



Strategic disclosure of opinions on a social network

Umberto Grandi, Emiliano Lorini, Arianna Novaro, Laurent Perrussel

► To cite this version:

Umberto Grandi, Emiliano Lorini, Arianna Novaro, Laurent Perrussel. Strategic disclosure of opinions on a social network. 16th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), May 2017, Sao Paulo, Brazil. pp.1196-1204. hal-01736247

HAL Id: hal-01736247

<https://hal.science/hal-01736247>

Submitted on 16 Mar 2018

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.



Open Archive TOULOUSE Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible.

This is an author-deposited version published in : <http://oatao.univ-toulouse.fr/>
Eprints ID : 18887

The contribution was presented at AAMAS 2017 :
<http://www.aamas2017.org/>

To cite this version : Grandi, Umberto and Lorini, Emiliano and Arianna, Novaro and Perrussel, Laurent *Strategic disclosure of opinions on a social network*. (2017) In: 16th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), 8 May 2017 - 12 May 2017 (Sao Paulo, Brazil).

Any correspondence concerning this service should be sent to the repository administrator: staff-oatao@listes-diff.inp-toulouse.fr

Strategic Disclosure of Opinions on a Social Network

Umberto Grandi
IRIT, University of Toulouse
umberto.grandi@irit.fr

Emiliano Lorini
IRIT-CNRS, University of
Toulouse
emiliano.lorini@irit.fr

Arianna Novaro
IRIT, University of Toulouse
arianna.novaro@irit.fr

Laurent Perrussel
IRIT, University of Toulouse
laurent.perrussel@irit.fr

ABSTRACT

This paper starts from a simple model of strategic reasoning in situations of social influence. Agents express binary views on a set of propositions, and iteratively update their views by taking into account the expressed opinion of their influencers. We empower agents with the ability to disclose or hide their opinions, in order to attain a predetermined goal. We study classical game-theoretic solution concepts in the resulting games, observing a non-trivial interplay between the individual goals and the structure of the underlying network. By making use of different logics for strategic reasoning, we show how apparently simple problems in strategic opinion diffusion require a complex logical machinery to be properly formalized and handled.

Keywords

Opinion diffusion; Nash equilibria; strategic reasoning; social influence

1. INTRODUCTION

Social influence can be seen as a process where an agent forms her opinion on the basis of the opinions expressed by agents she trusts [3, 10, 18]. Recent work in multi-agent systems proposed formal models of opinion diffusion which combined methods and techniques from social network analysis with those of belief merging and judgment aggregation [22, 13, 5, 23]. These frameworks built on classical models of opinion diffusion on networks such as the De Groot or Lehrer-Wagner model [9, 16] and threshold models [14], adapting them to more complex representations of individual opinions: belief bases, preferences, and binary judgments over interconnected issues. The focus has been on studying how agents' opinions evolve over time due to the influence of other agents in the society. Namely, an agent's opinion at a given time instant results from aggregating the opinions (at the previous time instant) of the agents she trusts and the trust relationship is modeled via a network.

In the present paper, we build on this work by looking at social influence from a strategic perspective. We introduce a new class of infinite repeated games, the *games of influ-*

ence. At each stage an agent can choose whether to make her opinions public, and she updates her opinion according to the *public* opinions of the agents she trusts and to some aggregation procedure. Moreover, to each agent is associated a goal she wants to achieve — in line with previous work on representation of agent's motivations [15], expressed in a variant of Linear Temporal Logic LTL to consider agents' future opinions.

Influence games provide a simple abstraction to explore the effects of a trust network on agents' behavior, as well as allowing us to study game-theoretic solution concepts such as winning strategy, weak dominance and Nash equilibrium. First, we analyze the interplay between the network structure and a player's strategic ability, focussing on two goal schemas: that of reaching consensus among a group of agents and that of influencing individuals towards one's own opinion. Since most of these results are of a limitative nature, pointing at the difficulties of introducing a network into an apparently simple problem, we move to explore the existence of game-theoretic solutions as a computational problem. We study its complexity by devising translations into existing logical formalisms for strategic reasoning, such as Alternating-time Temporal Logic [1] and Strategy Logic [17].

Previous work on logic-based models for social influence include the Facebook logic [24] and the preference and reliability models (which do not, however, consider strategic aspects in the update) [11]. Other known models differentiating private and public information have not focussed on strategic aspects [7, 8]. Problems close to opinion diffusion are those of information cascades and knowledge diffusion, which have been given formal treatment in a logical setting [20, 4]. Moreover, influence games can be considered as a variation of iterated boolean games [15] in which individuals do not have direct control over all the variables, but they concurrently participate in changing their truth values. Iterated boolean games have recently been extended with a social network structure where agents choose actions depending on the actions of their neighbors [25].

The remainder of the paper is organized as follows. Section 2 presents the basic definitions of private and public opinions, as well as our model of opinion diffusion. Moreover, we introduce state transitions given by agents' actions and the variant of LTL with which we express agents' goals. Section 3 introduces influence games, and presents the main results about the effects of the network structure on solution concepts from game theory. Section 4 studies complexity problems related to the solution concepts for influence games, and Section 5 concludes the paper.

