



UvA-DARE (Digital Academic Repository)

Pairwise Diffusion of Preference Rankings in Social Networks

Brill, M.; Elkind, E.; Endriss, U.; Grandi, U.

Publication date

2016

Document Version

Author accepted manuscript

Published in

Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence

[Link to publication](#)

Citation for published version (APA):

Brill, M., Elkind, E., Endriss, U., & Grandi, U. (2016). Pairwise Diffusion of Preference Rankings in Social Networks. In S. Kambhampati (Ed.), *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence: New York, NY, USA, 9-15 July, 2016* (Vol. 1, pp. 130-136). AAAI Press. <http://www.ilc.uva.nl/~ulle/pubs/files/BrillEtAlIJCAI2016.pdf>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Pairwise Diffusion of Preference Rankings in Social Networks

Markus Brill

University of Oxford
United Kingdom
mbrill@cs.ox.ac.uk

Edith Elkind

University of Oxford
United Kingdom
elkind@cs.ox.ac.uk

Ulle Endriss

University of Amsterdam
The Netherlands
ulle.endriss@uva.nl

Umberto Grandi

University of Toulouse
France
umberto.grandi@irit.fr

Abstract

We introduce a model of preference diffusion in which agents in a social network update their preferences based on those of their influencers in the network, and we study the dynamics of this model. Preferences are modelled as ordinal rankings over a finite set of alternatives. At each time step, some of the agents update the relative ordering of two alternatives adjacent in their current ranking with the majority view of their influencers. We consider both a synchronous and an asynchronous variant of this model. Our results show how the graph-theoretic structure of the social network and the structure of the agents' preferences affect the termination of the diffusion process and the properties of the preference profile at the time of termination.

1 Introduction

The development of models of opinion diffusion in a networked society is important for many reasons: such models enable us to obtain better predictions of electoral results, to maximise the effect of marketing campaigns, or to advance our understanding of the cognitive processes behind social influence. Social scientists have been the first to propose models of belief and opinion diffusion and their simple linear models have been deeply influential [de Groot, 1974; Lehrer and Wagner, 1981]. More complex models developed later on have borrowed heavily from related work in physics and biology [see, e.g., Jackson and Yariv, 2011].

In most of this work, individual opinions are modelled as real numbers or binary views on a given issue. For instance, models initially defined for the study of epidemics have been applied to the propagation of beliefs, by interpreting the likelihood of having a disease as an agent's degree of belief in some proposition. However, opinions can be complex objects that cannot always be accurately modelled by a single number, and this is particularly true for preferences. Indeed, the representation of preferences is a major research topic in Artificial Intelligence in its own right [see, e.g., Rossi *et al.*, 2011]. In this paper, we therefore ask how to model the diffusion of *structured* preference information across a network.

Our starting point is the intuition that agents update their preferences by aggregating the views of the agents they trust.

We model this trust relationship by means of a (directed) social network, where each agent trusts her influencers in the network. Furthermore, following standard conventions in voting theory [e.g., Taylor, 2005] we model preferences as linear orders over a set of alternatives. Preferences spread over the network by means of individual agents repeatedly reconsidering pairwise comparisons: an agent picks a pair of alternatives that are ranked one after the other in her preference order, and then compares her opinion on that pair to that of her influencers. If the agent's opinion differs from that of a majority of her influencers, then the pair is swapped. We call this process *pairwise preference diffusion*, and we study both its synchronous version, in which all agents make these updates simultaneously, and an asynchronous one, where only one agent at a time performs an update.

Our model extends that of Farnoud *et al.* [2013] who study iterated voting with agents updating their preferences on pairs of alternatives so as to agree with the *overall* majority. Their work can thus be viewed as a preference diffusion process on a complete network, where each agent is influenced by all members of the system. Our approach generalises their model to arbitrary social networks. We also draw inspiration from a recent line of work in multiagent systems, which studies diffusion processes based on complex representations of individual opinions, grounded in research on belief merging [Schwind *et al.*, 2015] and judgement aggregation [Grandi *et al.*, 2015]. A little further afield, related work modelling individually held points of view as numerical values includes the seminal contributions of Clifford and Sudbury [1973] and Holley and Liggett [1975], as well as more recent studies of linear threshold models [Kempe *et al.*, 2003; 2005] and refinements of the classic voter model [Raghavan *et al.*, 2007; Kearns and Tan, 2008; Hung *et al.*, 2013].

In this paper we identify conditions for a process of preference diffusion in our model to terminate, and we seek to characterise the profiles of preferences the process converges to in case of termination. Our results indicate that outcomes of a diffusion process depend both on the structure of the graph representing the social network and on the properties of the initial profile of preferences. After a review of basic definitions in Section 2, we formally introduce our model in Section 3. Section 4 presents our results on termination, Section 5 those on convergence, and Section 6 concludes. Some proofs are omitted due to space constraints.

2 Preliminaries

Let A be a finite set of *alternatives* and $N = \{1, \dots, n\}$ a finite set of *agents*. The preferences of agents over alternatives are recorded in a *preference profile* $\succ = (\succ_1, \dots, \succ_n)$, where each $\succ_i \subseteq A \times A$ is a (strict) *linear order* over A representing the preferences of agent i . For two distinct alternatives $a, b \in A$ we write $a \succ_i b$ if $(a, b) \in \succ_i$, and $a \not\succ_i b$ if $(a, b) \notin \succ_i$. Thus, $a \succ_i b$ signifies that agent i prefers alternative a to alternative b .

