standard LTL operators ‘next’ and ‘until’. As usual, we can define the temporal operators ‘eventually’ (\diamond) and ‘henceforth’ (\square) as $\diamond\varphi = \top\mathcal{U}\varphi$ and $\square\varphi = \neg\diamond\neg\varphi$.

The interpretation of $\mathcal{L}_{\text{LTL-1}}$ -formulas relative to histories is defined as follows.

Definition 12. *Let H be a history, φ a formula of $\mathcal{L}_{\text{LTL-1}}$ and $k, k', k'' \in \mathbb{N}$. Then:*

$$\begin{aligned} H_{\leq k} \models \text{op}(i, p) &\Leftrightarrow H_{i,k}^B(p) = 1 \\ H_{\leq k} \models \text{inf}(i, p) &\Leftrightarrow H_{i,k}^I(p) = 1 \\ H_{\leq k} \models \neg\varphi &\Leftrightarrow H_{\leq k} \not\models \varphi \\ H_{\leq k} \models \varphi_1 \wedge \varphi_2 &\Leftrightarrow H_{\leq k} \models \varphi_1 \text{ and } H_{\leq k} \models \varphi_2 \\ H_{\leq k} \models \bigcirc\varphi &\Leftrightarrow H_{\leq k+1} \models \varphi \\ H_{\leq k} \models \varphi_1\mathcal{U}\varphi_2 &\Leftrightarrow \exists k' : (k \leq k' \text{ and } H_{\leq k'} \models \varphi_2 \text{ and} \\ &\quad \forall k'' : \text{if } k \leq k'' < k' \text{ then } H_{\leq k''} \models \varphi_1) \end{aligned}$$

Formulas of $\mathcal{L}_{\text{LTL-1}}$ will be used to express agents’ goals on the iterative diffusion process. As individuals do not have any influence on the initial state of the history, we will consider only goals of the form $\bigcirc\varphi$ and $\varphi\mathcal{U}\psi$, for any φ and ψ in $\mathcal{L}_{\text{LTL-1}}$, which we denote as *goal formulas*.

For a subset of agents $C \subseteq \mathcal{N}$ and issues $J \subseteq \mathcal{I}$ consider the following goals on consensus and influence in situations of opinion diffusion:

$$\begin{aligned} \text{cons}(C, J) &:= \diamond\square(\text{pcons}(C, J) \vee \text{ncons}(C, J)) \\ \text{issue-cons}(C, J) &:= \diamond\square \bigwedge_{p \in J} (\text{pcons}(C, \{p\}) \vee \text{ncons}(C, \{p\})) \\ \text{influence}(i, C, J) &:= \diamond\square \bigwedge_{p \in J} ((\text{op}(i, p) \rightarrow \bigcirc\text{pcons}(C, p)) \\ &\quad \wedge (\neg\text{op}(i, p) \rightarrow \bigcirc\text{ncons}(C, p))) \end{aligned}$$

where:

$$\begin{aligned} \text{pcons}(C, J) &:= \bigwedge_{i \in C} \bigwedge_{p \in J} \text{op}(i, p) \\ \text{ncons}(C, J) &:= \bigwedge_{i \in C} \bigwedge_{p \in J} \neg\text{op}(i, p). \end{aligned}$$

Intuitively, an agent holding the goal $\text{cons}(C, J)$ wants at some point to reach a

stable consensus either for or against the issues in J with the agents in C . For the $\text{issue-cons}(C, J)$ goal, all agents in C must reach a consensus on the issues in C , which can be positive or negative depending on the issue. The $\text{influence}(i, C, J)$ goal expresses the idea that agent i wants to eventually gain a stable influence over the people in C on the issues in J (i.e., they will always hold her opinion at the next step).

Even if our model features an important epistemic aspect, formalized by the indistinguishability relations \sim_i and \sim_i^H on states and histories, we do not include an epistemic operator in the goal language, but rather use the classical notions of uniform winning strategies in the following sections to take this aspect into account.

3 Games of influence

We are now ready to combine all concepts introduced in the previous sections to give the definition of an influence game:

Definition 13. *An influence game is a tuple $IG = (\mathcal{N}, \mathcal{I}, E, S_0, F_i, \gamma_1, \dots, \gamma_n)$ where \mathcal{N} , \mathcal{I} , E and S_0 are, respectively, a set of agents, a set of issues, an influence network, and an initial state, F_i for $i \in \mathcal{N}$ is an aggregation procedure, and γ_i is agent i 's goal formula.*

Given an influence game, the agents build their strategies to attain their goals. We first introduce two kinds of strategies available to agents, namely memory-less and perfect-recall *uniform* strategies, taking into account the epistemic accessibility relation.

It is worth noting the similarity between stochastic games (Shapley, 1953) and influence games, in which the transition probabilities are degenerate (as captured by the transition function of Definition 8). A difference between the two game classes is that stochastic games have state-dependent utility functions, while in influence games agents have state-independent (possibly long-term) goals represented by LTL formulas.

3.1 Strategies

The first type of individual strategies to be considered are memory-less, i.e., they only depend on the equivalence class of the current state:

Definition 14. *A memory-less uniform strategy for player i is a function $Q_i : \mathcal{S} \rightarrow \mathcal{A}$ such that $Q_i(S) = Q_i(S')$ if $S \sim_i S'$, associating an action to every equivalence class of states for agent i .*

A strategy profile is a tuple $\mathbf{Q} = (Q_1, \dots, Q_n)$. For notational convenience, we also use \mathbf{Q} to denote the function $\mathbf{Q} : \mathcal{S} \rightarrow \mathcal{A}^n$ such that for all $S \in \mathcal{S}$ we have $\mathbf{Q}(S) = \mathbf{a}$ if and only if $Q_i(S) = a_i$, for all $i \in \mathcal{N}$.

The second definition we provide is the full-blown notion of perfect-recall uniform strategy, which assigns an action to each equivalence class of partial histories that has been observed by the player.

Definition 15. We call a perfect-recall strategy for player i a function $Q_i : \mathcal{H}^+ \rightarrow \mathcal{A}$ such that $Q_i(H') = Q_i(H'')$ for partial histories H' and H'' if $H' \sim_i^H H''$, which associates an action to every equivalence class of finite partial histories for agent i .

As the following definition highlights, every strategy profile (of either type) induces a different set of possible histories for each agent, when combined with an initial state.

Definition 16. Let S_0 be an initial state and let \mathbf{Q} be a memory-less uniform strategy profile. The induced history $H^{S_0, \mathbf{Q}} \in \mathcal{H}$ is defined as follows:

$$\begin{aligned} H_0^{S_0, \mathbf{Q}} &= S_0 \\ H_{t+1}^{S_0, \mathbf{Q}} &= \text{succ}(H_t^{S_0, \mathbf{Q}}, \mathbf{Q}(H_t^{S_0, \mathbf{Q}})) \end{aligned}$$

For uniform perfect-recall strategies, the condition is the following:

$$\begin{aligned} H_0^{S_0, \mathbf{Q}} &= S_0 \\ H_{t+1}^{S_0, \mathbf{Q}} &= \text{succ}(H_t^{S_0, \mathbf{Q}}, \mathbf{Q}(H_0^{S_0, \mathbf{Q}}, \mathbf{a}_0, \dots, H_t^{S_0, \mathbf{Q}})) \end{aligned}$$

where the successor function succ is as per Definition 8.

The two types of strategies let us differentiate between more or less sophisticated agents, depending on the information they take into account to choose their actions.

3.2 Solution concepts

The first solution concept we study is that of winning strategy. Intuitively, Q_i is a winning strategy for player i if and only if i knows that, by playing this strategy, she will achieve her goal no matter what the other players do. In this respect, we consider the winning condition from the *subjective* point of view of an agent:

Definition 17. Let IG be an influence game, S_0 its initial state, and Q_i a strategy for player i . We say that Q_i is a subjective winning strategy for player i in state S_0 if for all profiles \mathbf{Q}_{-i} of strategies of players other than i , and for all states S such that $S \sim_i S_0$, we have that $H^{S, \mathbf{Q}} \models \gamma_i$.

