

Propositionwise Opinion Diffusion with Constraints

Sirin Botan
ILLC, University of Amsterdam
sbotan@uva.nl

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

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

ABSTRACT

We study the diffusion of opinions on a social network as an iterated process of aggregating neighbouring opinions. Individual views are modelled as vectors of yes/no answers to a number of propositions connected by an integrity constraint, and each individual updates her opinion by looking at the aggregated opinion of her influencers. We propose and compare two alternative methods for such a process. The first simply ignores the inconsistent aggregated opinion, while the second performs propositionwise revisions whilst maintaining consistency. We characterise the set of integrity constraints that allow individuals to reach the aggregated opinion of their influencers by means of propositionwise updates, and we study under what conditions the termination of the two proposed processes can be guaranteed.

CCS Concepts

•Computing methodologies → Multi-agent systems;

Keywords

Social networks, judgment aggregation, opinion transformation, axiomatic method, strategy-proofness

1. INTRODUCTION

When faced with the opinions of others over multiple issues, people will often be influenced to change their own opinions in line with the opinions of their influencers. Social influence describes the effect one person’s opinion can have on the opinions of those around her, and in formal models of social influence, is often represented by means of an influence or trust network. Taking inspiration from game-theoretic models of social influence on networks [15, 14], a number of diffusion methods have been proposed for complex representations of individual opinions, such as preferences [3], beliefs [29], or judgments [17]. A common characteristic of all these settings is that changes in agents opinions are driven by the aggregate opinions of their influencers in the trust network. Thus, each agent has some initial opinion that might change if the agent is connected to others in the network who disagree with her on one or more issues.

In complex settings, when individual opinions are formed on multiple interconnected issues, a diffusion process needs to carefully account for the constraints that relate the issues at stake. Recent work in the diffusion of beliefs [12] realised the importance of considering integrity constraints on the set of opinions held by individuals. One of the central problems in this setting is the design of individual updates when there are certain dependencies between issues, and aggregating opinions might not always result in a rational opinion. Consider a voter in the U.S. that will vote for Donald Trump unless at least one of the following three conditions are true: (a) Trump will actually build a wall separating the U.S. from Mexico, (b) Trump was supported by Russia in his campaign, and (c) Trump will continue to manage his company during his presidency. Such a voter may be faced with a third of her influencers believing only (a), a different third believing only (b), and the remaining third believing only (c). Hence, the voter can safely go ahead and vote for Trump, since for each condition there is a two-third majority of her influencers that is of the opinion that the statement is false. However, none of her influencers vote for Trump, since they all believe in at least one of the three conditions. How should such an agent update her opinion, and what kind of opinion updates might lead the agent to change opinion on her vote?

In this paper we represent opinions as binary vectors, building on work by Grandi et al. [17], generalizing this setting to take into account possible correlations among the issues at stake. We represent such correlations by means of an integrity constraint, which prevents agents from holding certain opinions over the set of issues. The addition of the constraint to the framework is problematic in some cases, as it prevents certain changes in opinion.

Related work.

Diffusion on networks has been extensively studied in the field of social network analysis, be it diffusion of diseases, information, or opinions [20, 8]. Some of these models are developed to deal with the diffusion of individual opinions, in which individual views are updated by averaging the views of neighboring individuals. Two classical such examples are threshold models [18], with its more recent generalisations [22, 23], and the De Groot or Lehrer-Wagner model [5, 25], which are however based on a simple representation of opinions as a binary view on a single issue, or a real-valued view in the interval $[0, 1]$. Of particular interest is the recent work of Friedkin et al.[12], which studies the propagation of real-valued beliefs over multiple issues interconnected by logical

Appears at: 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE 2017). Held as part of the Workshops at the 16th International Conference on Autonomous Agents and Multiagent Systems. May 8th-9th, 2017. Sao Paulo, Brazil.

integrity constraints.

Building on this literature, a recent stream of papers have adapted averaging models to more complex and realistic representations of opinions: knowledge bases [29, 30], preferences over alternatives [13, 3], and binary evaluations [17]. The latter paper used techniques from judgment aggregation (see, e.g., [19, 10]) and binary aggregation [16], representing opinions as vectors of binary views on *uncorrelated* issues, obtaining a diffusion model assimilable to that of threshold models. We build on this latter paper by adding a constraint connecting the multiple issues at stake, tackling the non-trivial problem of updating individual opinions towards one that does not satisfy the constraint.

We mention in passing that a logical perspective on diffusion in social networks has been explored in a number of papers (see, e.g., [2, 4, 31]). The strategic aspects of diffusion have been studied in the related setting of product adoption [32, 1], which however mostly focuses on threshold models on uncorrelated products. Finally, our axiomatic analysis draws inspiration from the work of Miller and Os-heron [27], Knight and Johnson [24], and Dryzek and List [7], who use the notion of inter-agent communication as a mean of reconciling ideas from deliberative democracy with those from social choice.

Paper overview.

The paper is organised as follows. In Section 2 we define two mechanisms for opinion diffusion with constraints on a network. In Section 3 we examine the characteristics of the opinion profiles which result from each of the two mechanisms and in Section 4 we study the termination of the iterative diffusion processes we defined. In Section 5 we study strategic aspects of the diffusion processes, and in Section 6 we propose an axiomatic study of diffusion as a judgment transformation function. Section 7 concludes the paper and points at directions for future work.

2. THE GENERAL FRAMEWORK

In this section we present two models of opinion diffusion in the presence of an integrity constraint. The first is a straightforward generalization of the process of propositional opinion diffusion [17]. The second is instead based on updates to single issues. We represent agents' opinions as answers to a set of yes/no questions which are possibly connected by means of an integrity constraint. We model the social influence network as a directed graph.

2.1 Individual Opinions

Let $\mathcal{I} = \{p_1, \dots, p_m\}$ be a finite set of m issues (or propositions), where each issue represents a binary choice. We call $\mathcal{D} = \{0, 1\}^{\mathcal{I}}$ the *domain* associated with this set of issues. For a finite set of agents, $\mathcal{N} = \{1, \dots, n\}$, we say $B_i \in \mathcal{D}$ is the *opinion* of agent $i \in \mathcal{N}$ over all issues in \mathcal{I} . A vector $\mathbf{B} \in \mathcal{D}^{\mathcal{N}}$ of all opinions of agents in \mathcal{N} is a *profile*. Each opinion B represents an agent's acceptance/rejection of each of the issues in \mathcal{I} . For example, if our set of issues is $\mathcal{I} = \{p, q, r\}$, then $B = (110)$ is the opinion which accepts the first two issues p and q and rejects the third issue r . We call $flip(B, p)$ the opinion resulting from changing the judgment on p in the opinion B . In our example above where $B = (110)$, $flip(B, p) = (010)$.