2. OPINION DIFFUSION

In this section we present the model of opinion diffusion which is the starting point of our analysis. We generalize the model of propositional opinion diffusion introduced in related work [13] by separating private and public opinions with a notion of visibility, and we adapt the diffusion process through aggregation to this more complex setting.

2.1 Private and public opinions

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

DEFINITION 1 (PRIVATE OPINION). *The private 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 the agent's opinion that p is false.*

Let $\mathbf{B} = (B_1, \dots, B_n)$ denote the profile of private opinions of agents in $\mathcal{N}$. Propositional evaluations can be used to represent ballots in a multiple referendum, expressions of preference over alternatives, or judgements over correlated issues [6, 12]. Depending on the application at hand, an integrity constraint can be introduced to model the propositional correlation among issues. For the sake of simplicity we do not introduce any such constraint in this paper.

We also assume that each agent has the possibility of declaring her private opinion on each issue.

DEFINITION 2 (VISIBILITY FUNCTION). *We call visibility function of agent i any map $V_i : \mathcal{I} \rightarrow \{1, 0\}$ where $V_i(p) = 1$ expresses that agent i 's opinion on p is visible.*

We denote by $\mathbf{V} = (V_1, \dots, V_n)$ the profile composed of the agents' visibility functions. By combining the private opinion with the visibility function of an agent, we can build her public opinion as a three-valued function on the issues.

DEFINITION 3 (PUBLIC OPINION). *Let B_i be the opinion of agent i and V_i her visibility function. The public opinion of i is a function $P_i : \mathcal{I} \rightarrow \{1, 0, ?\}$ such that*

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

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

DEFINITION 4 (STATE). *A state is a tuple $S = (\mathbf{B}, \mathbf{V})$ where $\mathbf{B}$ is a profile of private opinions and $\mathbf{V}$ is a profile of visibility functions. The set of all states is denoted by $\mathcal{S}$.*

States are the building blocks of the model of strategic reasoning in opinion dynamics that we present in Section 2.3.

2.2 Unanimous opinion diffusion

We here define the influence process that is at the heart of our model. Our definition generalizes the model of propositional opinion diffusion [13] to take into account the visibility function. First, we assume that individuals are linked by an *influence network* modeled as a directed graph:

DEFINITION 5 (INFLUENCE NETWORK). *We call an influence network a directed irreflexive graph $E \subseteq \mathcal{N} \times \mathcal{N}$, where $(i, j) \in E$ reads "agent j is influenced by agent i ".*

We also refer to E as the influence graph and to 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 . Given a state S , this definition can be refined by considering $\text{Inf}^S(i, p) = \{k \in \mathcal{N} \mid (k, i) \in E \text{ and } P_k(p) \neq ?\}$ to be the subset of i 's influencers that are actually showing their private opinion about issue p .

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

DEFINITION 6 (AGGREGATION PROCEDURE). *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 individual opinion and the public opinions of a set of agents C to agent i 's individual opinion.

We drop C from the subscript when clear from the context, to simplify notation. Many aggregation procedures have been considered in the literature on judgment aggregation, 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 (cfr. the class of threshold models studied in the literature on opinion diffusion [14, 21]). Unanimity is another instance of a quota rule, which we adapt here:

DEFINITION 7. *The unanimous issue-by-issue aggregation procedure is defined as follows:*

$$F_i^U(B_i, \mathbf{P}_C)(p) = \begin{cases} B_i(p) & \text{if } C = \emptyset \\ x \in \{0, 1\} & \text{if } P_k(p) = x \text{ for all } k \in C \\ B_i(p) & \text{otherwise} \end{cases}$$

That is, an individual will change her private opinion about issue p if and only if all agents in C (usually among her influencers) publicly expressing their opinion are unanimous in disagreeing with her own. For the sake of simplicity, in the remainder of the paper we will consider that all agents use the unanimous aggregation procedure.

2.3 Strategic actions and state transitions

In our model, agents can make their opinions visible or invisible by specific actions of type **reveal**(J) — i.e., action of showing the opinion on issues in J , and **hide**(J) — i.e., action of hiding the opinion on issues in J . We allow for simultaneous disclosure on multiple propositions. Let thus:

$$\mathcal{A} = \{(\text{reveal}(J), \text{hide}(J')) \mid J, J' \subseteq \mathcal{I} \text{ and } J \cap J' = \emptyset\}$$

be the set of individual actions. Each joint action $\mathbf{a} = (a_1, \dots, a_n) \in \mathcal{A}^n$ induces a deterministic transition function between states:

DEFINITION 8 (TRANSITION FUNCTION). *The transition function $\text{succ} : \mathcal{S} \times \mathcal{A}^n \longrightarrow \mathcal{S}$ associates to each state S and joint action $\mathbf{a}$ a new state $S' = (\mathbf{B}', \mathbf{V}')$ as follows, for all $i \in \mathcal{N}$ and $p \in \mathcal{I}$. For $a_i = (\text{reveal}(J), \text{hide}(J')) \in \mathcal{A}$:*

- $V'_i(p) = \begin{cases} 1 & \text{if } p \in J \\ 0 & \text{if } p \in J' \\ V_i(p) & \text{otherwise} \end{cases}$
- $B'_i = F_i^U(B_i, \mathbf{P}'_{\text{Inf}(i)S})$

Where $\mathbf{P}'$ is the public profile obtained from private profile $\mathbf{B}$ and visibility profile $\mathbf{V}'$.