Namely, Q_i is a subjective winning strategy if, starting from any state indistinguishable from S_0 from i 's perspective, its combination with any profile of strategies by the other players leads to a history that verifies the goal of i . The definition applies to both memory-less and perfect-recall strategies.

We distinguish the *subjective* winning strategy of Definition 17 from an *objective* winning strategy: i.e., Q_i is an *objective winning strategy* for player i in state S_0 if for all profiles \mathbf{Q}_{-i} of strategies of players other than i , we have that $H^{S_0, \mathbf{Q}} \models \gamma_i$. Clearly, a subjective winning strategy is also objective since \sim_i is an equivalence relation.

A subjective winning strategy corresponds to the notion of power studied by sociologists (Barnes, 1988; Castelfranchi, 2003) according to which, for an agent to have the power of achieving γ_i , she must have both the objective capability of achieving γ_i and the knowledge about her capability (see Herzig and Troquard (2006) and Lorini et al. (2013) for a formalization of this concept of power in a logic of actions).

Example 5. *In our running example, we represent with a Boolean variable p the options at stake: $\text{op}(J, p)$ is true if Jesse's opinion is to go to the cinema, and false if she would rather go to the excursion. Let us start our analysis from when only Cathy pushes her opinion $H_0 = ((0, 1, 0, 1), (0, 0, 1, 0))$, and assume that the goal of Ann is $\gamma_A = \diamond \square \neg \text{op}(J, p)$: that is, Ann wants to influence Jesse so that eventually her opinion will be to go to the excursion. There are four possible states for Ann, since the opinion of Bob and Jesse has not been revealed yet. However, in any history starting from each of these initial states, if Ann plays $\text{exert}(p)$ then eventually a state in which $\neg \text{op}(J, p)$ holds is reached (see the figure in Example 4 for an example of two such histories). Therefore, playing $\text{exert}(p)$ in the equivalence class of state H_0 and skip otherwise is a memory-less subjective winning strategy for Ann.*

As we will show in further sections, the concept of winning strategy is rather strong for our setting. Let us then define the less demanding notion of weak dominance:

Definition 18. *Let IG be an influence game and Q_i a strategy for player i . We say that Q_i is a subjective weakly dominant strategy for player i and initial state S_0 if and only if for all profiles \mathbf{Q}_{-i} of strategies of players other than i , for all strategies Q'_i for player i , and for all states $S \sim_i S_0$, we have that*

$$H^{S, (Q'_i, \mathbf{Q}_{-i})} \models \gamma_i \Rightarrow H^{S, (Q_i, \mathbf{Q}_{-i})} \models \gamma_i.$$

Namely, Q_i is a subjective weakly dominant strategy, if it gives at least the same result (or better) as any other strategy Q'_i in all possible profiles and induced histories.

As for the concept of winning strategy, we also distinguish *subjective* weak dominance from *objective*. In particular, we say that Q_i is an *objective weakly dominant strategy* for player i and initial state S_0 if and only if for all profiles \mathbf{Q}_{-i} of strategies of players other than i , for all strategies Q'_i for player i , we have that $H^{S_0, (Q'_i, \mathbf{Q}_{-i})} \models \gamma_i \Rightarrow H^{S_0, (Q_i, \mathbf{Q}_{-i})} \models \gamma_i$.

Example 6. *Let us now revisit our running example and assume a different initial state in which no agent pushes anything. We are thus in state $H_0 = ((0, 1, 0, 1)(0, 0, 0, 0))$. There*

are eight indistinguishable states for Ann, since all other opinions are not expressed. In particular, without knowing the tacit opinion of Cathy, Ann cannot know if she can influence Jesse to go to the excursion: this is possible if also Cathy wants to go to the excursion, or if she wants to go to the cinema but does not push it. Otherwise, Jesse might still stick to her initial opinion since her influencers would not be unanimous. Still, if Ann's goal is $\gamma_A = \diamond \square \neg \text{op}(J, p)$, the action of revealing her opinion in any indistinguishable initial state (and do nothing afterwards) is a subjective weakly dominant strategy for Ann: if her goal is satisfied in a possible history, then it is also satisfied by playing the $\text{exert}(p)$ action in the initial state.

Finally, we introduce the concept of Nash equilibrium for influence games:

Definition 19. Let IG be an influence game and let \mathbf{Q} be a strategy profile. Agent i has a subjective profitable deviation Q'_i at initial state S_0 if for all states $S \sim_i S_0$ we have

$$H^{S, (Q'_i, \mathbf{Q}_{-i})} \models \gamma_i \text{ and } H^{S, (\mathbf{Q}_i, \mathbf{Q}_{-i})} \not\models \gamma_i.$$

We say that \mathbf{Q} is a subjective Nash equilibrium for initial state S_0 if and only if no agent has a profitable deviation at S_0 .

In words, in a subjective Nash equilibrium no agent wants to unilaterally deviate from the chosen strategy. Note that we are making a further assumption here, i.e., that a deviation is profitable only if it leads to the satisfaction of the agent's goal *in all states* indistinguishable from the initial one. This corresponds to a *qualitative* variant of the notion of Bayesian Nash equilibrium (BNE): a similar qualitative variant of BNE is studied by Amor et al. (2019) in the context of possibility theory.

We say that agent i has an *objective profitable deviation* at initial state S_0 if we have that $H^{S_0, (Q'_i, \mathbf{Q}_{-i})} \models \gamma_i$ and $H^{S_0, (\mathbf{Q}_i, \mathbf{Q}_{-i})} \not\models \gamma_i$. The concept of *objective* Nash equilibrium is defined in the usual way from the concept of *objective* profitable deviation.

The notion of subjective Nash equilibrium given in Definition 19 is quite demanding from an epistemic point view. Indeed, it requires that every agent knows the opponents' actual equilibrium strategy. In order to weaken it, we can draw inspiration from Kalai & Lehrer's notion of subjective equilibrium Kalai and Lehrer (1993) according to which (i) each agent is assumed to best respond to her beliefs (or conjectures) about the opponents' strategies, and (ii) the play induced by the agents' actual equilibrium strategy coincides with the plays induced by the beliefs of each agent.⁵ A crucial aspect of Kalai & Lehrer's notion is that an agent's beliefs about the opponents' strategies may be wrong and differ from the other agents' beliefs.

In our model, agents do not form conjectures about the opponents' strategies and have no uncertainty about future histories. They only have knowledge and uncertainty

⁵This notion of equilibrium has also been called conjectural equilibrium (Battigalli and Guaitoli, 1997) and self-confirming equilibrium (Fudenberg and Levine, 1993).

about the actual state which is solely determined either by the agents' present influencing opinions (in the memory-less case) or by the agents' present and past influencing opinions and the agents' past actions (in the perfect-recall case). For this reason, Kalai & Lehrer's notion of subjective equilibrium is not expressible in our model.

However, there is an alternative way to weaken the notion of subjective equilibrium of Definition 19 and to make it "more subjective". First of all, we could define, for every strategy profile $\mathbf{Q} = (Q_i, Q_{-i})$ and state S , the set

$$\mathcal{Q}(S, Q_i, Q_{-i}) = \{Q'_{-i} : H^{S, (Q_i, Q_{-i})} = H^{S, (Q_i, Q'_{-i})}\}.$$

This is the set of strategies of i 's opponents that with state S and i 's strategy Q_i are history-equivalent to their strategy Q_{-i} . Secondly, we could say that the strategy profile \mathbf{Q} is a *truly* subjective equilibrium for the initial state S_0 if and only if, for every agent i and every possible deviation Q'_i , there exists $Q'_{-i} \in \mathcal{Q}(S_0, Q_i, Q_{-i})$ such that, if $H^{S, (Q'_i, Q'_{-i})} \models \gamma_i$ then $H^{S, (Q_i, Q'_{-i})} \models \gamma_i$ for some $S \sim_i S_0$.