We write $B_i(p)$ to mean agent i 's judgment on $p \in \mathcal{I}$ in the profile \mathbf{B} . Thus if $B = (110)$, then $B(p) = B(q) = 1$

and $B(r) = 0$. We write $\mathbf{B} =_{-i} \mathbf{B}'$ to mean two profiles \mathbf{B} and \mathbf{B}' are identical if we ignore agent i 's opinion.

An *integrity constraint* $\text{IC} \subseteq \mathcal{D}$ defines a domain of feasible opinions. For instance, if we have three issues, p, q and r , and each agent can only accept at most two of the three, then $\text{IC} = \{(110), (011), (101), (100), (010), (001), (000)\}$. Integrity constraints are often represented compactly by means of a formula of propositional logic, such as $(\neg p \vee \neg q \vee \neg r)$ for the previous example. Issues of succinctness and computational complexity are out of the scope of this paper, hence we assume a set-theoretic representation of feasible opinions. For each agent i , we assume that $B_i \in \text{IC}$, meaning each individual opinion must satisfy the given rationality constraint.

2.2 The Social Influence Process

We assume that agents are connected by a *social influence network* $G = (\mathcal{N}, E)$ where $(i, j) \in E$ means agent i influences agent j and $\text{Inf}(i)_G = \{j \in \mathcal{N} \mid (j, i) \in E\}$ is the set of influencers of agent i in the network G . Observe that we do not make any assumption on whether $i \in \text{Inf}(i)$, defining the framework in full generality.

The first update procedure we propose is a direct generalisation of the *propositional opinion diffusion* (F -POD) proposed by Grandi et al. [17], in which agents simply aggregate the opinions of their influencers using some judgment aggregation functions $F = (F_1, \dots, F_n)$ —where $F_i : \text{IC}^{\text{Inf}(i)} \rightarrow \mathcal{D}$ is the rule agent i uses—and copy this aggregate opinion only if it satisfies the integrity constraint IC. Each step of the process is done according to the function POD. This function takes as input a network G , a profile of opinions $\mathbf{B} \in \mathcal{D}^{\mathcal{N}}$ and an agent $i \in \mathcal{N}$. Note that we assume F_i is resolute, meaning the function outputs a single opinion, and we do not require the outcome of F_i to be a model of the integrity constraint. The POD function returns the updated opinion of i according to her aggregation rule:


$$\text{POD}(G, \mathbf{B}, i) = \begin{cases} F_i(\mathbf{B}_{\text{Inf}(i)}) & \text{if } F_i(\mathbf{B}_{\text{Inf}(i)}) \in \text{IC} \\ B_i & \text{otherwise.} \end{cases}$$

We propose an alternative update, *propositionwise opinion diffusion* (F -PWOD), in which each agent updates on one issue at the time, provided that the updated opinion is consistent with the constraint. The function takes as additional argument the issue $p \in \mathcal{I}$ which agent i updates.

$$\text{PWOD}(G, \mathbf{B}, i, p) = \begin{cases} flip(B_i, p) & \text{if } F_i(\mathbf{B}_{\text{Inf}(i)})(p) \neq B_i(p) \\ & \text{and } flip(B_i, p) \in \text{IC} \\ B_i & \text{otherwise.} \end{cases}$$

The two processes can result in two different updates, as the following example shows:

EXAMPLE 1. *Three agents A, B, C, are voting on multiple referenda. They need to give opinions on three proposals; more parks in the city center, a homeless shelter and road repairs. Because of budget constraints, they can approve at most two of the proposals. Suppose the individuals are connected in the following social influence network, where the initial profile is $\mathbf{B} = (110, 011, 101)$, meaning A wants the parks and a homeless shelter, B wants the homeless shelter and road repairs, and C wants more parks and road repairs.*



Assume for each agent i , that F_i is the strict majority rule, accepting an issue only if a strict majority of the individuals accept it. If all agents update using the F -POD function, the resulting profile after one update will be $\mathbf{B}' = (110, 011, 011)$. If each agent uses the F -PWOD function instead, and we suppose they all update on the first issue, we get a different outcome after the first iteration— $\mathbf{B}' = (110, 011, 001)$.

2.3 The Iterative Process

A permissible transformation associates a profile of individual opinions with one of the possible outcomes of either a F -POD or an F -PWOD update:

DEFINITION 1. Let a network G , an integrity constraint IC , and a set of aggregators F_i for $i \in \mathcal{N}$ be given. We say there is a permissible F -POD transformation from profile \mathbf{B} to profile \mathbf{B}' if there exists $I \subseteq \mathcal{N}$ such that $B'_i = \text{POD}(G, \mathbf{B}, i)$ for all $i \in I$, and $B'_j = B_j$ for all $j \notin I$. Analogously, there is a permissible F -PWOD transformation from \mathbf{B} to \mathbf{B}' if there exists $I \subseteq \mathcal{N}$ and $p_i \in \mathcal{I}$ for each $i \in I$ such that $B'_i = \text{PWOD}(G, \mathbf{B}, i, p)$, and $B'_j = B_j$ for all $j \notin I$.

We say that a permissible transformation is *effective* if there is some $i \in \mathcal{N}$ such that $B_i \neq B'_i$. We further say that a profile \mathbf{B} is a *termination profile* if no effective transformation exists. An agent's opinion B_i is called *stable* on a network G (wrt. F) in profile \mathbf{B} if for any p , F -PWOD(G, \mathbf{B}, i, p) = B_i . Thus, a termination profile is a profile in which all individual opinions are stable.

Both F -POD and F -PWOD update functions can be used to define diffusion processes with discrete time. Let $\text{turn} : \mathbb{N} \rightarrow 2^{\mathcal{N}}$ be a turn function, indicating at each point in time t what are the agents that are updating their opinions. Let $\mathbf{B}^t = (B_1^t, \dots, B_n^t)$ be the profile of opinions at time t . At time $t + 1$, all and only the agents in $\text{turn}(t + 1)$ will perform an F -POD update aggregating the opinions of their influencers:

$$B_i^{t+1} = \text{POD}(G, \mathbf{B}^t, i).$$

An additional ingredient is required for F -PWOD. Let $\text{prop}_i : \mathbb{N} \rightarrow \mathcal{I}$ for each $i \in \mathcal{N}$ be a function which tells us which proposition agent i is allowed to update at any time t . At each time t all agents in $\text{turn}(t)$ update their opinion according to the aggregated opinion of their influencers at time $t - 1$ on issues $\text{prop}_i(t)$:

$$B_i^{t+1} = \text{PWOD}(G, \mathbf{B}^t, i, \text{prop}_i(t + 1))$$

for all $i \in \text{turn}(t + 1)$. Following Definition 1, there is a permissible transformation between each step of the diffusion process, i.e., each pair of profiles \mathbf{B}^t and \mathbf{B}^{t+1} . Observe that if \mathbf{B} is a termination profile and at time T we have that $\mathbf{B}^T = \mathbf{B}$ then it is the case that $\mathbf{B}^t = \mathbf{B}^T$ for all $t \geq T$.