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. We also use the following abbreviations: $\text{skip} = (\text{reveal}(\emptyset), \text{hide}(\emptyset))$ for doing nothing, $\text{reveal}(J) = (\text{reveal}(J), \text{hide}(\emptyset))$, $\text{hide}(J) = (\text{reveal}(\emptyset), \text{hide}(J))$, and we drop curly parentheses in $\text{reveal}(\{p\})$ and $\text{hide}(\{p\})$. Our definition assumes that the influence process occurs after the actions have changed the visibility of the agents' opinions. Specifically, first, actions affect the visibility of opinions, and then each agent modifies her private opinion on the basis of those opinions of her influencers that have become public.

We are now ready to define the concept of *history*, describing the temporal aspect of agents' opinion dynamic:

DEFINITION 9 (HISTORY). 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 , an history is an infinite sequence of states $H = (H_0, H_1, \dots)$ such that for all $t \in \mathbb{N}$ there exists a joint action $\mathbf{a}_t \in \mathcal{A}^n$ such that $H_{t+1} = \mathbf{a}_t(H_t)$.

Let $H = (H_0, H_1, \dots)$ be an history: the set of all histories is denoted by $\mathcal{H}$. Observe that our definition restricts the set of all possible histories to those that correspond to a run of the influence dynamic described above. 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 private opinion in state H_t and with $H_{i,t}^V$ agent i 's visibility function in state H_t .

EXAMPLE 1. Consider the example in Figure 1, where the initial state is H_0 and the agents are $\mathcal{N} = \{i, j, k\}$ such that $\text{Inf}(i) = \{j, k\}$. Let $\mathbf{a}_0 = (\text{skip}, \text{skip}, \text{reveal}(p))$ and $\mathbf{a}_1 = (\text{skip}, \text{hide}(p), \text{skip})$ be the joint actions of the agents at the first two states. Namely, agent k reveals her opinion on p , and at the next step j hides hers. If all individuals are using the unanimous aggregation procedure, then states H_1 and H_2 result from applying the joint actions from state H_0 . In state H_1 , agent i 's private opinion about p has changed to 1, as all her influencers are publicly unanimous on p , while in H_2 no opinion is updated.

$$\begin{array}{ccccc} ((0, 1, 1), (1, 1, 0)) & \xrightarrow{\mathbf{a}_0} & ((1, 1, 1), (1, 1, 1)) & \xrightarrow{\mathbf{a}_1} & ((1, 1, 1), (1, 0, 1)) \\ H_0 & & H_1 & & H_2 \end{array}$$

Figure 1: The first states of a history.

2.4 Individual goals

By revealing or hiding her opinion, an agent influences others towards the satisfaction of her goal. To account for the temporal aspect of our model, we follow recent work on iterated boolean games [15] and we define a language $\mathcal{L}_{\text{LTL-I}}$ to express individual goals using Linear Temporal Logic LTL.

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

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

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 $\neg\text{op}(i, p)$ has to be read “agent i 's opinion is that p is not true” (since agents have binary opinions). Moreover, $\text{vis}(i, p)$ has to be read “agent i 's opinion about p is visible”.

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’ ($\Box$) as $\Diamond\varphi = \top\mathcal{U}\varphi$ and $\Box\varphi = \neg\Diamond\neg\varphi$.

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

DEFINITION 10 (TRUTH CONDITIONS). Let H be a history, let φ be a formula of $\mathcal{L}_{\text{LTL-I}}$ and let $k, k', k'' \in \mathbb{N}$. Then:

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

Formulas of $\mathcal{L}_{\text{LTL-I}}$ 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 $\mathcal{L}_{\text{LTL-I}}$, which we denote as *goal formulas*.

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

$$\begin{aligned} \text{cons}(C, J) &:= \Diamond\Box(\text{pcons}(C, J) \vee \text{ncons}(C, J)) \\ \text{influ}(i, C, J) &:= \Diamond\Box \bigwedge_{p \in J} ((\text{op}(i, p) \rightarrow \bigcirc\text{pcons}(C, p)) \\ &\quad \wedge (\neg\text{op}(i, p) \rightarrow \bigcirc\text{ncons}(C, p))) \end{aligned}$$

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 first goal wants at some point in the history to reach a stable consensus either for or against the issues in J with the agents in C . The second goal expresses instead the idea that agent i wants to eventually gain a stable influence over the people in C about the issues in J (i.e., they will always hold her opinion at the next step).

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 11 (INFLUENCE GAME). An influence game is a tuple $IG = (\mathcal{N}, \mathcal{I}, E, F_i, S_0, \gamma_1, \dots, \gamma_n)$ 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, our agents will build their strategies in order to attain their goals. We first introduce two kinds of strategies available to agents, namely memory-less and perfect-recall strategies, then we study the existence of basic solution concepts from game theory when agents use influence or consensus goals (see Section 2.4).

3.1 Strategies

The first type of simple individual strategies we define is based on states, and thus memory-less.

DEFINITION 12 (MEMORY-LESS STRATEGY). A memory-less strategy for player i is a function $Q_i : \mathcal{S} \rightarrow \mathcal{A}$ that associates an action to every state.