This notion of equilibrium is weaker than the one of Definition 19 since $Q_{-i} \in \mathcal{Q}(S_0, Q_i, Q_{-i})$. This means that every subjective Nash equilibrium in the sense of Definition 19 is a truly subjective equilibrium as well. It is also less demanding in epistemic terms. It only requires that an agent believes that the opponents will play some strategy that, given the agent's actual equilibrium strategy, is history-equivalent to the opponents' actual equilibrium strategy. We leave for future work an in-depth analysis of this notion of *truly* subjective equilibrium.

3.3 Results

We begin by showing that joint observable actions in the incomplete information model can only reduce the uncertainty of the agents.

Proposition 3. *For history H , agent i and $t \geq 0$, we have that $|[H_{t+1}]_i| \leq |[H_t]_i|$.*

Proof. We provide the proof for the case of a single issue, whose generalization to multiple issues is straightforward. We begin by proving the following fact. If S is a state, let $U_i([S]_i) = \{j \in \mathcal{N} \mid \exists S, S' \in [S]_i \text{ s.t. } S_j^B(p) \neq S_j^{B'}(p)\}$ be the set of agents j on whose opinions agent i is uncertain about in state S . From the definition we obtain that $|[H_t]_i| = 2^{U_i([H_t]_i)}$. We now show that for all $t \geq 0$ we have that $U_i([H_{t+1}]_i) \subseteq U_i([H_t]_i)$. Since $\text{exert}(p)$ actions remove the uncertainty about the agent playing it, the only thing we need to show is that it is not possible to have an agent j such that $j \in U_i([H_{t+1}]_i)$ but $j \notin U_i([H_t]_i)$. Suppose this is the case. If $a_{t,j}$ (i.e., j 's action at time t) is $\text{exert}(p)$ or skip , then agent j 's opinion is known at step $t+1$. Suppose $a_{t,j} = \text{quiet}(p)$. We need to check that there are no $S, S' \in [H_{t+1}]_i$ where j has different opinions. But this is not possible, since j 's opinion is either the same at time t , or is the result of influence updates starting

from *known* opinions. Hence, the set of agents whose opinion i is uncertain about is decreasing with the play of each action profile. \square

The following result is a corollary of the previous proof, and shows the importance of the complete information model that will be studied in later sections:

Proposition 4. *Let H be a history. If for each agent $i \in \mathcal{N}$ and issue $p \in \mathcal{I}$ there is a time $t \geq 0$ such that $a_{t,i} = \text{exert}(p)$, then there is a time $T \geq 0$ such that $|[H_T]_i| = 1$ for all i (i.e., complete information is reached).*

Proof. (sketch). A careful analysis of the proof of Proposition 3 reveals that the strict inequality $|[H_{t+1}]_i| < |[H_t]_i|$ is obtained whenever an agent different than i reveals her tacit opinion on a previously unknown variable p . Equivalence classes have cardinality 2^U where U is the number of opinions that are uncertain from an agent i 's viewpoint. Since every $\text{exert}(p)$ action divides the number of models by a factor of 2, we eventually reach complete information, i.e., cardinality of 1. \square

4 Influence under complete information

In this section we focus on the special case of complete information, simplifying the model. While the notion of states remains the same, i.e., a pair composed of an opinion profile \mathbf{B} and an influence profile \mathbf{I} , no two distinct states are indistinguishable, as opinions are common knowledge among the agents from the start. A unique history is generated by a sequence of profiles of $\text{exert}(p)$ or $\text{quiet}(p)$ actions starting from an initial state S_0 .

Given a game of influence, the two notions of memory-less and perfect-recall strategies do not therefore need to take into consideration the indistinguishability of states. Definition 16 still applies, defining the unique history generated by a profile of strategies and an initial state.

4.1 Game-theoretic results under complete information

In this section we analyze the interplay between network structure and existence of solutions concepts for the goals defined in Section 2.4. We assume memory-less strategies and complete information.

We begin by restricting the class of graphs to *delegation graphs*,⁶ i.e., networks where each agent has at most one influencer:

⁶Similar networks are created in the study of *liquid democracy* or *transitive proxy voting*, where a set of voters can either vote directly on an issue at stake, or delegate their vote to a different member of the electorate (Christoff and Grossi, 2017). Note that agent i being influenced by agent j would correspond to i delegating her vote to j .

Proposition 5. *If E is acyclic and $|Inf(i)| \leq 1$ for all agents $i \in \mathcal{N}$, and if agent a has goal $\gamma_a := \text{cons}(C_a, J)$ where $J \subseteq \mathcal{I}$ and $C_a := \{k \in \mathcal{N} \mid a \in Inf(k)\} \cup \{a\}$, then agent a has an objective winning strategy.*

Proof (sketch). Consider an acyclic E and an agent a with goal γ_a . Let Q_a be the strategy associating to every state S action $\text{exert}(J)$. We want to show that $H^{S_0, (Q_a, \mathbf{Q}_{-a})} \models \gamma_a$ holds for all S_0 and \mathbf{Q}_{-a} . Consider the position of agent a in the graph for arbitrary S_0 . In case there is no agent b such that $a \in Inf(b)$, the goal reduces to $\text{cons}(\{a\}, J)$ which is always trivially satisfied. In case $Inf(a) = \emptyset$, by playing $\text{exert}(J)$ in S_0 and since every agent uses unanimous aggregation, at stage 1 all child nodes of a will update their beliefs on J by copying a 's opinion (she is their only influencer). Moreover, they can't change their opinions on J later on in the history.

On the other hand, suppose there is some agent b such that $a \in Inf(b)$ and some agent $c \in Inf(a)$. By assumption on E we thus have that $Inf(a) = \{c\}$ and $Inf(b) = \{a\}$. Hence, either at some point k in the history all ancestors of a will have reached consensus, such that by playing $\text{exert}(J)$ from point $k + 1$ onwards the consensus among a and her child nodes will be maintained, or there is no such k . Since there is a unique path linking a to one of the source nodes of E , if her ancestors always disagree in the history it means that there is some agent among them who has a different opinion and who will never play $\text{exert}(J)$. Thus, the opinion of a will nonetheless be stable and γ_a will be attained. \square

The assumption of acyclicity in the above result rules out the situation where all nodes in a cycle play $\text{exert}(J)$ and they start in S_0 by having alternating positive and negative opinions on the issues in J . Moreover, having at most one influencer per agent ensures each agent to have full control over their child nodes. Observe that the type of graphs covered by Proposition 5 include, among others, trees and chains. In general, however, we can see how winning strategy is too strong of a solution concept, since in order to attain it the scope of goals and network structures have to be narrow.

If we move to the less demanding concept of (objective) weak dominance, we may intuitively think that a strategy associating action $\text{exert}(J)$ to all states is weakly dominant for an agent a having goal $\gamma_a := \text{influence}(a, C, J)$ for $C \subseteq \mathcal{N}$, regardless of the network E or the initial state S_0 : in fact, all agents use the monotonic aggregation rule F_i^U . However, the following example shows that to satisfy goals of type γ_a as described, an agent could sometimes benefit from hiding her opinion.