3. TERMINATION PROFILES

In this section we focus on the properties of termination profiles and under which conditions F -POD and F -PWOD result in the same termination profiles. We characterise the set of integrity constraints for which F -PWOD termination profiles agree with the outcome of the respective aggregation functions, and we show that for both F -POD and F -PWOD, an agent's opinion at termination may be very distant from the opinions of her influencers.


3.1 Integrity Constraints with Open Structure

Given two opinions B and $B' \in \mathcal{D}$, recall that the *Hamming distance* between them is $H(B, B') = \sum_{p \in \mathcal{I}} |B(p) - B'(p)|$.

DEFINITION 2. An integrity constraint IC has an *open structure* if for any two opinions $B, B' \in IC$ where $H(B, B') = k$, there is some sequence of distinct opinions B_1, \dots, B_{k+1} —all in IC —such that $B_1 = B$, $B_{k+1} = B'$, and $H(B_i, B_{i+1}) = 1$ for all $1 \leq i \leq k$.

To visualise the idea underlining the above definition, we model the opinions on a hypercube. An edge between any two nodes means the Hamming distance between them is 1. An integrity constraint has an open structure if any two nodes at distance k are connected by a path of length exactly k .

EXAMPLE 2. We represent all opinions in the graph below, connecting only those that satisfy IC with a continuous edge. Let $IC = \{(000), (001), (010), (100), (011), (111)\}$. This integrity constraint does not have an open structure, and this can be visualised on the figure: the shortest path available between (100) and (111) is of length 4, which is strictly greater than the Hamming distance between the two models $H(100, 111) = 2$.



An important class of integrity constraints that has an open structure is the one used to represent preferences as linear orders over a set of alternatives (see, e.g., [28]). Let us see this example in details. Let A be a set of alternatives, a linear order is an irreflexive, transitive and complete binary relation over A . A linear order \succ can be represented as a binary evaluation over a set of issues $\mathcal{I}_A = \{p_{ab} \mid (a, b) \in A \times A \text{ and } a \neq b\}$, such that $B(p_{ab}) = 1$ if and only if $a \succ b$. For each pair (a, b) , we only include one of p_{ab} and p_{ba} in the issues as rejecting p_{ab} in a linear order is equivalent to accepting p_{ba} and vice versa. The integrity constraint IC_{\succ} therefore contains all opinions over \mathcal{I}_A corresponding to linear orders over A .¹ We can now prove the following:

PROPOSITION 1. IC_{\succ} has an open structure.

¹Representing preferences with binary evaluations is an idea that can be traced back to the work of Wilson [33].

PROOF. Let B and B' be two distinct opinions in IC_{\succ} such that $H(B, B') = k$, and let \prec and \prec' be the corresponding linear orders. Since the two orders \prec and \prec' are different, they also differ on pair which is adjacent in one of them, i.e., there exists a pair ab such that $B(p_{ab}) \neq B'(p_{ab})$ and there is no $c \in A$ such that $a \succ_i c \succ_i b$ or $b \succ_i c \succ_i a$.² Swapping an adjacent pair in a linear order results in a binary relation that still is a linear order, hence $\text{flip}(B, p_{ab}) \in \text{IC}_{\succ}$. By repeating updates on adjacent pairs we can therefore build a sequence of propositionwise updates of length k from B into B' . \square

3.2 F -Consistent Termination

We now give a formal definition that we will use to characterise integrity constraints on which the outcome of the propositionwise diffusion process matches the outcome of the rule F , if this outcome satisfies integrity constraint.

DEFINITION 3. *An opinion diffusion process is said to be F -consistent on a network G if for all termination profiles \mathbf{B} it is the case that for any $i \in \mathcal{N}$: if $F(\mathbf{B}_{\text{Inf}(i)}) \in \text{IC}$, then $B_i = F(\mathbf{B}_{\text{Inf}(i)})$.*

Clearly F -POD is F -consistent. We show that the same holds for F -PWOD if and only if IC has an open structure.

PROPOSITION 2. *F -PWOD is F -consistent if and only if IC has an open structure.*

PROOF. For the right to left direction, we first assume that IC has an open structure. Suppose further F -PWOD terminates on a profile \mathbf{B} and $F(\mathbf{B}_{\text{Inf}(i)}) \in \text{IC}$, and suppose for contradiction that F -PWOD is not F -consistent. That is, there is some agent $i \in \mathcal{N}$ such that $B'_i \neq F_i(\mathbf{B}_{\text{Inf}(i)})$. If $F_i(\mathbf{B}_{\text{Inf}(i)}) \in \text{IC}$, then since IC has an open structure, there must be some $p \in \mathcal{I}$ s.t. $F_i(\mathbf{B}_{\text{Inf}(i)})(p) \neq B_i(p)$ and $\text{flip}(B_i, p) \in \text{IC}$. By Definition 1, this implies the existence of a permissible and effective transformation from \mathbf{B} to a second profile \mathbf{B}' by having i updating on p , against the assumption that \mathbf{B} is a termination profile.

For the left to right direction, suppose that IC does not have an open structure. Then it must be the case that there are two opinions B, B' such that $H(B, B') = k$ and all paths of opinions in IC connecting them has length at least $k + 1$. By the pigeonhole principle, this implies the existence of two distinct opinions B'' and B''' , possibly equal to B and B' , such that there is no $p \in \mathcal{I}$ where $B''(p) \neq B'''(p)$ and $\text{flip}(B'', p) \in \text{IC}$. Let now $\mathcal{N} = \{1, 2\}$, $E = \{(1, 2)\}$, and $\mathbf{B} = (B''', B'')$. Observe that \mathbf{B} is a termination profile, since $F(\mathbf{B}_{\text{Inf}(2)}) = B'''$ and by construction there is no p such that $B''(p) \neq B'''(p)$ and $\text{flip}(B'', p) \in \text{IC}$. But $B'' \neq B'''$ and therefore F -PWOD is not F -consistent. \square

Proposition 2 shows that if aggregating an agent's influencers using F gives an opinion in the set IC, F -PWOD will eventually reach a state where each agent's opinion is equivalent to the outcome of F . We can in fact make a stronger claim if we know an agent's sources have stable opinions:

PROPOSITION 3. *Let $i \in \mathcal{N}$ and \mathbf{B} be a profile on G . If all $j \in \text{Inf}(i)$ have stable opinions in \mathbf{B} and IC has an open structure, then for any F -PWOD termination profile \mathbf{B}'' and F -POD termination profile \mathbf{B}' , both resulting from \mathbf{B} , we have that $H(B'_i, F(\mathbf{B}''_{\text{Inf}(i)})) \leq H(B'_i, F(\mathbf{B}'_{\text{Inf}(i)}))$.*

²This result is folklore, a formal proof can be found in [9].