Furthermore, for $a \neq b$, we let ab denote the unordered pair $\{a, b\}$. A pair ab is *adjacent in* \succ_i if there is no $c \in A$ with $a \succ_i c \succ_i b$ or $b \succ_i c \succ_i a$. We will make frequent, often implicit, use of the following folklore result (for a proof, see e.g., the paper of Elkind *et al.* [2009]).

Lemma 1. *Let \succ and \succ' be two distinct linear orders on A . Then \succ and \succ' must disagree on at least one pair of alternatives that are adjacent in \succ .*

We introduce some further notation. In case ab is an adjacent pair in \succ_i , let $\text{swap}(\succ_i, ab)$ be the linear order resulting from swapping alternatives a and b in \succ_i . For instance, if $a \succ_i b \succ_i c$ and $\succ'_i = \text{swap}(\succ_i, bc)$, then $a \succ'_i c \succ'_i b$. If R is a binary relation on A and $a, b \in A$, the *restriction* $R|_{ab}$ of R on ab is defined as $R|_{ab} = R \cap \{(a, b), (b, a)\}$. If \succ is a linear order, then $\succ|_{ab}$ is always a singleton, and if R is asymmetric, then $R|_{ab}$ is either a singleton or the empty set.

Agents are connected by an *influence network* represented by a directed graph $E \subseteq N \times N$. An edge from j to i signifies that agent i is influenced by agent j . We let $\text{Inf}(i) = \{j \in N : (j, i) \in E\}$ denote the set of *influencers* of agent i . We often refer to agents in N as the *nodes* of the influence network E . A node i such that $\text{Inf}(i) = \emptyset$ is called a *source*, and S_E denotes the set of all sources of network E .

We restrict attention to graphs without self-loops, i.e., we assume $(i, i) \notin E$ for all $i \in N$.¹ Some of our results apply to specific classes of graphs: *directed acyclic graphs* (DAGs), the *complete graph* $E = \{(N \times N) \setminus \{(i, i) : i \in N\}\}$, and *simple cycles* $E = \{(1, 2), (2, 3), \dots, (n, 1)\}$.

3 The Model

In this section, we introduce our model and formally establish some of its basic properties.

3.1 Pairwise Preference Diffusion

Our diffusion model is based on the familiar *majority rule*. If \succ is a preference profile and $C \subseteq N$ a subset of agents, the (strict) *majority relation with respect to* C is defined by

$$a \succ_C b \iff |\{i \in C : a \succ_i b\}| > |\{i \in C : b \succ_i a\}|.$$

Thus, $a \succ_C b$ if a strict majority of agents in C prefer a to b . Majority relations are asymmetric ($a \succ_C b$ implies $b \not\succ_C a$), but, in contrast to linear orders, they may have ties and cycles. In the case of a *majority tie*, i.e., $|\{i \in C : a \succ_i b\}| = |C|/2$, we have $\succ_C|_{ab} = \emptyset$. This can only happen if $|C|$ is even.

Next, we define the basic update function underlying our diffusion process. Let \succ be a preference profile, $i \in N$, and

¹Thus, we assume that agents do not influence themselves. This simplifies exposition and does not significantly affect our results.

$a, b \in A$. The function PPD (“pairwise preference diffusion”) is defined as

$$\text{PPD}(\succ, i, ab) = \begin{cases} \text{swap}(\succ_i, ab) & \text{if } ab \text{ is adjacent in } \succ_i \\ & \text{and } \succ_{\text{Inf}(i)}|_{ab} \not\subseteq \succ_i|_{ab} \\ \succ_i & \text{otherwise.} \end{cases}$$

Given a profile, an agent i , and a pair of alternatives ab , the PPD update swaps a and b in \succ_i if ab is adjacent in \succ_i and a majority of i ’s influencers disagrees with i on ab . In case of a tie, i.e., when $\succ_{\text{Inf}(i)}|_{ab} = \emptyset$, i ’s preferences remain unchanged.² Note that PPD takes a profile as its (first) input, but returns only a single preference relation. Nevertheless, we will often speak of the new *profile* obtained as a result of an update. This is the profile that results from \succ by replacing \succ_i with $\text{PPD}(\succ, i, ab)$.

Let now $\text{turn} : \mathbb{N} \rightarrow 2^N \setminus \{\emptyset\}$ be a *turn function*, indicating which individuals in N are updating their preferences at each point in time, and let *pair* be a randomised procedure that, given a voter in N , outputs a pair of alternatives sampled from $A \times A$ uniformly at random. Starting from a preference profile \succ^0 , *pairwise preference diffusion* (PPD) is the following discrete-time iterative process:

$$\succ_i^{t+1} = \begin{cases} \text{PPD}(\succ^t, i, \text{pair}(i)) & \text{if } i \in \text{turn}(t) \\ \succ_i^t & \text{otherwise.} \end{cases}$$

An update from \succ^t to \succ^{t+1} is called *effective* if at least one agent performs a swap, i.e., if $\succ^{t+1} \neq \succ^t$. We define two diffusion processes. Under *synchronous* PPD the turn function is such that $\text{turn}(t) = N$ for all $t \in \mathbb{N}$, i.e., at every step all agents sample (possibly different) pairs of alternatives and update accordingly. Under *asynchronous* PPD the turn function selects only one agent uniformly at random at every time point $t \in \mathbb{N}$, i.e., one agent at a time samples a pair of alternatives and updates accordingly. See Figure 1 for an example.

