

Iterative Judgment Aggregation

Marija Slavkovik¹ and Wojciech Jamroga²

Abstract. Judgment aggregation problems form a class of collective decision-making problems represented in an abstract way, subsuming some well known problems such as voting. A collective decision can be reached in many ways, but a direct one-step aggregation of individual decisions is arguably most studied. Another way to reach collective decisions is by iterative consensus building – allowing each decision-maker to change their individual decision in response to the choices of the other agents until a consensus is reached. Iterative consensus building has so far only been studied for voting problems. Here we propose an iterative judgment aggregation algorithm, based on movements in an undirected graph, and we study for which instances it terminates with a consensus. We also compare the computational complexity of our iterative procedure with that of related judgment aggregation operators.

1 Introduction

Social choice aggregation methods, such as voting [32], are of interest to artificial intelligence as methods for collective decision-making among humans and automated agents alike [3]. Judgment aggregation problems [26] are problems of aggregating individual judgments on a fixed set of logically related issues, called an agenda. Intuitively, an issue is a question that can be answered “yes” or “no.” Alternatively, an issue is a proposition that can be accepted or rejected. A judgment is a consistent collection of “yes” and “no” answers, one for each issue. Judgment aggregation has been used to model collective decision-making in multi-agent systems [2, 39, 41]. It is also interesting because it generalises voting problems, i.e., problems of choosing one option from a set of available options by aggregating agents’ individual preferences over these options. A voting problem can be represented as a judgment aggregation problem under some mild conditions, see e.g. [6, 23].

Aggregation methods produce a joint decision for a group of agents by aggregating the set of their individual decisions, called a profile, using an aggregation operator. Another approach to collective decision-making is deliberation, when agents negotiate which decisions to make. In multi-agent systems, deliberation procedures are constructed using an abstract argumentation framework to model relations between decisions, cf. e.g. [35, 20]. A third, comparatively less explored method to reach collective decisions is by *iterative consensus building*: each agent starts with an individual decision which she then iteratively changes in response to the individual decisions of the other agents until all agents end up supporting the same decision, i.e., until a *consensus* is reached. While in standard aggregation all individual decisions are elicited once, forming a profile, and after the

elicitation the agents can no longer change the profile, an iterative procedure allows agents to change their decisions many times, even back and forth.

The existence of judgment transformation functions, i.e., functions that transform one profile of individual judgments into another profile of individual judgments (possibly towards consensus) has been considered by List [25]. It was shown that under a set of reasonable and minimal desirable conditions no transformation function can exist. Social choice aggregation theory is rife with impossibility results such as this, yet few specific aggregation operators (that violate some of the desirable conditions) are proposed. There are more voting operators than judgment aggregation operators, which is unsurprising since voting is a much older discipline, but the number of judgment aggregation operators is also on the rise. Those include: quota-based rules [7, 9], distance-based rules [11, 13, 14, 28, 37, 40], generalisations of Condorcet-consistent voting rules [21, 29, 30], and rules based on the maximisation of some scoring function [5, 21, 42]. Deliberation and iterative consensus reaching procedures for voting problems have been explored, e.g., in [27, 24, 17, 18, 34]. However, to the best of our knowledge, there are no iterative procedures for aggregating judgments. With this work we aim to fill in the gap.

We consider all possible judgments for an agenda as vertices in a graph. The existence of an edge between judgments in the graph depends on the relations between the truth-value assignments on the same issue in the connected judgments. We define three intuitive agenda graphs. We design an iterative consensus building algorithm which reaches consensus in the following way: In the first step of the algorithm, each agent chooses a vertex and lets the other agents know what she has chosen. In each subsequent step each agent independently from the other agents moves to an adjacent vertex if this move reduces her path distance to the other vertices in the profile. The agents are only allowed to move along a shortest path towards some other agent. The moving continues until a consensus is reached (i.e., when all agents end up on the same vertex). We then exploit properties of graphs to study for which subgraphs corresponding to a profile of judgments the algorithm terminates with a consensus.

Judgment aggregation operators suffer from two shortcomings. First, they are often irresolute, i.e., more than one collective decision is produced. Unlike in voting, tie-breaking in judgment aggregation is not straightforward and virtually unexplored. Secondly, deciding if a given judgment is among the possible “winners” of the aggregation is often intractable [23, 12]. The set of tractable aggregation functions is very limited, exceptions being [7, 9, 13]. An iterative procedure clearly avoids ties when it reaches a consensus, and this is one advantage of our proposal. We also show that our consensus-oriented procedure may offer some computational benefits when compared to standard judgment aggregation rules.