PROOF. If all agents in $\text{Inf}(i)$ have stable opinions on G , then $F(\mathbf{B}''_{\text{Inf}(i)}) = F(\mathbf{B}'_{\text{Inf}(i)}) = F(\mathbf{B}_{\text{Inf}(i)})$. Suppose $F(\mathbf{B}_{\text{Inf}(i)}) \in \text{IC}$. By Proposition 2, F -PWOD ensures that at termination $B''_i = F(\mathbf{B}_{\text{Inf}(i)})$, and the same will hold for F -POD. Suppose $F(\mathbf{B}_{\text{Inf}(i)}) \notin \text{IC}$. Then $B'_i = B_i$. If F -PWOD is not able to perform any updates, $B''_i = B_i$ as well, but if even one update is performed, $H(F(\mathbf{B}''_{\text{Inf}(i)}), B'_i) < H(F(\mathbf{B}'_{\text{Inf}(i)}), B'_i)$. \square

Although the assumption of IC having an open structure guarantees that F -PWOD will be able to make at least as many updates as F -POD when faced with an outcome which does not satisfy the constraint, this might in some cases simply mean that neither F -PWOD nor F -POD will be able to update. In the worst case, this means that an agent will end up with an opinion that is very distant from the opinions of her influencers.

PROPOSITION 4. *For any number of issues m , there is always some IC with open structure such that we can construct a network G and a F -PWOD termination profile \mathbf{B} where an agent i is at distance $m - 2$ from $F(\mathbf{B}_{\text{Inf}(i)})$.*

PROOF. Let $\mathcal{I} = \{p_1, \dots, p_m\}$ and $\text{IC} = (\neg p_1 \wedge \neg p_m) \rightarrow (p_2 \wedge \dots \wedge p_{m-1})$, i.e. IC allows all opinions except those which reject the first and last issue and at least one other issue. We slightly abuse our notation and say that B_0 is the opinion which rejects all issues, B_k only accepts the k th issue, and B_* rejects the first and last issue, but accepts all others.

Let F_i be the strict majority rule. We first show that we can construct a network and termination profile such that there is an agent who is at distance $m - 2$ from the outcome of F_i over her influencers. Take the following network G and profile \mathbf{B} :



Here $F_i(\mathbf{B}_{\text{Inf}(i)}) = B_0 \notin \text{IC}$. However, for any issue on which agent i does not agree with the majority, namely p_2 to p_{m-1} , she cannot update her opinion without ending up in one of the opinions prohibited by IC.

We now show such an IC must have an open structure. Let $B, B' \models \text{IC}$. Suppose both B and B' accept the first (last) issue. Then since all opinions accepting the first (last) issue satisfy IC, we can freely move between the two by performing updates to single propositions. If both B and B' reject both the first and the last issue, then $B = B' = B_*$ as this is the only opinion which satisfies IC.

Thus, we only need to check if there is a required sequence of opinions between B and B' if they disagree on either the first issue or the last. W.l.o.g., suppose they disagree on the first issue and B rejects the first issue and B' accepts it. Then B can update on the first issue before performing any other updates, as $\text{flip}(B, p_1)$ satisfies IC. Now since, $\text{flip}(B, p_1)$ and B' both accept the first issue, there must be a sequence of opinions from $\text{flip}(B, p_1)$ to B' where the distance between any two successive opinions is 1 and each satisfies IC. Further, since $H(B, \text{flip}(B, p_1)) = 1$, we can conclude that the constructed sequence has length exactly $H(B, B') + 1$. \square

Note that in the construction in Proposition 4, F -POD would result in the same termination profile.

4. TERMINATION OF ITERATIVE OPINION DIFFUSION

In this section we compare the two proposed diffusion models with respect to the termination of the associated iterative process. We first need to introduce a number of definitions.

Recall that by fixing a $turn$ function and functions $prop_i$ for $i \in \mathcal{N}$, deciding which agents are updating and on which issues, we can define iterative processes associated to F -POD and F -PWOD. The following definitions are straightforward adaptations of those proposed by Brill et al. [3]. We call an iterative process *asynchronous* if $|turn(t)| = 1$ for all $t \in \mathbb{N}$, and *synchronous* if $turn(t) = \mathcal{N}$ for all $t \in \mathbb{N}$. We say that the iterative process F -POD or F -PWOD *universally terminate* on a class of graphs \mathcal{E} if for all $G \in \mathcal{E}$ and each initial opinion profile \mathbf{B} there does not exist an infinite sequence of effective transformations starting from \mathbf{B} . We say that F -POD or F -PWOD *asymptotically terminate* on a class of graphs \mathcal{E} if for all $E \in \mathcal{E}$ and profiles \mathbf{B} the following condition holds: from all profiles \mathbf{B}' reachable from \mathbf{B} there exists a path of permissible transformations leading to a termination profile. When both the $turn$ and $prop_i$ functions select an agent and an issue uniformly at random, asymptotic termination implies that the probability of eventually reaching a termination state tends to 1 as t goes to infinity. Finally, a *consensual termination profile* is a profile \mathbf{B} such that for all $i, j \in \mathcal{N}$ we have that $B_i = B_j$.

Aggregation functions F_i are typically classified by means of axiomatic properties. A full-blown analysis of the influence of these properties on termination is out of the scope of this paper, but we still need one such definition. We say that an aggregator F_i is *unanimous* for agent i if, whenever $B_j = B^*$ for all $j \neq i$ then $F_i(\mathbf{B}) = B^*$. In words, whenever all influencers (excluding the updating agent) are unanimous, F updates according to the influencers.

4.1 Simple cycles

A simple cycle is a finite network E such that every agent has exactly one outgoing edge and exactly one incoming edge.