Example 7. *For four agents $\mathcal{N} = \{1, 2, 3, 4\}$ and one issue $\mathcal{I} = \{p\}$ consider the network $E = \{(1, 2), (2, 3), (3, 4)\}$. Suppose that agents 1 and 2 associate action $\text{exert}(p)$ to all states, agent 3 associates action $\text{quiet}(p)$ only to those states where 1, 2 and 3 agree on p , while agent 4 always skips. Let the goal of agent 2 be $\gamma_2 = \text{influence}(2, \{4\}, \{p\})$ and consider the history below for these strategies, where goal γ_2 is not attained (we only represent \mathbf{B}):*

$$\begin{array}{ccccccc}
(0, 1, 0, 1) & \xrightarrow{\quad\quad\quad} & (0, 0, 1, 0) & \xrightarrow{\quad\quad\quad} & (0, 0, 0, 1) & \xrightarrow{\quad\quad\quad} & \dots \\
H_0 & \xrightarrow{\quad\quad\quad} & H_1 & \xrightarrow{\quad\quad\quad} & H_2 & &
\end{array}$$

$((\text{exert}(p), \text{exert}(p), \text{exert}(p), \text{skip}))$

From state H_2 onwards, given the strategies of the agents, the profile of opinions $\mathbf{B} = (0, 0, 0, 1)$ won't change. Consider now a strategy for agent 2 identical to the previous one, but for the fact that it associates action $\text{quiet}(p)$ to state H_0 :

$$\begin{array}{ccc}
(0, 1, 0, 1) & \xrightarrow{\quad\quad\quad} & (0, 0, 0, 0) & \xrightarrow{\quad\quad\quad} & \dots \\
H_0 & \xrightarrow{\quad\quad\quad} & H_1 & &
\end{array}$$

$((\text{exert}(p), \text{quiet}(p), \text{exert}(p), \text{skip}))$

From state H_1 onwards, given the strategies of the agents, the profile of opinions won't change. Thus, we found a network, an initial state H_0 , and strategies for the other agents, such that agent 2 is better off by hiding her opinion on p to satisfy her influence goal γ_2 .

We can now see an example of how the network structure and the agents' goals can yield a Nash equilibrium in the spirit of anti-coordination games (Rapoport and Chamamah, 1966).

Proposition 6. *Let IG be an influence game. If E is a graph such that $\text{Inf}(i) \neq \emptyset$ for all $i \in \mathcal{N}$ and $\gamma_i = \text{issue-cons}(\mathcal{N}, J)$ for $i \in \mathcal{N}$ and $J \subseteq \mathcal{I}$, then there exists an objective Nash equilibrium \mathbf{Q} for any initial state S_0 .*

Proof (sketch). We construct a strategy profile \mathbf{Q} such that the issue-wise consensus goal $\gamma_i = \text{issue-cons}(\mathcal{N}, J)$ of all agents in \mathcal{N} is attained, thus making \mathbf{Q} a Nash equilibrium. Intuitively, the agents must coordinate on the issues in J about which they disagree in S_0 . We denote S_0 with H_0 for clarity. Recall that we are considering IG s where all agents use the unanimous aggregation of Definition 6.

Observe that the agents want to reach a consensus on each issue in J , regardless of whether the consensus is in favour or against it. Let $J^c \subseteq J$ be the set of issues on which the agents agree already in S_0 : for these issues, the strategy of the agents can associate any action to any state (as, by definition, they already have a consensus).

For the issues in $J \setminus J^c$, it suffices to construct a strategy which progressively propagates, w.l.o.g., the positive opinion on all the issues in J in the network (an equivalent strategy would propagate the negative opinion on the issues, or a mix of both). The assumption of $\text{Inf}(i) \neq \emptyset$ for all $i \in \mathcal{N}$ prevents the case where two sources have opposite opinions on the issues in J , which would make the consensus among all agents impossible (as they can never change their opinions). Hence, construct a strategy profile \mathbf{Q} such that $Q_i(H_0) = \text{exert}(p)$ for all i and p such that $H_{i,0}^B(p) = 1$ and $Q_k(H_0) = \text{quiet}(q)$ for all k and q such that $H_{k,0}^B(q) = 0$. The strategy is defined analogously for all other states until consensus is reached, and then it can just associate action skip to the issues in J .

By this coordination effort of the agents, all the $\gamma_i = \text{issue-cons}(\mathcal{N}, J)$ will be satisfied, thus making Q a Nash equilibrium. \square

The above proof can easily be adapted to the case where the agents want to reach specifically a positive or negative consensus (i.e., the goals $\diamond\Box\text{pcons}(\mathcal{N}, J)$ and $\diamond\Box\text{ncons}(\mathcal{N}, J)$), provided that there is at least one agent holding a positive or negative opinion on the issues in S_0 ; as well as to the case where the agents want to reach a global consensus on all issues (i.e., the goal $\text{cons}(\mathcal{N}, J)$), provided that they are currently disagreeing on all issues in J (i.e., $J^c = \emptyset$).

For the complete information model, we thus see how limited (and sometimes counter-intuitive) results can be obtained on the existence of classical solution concepts, depending on the network structure and the goal of the agents. In the next section we thus move to the study of computational complexity problems related to existence of solution concepts in influence games.

5 Computational complexity and solution concepts

We now study the computational complexity of reasoning about solution concepts in influence games. Our results will rely on reductions to well-known logical formalisms, such as ATL, LTL, and graded strategy logic. The results presented in this section show upper bounds on the main strategic reasoning tasks associated to influence games: deciding if a game admits a winning strategy or a unique Nash equilibrium, and testing if a given profile of strategies is a Nash equilibrium. We will consider three complexity classes: PSPACE, the class of problems solvable by a Turing machine in polynomial space, 2-EXPTIME and 3-EXPTIME, the class of problems solvable in time that is doubly (triply, respectively) exponential in the size of the input.⁷

5.1 Problem definitions

We provide the formal definitions of the complexity problems we are interested in studying for influence games for both information models, in line with the solution concepts we introduced in Section 3.2, and we provide complexity results.

Let $x \in \{\text{subjective}, \text{objective}\}$. The first problem we define takes as input an influence game and an agent, and it asks whether there exists a subjective/objective winning strategy for that agent in the game:

INPUT: $IG = (\mathcal{N}, \mathcal{I}, E, F, S_0, \gamma_1, \dots, \gamma_n), i \in \mathcal{N}$.
E-WINNING $_i^x(F)$: Is there a x winning strategy Q_i for agent i in IG ?

⁷For a general reference on complexity theory see, e.g., Arora and Barak (2009).

Next, we define the problem of checking if a strategy profile is a subjective/objective Nash equilibrium in a given influence game:

INPUT: $IG = (\mathcal{N}, \mathcal{I}, E, F, S_0, \gamma_1, \dots, \gamma_n), \mathbf{Q}$.
M-NASH $^x(F)$: Is \mathbf{Q} a x Nash equilibrium of IG ?

Then, given an influence game, we define the problems of existence and uniqueness of a subjective/objective Nash equilibrium:

INPUT: $IG = (\mathcal{N}, \mathcal{I}, E, F, S_0, \gamma_1, \dots, \gamma_n)$.
E-NASH $^x(F)$: Is there some x Nash equilibrium \mathbf{Q} of IG ?
U-NASH $^x(F)$: Is there a unique x Nash equilibrium \mathbf{Q} of IG ?

Observe that F is not part of the input but a parameter of the problems defined above, since, as we shall see later, different aggregation functions may give rise to computational problems in different complexity classes.

5.2 Existence of winning strategy

We recall here the syntax and semantics of Alternating-time Temporal Logic ATL. The language of ATL is defined by the following BNF:

$$\varphi ::= p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \langle\langle C \rangle\rangle \bigcirc \varphi \mid \langle\langle C \rangle\rangle (\varphi_1 \mathcal{U} \varphi_2)$$

where C ranges over $2^{\mathcal{N}}$ and p ranges over a set of atomic propositions $\mathcal{A}tm$. The formula $\langle\langle C \rangle\rangle \bigcirc \varphi$ is read “coalition C has the capability to ensure that φ is going to be true in the next state, regardless of what the agents outside C do”, and the formula $\langle\langle C \rangle\rangle (\varphi_1 \mathcal{U} \varphi_2)$ is read “coalition C has the capability to ensure that φ_1 will be true until φ_2 is true, regardless of what the agents outside C do”.