The motivation for our iterative procedure is both descriptive and prescriptive. On one hand, our algorithm is meant to approximate

¹ Department of Information Science and Media Studies, University of Bergen, Norway, email: Marija.Slavkovik@infomedia.uib.no

² Institute of Computer Science, Polish Academy of Sciences, email: w.jamroga@ipipan.waw.pl

	φ_1	φ_2	φ_3
agent 1	0	1	0
agent 2	1	0	0
agent 3	1	1	1
majority	1	1	0

Figure 1. Doctrinal paradox

consensus formation that happens in human societies. On the other hand, our procedure can be useful for implementing artificial agents, as producing a consensual judgment is in some cases distinctly cheaper than computing the collective opinion in one step elicitation by a standard judgment aggregation procedure.

The paper is structured as follows. In Section 2 we introduce the judgment aggregation framework. In Section 3 we define three relevant agenda graphs and recall some useful concepts from graph theory. In Section 4 we present the algorithm for iterative judgment aggregation, and discuss necessary conditions for its termination with a consensus. In Section 5 we investigate sufficient termination conditions for each of the three agenda graphs. In Section 6 we briefly discuss the quality of the reached consensus with respect to some judgment aggregation operators, and study the computational complexity of the algorithm. In Section 7 we discuss the related work. In Section 8 we present our conclusions and discuss future work.

2 Preliminaries

We first introduce the basic definitions of judgment aggregation.

Judgments. Let \mathcal{L} be a set of propositional variables. An *agenda* $\mathcal{A} = \{\varphi_1, \dots, \varphi_m\}$ is a finite set $\mathcal{A} \subseteq \mathcal{L}$. The elements of \mathcal{A} are called *issues*. A *judgment* is a (truth assignment) function $J : \mathcal{A} \rightarrow \{0, 1\}$ mapping each issue to either 0 (reject) or 1 (accept). We write $\{0, 1\}^{\mathcal{A}}$ as a shorthand for $\mathcal{A} \rightarrow \{0, 1\}$, the space of all possible judgments for m issues, i.e., all sequences of length m comprised of 0s and 1s. We use $J(\varphi)$ to denote the value assigned to $\varphi \in \mathcal{A}$. The Hamming distance between two judgments is defined as the number of issues on which the judgments differ, i.e., $d_h(J^a, J^b) = \#\{\varphi \in \mathcal{A} \mid J^a(\varphi) \neq J^b(\varphi)\}$.

With each agenda we associate a *constraint* $\Gamma \in \mathcal{L}_{\mathcal{A}}$, where $\mathcal{L}_{\mathcal{A}}$ is the set of well formed formulas constructed with variables from \mathcal{A} and the logical connectives \neg (negation), \wedge (conjunction), \vee (disjunction), and \rightarrow (implication). The formula Γ is assumed not to be a contradiction. A judgment from $\{0, 1\}^{\mathcal{A}}$ is *rational* for Γ if and only if it is a model for Γ in the sense of classical propositional logic. For a given $\Gamma \in \mathcal{L}_{\mathcal{A}}$, we define $\mathcal{J}_{\mathcal{A}, \Gamma} \subseteq \{0, 1\}^{\mathcal{A}}$ to be the set of all *rational judgments* for \mathcal{A} and Γ .

Agents and profiles. Let $N = \{1, 2, \dots, n\}$ be a finite set of *agents*. A profile $P = (J_1, \dots, J_i, \dots, J_n) \in \mathcal{J}_{\mathcal{A}, \Gamma}^n$ is a list of rational judgments, one for each agent. We denote $P[i] = J_i$ and $P_{-i} = (J_1, \dots, J_{i-1}, J_{i+1}, \dots, J_n)$. Further let $\{P\}$ be the set of all distinct judgments that are in P . We often abuse notation and write $J_i \in P$ when $P[i] = J_i$. We reserve subscripted judgments, e.g., J_i , to denote judgments that belong to some profile and the superscripted judgments, e.g., J^a, J^b , to denote rational judgments that may not belong to some profile. A profile is *unanimous* if $\{P\} = \{J\}$, for some $J \in \mathcal{J}_{\mathcal{A}, \Gamma}$. A judgment J is a *plurality judgment* in P if and only if $\#\{i \mid P[i] = J\} \geq \#\{i \mid P[i] \neq J\}$.

Example 1 The quintessential example in judgment aggregation is the “doctrinal paradox” which is described with $\mathcal{A} = \{\varphi_1, \varphi_2, \varphi_3\}$ and $\Gamma = (\varphi_1 \wedge \varphi_2) \leftrightarrow \varphi_3$. The $\mathcal{J}_{\mathcal{A}, \Gamma} =$




Figure 2. Agenda graphs for the doctrinal paradox: (a) the Hamming graph $G_{\mathcal{A}}^h$, (b) the Model graph $G_{\mathcal{A}, \Gamma}^m$, (c) the Complete graph $G_{\mathcal{A}, \Gamma}^c$ for \mathcal{A} and Γ in Example 1.