3.2 Termination and Convergence

Both synchronous and asynchronous PPD are discrete-time dynamical processes, and a natural question is whether they terminate. A preference profile \succ is called *stable for* E if $\text{PPD}(\succ^t, i, ab) = \succ_i^t$ for all $i \in N$ and $a, b \in A$. Note that this definition of stability is independent of the turn function and thus not specific to either synchronous or asynchronous diffusion. Let $\mathbb{P}[X]$ denote the probability of event X under PPD. We can now define asymptotic termination.

Definition 2. *We say that PPD asymptotically terminates on a class of graphs $\mathcal{E} \subseteq 2^{N^2}$ if for each graph $E \in \mathcal{E}$ and for each profile of initial preferences \succ^0 we have*

$$\lim_{t \rightarrow +\infty} \mathbb{P}[\succ^{t+1} \neq \succ^t] = 0.$$

It will be useful to view PPD as a simple Markov chain, in which the states are all possible preference profiles, and PPD

²According to the definitions of Farnoud *et al.* [2013], an adjacent pair ab with $a \succ_i b$ is swapped if and only if $|\{i \in N : b \succ_i a\}| \geq |\{i \in N : a \succ_i b\}|$. If $\text{Inf}(i) = N \setminus \{i\}$, the latter is equivalent to $b \succ_{\text{Inf}(i)} a$. Thus, the setting of Farnoud *et al.* [2013] exactly corresponds to our setting for the special case of a complete graph.

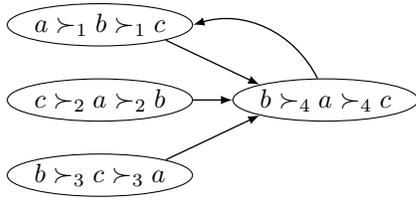


Figure 1: A simple influence network with $n = 4$ agents and $|A| = 3$ alternatives. Observe that the preferences of agents 1, 2, and 3 form a *Condorcet cycle*, i.e., the majority relation of their preferences is cyclic. In one possible sequence of updates for asynchronous PPD, agent 4 updates on pair ab , moving to preference $a \succ_4 b \succ_4 c$. After that, no further updates are possible: even though agent 4 disagrees with its influencers on pair ac , this pair cannot be swapped since it is no longer adjacent in \succ_4 . If we now consider synchronous PPD, and let agents 1 and 4 update repeatedly on pair ab , we obtain an infinite update sequence. Agents 2 and 3 are sources, and thus never update their preferences.

updates define transitions from one state to another. That is, in the case of asynchronous PPD, it is possible to move from state \succ to a state \succ' if and only if there exists an agent i and a pair ab such that $\succ'_i = \text{PPD}(\succ, i, ab)$ and $\succ'_j = \succ_j$ for all agents $j \neq i$. For synchronous PPD, we require that for each agent i there is a (possibly different) pair $a_i b_i$ such that $\succ'_i = \text{PPD}(\succ, i, a_i b_i)$. Transitions between *distinct* states correspond to effective PPD updates, and absorbing states, i.e., states without effective PPD updates, correspond to stable profiles. In Section 4 we will use the fact that PPD asymptotically terminates if and only if the respective Markov chain is *absorbing*, i.e., from each state there exists a sequence of PPD updates leading to a stable state [see, e.g., Kemeny and Snell, 1976].

We now define a second, stronger notion of termination.

Definition 3. We say that PPD universally terminates on a class of graphs \mathcal{E} if for each graph $E \in \mathcal{E}$ and each preference profile \succ on E there does not exist an infinite sequence of effective PPD updates starting from \succ .

Intuitively, Definition 3 says that, as long as we keep moving, we are guaranteed to reach a stable state, irrespective of the random choices made by *turn* and *pair*. Clearly, universal termination implies asymptotic termination. Next, we establish some further connections between the two concepts.

Lemma 4. If $|A| \geq 3$, then every sequence of asynchronous updates is also a sequence of synchronous updates.

Proof. We show that any asynchronous update can be simulated with a synchronous update. Let ab be the pair that is being swapped during the asynchronous update in \succ_i . For each agent $j \neq i$, we choose a pair that is nonadjacent in \succ_j (which is possible since $|A| \geq 3$), and we keep ab as the selected pair for i . Now we have $\succ_i^{t+1} = \text{PPD}(\succ^t, i, ab)$ and $\succ_j^{t+1} = \succ_j^t$ for all $j \neq i$. \square

As a direct consequence of Lemma 4, we have the following.

Proposition 5. Let $|A| \geq 3$. If synchronous PPD universally terminates on graph E , then so does asynchronous PPD. If

asynchronous PPD asymptotically terminates on E , then so does synchronous PPD.