We consider the standard ATL semantics in terms of concurrent game structures:

Definition 20. A concurrent game structure is a tuple $\mathbf{G} = (\mathcal{W}, \mathcal{M}, R, T, Val)$ where \mathcal{W} is a set of worlds or states, \mathcal{M} is a set of moves, function $R : \mathcal{N} \times \mathcal{W} \rightarrow 2^{\mathcal{M}} \setminus \emptyset$ defines a nonempty repertoire of moves for each agent at each world, $T : \mathcal{W} \times \mathcal{M}^n \rightarrow \mathcal{W}$ is a transition function mapping a world w and a move profile $m = (m_1, \dots, m_n)$ to the successor world $T(w, m)$, and $Val : \mathcal{W} \rightarrow 2^{\mathcal{A}tm}$ is a valuation function.

In ATL, a strategy for player i is a function f_i that maps every finite sequence of worlds $\pi = w_0 \dots w_n$ in \mathcal{W}^+ (i.e., a *path*) to a move $f_i(\pi) \in R(i, w_n)$ available to agent i at the end of path π .⁸ A strategy for coalition C is a function G_C that maps every agent

⁸Observe that ATL does not distinguish a semantics based on perfect-recall strategies from a semantics based on memory-less strategies, i.e., a function mapping a world to the set of moves available in this world. Specifically, the set of ATL validities with a semantics based on perfect-recall strategies and the set of ATL validities with a semantics based on memory-less strategies are the same.

$i \in C$ to a strategy $G_C(i)$ for i . The set of strategies for coalition C is denoted by Str_C . We write G instead of $G_{\mathcal{N}}$, and Str instead of $Str_{\mathcal{N}}$.

A move profile is used to determine a successor of a world using the transition function T . We define the set of available move profiles at world w as follows:

$$\mathcal{P}(w) = \{(m_1, \dots, m_n) \in \mathcal{M}^n \mid m_i \in R(i, w) \text{ for all } i \in \mathcal{N}\}$$

The set of *possible successors* of w is defined as $\mathcal{Succ}(w) = \{T(w, m) \mid m \in \mathcal{P}(w)\}$. An infinite sequence $\lambda = w_0 w_1 w_2 \dots$ of worlds from \mathcal{W} is called a *computation* if $w_{k+1} \in \mathcal{Succ}(w_k)$ for all $k \geq 0$. The k -th component w_k in λ is denoted by $\lambda[k]$. Moreover, for every computation $\lambda = w_0 w_1 w_2 \dots$ and for every positive integer k , $Prf(\lambda, k) = w_0 \dots w_k$ denotes the prefix of λ of length k .

The set $\mathcal{O}(w, G_C)$ denotes all computations $\lambda = w_0 w_1 w_2 \dots$ such that $w_0 = w$ and, for every $k \geq 0$, there is $m = (m_1, \dots, m_n) \in \mathcal{P}(w_k)$ such that $G_C(i)(w_0 \dots w_k) = m_i$ for all $i \in C$, and $T(w, m) = w_{k+1}$. Notice that $\mathcal{O}(w, G)$ is a singleton.

Truth conditions of ATL are defined relative to a CGS $\mathbf{G} = (\mathcal{W}, \mathcal{M}, R, T, Val)$ and a world $w \in \mathcal{W}$ —we omit the standard truth conditions for Boolean formulas:

$$\begin{aligned} \mathbf{G}, w \models \langle\langle C \rangle\rangle \bigcirc \varphi &\Leftrightarrow \text{there exists } G_C \in Str_C \text{ such that} \\ &\quad \mathbf{G}, \lambda[1] \models \varphi \text{ for all } \lambda \in \mathcal{O}(w, G_C) \\ \mathbf{G}, w \models \langle\langle C \rangle\rangle (\varphi_1 \mathcal{U} \varphi_2) &\Leftrightarrow \text{there is } G_C \in Str_C \text{ such that for} \\ &\quad \text{all } \lambda \in \mathcal{O}(w, G_C) \text{ there is } k \geq 0 \\ &\quad \text{such that } \mathbf{G}, \lambda[k] \models \varphi_2 \text{ and} \\ &\quad \mathbf{G}, \lambda[h] \models \varphi_1 \text{ for all } 0 \leq h < k. \end{aligned}$$

The following theorem gives an upper bound on the complexity of the problem of checking the existence of an objective winning strategy for perfect-recall strategies and for the unanimous aggregation procedure under complete information.

Theorem 1. *Let $x = \text{objective}$. Then, $\text{E-WINNING}_i^x(F_i^U)$ for perfect-recall strategies is in EXPTIME.*

Proof (sketch). We reduce $\text{E-WINNING}_i^x(F_i^U)$ to the model checking problem of ATL where $Atm = \{\text{op}(i, p) \mid i \in \mathcal{N} \text{ and } p \in \mathcal{I}\} \cup \{\text{inf}(i, p) \mid i \in \mathcal{N} \text{ and } p \in \mathcal{I}\}$. First of all, recall our assumption that *goal formulas* in influence games are of the form $\bigcirc \varphi$ and $\varphi \mathcal{U} \psi$, for any φ and ψ in $\mathcal{L}_{\text{LTL-}\perp}$. Moreover, observe that we can generate the CGS $\mathbf{G}_{IG} = (\mathcal{W}, \mathcal{M}, R, T, Val)$ corresponding to the influence game IG as follows:

- $\mathcal{W} = \mathcal{S}$,
- $\mathcal{M} = \mathcal{A}$,

- for all $i \in \mathcal{N}$ and $w \in \mathcal{W}$, $R(i, w) = \mathcal{M}$,
- for all $w \in \mathcal{W}$ and $\mathbf{a} \in \mathcal{A}^n$, $T(w, \mathbf{a}) = \text{succ}(w, \mathbf{a})$,
- for all $\text{op}(i, p), \text{inf}(i, p) \in \mathcal{A}tm$ and $w \in \mathcal{W}$, $\text{op}(i, p) \in \text{Val}(w)$ iff $B_i(p) = 1$, $\text{inf}(i, p) \in \text{Val}(w)$ iff $V_i(p) = 1$.

Secondly, it is easy to verify that the answer to the problem $\text{E-WINNING}_i^x(F_i^U)$ is positive if and only if $\mathbf{G}_{IG}, w_0 \models \langle\langle \{i\} \rangle\rangle \gamma_i$ with $w_0 = S_0$. Alur et al. (2002, Theorem 5.2) show that the model checking problem for ATL can be solved in time $O(m \cdot l)$ where m is the number of transitions in the CGS (which is polynomial in the size of \mathcal{W} and \mathcal{M}) and l is the size of the formula. Since the number of states S is exponential in the size of an influence game (more precisely, it is $O(2^{|\mathcal{I}|} \times |\mathcal{N}|)$ where \mathcal{I} is the set of issues and \mathcal{N} is the set of agents), we obtain the upper bound. \square

5.3 Nash membership

We study here the complexity of the $\text{M-NASH}^x(F)$ problem, for memory-less strategies.

We begin by translating a memory-less strategy in the language of LTL-I defined in Section 2.4. A conjunction of literals $\alpha(S)$ can be defined to uniquely identify a state S : $\alpha(S)$ will specify the tacit opinion of all individuals and their influence function. For an action $a_i = (\text{exert}(J), \text{quiet}(J'))$ for agent i , let $\beta_i(a)$ be the following formula:

$$\beta_i(a) = \bigwedge_{p \in J} \bigcirc \text{inf}(i, p) \wedge \bigwedge_{q \in J'} \bigcirc \neg \text{inf}(i, q).$$

In case $a = \text{skip}$ we let $\beta_i(a) = \top$. Given a memory-less strategy Q_i , we construct the following formula:

$$\tau_i(Q_i) = \bigwedge_{S \in \mathcal{S}} \alpha(S) \rightarrow \beta_i(Q_i(S)).$$

If \mathbf{Q} is a strategy profile, let $\tau(\mathbf{Q}) = \bigwedge_{i \in \mathcal{N}} \tau_i(Q_i)$. We now need to encode the unanimous aggregation function into a formula as well.