PROPOSITION 5. *If G is a simple cycle and F_i are unanimous, then asynchronous F -POD terminates asymptotically to a consensual termination profile.*

PROOF. Let \mathbf{B}^0 be a profile on the simple cycle G , where $E = \{(1, 2), \dots, (i, i+1), \dots, (n, 1)\}$. Let $i^* \in \mathcal{N}$ be such that $B_{i^*}^0 \neq B_{i^*+1}^0$. If such an agent does not exist then the profile \mathbf{B} is already a consensual termination profile. Let us now define the following turn function. Let $turn(t) = i^* + t + 1$, for $t = 0, \dots, n-1$. Since $B_{i^*}^0$ satisfy the integrity constraint by assumption, and all F_i are unanimous aggregators, then at each iteration step t agent $i^* + t + 1$ will copy the opinion of agent $i^* + t$, obtaining a consensual termination profile at $t = n - 1$ in which all agents have the same opinion $B_{i^*}^0$. \square

The same result holds for F -PWOD, albeit under additional assumptions on the integrity constraint:

PROPOSITION 6. *If G is a simple cycle, F is unanimous, and IC has an open structure, then asynchronous F -PWOD terminates asymptotically to a consensual termination profile. The same holds for synchronous F -PWOD if $|\mathcal{I}| \geq 2$.*

PROOF SKETCH. Let \mathbf{B}^0 be a profile on G , and let $i^* \in \mathcal{N}$ be such that $B_{i^*}^0 \neq B_{i^*+1}^0$. Since IC has an open structure, there is a sequence of propositionwise updates of length $k = H(B_{i^*}^0, B_{i^*+1}^0)$ that transforms the latter opinion into the former. By defining $turn(t) = i^* + 1$ for $t = 0, \dots, k$, and $prop_i$ according to the sequence above, we obtain a resulting profile \mathbf{B}^k such that $B_{i^*+1}^k = B_{i^*}^0$ and $B_j^k = B_j^0$ for all $j \neq i^* + 1$. The process can then be repeated for $i^* + 2$, and sequentially until reaching again agent i^* , to obtain a consensual termination profile in which all agents have the same opinion $B_{i^*}^0$.

The proof for synchronous F -PWOD uses the same construction as above, setting the $prop_i$ functions to update on irrelevant issues for the non updating agents. \square

Observe that the set of termination profiles that can be reached starting from the same profile of initial opinions can be different depending on whether we are using F -POD or F -PWOD. In particular, while the former leads to profiles that are consensual on opinions that are already present in the initial profile, the second can result in consensual profiles on opinions that are a combination of the initial ones.

4.2 Directed acyclic graphs

A directed acyclic graph (DAG) is a directed graph that contains no cycle involving two or more vertices. A simple argument of propagation allows us to prove the following:

PROPOSITION 7. *If G is a DAG, then both synchronous and asynchronous F -POD and F -PWOD converge universally.*

PROOF SKETCH. We define potential functions h_i for each node i , as follows: $h_i(t) = H(B_i^t, F_i(\mathbf{B}_{Inf(i)}))$, measuring the distance between an individual's opinion and the aggregated opinion of its influencers. Each effective transformation under both F -POD and F -PWOD decreases one such function, the one of the updating agent, and possibly increases others, those of the agents influenced by the one updating. By ordering such potential functions based on the distance from a node to a source, which is possible given the assumption that G is a DAG, we obtain a lexicographic ordering of all functions h_i that decreases strictly with each effective transformation. Therefore, for any set of aggregators F_i and any DAG it is impossible to build an infinite sequence of F -POD or F -PWOD effective transformations. \square

4.3 Complete graphs

Let a complete graph be such that $E = \mathcal{N} \times \mathcal{N}$. Observe in particular that this means $i \in Inf(i)$ for each $i \in \mathcal{N}$. Using an idea from Farnoud et al. [11], we are able to show the following:

PROPOSITION 8. *If G is the complete graph, then both synchronous and asynchronous F -POD and F -PWOD converge universally.*

PROOF. On a complete graph the set of influencers $Inf(i) = \mathcal{N}$ for all i . Let therefore $h(t) = \sum_i H(B_i, F(\mathbf{B}))$ be a potential function that measures the overall distance of the individual opinions from the overall aggregated one. Every effective transformation for both F -POD and F -PWOD decreases the value of h , hence obtaining the desired result. \square


A general result on the asymptotic convergence of F -POD or F -PWOD is an open problem. A proof similar to the

one used by [3] could be adapted to show that F -PWOD asymptotically converges on any graph, provided that at any point in time the aggregated opinion of any set of influencers satisfy the integrity constraint. This assumption seems however too restrictive for diffusion processes that are designed to deal with integrity constraints. Universal convergence cannot be guaranteed even on simple cycles, for both F -POD and F -PWOD, at least when more than two issues are present. To see this it is sufficient to consider a simple cycle with only one agent having opinion 11 and all others 00, and devise $turn$ and $prop_i$ functions that make the 11 opinion turn in the cycle whilst keeping all other opinions at 00.

4.4 Update Order Dependence

When updating on single propositions at a time, even with all agents updating synchronously, the order in which each agent updates their opinions matters in determining what the possible termination profiles look like. Consider for instance the following example:

EXAMPLE 3. Let a network and a profile of opinions be as in the figure below and let $IC = \mathcal{D} \setminus \{(111)\}$.



Two agents with the same initial opinion can have the same set of influencers yet end up with different opinions in a termination profile, depending on the order in which they update their opinions on the issues. We can see this with agents D and E who have the same initial opinion. In our case if D updates the issues in the order p, q, r , obtaining 110, and E in the order r, q, p , obtaining 011, these will be their opinions in the termination profile.

Similar situations occur when an integrity constraint blocks the update on a certain set of issues, even though the result of the majority rule is consistent. This does not happen with IC with open structure, as can be shown in the following proposition. Recall from Section 2 that by fixing a function $prop_i$ for each individual i we obtain an iterative diffusion process. We say that the $prop_i$ functions are *balanced* if all profiles at which the iterative process stabilizes, i.e., when there is a T such that $\mathbf{B}^t = \mathbf{B}^T$ for all $t \geq T$, then \mathbf{B}^T is a termination profile.

PROPOSITION 9. If IC has an open structure, that IC is guaranteed to be satisfied by the outcome of F , and that no ties will occur at any iteration step, then on any directed acyclic graph any choice of balanced $prop_i$ functions results in the same termination profile as F -POD.

PROOF SKETCH. By Proposition 2, if IC has an open structure and the outcome of F is guaranteed to satisfy the integrity constraint, then the process will converge to the result of aggregating the influencers' opinions via F . By Proposition 7 we know that the iterative process on DAGs converge, and by a simple algorithm of propagation from the sources we can also show that the iterative F -PWOD process stabilises on a termination profile that is uniquely determined by the initial profile \mathbf{B}^0 . \square

Proposition 9 does not generalize to arbitrary network containing cycles. Consider the following example:

EXAMPLE 4. Let G be as depicted in the figure below, and let there be no integrity constraint, i.e., $IC = \mathcal{D}$.