Proof. Universal termination requires the *absence* of infinite sequences of effective PPD updates. The claim regarding universal termination thus follows immediately from Lemma 4. Asymptotic termination requires, for every profile \succ , the *presence* of a sequence of PPD updates from \succ to a stable state. Lemma 4 thus also establishes the claim regarding asymptotic termination. \square

We will now introduce further terminology to characterise the profiles encountered at termination. Let a *termination profile* for \succ^0 be a stable profile that is reachable from \succ^0 . We say that (asynchronous or synchronous) PPD *converges to a unique profile* for \succ^0 if it terminates when starting from \succ^0 and all termination profiles for \succ^0 coincide. We say that PPD *converges to consensus* for \succ^0 if it terminates when starting from \succ^0 and all termination profiles \succ^* are unanimous, i.e., $\succ_i^* = \succ_j^*$ for all $i, j \in N$. Observe that we do not require all termination profiles to be equal, thereby allowing for convergence to different consensus profiles.

4 Termination Results

In this section, we investigate how the structure of the influence network affects the termination of PPD.

4.1 Universal Termination

Universal termination is a very strong property, and we show that it can only be obtained by heavily restricting the class of graphs. As noted in Footnote 2, in the special case where the influence network E is a complete graph (without self-loops), our model coincides with the setting of Farnoud *et al.* [2013], who (implicitly) prove the following.

Proposition 6 (Farnoud *et al.*, 2013). *Asynchronous PPD universally terminates on complete graphs.*

Simple examples, such as that of two connected agents with opposing preferences $a \succ_1 b$ and $b \succ_2 a$, are sufficient to show that synchronous PPD does *not* universally terminate on complete graphs. However, if we restrict attention to acyclic networks (DAGs), we obtain universal termination for both synchronous and asynchronous PPD.

Proposition 7. *Synchronous PPD universally terminates on DAGs.*

The proof of Proposition 7 uses a potential function argument similar to that used by Farnoud *et al.* [2013]. Proposition 7 is tight, because—as we have seen already—synchronous PPD does not even terminate universally on a simple cycle with just two alternatives. Combining Proposition 7 with Proposition 5 we directly obtain the following result.

Corollary 8. *Asynchronous PPD universally terminates on DAGs.*

4.2 Asymptotic Termination

We now prove an asymptotic termination result for arbitrary networks. We need the following definition.

Definition 9. Let \succ^0 be a preference profile and E an influence network. The pair (\succ^0, E) satisfies the linear local majority property (LLM) if for all profiles \succ reachable from \succ^0 and for all $i \in N$ the relation $\succ_{\text{Inf}(i)}$ is a linear order.

Thus, we require the influence exercised over any agent in the network after any update sequence to be based on a linear order. Profiles satisfying domain restrictions such as those studied in Section 5.1 satisfy LLM on any network.

Theorem 10. Let \succ^0 denote the initial preference profile. For $|A| = 2$, asynchronous PPD asymptotically terminates on any graph E . For $|A| \geq 3$, asynchronous PPD asymptotically terminates on E when (\succ^0, E) satisfies LLM.

Proof sketch. For $A = \{a, b\}$, we define a two-phase process inspired by an idea of Chierichetti *et al.* [2013]. In the first phase, we consider all agents i with $b \succ_i a$ and $a \succ_{\text{Inf}(i)} b$, and we perform a PPD update on them. We iterate this process until all remaining agents with $b \succ_i a$ agree with their influencers. In the second phase, we consider agents j with $a \succ_j b$, and we iteratively update all those that disagree with a majority of their influencers. We claim that this process is guaranteed to reach a stable profile. To see this, let \succ' and \succ'' denote the profile after the first and after the second phase, respectively. We show that \succ'' is stable.

Assume for contradiction that there is an agent i such that $\succ''_{\text{Inf}(i)} \not\subseteq \succ''_i$. We distinguish two cases. If $a \succ''_i b$ (and thus $b \succ''_{\text{Inf}(i)} a$), then \succ_i would have been updated at the end of the second phase, a contradiction. If, on the other hand, $b \succ''_i a$ (and thus $a \succ''_{\text{Inf}(i)} b$), then at every time step during the second phase a majority of i 's influencers preferred a to b . Consequently, at the end of the first phase we already had $b \succ'_i a$ and $a \succ'_{\text{Inf}(i)} b$. But then \succ_i would have been updated to $a \succ_i b$, a contradiction. This proves that \succ'' is stable.

Now let $|A| \geq 3$, and assume that (\succ^0, E) satisfies LLM. Fix an arbitrary enumeration $p_1, \dots, p_{\lfloor |A|/2 \rfloor}$ of all unordered pairs of alternatives in A and consider the pairs in this order. Say the first pair is $p_1 = ab$. For this pair, we execute a generalised variant of the two-phase process defined above: In the first phase, we consider all agents i with $b \succ_i a$ and $a \succ_{\text{Inf}(i)} b$, and we repeatedly perform PPD updates on agent i until \succ_i is identical to $\succ_{\text{Inf}(i)}$ (we ‘‘copy’’ the entire relation $\succ_{\text{Inf}(i)}$ to agent i). We can do this because $\succ_{\text{Inf}(i)}$ is guaranteed to be a linear order by the LLM property. In particular, agent i prefers a to b after such an update. We iterate this process until all remaining agents with $b \succ_i a$ agree with their influencers on ab . In the second phase, we consider agents j with $a \succ_j b$ and iteratively copy the linear order $\succ_{\text{Inf}(j)}$ to agent j whenever $b \succ_{\text{Inf}(j)} a$.