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 $\mathbf{Q}(S) = \mathbf{a}$ if and only if $Q_i(S) = a_i$, for all $S \in \mathcal{S}$ and $i \in \mathcal{N}$.

The second definition we provide is the full-blown notion of perfect-recall strategy, which assigns an action to each partial history that has been observed by the player. Let $\mathcal{H}^+$ denote the set of all partial histories — i.e., finite sequences of states satisfying Definition 9.

DEFINITION 13 (PERFECT-RECALL STRATEGY). We call a perfect-recall strategy for player i a function $Q_i : \mathcal{H}^+ \rightarrow \mathcal{A}$ that associates an action to every finite history.

As the following definition highlights, every strategy profile (of either type) induces a history if combined with an initial state:

DEFINITION 14 (INDUCED HISTORY). Let S_0 be an initial state and let $\mathbf{Q}$ be a strategy profile. The induced history $H_{S_0, \mathbf{Q}} \in \mathcal{H}$ is defined as follows:

$$\begin{aligned} H_0(S_0, \mathbf{Q}) &= S_0 \\ H_{n+1}(S_0, \mathbf{Q}) &= \text{succ}(S_n, \mathbf{Q}(S_n)) \text{ for all } n \in \mathbb{N} \end{aligned}$$

The two definitions allow us to differentiate the study among more or less sophisticated agents. It is important to observe, however, that we study influence games as repeated games of *complete information*. This is to be considered as a simplifying assumption — a more realistic model would require the use of *uniform strategies* and the corresponding indistinguishability relation. The hide action can alternatively be interpreted in our model as “not exercising one’s own influence”, without any epistemic interpretation, and the reveal action as “persuade”.

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.

DEFINITION 15 (WINNING STRATEGY). Let IG be an influence game and let Q_i be a strategy for player i . We say that Q_i is a winning strategy for player i in state S_0 if

$$H_{S_0, (Q_i, \mathbf{Q}_{-i})} \models \gamma_i$$

for all profiles $\mathbf{Q}_{-i}$ of strategies of players other than i .

EXAMPLE 2. Let Ann, Bob and Jesse be three agents and suppose that $B_{\text{Ann}}(p) = 1$, $B_{\text{Bob}}(p) = 0$, $B_{\text{Jesse}}(p) = 0$ for $p \in \mathcal{I}$. Their influence network is as in this picture:



Suppose Ann’s goal is $\Diamond \Box \text{op}(\text{Jesse}, p)$. Her winning memory-less strategy is to play $\text{reveal}(p)$ in all states. Bob will be influenced to believe p at the second stage in the history and similarly for Jesse at the third stage, since her influencers are unanimous even if Bob plays $\text{hide}(p)$.

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

DEFINITION 16 (WEAKLY DOMINANT STRATEGY). Let IG be an influence game and Q_i a strategy for player i . We say that Q_i is a 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 and for all strategies Q'_i we have:

$$H_{S_0, (Q'_i, \mathbf{Q}_{-i})} \models \gamma_i \Rightarrow H_{S_0, (Q_i, \mathbf{Q}_{-i})} \models \gamma_i$$

EXAMPLE 3. Suppose that in the previous example Ann still believes p , but does not influence Jesse any longer. Here, Ann does not have a winning strategy: if neither Bob nor Jesse believe p , it is sufficient for Bob to play $\text{hide}(p)$ to make sure that she will never satisfy her goal. However, playing action $\text{reveal}(p)$ is weakly dominant for Ann.

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

DEFINITION 17 (NASH EQUILIBRIUM). Let IG be an influence game and let $\mathbf{Q}$ be a strategy profile. We say that $\mathbf{Q}$ is a Nash equilibrium for initial state S_0 if and only if for all $i \in \mathcal{N}$ and for all $Q'_i \in \mathcal{Q}_i$:

$$H_{S_0, (Q_i, \mathbf{Q}_{-i})} \models \gamma_i \text{ or } H_{S_0, (Q'_i, \mathbf{Q}_{-i})} \not\models \gamma_i.$$

Problems of membership, existence and uniqueness of Nash equilibria will be studied in Section 4.

3.3 Influence network and solution concepts

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.

PROPOSITION 1. If E is a directed acyclic graph (DAG) such that $|\text{Inf}(i)| \leq 1$ 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 $C_a := \{k \in \mathcal{N} \mid a \in \text{Inf}(k)\} \cup \{a\}$, then agent a has a winning strategy.

PROOF. (sketch) Consider a DAG E and an agent a with goal γ_a . Let Q_a be the strategy associating to every state S action $\text{reveal}(J)$. 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 $a \in \text{Inf}(b)$, the goal reduces to $\text{cons}(\{a\}, J)$ which is always trivially satisfied. In case $\text{Inf}(a) = \emptyset$, by playing $\text{reveal}(J)$ in S_0 and since every agent uses the unanimous aggregation rule, 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 $a \in \text{Inf}(b)$ and some agent $c \in \text{Inf}(a)$. By assumption on E we thus have that $\text{Inf}(a) = \{c\}$ and $\text{Inf}(b) = \{a\}$. Hence, either at some point k in the history all ancestors of a will have reached consensus, such that by playing $\text{reveal}(J)$ 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 $\text{reveal}(J)$. Therefore, 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 $\text{reveal}(J)$ and they start in S_0 by having alternating positive and negative opinions on the issues in J . In the second place, having at most one influencer per agent ensures each agent to have full control over their child nodes.