Suppose in the first round of updates, agent A updates on the first issue, and B on the second issue. In the second round, A updates on the second issue, and B on the first issue. PWOD will then terminate on the unanimous profile (01), (01). However if the agents switch the order of updates (A updates the second issue first, then the first issue, similarly for B), we arrive at the profile (10), (10).


5. STRATEGIC MANIPULATION

In this section we examine the possibility of a strategic agent guiding the outcome of a diffusion process by misreporting her initial opinion. We limit our attention to the source agents when considering possible cases of manipulation, as these are the only agents in the network who are in a sense, sure about their opinion and will only change it for strategic reasons.

We assume agents' preferences are defined by means of the Hamming distance wrt. their initial opinion (this is one of many possible choices, see, e.g., [6]). Each agent i with initial opinion B_i is associated with a weak ordering \succeq_i defined as follows: $B \succeq_i B'$ if and only if $H(B, B_i) \leq H(B', B_i)$, i.e., when the Hamming distance between her truthful opinion and B is less than or equal to the distance between her truthful opinion and the opinion B' . In what follows we provide two examples in which a source agent is able to guide the influence process to obtain a resulting opinion profile where agents influenced by her have opinions closer to hers if compared to the outcome of the diffusion process had she been honest about her opinion.

We begin by showing how F -POD can be manipulated in presence of an integrity constraint.


EXAMPLE 5. Let there be four agents, and let $\mathcal{D} \setminus IC = \{111\}$, i.e. let (111) be the only forbidden opinion. Let the network and the profile be defined as below:



Suppose F_D is the strict majority rule, resulting in an aggregated result of (111), and no update for agent D . Then agent A will benefit by reporting (010) instead of her truthful opinion above: in the truthful profile D does not update, keeping her opinion which is at distance 3 from A 's opinion, while in the second profile D updates to (110), which is at distance 2 from A 's truthful opinion (011).

The situation is similar for F -PWOD, except that a potential manipulator needs to know the order of updates on the issues in advance to be sure of the effect of her manipulation.

EXAMPLE 6. Let $IC = \{111, 100, 010, 001, 011, 000\}$. Let the network and the profile be defined as follows:



If agent D updates first her opinion on p_1 then p_2 then p_3 , then her opinion at termination will be (100). If agent A knows this is the order in which D will perform the updates, it is in her interest to report the opinion (011) instead of her truthful opinion (111), as this results in the opinion (011) for agent D at termination. Note however that if agent D updates in a different order, say first p_3 then p_2 then p_1 , then D 's opinion at termination is (111) and it is in agent A 's interest to keep her initial opinion.

6. TRANSFORMATION FUNCTIONS

A transformation function is one way of representing opinion change among a group of agents. Following the definition of List [26], such a function takes as input a profile of opinions \mathbf{B} and outputs a second profile \mathbf{B}' , representing the influenced or updated opinions. One example of such a transformation function T is *deference to unanimity*, where $T_i(\mathbf{B})$ – the transformed opinion of agent i – accepts only those issues unanimously accepted by all agents in \mathbf{B} . In this section, we adapt this definition to take into account the network relating the individuals, and we adapt the axioms initially proposed by List [26], and subsequently formalized by Grossi and Pigozzi [19], to this setting.

6.1 Opinion Transformation on a Network

Given a social influence network and a profile corresponding to the opinions of the agents in the network, a *network-based opinion transformation function* returns a profile which comprises the updated opinions of each agent in the network. Formally (recall that \mathcal{D} is the set of all individual opinions):

$$T : \mathcal{D}^{\mathcal{N}} \times 2^{(\mathcal{N} \times \mathcal{N})} \rightarrow \mathcal{D}^{\mathcal{N}}.$$

The propositionwise opinion diffusion mechanisms defined in Section 2 can be viewed as network-based opinion transformation functions. Given a set of issues \mathcal{I} , agents \mathcal{N} , an integrity constraints $\text{IC} \subseteq \mathcal{D}$, and an influence network $G = (\mathcal{N}, E)$, for any $p \in \mathcal{I}$ we can define a transformation function T where:

$$T_i(\mathbf{B}, G) = F\text{-PWOD}(G, \mathbf{B}, i, p).$$

With p corresponding to $\text{prop}_i(t)$, at any time t of the iterative diffusion process.

6.2 Axioms for Opinion Transformations

In this section we adapt some of the axioms proposed by List [26] to the current setting, and we propose novel network-specific properties. We begin with the following straightforward adaptation of some classical axioms. Note that by $T_{i,p}(\mathbf{B}, G)$ we mean the opinion of agent i on p in the transformed profile.

Rationality: for all networks $G \in \mathcal{G}$, agents $i \in \mathcal{N}$, profiles $\mathbf{B} \in \text{IC}^{\mathcal{N}}$ we have that $T_i(\mathbf{B}, G) \in \text{IC}$.³

Unanimity: for all networks $G \in \mathcal{G}$ and opinions $B^* \in \mathcal{D}$, if it is the case that $B_i = B^*$ for all agents $i \in \mathcal{N}$, then $T_i(\mathbf{B}, G) = B^*$ for all $i \in \mathcal{N}$.

³Observe that IC is a parameter of this axiom.

Responsiveness: for all networks $G \in \mathcal{G}$ and agents $i \in \mathcal{N}$, there exist two profiles $\mathbf{B}, \mathbf{B}' \in \mathcal{D}^{\mathcal{N}}$ such that $\mathbf{B} =_{-i} \mathbf{B}'$, $B_i \neq B'_i$ and $T_i(\mathbf{B}, G) \neq T_i(\mathbf{B}', G)$.

Independence: for all networks $G \in \mathcal{G}$, issues $p \in \mathcal{I}$, and pairs of profiles $\mathbf{B}, \mathbf{B}' \in \mathcal{D}^{\mathcal{N}}$, if it is the case that $B_i(p) = B'_i(p)$ for all $i \in \mathcal{N}$ then $T_{i,p}(\mathbf{B}, G) = T_{i,p}(\mathbf{B}', G)$ for all $i \in \mathcal{N}$.

Monotonicity: for all networks $G \in \mathcal{G}$, issues $p \in \mathcal{I}$, and pairs of profiles $\mathbf{B}, \mathbf{B}' \in \mathcal{D}^{\mathcal{N}}$, if for $i, j \in \mathcal{N}$, $\mathbf{B} =_{-j} \mathbf{B}'$, $B_j =_{-p} B'_j$, $B(p) = 0$ and $B'(p) = 1$, then $T_i(\mathbf{B}, G)(p) = 1 \Rightarrow T_i(\mathbf{B}', G)(p) = 1$.