After completing this two-phase process for $p_1 = ab$, we execute the two-phase process for p_2 , then for p_3 , and so on. We claim that, after all pairs have been considered, we have reached a stable profile. Say that a pair xy is *stable* if $\succ_i \upharpoonright_{xy} = \succ_{\text{Inf}(i)} \upharpoonright_{xy}$ for every agent i , and *unstable* otherwise. If all pairs are stable, the profile itself is stable (but note that the converse does not hold, since there may be unstable pairs in a stable profile, notably non-adjacent ones).

Now let xy be an arbitrary pair. After the two-phase process has completed for xy , we know (by the argument for

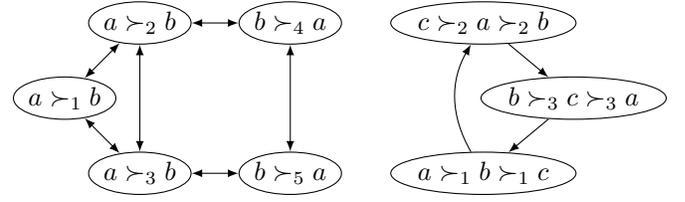


Figure 2: Left: A strongly connected influence network and a stable profile. PPD does not converge to consensus despite the fact that the majority relation is (trivially) acyclic. Right: A simple cycle influence network and a preference profile with a cyclic majority relation. PPD converges to consensus (see Section 5.2).

$|A| = 2$) that xy is stable. It remains to show that executing the two-phase process for subsequent pairs does not make xy unstable. But this is true simply because, once a pair is stable, *no* sequence of asynchronous PPD updates will lead to an agent performing an xy swap. Indeed, suppose that xy is stable and assume for contradiction that there is a sequence of PPD updates that swaps xy for some agent. Let t be the first time that happens, and let i be the agent whose preferences are swapped at time t . The swap can only be caused by some agent in $\text{Inf}(i)$ having swapped xy at an earlier time $t' < t$, contradicting the minimality of t . Therefore, at the end of the process all pairs are stable and so is the profile. \square

Proposition 5 implies asymptotic termination of *synchronous* PPD under the same conditions.

It is not clear whether Theorem 10 remains true when LLM is not satisfied. We leave this as an open problem for future work. A useful corollary of Theorem 10 concerns simple cycles, where LLM trivially holds.

Corollary 11. Asynchronous PPD asymptotically terminates on the class of simple cycles.

Asynchronous PPD does not *universally* terminate on simple cycles. To see this, take two alternatives a and b , and let all agents on a simple cycle of length $n \geq 3$ have preference $a \succ_i b$ except for one agent who prefers b to a . This situation allows for an infinite sequence of effective PPD updates, with the preference of the outlier moving around the cycle.

5 Convergence Results

In this section we study the properties of termination profiles. We consider two graph classes: directed acyclic graphs (Section 5.1) and simple cycles (Section 5.2).

Recall that when the influence network is complete, our model is equivalent to the one studied by Farnoud *et al.* [2013]. They show that, if n is odd, PPD converges to consensus if and only if the majority relation \succ_N is acyclic [Farnoud *et al.*, 2013, Theorem 8]. It is straightforward to check that acyclicity of \succ_N is neither necessary nor sufficient for convergence to consensus if the network is *incomplete* (see Figure 2). We first consider DAGs.

5.1 Directed Acyclic Graphs (DAGs)

For an acyclic network E , the sources \mathcal{S}_E play an important role because their preferences never change. We now study the influence that the preferences of the sources have on the termination profile. In particular, we consider the case where the preferences of the sources are in some sense “aligned,” and study whether at termination this will be reflected across the whole network. Such results are important when the diffusion process precedes the aggregation of individual preferences into a collective preference ordering, e.g., when taking a collective decision in an election. We identify conditions under which the asynchronous PPD diffusion process on a DAG with aligned sources leads to termination profiles that can be safely aggregated.³

The idea of aligned preferences can be formalised by restricting the set of preference relations that can co-occur in a profile, i.e., imposing a *domain restriction*. A very well-known example of a domain restriction is single-peakedness. A profile $\succ = (\succ_1, \dots, \succ_n)$ is *single-peaked* if there exists a linear order \triangleleft on A , called the *axis*, such that, for each $i \in N$, $a \triangleleft b \triangleleft c$ implies that $b \succ_i a$ or $b \succ_i c$. Single-peaked preference profiles have many desirable properties. For instance, it is well known that the majority relation \succ_N is acyclic (and thus gives rise to weak Condorcet winners) whenever preferences are single-peaked [Black, 1948]. Other well-known restricted domains are those of *single-crossing* preferences [Mirrlees, 1971; Roberts, 1977] and *value-restricted* preferences [Sen, 1966; Sen and Pattanaik, 1969] (we omit the definitions due to space constraints; see, e.g., [Gaertner, 2009]); a further (trivial) example is full consensus.