Consider the following formula $\text{unan}(i, p)$:

$$\begin{aligned} & \bigcirc \text{op}(i, p) \leftrightarrow \\ & \left(\left[\bigwedge_{j \in \text{Inf}(i)} \bigcirc \neg \text{inf}(j, p) \wedge \text{op}(i, p) \right] \vee \right. \\ & \left[\bigvee_{j \in \text{Inf}(i)} \bigcirc \text{inf}(j, p) \wedge \bigwedge_{j \in \text{Inf}(i)} (\bigcirc \text{inf}(j, p) \rightarrow \text{op}(j, p)) \right] \vee \\ & \left. \left[\bigvee_{j, z \in \text{Inf}(i):} (\bigcirc \text{inf}(j, p) \wedge \bigcirc \text{inf}(z, p) \wedge \right. \right. \\ & \left. \left. \text{op}(j, p) \wedge \neg \text{op}(z, p) \right) \wedge \text{op}(i, p) \right] \right). \end{aligned}$$

as well as the following formula $unan(i, \neg p)$:

$$\begin{aligned} & \bigcirc \neg \text{op}(i, p) \leftrightarrow \\ & \left(\left[\bigwedge_{j \in \text{Inf}(i)} \bigcirc \neg \text{inf}(j, p) \wedge \neg \text{op}(i, p) \right] \vee \right. \\ & \left[\bigvee_{j \in \text{Inf}(i)} \bigcirc \text{inf}(j, p) \wedge \bigwedge_{j \in \text{Inf}(i)} (\bigcirc \text{inf}(j, p) \rightarrow \neg \text{op}(j, p)) \right] \vee \\ & \left. \left[\bigvee_{j, z \in \text{Inf}(i):} (\bigcirc \text{inf}(j, p) \wedge \bigcirc \text{inf}(z, p) \wedge \right. \right. \\ & \left. \left. \neg \text{op}(j, p) \wedge \text{op}(z, p) \right) \wedge \neg \text{op}(i, p) \right] \right). \end{aligned}$$

If the influencers of agent i are unanimous, then agent i 's opinion should be defined according to the three cases described in Definition 6. Recall that, while actions take one time unit to be performed—hence the \bigcirc operator in front of $\text{inf}(j, p)$, the diffusion of opinions is simultaneous. Let now:

$$\tau(F_i^U) = \bigwedge_{\{i \in \mathcal{N} \mid \text{Inf}(i) \neq \emptyset\}} \bigwedge_{p \in \mathcal{I}} (unan(i, p) \wedge unan(i, \neg p))$$

be the formula encoding the transition process defined by the opinion diffusion. The formula $\tau(F_i^U)$ is polynomial in both the number of individuals and the number of issues—in the worst case it is quadratic in the number of agents and linear in the number of issues. We are now ready to prove the following:

Theorem 2. *Let $x = \{\text{objective}, \text{subjective}\}$. M-NASH $^x(F_i^U)$ for memory-less strategies is in PSPACE.*

Proof. We give the proof for the subjective case, thus under incomplete information. The case for complete information is straightforward from the given proof. Let \mathbf{Q} be a memory-less strategy profile for game IG . The following algorithm can be used to check whether \mathbf{Q} is a subjective Nash equilibrium. For all individuals $i \in \mathcal{N}$, and for all states $S \in [S_0]_i$, we check the following entailment:

$$\alpha(S) \wedge \tau(\mathbf{Q}) \wedge \tau(F_i^U) \models_{\text{LTL}} \gamma_i$$

in the language $\mathcal{L}_{\text{LTL-1}}$ of LTL built out of the set of propositions $\{\text{op}(i, p) \mid i \in \mathcal{N} \text{ and } p \in \mathcal{I}\} \cup \{\text{inf}(i, p) \mid i \in \mathcal{N} \text{ and } p \in \mathcal{I}\}$.

If for a given agent i and state S the previous entailment is not verified, we consider all the possible strategies $Q'_i \neq Q_i$ for agent i —there are exponentially many of them, but each one can be specified in space polynomial in the size of the input—and check the

following entailment:

$$\alpha(S) \wedge \tau(Q_{-i}, Q'_i) \wedge \tau(F_i^U) \models_{\text{LTL}} \gamma_i$$

If the answer is positive we output NO; otherwise, we proceed until all strategies and all individuals have been considered (as well as the associated initial states).

The entailment for LTL can be reduced to the problem of checking validity in LTL. Indeed, the following holds:

$$\psi \models_{\text{LTL}} \varphi \text{ iff } \models_{\text{LTL}} \Box\psi \rightarrow \varphi.$$

Since the problem of checking validity in LTL can be solved in PSPACE (Sistla and Clarke, 1985), we obtain that M-NASH^x is in PSPACE. \square

Theorem 2 can be easily generalized to all aggregation procedures that can be axiomatized by means of polynomially many formulas in $\mathcal{L}_{\text{LTL}-1}$, such as quota rules. This is not the case for all aggregation procedures: the majority rule would for instance require an exponential number of formulas, one for each subset of influencers forming a relative majority, against the constraint of using polynomial space.

Observe also that the reduction described in this setting can be used to check the existence of Nash equilibria (and uniqueness, or existence of winning strategies), by model-checking an LTL formula for each strategy profile. However, such LTL formulas would be doubly exponential in the size of the game description.

5.4 Nash existence and uniqueness

Strategy logic SL is a logic of strategic reasoning that embeds ATL (Mogavero et al., 2010), but that allows to quantify over strategies in a more flexible way. The flexibility of SL is mainly due to the fact that it has variables for strategies that are associated to specific agents with a binding operator.

An extension of SL by *graded* quantifiers over tuples of strategy variables has also been presented (Aminof et al., 2016). We here focus on this SL extension, denoted by G-SL, whose language is defined by the following BNF:

$$\begin{aligned} \varphi ::= & p \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \bigcirc\varphi \mid \varphi_1 \mathcal{U}\varphi_2 \mid \langle\langle x_1, \dots, x_\ell \rangle\rangle^{\geq k} \varphi \mid \\ & [i \mapsto x]\varphi \end{aligned}$$

where p ranges over $\mathcal{A}tm$, i ranges over \mathcal{N} , k and ℓ range over the set of positive integers, x_1, \dots, x_ℓ range over a countable infinite set of variables $\mathcal{V}ar$ with the additional constraint that $x_h \neq x_{h'}$ for all $1 \leq h, h' \leq \ell$ such that $h \neq h'$. We read $\langle\langle x_1, \dots, x_\ell \rangle\rangle^{\geq k} \varphi$ as “there exist at least k different ℓ -tuples of strategies ensuring φ ” while $[i \mapsto x]\varphi$ is read

“if agent i selects strategy x then φ will hold”. Formulas $\bigcirc\varphi$ and $\varphi_1\mathcal{U}\varphi_2$ have the same reading as in LTL. Moreover, the basic SL operator $\langle\langle x \rangle\rangle$ abbreviates $\langle\langle x \rangle\rangle^{\geq 1}$ in G–SL.

As for ATL, SL semantics is based on concurrent game structures. However, to provide an interpretation of SL some extra components are needed. We denote with X an *assignment* function mapping agents and variables in $\mathcal{N} \cup \mathcal{V}ar$ to strategies in Str . For every assignment function X , strategy $G \in Str$ and $e \in \mathcal{N} \cup \mathcal{V}ar$, we write $X[e \mapsto G]$ to denote the assignment function that differs from X only in the fact that e maps to G . We extend this definition to tuples of agents or variables $\vec{e} = (e_1, \dots, e_\ell)$ and tuples of strategies $\vec{G} = (G_1, \dots, G_\ell)$ with $e_h \neq e_{h'}$ for $h \neq h'$, by denoting with $X[\vec{e} \mapsto \vec{G}]$ the assignment function differing from X only in that e_h maps to G_h for each $1 \leq h \leq \ell$.