Rationality states that if the input to the transformation function is a profile of rational opinions, then the outcome of the transformation should be a profile of rational opinions. Unanimity states that if every opinion in the input profile is the same, then the function simply outputs this same profile. A transformation function is Responsive if there are two profiles in which only agent i changes her opinion, and her opinion in the outcome is different for the two profiles. Independence states that the opinion an agent has on a proposition p in the outcome of the transformation function depends only on agents' opinions on p in the input profile. Monotonicity requires that for any agent i , if they accepted a proposition p in the outcome of a transformation function T applied to a profile \mathbf{B} , then added support to this proposition in a profile \mathbf{B}' should imply that p remains accepted by agent i in the outcome of T .

Several of the axioms for transformation function have counterparts in judgment aggregation. For example, the Monotonicity Axiom for network-based transformation functions simply states that the aggregation function each agent uses must satisfy Monotonicity as defined for aggregation functions.

We now give three axioms that are specifically defined for transformations on a social network.

Influencer-Unanimity: for all networks $G \in \mathcal{G}$, opinions $B^* \in \mathcal{D}$, and agents $i \in \mathcal{N}$, if for all agents $j \neq i \in \text{Inf}(i)$ we have that $B_j = B^*$ then $T_i(\mathbf{B}, G) = B^*$.

Influencer-Independence: for all networks $G \in \mathcal{G}$, issues $p \in \mathcal{I}$, agents $i \in \mathcal{N}$, and profiles $\mathbf{B}, \mathbf{B}' \in \mathcal{D}^{\mathcal{N}}$, if it is the case that $B_j(p) = B'_j(p)$ for all $j \in \text{Inf}(i)$ then $T_{i,p}(\mathbf{B}, G) = T_{i,p}(\mathbf{B}', G)$.

Exclusiveness: for all networks $G \in \mathcal{G}$, agents $i \in \mathcal{N}$, and profiles $\mathbf{B}, \mathbf{B}' \in \mathcal{D}^{\mathcal{N}}$, if $[\forall j \in \text{Inf}(i) \cup \{i\} : B_j = B'_j]$, then $[T_i(\mathbf{B}, G) = T_i(\mathbf{B}', G)]$.

Influencer-Unanimity states that if all influencers of an agent submit the same opinion in the input to the transformation function, then the agent submits that same opinion in the output profile. Influencer-Independence states that a transformation function is Independent with respect to the opinions of an agent's influencers. For the complete network where $\text{Inf}(i) = \mathcal{N}$ for all agents i , Influencer-Independence corresponds to Independence. Finally, a transformation function is Exclusive if an agent's opinion in the output of the transformation function depends only on her own opinion and the opinions of her influencers in the input profile. This means that someone who is not an influencer of an agent cannot play any role in the opinion update of this agent at any step in the diffusion.

6.3 Majority PWOD


The aggregation rule F_{Maj} , which accepts only the issues accepted by a (strict) majority of agents, is the *strict majority rule*, and is defined such that for any proposition p , $F_{Maj}(\mathbf{B})(p) = 1$ if and only if $|N_p^{\mathbf{B}}| > \frac{n}{2}$, where $N_p^{\mathbf{B}}$ is the set of agents who accept p in the profile \mathbf{B} .⁴ Let MAJ-PWOD be the propositionwise opinion diffusion model in which each agent uses the strict majority rule to update, i.e., where $F_i = F_{Maj}$ for all $i \in \mathcal{N}$. By definition, F -POD and F -PWOD for any F satisfy Rationality and Exclusiveness. The same holds for Responsiveness as agents must always take into account their own opinion to ensure that changes can be made on a subset of \mathcal{I} while still satisfying the constraint. Though the majority judgment aggregation rule satisfies Independence and Unanimity, the propositionwise updates lead to a violation of the corresponding axioms for transformation functions.

PROPOSITION 10. *The MAJ-PWOD transformation function is rational, unanimous, responsive and monotonic. It does not satisfy independence, influencer-independence, or influencer-unanimity.*

PROOF. As noted above, MAJ-PWOD satisfies Rationality and Exclusiveness by definition. Moreover, it is straightforward to observe that Unanimity is also satisfied, as if every agent submits the same ballot B , then any agent i will agree with her influencers on any proposition p and will never change her opinion.

For Monotonicity, suppose for profiles \mathbf{B} and \mathbf{B}' that for $i, j \in \mathcal{N}$, $\mathbf{B} =_{-j} \mathbf{B}'$, $B_j =_{-p} B'_j$, $B_j(p) = 0$ and $B'_j(p) = 1$, and further, that $T_{i,p}(\mathbf{B}, G) = 1$. If $j \neq i$ and $j \notin \text{Inf}(i)$ we know $T_{i,p}(\mathbf{B}', G) = T_{i,p}(\mathbf{B}, G)$ because MAJ-PWOD is Exclusive. If $j \in \text{Inf}(i)$, then $T_{i,p}(\mathbf{B}, G) = 1$ means there was a majority of acceptances for p among agent i 's influencers, or there was a majority of rejections but a change in opinion was blocked by IC. If it is the former, we know that an additional acceptance for p in \mathbf{B}' means it remains the case that $T_{i,p}(\mathbf{B}', G) = 1$. If it is the latter, then it must have been the case that $B_i(p) = 1$ and thus $B'_i(p) = T_{i,p}(\mathbf{B}', G) = 1$. Now suppose $i = j$. If $j \in \text{Inf}(i)$, then the only way $T_{i,p}(\mathbf{B}', G) = 0$ is if there is a majority of rejections for p among agent i 's influencers, but since no agent but i changes her opinion, this cannot be the case. If $i \notin \text{Inf}(i)$ we fall back into the first case we analysed.

We provide a counterexample to show that Influencer-Unanimity fails. Let there be two issues $\mathcal{I} = \{p, q\}$ and suppose $\text{IC} = p \rightarrow q$. Let G be the following network and \mathbf{B} the profile shown in the network below:



Let p be the issue agent i is updating. Then $T_i(\mathbf{B}, G) = B_i$ as an update to p would lead to an opinion which does not satisfy the constraint, falsifying Influencer-Unanimity.

Take now a second profile \mathbf{B}' that coincides with the one described above, with the exception that $B'_i = (01)$,

⁴Here we take n to be the number of opinions in the input to the aggregation rule and not the total number of agents in the network.

hence such that $\mathbf{B} =_{-i} \mathbf{B}'$. We have that $T_{i,p}(\mathbf{B}, G) = 0$, since the update is blocked by the integrity constraint, while $T_{i,p}(\mathbf{B}', G) = 1$, falsifying Independence.