Now, it is straightforward to observe that if the sources form a consensus, i.e., if $\succ_i = \succ_j$ for all $i, j \in \mathcal{S}_E$, then PPD converges to consensus. What we want to show next is that more complex domain restrictions such as single-peakedness similarly spread to the whole network whenever they hold at the source nodes. A simple example shows that this is not possible if there can be majority ties among the influencers of a node (see Figure 3). For this reason, we need to assume that no majority ties occur during the PPD process. This assumption is formalised in the no-tie property.

Definition 12. Let \succ^0 be a preference profile and E an influence network. The pair (\succ^0, E) satisfies the no-tie property if, for all profiles \succ reachable from \succ^0 and for all $i \in N$, $\succ|_{\text{Inf}(i)}$ does not have ties, i.e., $\succ|_{\text{Inf}(i)}|_{ab} \neq \emptyset$ for all $a \neq b$.

The no-tie property generalises a common assumption in social choice theory—that of assuming that the number of decision-makers is odd—and can be regarded as natural when considering large enough networks.

Using this assumption, we will show that PPD propagates a number of domain restrictions, including, in particular, single-peakedness, single-crossingness, and value-restrictedness, to the whole network. We now define three properties a domain restriction may possess.

³This argument has been used in the literature on deliberative democracy to suggest that deliberation can be an effective process to obtain an “aggregatable” preference profile [see, e.g., Miller, 1992; Knight and Johnson, 1994; Dryzek and List, 2004].

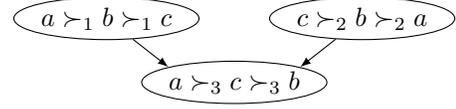


Figure 3: A stable profile \succ showing that single-peakedness is not propagated when there are ties in $\succ|_{\text{Inf}(i)}$. Profile \succ is not single-peaked, even though $\succ|_{\mathcal{S}_E}$ is.

Definition 13. Let D be a domain restriction.

- D is closed under subprofiles if, for each profile \succ satisfying D and for each $I \subseteq N$, the profile $\succ|_I = (\succ_i)_{i \in I}$ satisfies D .
- D is closed under majority if, for each profile \succ satisfying D and for each $I \subseteq N$ such that \succ_I is a linear order, the profile $\succ \cup \{\succ_I\} = (\succ_1, \dots, \succ_n, \succ_I)$ satisfies D .⁴
- D satisfies the Condorcet condition if, for each profile \succ satisfying D , the majority relation \succ_N is acyclic.

We are now ready to show that these properties are sufficient to guarantee PPD propagation from the sources.

Proposition 14. Let D be a domain restriction that is closed under subprofiles, closed under majority, and satisfies the Condorcet condition. Let E be a DAG and let \succ^0 be a profile such that (\succ^0, E) satisfies the no-tie property. If $\succ^0|_{\mathcal{S}_E}$ satisfies D , then all profiles at termination satisfy D .

Proof. Let the *depth* of node i be the length of a longest path from a source to i (this number is finite, as E is a DAG). We proceed by induction on the depth of a node. Nodes of depth zero are sources, so for them our claim is immediate.

Now, consider some $d > 0$ and assume that our claim is true for profiles formed by preferences of agents in nodes of depth $d - 1$ or less; we will show that it holds for profiles formed by preferences of agents in nodes of depth d or less. Consider an agent i at depth d and assume that her influencers $\text{Inf}(i)$ have stabilised. Note that each node in $\text{Inf}(i)$ has depth at most $d - 1$. Since D is closed under subprofiles and because of the induction hypothesis, the profile $\succ|_{\text{Inf}(i)}$ satisfies D . This in turn implies that the majority relation $\succ|_{\text{Inf}(i)}$ is a linear order, because (1) D satisfies the Condorcet condition and (2) $\succ|_{\text{Inf}(i)}$ does not have majority ties due to (\succ^0, E) satisfying the no-tie property. Thus, Lemma 1 implies that agent i ’s preference relation will eventually stabilise to $\succ|_{\text{Inf}(i)}$: as long as agent i disagrees with $\succ|_{\text{Inf}(i)}$, she will have an adjacent pair she can update. Our claim now follows from the assumption that D is closed under majority. \square

The requirements of Proposition 14 are satisfied by the three aforementioned domain restrictions: single-peakedness (SP), single-crossingness (SC), and value-restrictedness (VR).

Corollary 15. Let $D \in \{SP, SC, VR\}$. Consider a DAG E and a profile \succ^0 such that (\succ^0, E) satisfies the no-tie property. If $\succ^0|_{\mathcal{S}_E}$ satisfies D , then all profiles at termination satisfy D .

⁴Being closed under majority is a slightly more demanding requirement than the *closedness* property that is studied by Puppe and Slinko [2015] in the context of Condorcet domains.

Proof sketch. According to Proposition 14, it suffices to show that all three domain restrictions satisfy the Condorcet condition and are closed under subprofiles and majority.

It is well known that SP , SC , and VR satisfy the Condorcet condition and are closed under subprofiles. The proof that SP is closed under majority can be found, e.g., in Moulin’s textbook [Moulin, 2003, Exercise 4.6]. SC is closed under majority because the preferences of the median voter(s) agree with the majority relation (assuming the latter has no ties). Finally, we can show that VR is closed under majority by straightforward case analysis. \square

Recall that PPD universally terminates on a DAG. A consequence of the results shown above is that if the sources are aligned then PPD *converges to a unique profile*, which can be computed in polynomial time with a sequence of updates that stabilises each node in increasing order of their depth.