Observe also that Proposition 1 implies the same result to hold for $\gamma_a := \square \diamond (\text{pcons}(C_a, J) \vee \text{ncons}(C_a, J))$. In fact, since eventually a will reach a stable consensus with her child nodes, this implies that it is always true that we can find some later point in the history where consensus holds. In general, however, Proposition 1 shows us how winning strategies are too strong of a solution concept, so that the type of goals which can be attained have a narrow scope.

If we move to the less demanding concept of weak dominance, we may intuitively think that a strategy associating action $\text{reveal}(J)$ to all states is weakly dominant for an agent a having goal $\gamma_a := \text{influ}(a, C, J)$ 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 . Yet, we show in the following example that to satisfy goals of type γ_a as described, an agent could sometimes benefit from hiding her opinion.

EXAMPLE 4. 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 agent 1 and 2 associate action $\text{reveal}(p)$ to all states, and agent 3 associates action $\text{hide}(p)$ only to those states where 1, 2 and 3 agree on p . Suppose that the goal of agent 2 is $\gamma_2 = \text{influ}(2, \{4\}, \{p\})$. Consider the history below for these strategies, where goal γ_2 is not attained (we only represent $\mathbf{B}$):

$$\begin{array}{ccccc} (0, 1, 0, 1) & \xrightarrow{\mathbf{a}_0} & (0, 0, 1, 0) & \xrightarrow{\mathbf{a}_1} & (0, 0, 0, 1) \\ H_0 & & H_1 & & H_2 \end{array}$$

From state H_2 onwards, given the strategies of the agents, the profile of opinions $\mathbf{B} = (0, 0, 0, 1)$ won't change. On the other hand, consider a strategy for agent 2 identical to the previous one, but for the fact that it associates to state H_0 action $\text{hide}(p)$. This is what would happen:

$$\begin{array}{ccccc} (0, 1, 0, 1) & \xrightarrow{\mathbf{a}_0} & (0, 0, 0, 0) \\ H_0 & & H_1 \end{array}$$

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 sometimes better off by hiding her opinion on p to satisfy her influence goal γ_2 .

We can now see an easy example of how the network structure and the agents' goals can yield a Nash equilibrium in the spirit of anti-coordination games [19].

PROPOSITION 2. Let E be a cycle for $\mathcal{N} = \{1, 2\}$. If $\gamma_1 = \gamma_2 = \text{cons}(\mathcal{N}, J)$, where $J \subseteq \mathcal{I}$, then there exists a Nash equilibrium for any initial state S_0 .

PROOF. (sketch) To attain their goal the agents must coordinate on the issues in J on which they disagree in S_0 . In fact, in case at some stage k of the history they both play $\text{hide}(p)$ for $p \in J$ their private opinion would stay the same at stage $k + 1$. If they both play $\text{reveal}(p)$, at the next stage they would just swap their opinions on p (since they are each other's only influencers and they both use the unanimous rule). Hence, agent 1 has to play $\text{reveal}(p)$ whenever the other agent is playing $\text{hide}(p)$ so that at the next stage in the history he will have copied her opinion, while she would have not changed hers — and similarly if the other agent is playing $\text{reveal}(p)$.

Consider thus an arbitrary strategy Q_1 for agent 1 and some initial state S_0 . Construct now strategy Q_2 for agent 2 associating action $\text{reveal}(J^S)$ to all states where strategy Q_1 associates action $\text{hide}(J^S)$ for $J^S = \{p \in J \mid b_{1p} = 1 - b_{2p} \text{ in } S\}$, and viceversa for action $\text{reveal}(J^S)$. By the above reasoning, the strategy profile $\mathbf{Q} = (Q_1, Q_2)$ generates a history that satisfies both γ_1 and γ_2 , and therefore is a Nash equilibrium. The same construction can be done for an arbitrary strategy of agent 2 as well. $\square$

4. COMPUTATIONAL COMPLEXITY AND SOLUTION CONCEPTS

In this section we provide complexity results for a variety of computational problems for the solution concepts introduced in Section 3.2. We start by studying the complexity of deciding whether a certain memory-less strategy profile of an influence game is a Nash equilibrium.

Then, we present two logics of strategic reasoning that have been studied in the area of logics for multi-agent systems: Alternating-time Temporal Logic ATL [1] and the graded variant of Strategy Logic SL [17, 2]. We then focus on perfect-recall strategies for influence games and we reduce the problem of checking existence of a winning strategy to the model checking problem of the former, and the problem of checking existence or uniqueness of a Nash equilibrium to the model checking problem of the latter.

4.1 Problem definitions

We now provide the formal definitions of the complexity problems we are interested in studying for influence games.

INPUT: $IG = (\mathcal{N}, \mathcal{I}, E, F, S_0, \gamma_1, \dots, \gamma_n)$, $\mathbf{Q}$.

M-NASH(F): Is $\mathbf{Q}$ a Nash equilibrium of IG ?