Furthermore, for every assignment function X , we write G_X to denote the strategy in Str generated by the assignment function X . That is, G_X is the strategy in Str such that, for all $i \in \mathcal{N}$, $G_X(i) = X(i)(i)$. For every world $w \in W$, we write $\lambda_{w,X}$ to denote the computation starting in w generated by the assignment function X . More precisely, $\lambda_{w,X}$ denotes the computation $w_0w_1\dots$ such that $\mathcal{O}(w, G_X) = \{(w_0w_1\dots)\}$. For every path π , we write X_π to denote the assignment function obtained by shifting all strategies in the image of X by π . For all $e \in \mathcal{N} \cup \mathcal{V}ar$, $i \in \mathcal{N}$ and $\pi, \pi' \in W^+$, we thus have:

$$X_\pi(e)(i)(\pi') = X(e)(i)(\pi \cdot \pi')$$

Any G–SL formula φ is evaluated relative to a CGS \mathbf{G} , a world $w \in \mathcal{W}$ and assignment function X . Let $\vec{x} = (x_1, \dots, x_\ell)$. Then:

$$\begin{aligned} \mathbf{G}, w, X \models [i \mapsto x]\varphi &\Leftrightarrow \mathbf{G}, w, X[i \mapsto X(x)] \models \varphi \\ \mathbf{G}, w, X \models \langle\langle \vec{x} \rangle\rangle^{\geq k}\varphi &\Leftrightarrow \text{there exists } k \text{ many } \ell\text{-tuples} \\ &\quad \vec{G}_1, \dots, \vec{G}_k \text{ of strategies} \\ &\quad \text{such that for all } 1 \leq h, h' \leq k \\ &\quad \text{if } \vec{G}_h \neq \vec{G}_{h'} \text{ then } h \neq h' \\ &\quad \text{and } \mathbf{G}, w, X[\vec{x} \mapsto \vec{G}_h] \models \varphi \\ \mathbf{G}, w, X \models \bigcirc\varphi &\Leftrightarrow \mathbf{G}, \lambda_{w,X}[1], X_{Prf(\lambda_{w,X},1)} \models \varphi \\ \mathbf{G}, w, X \models \varphi_1\mathcal{U}\varphi_2 &\Leftrightarrow \text{there exists } k \geq 0 \text{ such that} \\ &\quad \mathbf{G}, \lambda_{w,X}[k], X_{Prf(\lambda_{w,X},k)} \models \varphi_2 \\ &\quad \text{and for all } 0 \leq h < k \text{ we have} \\ &\quad \mathbf{G}, \lambda_{w,X}[h], X_{Prf(\lambda_{w,X},h)} \models \varphi_1 \end{aligned}$$

The following theorem highlights the high complexity of checking existence and uniqueness of Nash equilibria in influence games.

Theorem 3. *Let $x = \text{objective}$. E-NASH $^x(F_i^U)$ and U-NASH $^x(F_i^U)$ for perfect-recall*

strategies are both in 3EXPTIME.

Proof (sketch). As in the proof of Theorem 1, we generate the CGS \mathbf{G}_{IG} corresponding to the influence game IG . Following Aminof et al. (2016), we reduce $E\text{-NASH}^x(F_i^U)$ and $U\text{-NASH}^x(F_i^U)$ to the problems of checking $\mathbf{G}_{IG}, w_0 \models E\text{-NASH}_{IG}$ and $\mathbf{G}_{IG}, w_0 \models U\text{-NASH}_{IG}$ respectively, with $w_0 = S_0$, where:

$$\begin{aligned} E\text{-NASH} &= \langle\langle x_1, \dots, x_n \rangle\rangle^{\geq 1} \psi_{Nash} \\ U\text{-NASH} &= E\text{-NASH} \wedge \neg \langle\langle x_1, \dots, x_n \rangle\rangle^{\geq 2} \psi_{Nash} \end{aligned}$$

with:

$$\psi_{Nash} = [1 \mapsto x_1] \dots [n \mapsto x_n] \bigwedge_{1 \leq i \leq n} (\langle\langle y \rangle\rangle [i \mapsto y] \gamma_i \rightarrow \gamma_i).$$

Aminof et al. (2016, Theorems 3.1) show that the model checking problems for the G–SL formulas $E\text{-NASH}$ and $U\text{-NASH}$ can be decided in 2EXPTIME (with respect to the size of the CGS and the size of the agents’ temporal goals $\gamma_1, \dots, \gamma_n$). As in the proof of Theorem 1, we obtain the desired bound by recalling that the size of the CGS generated by an influence game is exponential in the size of the game. \square

Note that existence and uniqueness of Nash equilibria have been proved to be in 2EXPTIME in the context of iterated Boolean games (Gutierrez et al., 2015). We believe that the discrepancy between these results and Theorem 3 is due to the different representations of strategies we use. Our perfect-recall strategies may not have a finite representation, while in the setting of iterated Boolean games strategies are deterministic finite state machines having a finite representation.

6 Conclusions and future work

We have presented here two models of opinion diffusion on a network, where the agents try to achieve their individual goals by deciding to enforce (or not) their opinions over the agents they can influence. The difference in the two models lies in what is common knowledge for the agents with respect to their opinions (at least until they decide to use their influence for the first time). In order to model the strategic dynamics of agents over such networks we introduced *influence games* and found agents to be greatly empowered by the basic action of deciding whether to enforce their opinion, as reflected by our results on the interaction between goals, network structure and solution concepts.

From a computational complexity point of view, we found that for memory-less strategies checking whether a given strategy profile is a Nash equilibrium is in PSPACE, for both the complete and the incomplete information model, as it is the model-checking problem

of LTL in which individual goals are expressed. Moreover, for perfect-recall strategies and complete information the problems of checking existence of a winning strategy is in EXPTIME, while existence and uniqueness of Nash equilibrium are both in 3EXPTIME.

An obvious extension of this work would be to endow agents with additional strategic actions, e.g., lying about their tacit opinions, in line with the work by Christoff and Hansen (2015). Secondly, changing the aggregation procedures used to update agents' opinions would have a great impact on the game-theoretic structures of our setting, as well as on the computational complexity of reasoning about solution concepts.

Acknowledgments

The authors acknowledge the support of the ANR project SCONE on “Social Choice and Social Networks” (ANR 18-CE23-0009-01). Support from the ANR-3IA Artificial and Natural Intelligence Toulouse Institute is also gratefully acknowledged. This work was significantly improved thanks to the comments of the anonymous reviewers of AAMAS-2017 and JLC.

References

- Ågotnes, T., Harrenstein, P., van der Hoek, W., and Wooldridge, M. J. (2013). Boolean Games with Epistemic Goals. In *Proceedings of the 4th International Workshop on Logic, Rationality, and Interaction (LORI-2013)*.
- Alur, R., Henzinger, T. A., and Kupferman, O. (2002). Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713.
- Aminof, B., Malvone, V., Murano, A., and Rubin, S. (2016). Graded Strategy Logic: Reasoning about Uniqueness of Nash Equilibria. In *Proceedings of the 15th International Joint Conference on Agents and Multiagent Systems (AAMAS-2016)*.
- Amor, N. B., Fargier, H., Sabbadin, R., and Trabelsi, M. (2019). Possibilistic Games with Incomplete Information. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-2019)*.
- Arora, S. and Barak, B. (2009). *Computational Complexity: A Modern Approach*. Cambridge University Press.
- Asch, S. E. (1956). Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological Monographs*, 70(9):1–70.