We conclude this discussion of the case of DAGs by stating a result that does not require the no-tie property. The following property is a strengthening of closedness under majority.

Definition 16. A domain restriction D is closed under majority influence if the following holds for all profiles \succ satisfying D and for all $I \subseteq N$. Let E^* be an influence network on $|I|$ nodes, consisting of $|I| - 1$ sources that all influence the single non-source node. Assign the preferences in $\succ|_I$ to the nodes in E^* in an arbitrary order. Consider the PPD process on E^* and let \succ_{sink} denote the preferences of the non-source node at termination. Then, the profile $\succ \cup \{\succ_{sink}\} = (\succ_1, \dots, \succ_n, \succ_{sink})$ also satisfies D .

Whenever the majority relation with respect to the source nodes in E^* is a linear order, \succ_{sink} coincides with this order and closedness under majority influence follows from closedness under majority (Definition 13). However, when there are majority ties (or even cycles), then closedness under majority influence is a stronger requirement.

The following result is analogous to Proposition 14, but instead of requiring the no-tie property, we need to start with a network in which the preferences of all nodes, not only those of the sources, satisfy the domain restriction.

Proposition 17. Let D be a restricted domain that satisfies the Condorcet condition and is closed under subprofiles and majority influence. Consider a DAG E and a profile \succ^0 . If \succ^0 satisfies D , then all profiles at termination satisfy D .

The proof of Proposition 17 is very similar to that of Proposition 14, with closedness under majority influence taking over the roles of closedness under majority and the no-tie property. Proposition 17 applies to single-peaked and to single-crossing preferences.

Lemma 18. SP and SC are closed under majority influence.

Corollary 19. Let $D \in \{SP, SC\}$ and let E be a DAG. If \succ^0 satisfies D , then all profiles at termination satisfy D .

5.2 Simple Cycles

We now investigate termination profiles for the special case of simple cycles. Lemma 1 implies that, if E is a simple cycle, then all stable profiles (and thus all termination profiles) are unanimous. We are going to characterise the set of linear orders that can be reached as the group consensus.

Clearly, this set contains all linear orders that occur in the initial profile \succ^0 , as we can simply “copy” preferences along an influence edge. Also observe that if \succ^0 is unanimous on a pair ab , then this unanimity will be preserved throughout the PPD process. We show that, for up to three alternatives, all linear orders that respect such initial unanimities can be obtained at termination.

Proposition 20. For $|A| \leq 3$, a simple cycle of arbitrary size can converge to any linear order that does not conflict with a unanimously accepted pair in \succ^0 (and only to those).

Proof sketch. The claim clearly holds for $|A| \leq 2$. For $A = \{a, b, c\}$, suppose w.l.o.g. that our goal is to reach consensus on $a \succ b \succ c$. Clearly, this is impossible if one of the comparisons $a \succ b$, $a \succ c$, or $b \succ c$ is unanimously rejected. It remains to show that convergence to $a \succ b \succ c$ is possible when none of the comparisons $a \succ b$, $b \succ c$, $a \succ c$ is unanimously rejected. Consider a greedy algorithm that repeatedly finds an agent i that disagrees with both $a \succ b \succ c$ and her influencer on some pair in $\{ab, bc, ac\}$ and updates i ’s preferences on that pair. This algorithm terminates after at most $3n$ updates; we show that it always results in consensus on $a \succ b \succ c$. To prove this, we assume for the sake of contradiction that the greedy algorithm gets stuck, and analyse what a “stuck” profile looks like. By enumerating all six possible preference orders an agent might have, we show that any stuck profile consists of the linear orders $c \succ a \succ b$ and $b \succ c \succ a$, in an alternating pattern. Both orders have $c \succ a$, and hence so do the initial preferences (since the greedy algorithm never reverses $a \succ c$), contradicting our assumption that $a \succ c$ is not unanimously rejected. \square

While it seems likely that Proposition 20 generalises to more than three alternatives, the proof technique employed above is not applicable to $|A| > 3$. We leave this as an open problem.

6 Conclusions

We have introduced a novel model of opinion diffusion on a social network that is tailored to the representation of individual preferences as linear orders over a finite set of alternatives. When prompted to reflect on the relative ordering of a pair of alternatives adjacent in her current preference ranking, an agent will swap them if a majority of her influencers disagree with the current ordering. At what point in time an agent is prompted to do so depends on the diffusion process, which may be synchronous or asynchronous. We have analysed two notions of termination for such processes, universal and asymptotic termination, and we have characterised the profiles of collective preferences at the time of termination for two classes of networks, namely directed acyclic graphs and simple cycles.

Modelling social influence as aggregation of the opinions of an agent’s influencers is a simple idea that opens up a number of interesting directions for future investigation. For instance, the field of voting theory provides a huge variety of aggregators that can be explored in this setting. Interesting questions include how to measure and compute the influence of a given agent, and how to efficiently infer properties of a system at the time of convergence when given its initial state.