This problem is out of reach for perfect-recall strategies, since the size of the input (i.e., the size of a perfect-recall strategy profile) would already be exponential in the size of the initial game. We will thus focus on the problems E-NASH(F) and U-NASH(F), defined below.

INPUT: $IG = (\mathcal{N}, \mathcal{I}, E, F, S_0, \gamma_1, \dots, \gamma_n)$.

E-NASH(F): Is there some Nash equilibrium $\mathbf{Q}$ of IG ?

U-NASH(F): Is there a unique 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.

4.2 Memory-less Nash equilibrium

We begin by translating a memory-less strategy in the language of the logic LTL-I (see Section 2.4). The algorithm presented in the setting of iterated boolean games [15] cannot be directly applied to ours for two reasons. First, our histories are generated by an aggregation function F that models opinion diffusion — i.e., agents have a concurrent control over a set of propositional variables. Second, in this section we focus on memory-less strategies only.

A conjunction of literals $\alpha(S)$ can be defined to uniquely identify a state S : $\alpha(S)$ will specify the private opinion of all individuals and their visibility function. For an action $a = (\text{reveal}(J), \text{hide}(J'))$, let $\beta_i(a)$ be the following formula:

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

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 aggregation function into a formula as well. Consider the following formula $\text{unan}(i, p)$:

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

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

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

This formula ensures that if the influencers of agent i are unanimous, then agent i 's opinion should be defined according to the three cases described in Definition 7. Recall that, while actions take one time unit to be performed — hence the $\bigcirc$ operator in front of $\text{vis}(j, p)$, the diffusion of opinions is simultaneous. Let now:

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

be the formula encoding the transition process defined by the opinion diffusion. $\tau(F_i^U)$ 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 1. *Let F_i^U be the unanimous issue-by-issue aggregator. M-NASH(F_i^U), E-NASH(F_i^U) and U-NASH(F_i^U) for memory-less strategies are in PSPACE.*

PROOF. We begin with M-NASH(F_i^U). 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 Nash equilibrium. For all individuals $i \in \mathcal{N}$, we first check the following entailment:

$$\tau(\mathbf{Q}) \wedge \tau(F_i) \models_{\text{LTL}} \gamma_i$$

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

If this is not the case, 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:

$$\tau(\mathbf{Q}_{-i}, Q'_i) \wedge \tau(F_i) \models_{\text{LTL}} \gamma_i$$

If the answer is positive we output NO, otherwise we proceed until all strategies and all individuals have been considered. 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 [26], we obtain the desired upper bound.

Similar algorithms can then be devised to solve E-NASH and U-NASH. To see this, it is sufficient to consider one strategy profile $\mathbf{Q}$ at the time, and use M-NASH with $\mathbf{Q}$ in the input as a PSPACE-oracle. For E-NASH the algorithm can stop at the first Nash equilibria that is found, while for U-NASH the entire set of strategy profiles eventually needs to be visited — again, there are exponentially many profiles, but each one can be specified in space polynomial in the size of the input. $\square$

As for completeness, a reduction from similar results in iterated boolean games [15] is likely to be obtained. We leave a more careful study of the relations between these two settings for future work. Observe also that Theorem 1 can be easily generalised to all aggregation procedures that can be axiomatized by means of polynomially many LTL-I formulas, 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. The study of the axiomatization of aggregation procedures for opinion diffusion also constitutes a promising direction for future work.

4.3 Alternating-time temporal logic

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

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

where C ranges over $2^{\mathcal{N}}$ and q ranges over a set of atomic propositions Atm . 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 decide to 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 decide to do”.

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

DEFINITION 18 (CONCURRENT GAME STRUCTURE). 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 $Val : \mathcal{W} \rightarrow 2^{Atm}$ 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 .¹ A strategy for coalition C is a function G_C that maps every agent $i \in C$ 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 the following:

$$Succ(w) = \{T(w, m) \mid m \in \mathcal{P}(w)\}.$$

An infinite sequence $\lambda = w_0 w_1 w_2 \dots$ of worlds from $\mathcal{W}$ is called a *computation* if $w_{k+1} \in Succ(w_k)$ 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 the set of 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 definition of truth conditions for Boolean formulas.

$$\mathbf{G}, w \models \langle\langle C \rangle\rangle \bigcirc \varphi \Leftrightarrow \text{there exists } G_C \in Str_C \text{ such that} \\ \mathbf{G}, \lambda[1] \models \varphi \text{ for all } \lambda \in \mathcal{O}(w, G_C)$$

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

¹Observe that ATL does not distinguish a semantics based on perfect-recall strategies from a semantics based on memory-less strategies in which a strategy is 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.

4.4 Graded strategy logic

Strategy logic SL is a logic of strategic reasoning that embeds ATL [17], but it 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 [2]. We here focus on this SL extension, denoted by $\mathbf{G} - \mathbf{SL}$, whose language is defined by the following BNF:

$$\varphi ::= q \mid \neg \varphi \mid \varphi \wedge \varphi \mid \bigcirc \varphi \mid \varphi \mathcal{U} \varphi \mid \langle\langle x_1, \dots, x_\ell \rangle\rangle^{\geq k} \varphi \mid [i \mapsto x] \varphi$$

where q ranges over Atm , i ranges over $\mathcal{N}$, k and ℓ range over the set of positive integers, $x_1, \dots, x_\ell$ range over a countable infinite set of variables $\mathcal{Var}$ with the additional constraint that $x_h \neq x_{h'}$ for all $1 \leq h, h' \leq \ell$ such that $h \neq h'$. $\langle\langle x_1, \dots, x_\ell \rangle\rangle^{\geq k} \varphi$ is read “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 $\mathbf{G} - \mathbf{SL}$.