Influencer-Independence also fails, as can be seen in the following example. Take two profiles $\mathbf{B}, \mathbf{B}' \in \mathcal{D}^{\mathcal{N}}$, and let $B_j(p) = B'_j(p) = 1$ for all $j \in \text{Inf}(i)$ for some agent i . Suppose $i \notin \text{Inf}(i)$ and let $B_i = (10)$, $B'_i = (01)$. Further, let $\text{IC} = \{(01), (10)\}$. Then, even if $F_{Maj}(\mathbf{B}_{\text{Inf}(i)})(p) = F_{Maj}(\mathbf{B}'_{\text{Inf}(i)})(p) = 1$, we still have that $T_i(\mathbf{B}, G) = (10)$ while $T_i(\mathbf{B}', G) = (01)$, contradicting the axiom of Influencer-Independence. \square

7. CONCLUSIONS AND FUTURE WORK

In this paper we have introduced and studied two models for opinion diffusion on multiple binary issues connected by an integrity constraint. Propositional opinion diffusion F -POD updates on all issues at the same time, provided that the aggregated opinion of one's influencers satisfy the integrity constraint. Propositionwise opinion diffusion F -PWOD, instead, updates on one issue at the time towards the aggregated opinion of the influencers, provided that this single change satisfies the integrity constraint. We have characterised the set of integrity constraints on which F -PWOD coincides with F -POD at termination of the diffusion process, and compared the two processes on the distance between an agent's opinion and the one of her influencers. We have given sufficient conditions for the termination of the iterated diffusion process, and provided initial results on the strategic abilities of source agents in the network. We also adapted axiomatic conditions for profile transformation functions, previously defined in judgment aggregation, to take into account a social network relating the individuals, and used these novel formulations to analyze the majoritarian propositionwise opinion diffusion method.

This paper poses a number of open questions, and suggests fascinating directions for future research. First, obtaining termination results for arbitrary integrity constraints, or characterising the set of constraints that guarantee termination on arbitrary networks, would be a major advancement. Techniques from finite Markov chains may be useful in such proofs (see, e.g., [21]). Second, the interplay between the properties of the aggregators, the structure of the integrity constraints, and the network, need to be investigated further. Third, issues of succinctness and computational complexity should be tackled. Once the integrity constraint is represented as a logical formula, a number of strategic questions related to influence maximisation may become intractable.

REFERENCES

- [1] K. R. Apt and E. Markakis. Diffusion in social networks with competing products. In *Algorithmic Game Theory*. Springer, 2011.
- [2] A. Baltag, Z. Christoff, J. U. Hansen, and S. Smets. Logical models of informational cascades. *Studies in Logic*, 47:405–432, 2013.
- [3] 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.
- [4] Z. Christoff and J. U. Hansen. A logic for diffusion in social networks. *Journal of Applied Logic*, 13(1):48–77, 2015.
- [5] M. H. DeGroot. Reaching a consensus. *Journal of the Americal Statistical Association*, 69(345):118–121, 1974.
- [6] F. Dietrich and C. List. Strategy-proof judgment aggregation. *Economics and Philosophy*, 23(3):269–300, 2007.
- [7] J. Dryzek and C. List. Social choice theory and deliberative democracy: a reconciliation. *British Journal of Political Science*, 2003.
- [8] D. Easley and J. Kleinberg. *Networks, Crowds, and Markets: Reasoning About a Highly Connected World*. Cambridge University Press, New York, NY, USA, 2010.
- [9] E. Elkind, P. Faliszewski, and A. Slinko. Swap bribery. In *Algorithmic Game Theory*, pages 299–310. Springer, 2009.
- [10] U. Endriss. Judgment aggregation. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 17. Cambridge University Press, 2016.
- [11] F. Farnoud, E. Yaakobi, B. Touri, O. Milenkovic, and J. Bruck. Building consensus via iterative voting. In *Proceedings of the 2013 IEEE International Symposium on Information Theory*, 2013.
- [12] N. E. Friedkin, A. V. Proskurnikov, R. Tempo, and S. Parsegov. Network science on belief system dynamics under logical constraints. *Science*, 354(6310):321–326, 2016.
- [13] 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.
- [14] M. Grabisch and A. Rusinowska. Different approaches to influence based on social networks and simple games. In A. Van Deemen and A. Rusinowska, editors, *Collective Decision Making*, volume 43 of *Theory and Decision Library C*, pages 185–209. Springer, 2010.
- [15] M. Grabisch and A. Rusinowska. A model of influence in a social network. *Theory and Decision*, 69(1):69–96, 2010.
- [16] U. Grandi and U. Endriss. Lifting integrity constraints in binary aggregation. *Artificial Intelligence*, 199-200:45–66, 2013.
- [17] U. Grandi, E. Lorini, and L. Perrussel. Propositional opinion diffusion. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2015.
- [18] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6):1420–1443, 1978.
- [19] D. Grossi and G. Pigozzi. *Judgment Aggregation: A Primer*. Morgan and Claypool, 2014.
- [20] M. O. Jackson and L. Yariv. Diffusion, strategic interaction, and social structure. In J. Benhabib, A. Bisin, and M. O. Jackson, editors, *Handbook of Social Economics*, volume 1, pages 645–678. North-Holland, 2011.
- [21] J. G. Kemeny and J. L. Snell. *Finite Markov Chains*. Springer-Verlag, 1976.
- [22] D. Kempe, J. M. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2003.
- [23] D. Kempe, J. M. Kleinberg, and E. Tardos. Influential nodes in a diffusion model for social networks. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP-2005)*, 2005.
- [24] J. Knight and J. Johnson. Aggregation and deliberation: On the possibility of democratic legitimacy. *Political Theory*, 22(2):277–296, 1994.
- [25] K. Lehrer and C. Wagner. *Rational Consensus in Science and Society*. Springer, 1981.
- [26] C. List. Group communication and the transformation of judgments: An impossibility result. *The Journal of Political Philosophy*, 19(1):1–27, 2011.
- [27] D. Miller. Deliberative democracy and social choice. *Political Studies*, 40(1):54–67, 1992.
- [28] F. Rossi, K. B. Venable, and T. Walsh. *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2011.
- [29] N. Schwind, K. Inoue, G. Bourgne, S. Konieczny, and P. Marquis. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [30] 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.
- [31] 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.
- [32] S. Simon and K. R. Apt. Social network games. *Journal of Logic and Computation*, 2015.
- [33] R. B. Wilson. On the theory of aggregation. *Journal of Economic Theory*, 10(1):89–99, 1975.