- Attanasi, G., Hopfensitz, A., Lorini, E., and Moisan, F. (2016). Social connectedness improves co-ordination on individually costly, efficient outcomes. *European Economic Review*, 90:86–106.
- Baltag, A., Christoff, Z., Hansen, J. U., and Smets, S. (2013). Logical Models of Informational Cascades. *Studies in Logic*, 47:405–432.
- Bandura, A. (1986). *Social foundations of thought and action: A social cognitive theory*. Englewood Cliffs.
- Barnes, B. (1988). *The Nature of Power*. Polity Press.
- Battigalli, P. and Guaitoli, D. (1997). Conjectural equilibria and rationalizability in a game with incomplete information. In Battigalli, P., Montesano, A., and Panunzi, F., editors, *Decisions, games and markets*, volume 43 of *Theory and Decision Library C*, pages 97–124. Kluwer.
- Belardinelli, F., Grandi, U., Herzig, A., Longin, D., Lorini, E., Novaro, A., and Perrussel, L. (2017). Relaxing Exclusive Control in Boolean Games. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2017)*.
- Bonzon, E., Lagasquie-Schiex, M., Lang, J., and Zanuttini, B. (2006). Boolean Games Revisited. In *Proceedings of the 7th European Conference on Artificial Intelligence (ECAI-2006)*.
- Botan, S., Grandi, U., and Perrussel, L. (2019). Multi-Issue Opinion Diffusion under Constraints. In *Proceedings of the 18th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2019)*.
- Brill, M., Elkind, E., Endriss, U., and Grandi, U. (2016). Pairwise Diffusion of Preference Rankings in Social Networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016)*.
- Castelfranchi, C. (2003). The Micro-Macro Constitution of Power. *Protosociology*, (18):208–265.
- Christian, R., Fellows, M., Rosamond, F., and Slinko, A. (2007). On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224.
- Christoff, Z. and Grossi, D. (2017). Binary Voting with Delegable Proxy: An Analysis of Liquid Democracy. In *Proceedings of the 16th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2017)*.

- Christoff, Z. and Hansen, J. U. (2013). A two-tiered formalization of social influence. In *Proceedings of the 4th International Workshop on Logic, Rationality, and Interaction - 4th International Workshop (LORI)*.
- Christoff, Z. and Hansen, J. U. (2015). A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1):48–77.
- Christoff, Z., Hansen, J. U., and Proietti, C. (2016). Reflecting on Social Influence in Networks. *Journal of Logic, Language, and Information*, 25(3):299–333.
- de Groot, M. H. (1974). Reaching a Consensus. *Journal of the American Statistical Association*, 69(345):118–121.
- Dimant, E. (2019). Contagion of pro- and anti-social behavior among peers and the role of social proximity. *Journal of Economic Psychology*, 73:66–88.
- Ellwood, C. A. (1901). The theory of imitation in social psychology. *American Journal of Sociology*, 6(6):721–741.
- Festinger, L. (1950). Informal social communication. *Psychological Review*, 57(5):271–282.
- Fudenberg, D. and Levine, D. K. (1993). Self-confirming equilibrium. *Econometrica*, 61(3):523–545.
- Ghosh, S. and Velázquez-Quesada, F. R. (2015). Agreeing to Agree: Reaching Unanimity via Preference Dynamics Based on Reliable Agents. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*.
- Grandi, U. and Endriss, U. (2011). Binary Aggregation with Integrity Constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*.
- Grandi, U., Lorini, E., Novaro, A., and Perrussel, L. (2017). Strategic Disclosure of Opinions on a Social Network. In *Proceedings of the 16th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2017)*.
- Grandi, U., Lorini, E., and Perrussel, L. (2015). Propositional Opinion Diffusion. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*.
- Granovetter, M. (1973). The Strength of Weak Ties. *American Journal of Sociology*, 78(6):1360–1380.

- Granovetter, M. (1978). Threshold Models of Collective Behavior. *American Journal of Sociology*, 83(6):1420–1443.
- Grant, J., Kraus, S., Wooldridge, M. J., and Zuckerman, I. (2011). Manipulating Boolean Games through Communication. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*.
- Gutierrez, J., Harrenstein, P., and Wooldridge, M. (2015). Iterated Boolean games. *Information and Computation*, 242:53–79.
- Harrenstein, P., van der Hoek, W., Meyer, J.-J., and Witteveen, C. (2001). Boolean games. In *Proceedings of the 8th Conference on Theoretical aspects of Rationality and Knowledge (TARK-2001)*.
- Herzig, A., Lorini, E., and Maffre, F. (2015). A Poor Man’s Epistemic Logic Based on Propositional Assignment and Higher-Order Observation. In *Proceedings of the 4th International Workshop on Logic, Rationality, and Interaction (LORI-2015)*.
- Herzig, A. and Troquard, N. (2006). Knowing how to play: uniform choices in logics of agency. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2006)*.
- Jamroga, W. and Ågotnes, T. (2016). What agents can achieve under incomplete information. In *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2016)*.
- Kalai, E. and Lehrer, E. (1993). Subjective equilibrium in repeated games. *Econometrica*, 61(5):1231–1240.
- Lang, J. and Slavkovik, M. (2013). Judgment aggregation rules and voting rules. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT-2013)*.
- Lehrer, K. and Wagner, C. (1981). *Rational Consensus in Science and Society*. Springer.
- Lorini, E., Troquard, N., Herzig, A., and Broersen, J. (2013). Grounding power on actions and mental attitudes. *Logic Journal of the IGPL*, 21(3):311–331.
- Mogavero, F., Murano, A., and Vardi, M. Y. (2010). Reasoning About Strategies. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS-2010)*.
- Moscovici, S. and Personnaz, B. (1980). Studies in social influence: V. Minority influence and conversion behavior in a perceptual task. *Journal of Experimental Social Psychology*, 16(3):270–282.

- Rapoport, A. and Chammah, A. M. (1966). The Game of Chicken. *American Behavioral Scientist*, 10(3):10–28.
- Rogers, E. M. (1983). *Diffusion of Innovations*. The Free Press, 3rd edition.
- Ruan, J. and Thielscher, M. (2011). A Logic for Knowledge Flow in Social Networks. In *Proceedings of the 24th Australasian Joint Conference on Artificial Intelligence (AI-2011)*.
- Schelling, T. (1978). *Micromotives and macrobehavior*. Norton.
- Schobbens, P. Y. (2004). Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93.
- Schwind, N., Inoue, K., Bourgne, G., Konieczny, S., and Marquis, P. (2015). Belief Revision Games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-2015)*.
- Schwind, N., Inoue, K., Bourgne, G., Konieczny, S., and Marquis, P. (2016). Is Promoting Beliefs Useful to Make Them Accepted in Networks of Agents? In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016)*.
- Seligman, J., Liu, F., and Girard, P. (2011). Logic in the community. In *Proceedings of the 4th Indian Conference on Logic and Its Applications*.
- Seligman, J., Liu, F., and Girard, P. (2013). Facebook and the Epistemic Logic of Friendship. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2013)*.
- Seligman, J. and Thompson, D. (2015). Boolean Network Games and Iterated Boolean Games. In *Proceedings of the 5th International Workshop on Logic, Rationality and Interaction (LORI-2015)*.
- Shapley, L. (1953). Stochastic games. *Proceedings of the National Academy of Sciences*, 39:1095–1100.
- Sistla, A. P. and Clarke, E. M. (1985). The complexity of propositional linear temporal logics. *Journal of the Association of Computing Machinery*, 32(3):733–749.
- van Benthem, J. (2001). Games in Dynamic-Epistemic Logic. *Bulletin of Economic Research*, 53(4):219–248.
- van Benthem, J., Gerbrandy, J., Hoshi, T., and Pacuit, E. (2009). Merging Frameworks for Interaction. *Journal of Philosophical Logic*, 38:491–526.

van der Hoek, W., Troquard, N., and Wooldridge, M. J. (2011). Knowledge and control. In *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2011)*.

Wooldridge, M., Endriss, U., Kraus, S., and Lang, J. (2013). Incentive engineering for Boolean games. *Artificial Intelligence*, 195:418–439.