Like ATL, SL semantics is based on the concept of concurrent game structure $\mathbf{G} = (\mathcal{W}, \mathcal{M}, R, T, Val)$. However, to provide an interpretation of SL some extra-components are needed. We denote with X a function mapping agents and variables in $\mathcal{N} \cup \mathcal{Var}$ to strategies in Str . We call X an *assignment* function. For every assignment function X , strategy $G \in Str$ and $e \in \mathcal{N} \cup \mathcal{Var}$, 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 that differs 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 $i \in \mathcal{N}$, $G_X(i) = X(i)(i)$.

Moreover, for every world $w \in \mathcal{W}$, 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 $w_0 w_1 \dots$ such that $\mathcal{O}(w, G_X) = \{(w_0 w_1 \dots)\}$. Finally, 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 $e \in \mathcal{N} \cup \mathcal{Var}$, $i \in \mathcal{N}$ and $\pi, \pi' \in \mathcal{W}^+$, we thus have:

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

Any $\mathbf{G} - \mathbf{SL}$ 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:

$$\mathbf{G}, w, X \models [i \mapsto x] \varphi \Leftrightarrow \mathbf{G}, w, X[i \mapsto X(x)] \models \varphi$$

$$\mathbf{G}, w, X \models \langle\langle \vec{x} \rangle\rangle^{\geq k} \varphi \Leftrightarrow \text{there exists } k \text{ many } \ell\text{-tuples} \\ \vec{G}_1, \dots, \vec{G}_k \text{ of strategies} \\ \text{such that for all } 1 \leq h, h' \leq k \\ \text{if } \vec{G}_h \neq \vec{G}_{h'} \text{ then } h \neq h' \\ \text{and } \mathbf{G}, w, X[\vec{x} \mapsto \vec{G}_h] \models \varphi$$

$$\begin{aligned}
\mathbf{G}, w, X \models \bigcirc \varphi &\Leftrightarrow \mathbf{G}, \lambda_{w,X}[1], X_{\text{Prf}(\lambda_{w,X},1)} \models \varphi \\
\mathbf{G}, w, X \models \varphi_1 \mathcal{U} \varphi_2 &\Leftrightarrow \text{there exists } k \geq 0 \text{ such that} \\
&\quad \mathbf{G}, \lambda_{w,X}[k], X_{\text{Prf}(\lambda_{w,X},k)} \models \varphi_2 \\
&\quad \text{and for all } 0 \leq h < k \text{ we have} \\
&\quad \mathbf{G}, \lambda_{w,X}[h], X_{\text{Prf}(\lambda_{w,X},h)} \models \varphi_1
\end{aligned}$$

4.5 Perfect-recall Nash equilibrium

We now introduce the problem of checking the existence of a winning strategy for an agent $i \in \mathcal{N}$ in an influence game IG as follows:

INPUT: $IG = (\mathcal{N}, \mathcal{I}, E, F, S_0, \gamma_1, \dots, \gamma_n)$, $i \in \mathcal{N}$.

E-WINNING $_i(F)$: Is there a winning strategy Q_i in IG ?

The following theorem studies the complexity of the problem E-WINNING $_i(F)$ for perfect-recall strategies and for the unanimous aggregation procedure.

THEOREM 2. E-WINNING $_i(F_i^U)$ for perfect-recall strategies is in EXPTIME.

PROOF. (sketch) We reduce E-WINNING $_i(F_i^U)$ to the model checking problem of ATL where $\mathcal{Atm} = \{\text{op}(i, p) \mid i \in \mathcal{N} \text{ and } p \in \mathcal{I}\} \cup \{\text{vis}(i, p) \mid i \in \mathcal{N} \text{ and } p \in \mathcal{I}\}$. First of all, observe that we can generate the CGS $\mathbf{G}_{IG} = (\mathcal{W}, \mathcal{M}, R, T, \text{Val})$ 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 $\text{op}(i, p), \text{vis}(i, p) \in \mathcal{Atm}$ and $w \in \mathcal{W}$, $\text{op}(i, p) \in \text{Val}(w)$ iff $B_i(p) = 1$, $\text{vis}(i, p) \in \text{Val}(w)$ iff $V_i(p) = 1$.

Secondly, it is easy to verify that the answer to the problem E-WINNING $_i(F_i^U)$ is YES if and only if $\mathbf{G}_{IG}, w_0 \models \langle\langle \{i\} \rangle\rangle \gamma_i$ with $w_0 = S_0$. In [1, Theorem 5.2] it is shown 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 $\mathcal{S}$ 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 above upper bound. $\square$

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

THEOREM 3. E-NASH(F_i^U) and U-NASH(F_i^U) for perfect-recall strategies are both in 3EXPTIME.