Acknowledgments

This work has been partly supported by COST Action IC1205 on Computational Social Choice, the Labex CIMI project “Social Choice on Networks” (ANR-11-LABX-0040-CIMI), a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation, and ERC-StG 639945 (ACCORD).

References

- [Black, 1948] Duncan Black. On the rationale of group decision-making. *The Journal of Political Economy*, 56(1):23–34, 1948.
- [Chierichetti *et al.*, 2013] Flavio Chierichetti, Jon Kleinberg, and Sigal Oren. On discrete preferences and coordination. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC-2013)*, 2013.
- [Clifford and Sudbury, 1973] Peter Clifford and Aidan Sudbury. A model for spatial conflict. *Biometrika*, 60(3):581–588, 1973.
- [de Groot, 1974] Morris H. de Groot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [Dryzek and List, 2004] John S. Dryzek and Christian List. Social choice theory and deliberative democracy: A reconciliation. *British Journal of Political Science*, 34(04):752–758, 2004.
- [Elkind *et al.*, 2009] Edith Elkind, Piotr Faliszewski, and Arkadii Slinko. Swap bribery. In *Proceedings of the 2nd International Symposium on Algorithmic Game Theory (SAGT-2009)*, 2009.
- [Farnoud *et al.*, 2013] Farzad Farnoud, Eitan Yaakobi, Behrouz Touri, Olgica Milenkovic, and Jehoshua Bruck. Building consensus via iterative voting. In *Proceedings of the 2013 IEEE International Symposium on Information Theory*, 2013.
- [Gaertner, 2009] Wulf Gaertner. *A Primer in Social Choice Theory*. OUP Oxford, 2009.
- [Grandi *et al.*, 2015] Umberto Grandi, Emiliano Lorini, and Laurent Perrussel. Propositional opinion diffusion. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-2015)*, 2015.
- [Holley and Liggett, 1975] Richard A. Holley and Thomas M. Liggett. Ergodic theorems for weakly interacting infinite systems and the voter model. *The Annals of Probability*, 3(4):643–663, 1975.
- [Hung *et al.*, 2013] San-Chuan Hung, Perng-Hwa Kung, Shou-De Lin, Jing-Kai Lou, Chin-Hua Tsai, and Fu-Min Wang. Modeling the diffusion of preferences on social networks. In *Proceedings of the 13th SIAM International Conference on Data Mining*, 2013.
- [Jackson and Yariv, 2011] Matthew O. Jackson and Leeat Yariv. Diffusion, strategic interaction, and social structure. In Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, editors, *Handbook of Social Economics*, volume 1, pages 645–678. North-Holland, 2011.
- [Kearns and Tan, 2008] Michael Kearns and Jinsong Tan. Biased voting and the democratic primary problem. In *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE-2008)*, 2008.
- [Kemeny and Snell, 1976] John G. Kemeny and J. Laurie Snell. *Finite Markov Chains*. Springer-Verlag, 1976.
- [Kempe *et al.*, 2003] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-2003)*, 2003.
- [Kempe *et al.*, 2005] David Kempe, Jon M. Kleinberg, and Éva 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.
- [Knight and Johnson, 1994] Jack Knight and James Johnson. Aggregation and deliberation: On the possibility of democratic legitimacy. *Political Theory*, 22(2):277–296, 1994.
- [Lehrer and Wagner, 1981] Keith Lehrer and Carl Wagner. *Rational Consensus in Science and Society*. D. Reidel Publishing Company, Dordrecht, Holland, 1981.
- [Miller, 1992] David Miller. Deliberative democracy and social choice. *Political Studies*, 40:54–67, 1992.
- [Mirrlees, 1971] James A. Mirrlees. An exploration in the theory of optimum income taxation. *Review of Economic Studies*, 38(2):175–208, 1971.
- [Moulin, 2003] Hervé Moulin. *Fair Division and Collective Welfare*. The MIT Press, 2003.
- [Puppe and Slinko, 2015] Clemens Puppe and Arkadii Slinko. Condorcet domains, median graphs and the single crossing property. Technical report, <http://arxiv.org/abs/1507.08219>, 2015.
- [Raghavan *et al.*, 2007] Usha Nandini Raghavan, Réka Albert, and Soundar Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Physical Reviews E*, 76:036106, 2007.
- [Roberts, 1977] Kevin W.S. Roberts. Voting over income tax schedules. *Journal of Public Economics*, 8(3):329–340, 1977.
- [Rossi *et al.*, 2011] Francesca Rossi, Kristen Brent Venable, and Toby Walsh. *A Short Introduction to Preferences: Between Artificial Intelligence and Social Choice*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers, 2011.
- [Schwind *et al.*, 2015] Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, and Pierre Marquis. Belief revision games. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI-2015)*, 2015.
- [Sen and Pattanaik, 1969] Amartya K. Sen and Prasanta K. Pattanaik. Necessary and sufficient conditions for rational choice under majority decision. *Journal of Economic Theory*, 1:178–202, 1969.
- [Sen, 1966] Amartya K. Sen. A possibility theorem on majority decisions. *Econometrica*, 34(2):491–499, 1966.
- [Taylor, 2005] Alan D. Taylor. *Social Choice and the Mathematics of Manipulation*. Cambridge University Press, 2005.