PROOF. (sketch) As in the proof of Theorem 2, we generate the CGS $\mathbf{G}_{IG}$ corresponding to the influence game IG . Following [2], we reduce E-NASH(F) and U-NASH(F) 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:

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

In [2, Theorems 3.1] it is shown that the model checking problems for the $\mathbf{G} - \mathbf{SL}$ 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 2, we obtain the desired bound by recalling that the size of the GCS 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 [15]. 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 which have a finite representation.

5. CONCLUSIONS AND FUTURE WORK

We have extended a simple model of opinion diffusion on networks [13] with a strategic component (i.e., the possibility for the agents to hide or show their opinions) and individual goals, which allowed us to define influence games. We found that agents were greatly empowered by these basic actions, as reflected by our results on the interaction between goals, network structure and solution concepts (see Section 3). Interestingly, we found that for influence goals and the unanimous rule, playing a strategy always revealing the agent's opinion was not necessarily weakly dominant.

We then approached the study of influence games and solution concepts from a computational point of view. Specifically, we found that for memory-less strategies the problems of checking whether a given strategy profile is a Nash equilibrium, checking whether a Nash equilibrium exists and is unique are all in PSPACE. On the other hand, for perfect-recall strategies the problems of checking existence of a winning strategy is in EXPTIME, while existence and uniqueness of Nash equilibrium are both in 3EXPTIME.

There are multiple directions to expand the present work. In the first place, following the work by Christoff and Hansen [7], we might provide our agents with additional actions to allow them to lie about their private opinions. This will give them more sophisticated strategies to attain their goals. Secondly, we plan to explore the use of different aggregation procedures to perform the update of opinions, studying the novel game-theoretic structures that may arise in this setting. Last but not least, we are interested in a deeper analysis of the relationship between influence games and iterated boolean games [15]. We conjecture that the latter can be embedded into the version of the former with perfect-recall strategies by only considering atomic formulas of type $\text{vis}(i, p)$, whose truth values are under the direct control of the agents and do not depend on the aggregation procedure.

Acknowledgments

This work was partially supported by the project ‘‘Social Choice on Networks’’ of Labex CIMI ANR-11-LABX-0040-CIMI within the program ANR-11-IDEX-0002-02, and by the CNRS PEPS project ‘‘ODIVO’’.

REFERENCES

- [1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [2] B. Aminof, V. Malvone, A. Murano, and S. Rubin. Graded strategy logic: Reasoning about uniqueness of nash equilibria. In *Proceedings of the 15th International Joint Conference on Agents and Multiagent Systems (AAMAS-2016)*, pages 698–706. ACM, 2016.
- [3] S. E. Asch. Studies of independence and conformity: a minority of one against a unanimous majority. *Psychological Monographs*, 70(9), 1956.
- [4] A. Baltag, Z. Christoff, J. U. Hansen, and S. Smets. Logical models of informational cascades. *Studies in Logic*, 47:405–432, 2013.
- [5] M. Brill, E. Elkind, U. Endriss, and U. Grandi. Pairwise diffusion of preference rankings in social networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-2016)*, 2016.
- [6] R. Christian, M. Fellows, F. Rosamond, and A. Slinko. On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224, 2007.
- [7] Z. Christoff and J. U. Hansen. A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1):48 – 77, 2015.
- [8] Z. Christoff, J. U. Hansen, and C. Proietti. Reflecting on social influence in networks. In *Proceedings of the Information Dynamics in Artificial Societies Workshop (IDAS-2014)*, 2014.
- [9] M. H. de Groot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [10] L. Festinger. Informal social communication. *Psychological Review*, 57(5):271–282, 1950.
- [11] S. Ghosh and F. R. Velázquez-Quesada. 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)*, 2015.
- [12] U. Grandi and U. Endriss. Binary aggregation with integrity constraints. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI-2011)*, 2011.
- [13] U. Grandi, E. Lorini, and L. Perrussel. Propositional opinion diffusion. In *Proceedings of the 14th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, 2015.
- [14] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.
- [15] J. Gutierrez, P. Harrenstein, and M. Wooldridge. Iterated boolean games. *Information and Computation*, 242:53–79, 2015.
- [16] K. Lehrer and C. Wagner. *Rational Consensus in Science and Society*. Springer, 1981.
- [17] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 8 of *LIPICs*, pages 133–144, 2010.
- [18] S. Moscovici and B. Personnaz. Studies in social influence. *Journal of Experimental Social Psychology*, 16:270–282, 1980.
- [19] A. Rapoport and A. M. Chammah. The game of chicken. *American Behavioral Scientist*, 10(3):10–28, 1966.
- [20] J. Ruan and M. Thielscher. A logic for knowledge flow in social networks. In *24th Australasian Joint Conference on Artificial Intelligence*, 2011.
- [21] T. Schelling. *Micromotives and macrobehavior*. Norton, 1978.
- [22] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-2015)*, 2015.
- [23] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. 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)*, 2016.
- [24] J. Seligman, F. Liu, and P. Girard. Facebook and the epistemic logic of friendship. In *Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2013)*, 2013.
- [25] J. Seligman and D. Thompson. Boolean network games and iterated boolean games. In *Proceedings of the 5th International Workshop on Logic, Rationality and Interaction (LORI-2015)*, pages 353–365. Springer-Verlag, 2015.
- [26] A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *Journal of the Association of Computing Machinery*, 32(3):733–749, 1985.