$\{(0, 0, 0), (0, 1, 0), (1, 0, 0), (1, 1, 1)\}$. The doctrinal paradox profile is $P = ((0, 1, 0), (1, 0, 0), (1, 1, 1))$, see also Figure 1. Note that all the three profile judgments are rational, but the collective judgment obtained by taking the value for each issue assigned by a strict majority of agents, the so called majority rule, is not rational.

3 Agenda Graphs

An *agenda graph* is a graph $G_{\mathcal{A}, \Gamma} = \langle V, E \rangle$ whose nodes are judgments for some agenda \mathcal{A} , namely $V \subseteq \{0, 1\}^{\mathcal{A}}$. Given an agenda \mathcal{A} and constraints Γ we define three agenda graphs. The *Hamming graph* $G_{\mathcal{A}}^h$ is the graph over all possible (not necessarily rational!) judgments, that connects vertices which differ on exactly one issue. The *Model graph* $G_{\mathcal{A}, \Gamma}^m$ is the graph over all *rational* judgments, where two judgments are adjacent iff no “compromise” exists between them. The *Complete graph* $G_{\mathcal{A}, \Gamma}^c$ is the fully connected graph over all rational judgments. Formally:

The **Hamming graph** is $G_{\mathcal{A}}^h = \langle \{0, 1\}^{\mathcal{A}}, E^h \rangle$ where $(J^a, J^b) \in E^h$ iff $d_h(J^a, J^b) = 1$.

The **Model graph** is $G_{\mathcal{A}, \Gamma}^m = \langle \mathcal{J}_{\mathcal{A}, \Gamma}, E^m \rangle$ where $(J^a, J^b) \in E^m$ iff there exists no judgment $J^c \in \mathcal{J}_{\mathcal{A}, \Gamma}$ between J^a and J^b . A judgment $J^c \in \mathcal{J}_{\mathcal{A}, \Gamma}$ is between judgments $J^a \in \mathcal{J}_{\mathcal{A}, \Gamma}$ and $J^b \in \mathcal{J}_{\mathcal{A}, \Gamma}$ when $J^c \neq J^a, J^c \neq J^b, J^a \neq J^b$ and for every $\varphi \in \mathcal{A}$ if $J^a(\varphi) = J^b(\varphi)$, then $J^c(\varphi) = J^a(\varphi) = J^b(\varphi)$.

The **Complete graph** is $G_{\mathcal{A}, \Gamma}^c = \langle \mathcal{J}_{\mathcal{A}, \Gamma}, E^c \rangle$ where $E^c = \mathcal{J}_{\mathcal{A}, \Gamma} \times \mathcal{J}_{\mathcal{A}, \Gamma}$.

The agenda graphs for the doctrinal paradox of Example 1 are shown in Figure 2.

We use d_x to denote the (shortest) path distance on an agenda graph $G_{\mathcal{A}, \Gamma}^x$. The path distance on $G_{\mathcal{A}, \Gamma}^c$ is also known as the *drastic distance*: the distance between two judgments is 0 iff they are the same on all issues and 1 iff they differ on at least one issue. The path distance on $G_{\mathcal{A}}^h$ is the Hamming distance, and the path distance on $G_{\mathcal{A}, \Gamma}^m$ is the Model distance introduced and formally characterized in [11]. Recall that a path distance on a graph $G_{\mathcal{A}, \Gamma}$, as on any graph, is a distance function in the true sense since for every $J^a, J^b, J^c \in V$ it satisfies: $d(J^a, J^b) = 0$ iff $J^a = J^b$, $d(J^a, J^b) = d(J^b, J^a)$, and $d(J^a, J^c) \leq d(J^a, J^b) + d(J^b, J^c)$ (triangle inequality).

A graph $G' = \langle V', E' \rangle$ is a subgraph of graph $G = \langle V, E \rangle$, denoted $G' \subseteq G$, if $V' \subseteq V$ and $E' \subseteq E$. The V' -induced subgraph of a graph G is the graph $G' \subseteq G$ with vertices V' and edges E' which satisfies that, for every pair of vertices in V' , they are adjacent in G if and only if they are adjacent in G' .

We consider profile-induced subgraphs of $G_{\mathcal{A}, \Gamma}$ and make use of their “geometry”. Therefore we define some useful concepts for a

given agenda graph $G_{\mathcal{A},\Gamma} = \langle V, E \rangle$ and agents $N = \{1, 2, \dots, n\}$, following the terminology from graph theory [36].

The interval between a pair of vertices $J^a, J^b \in V$, denoted $I_{\mathcal{A},\Gamma}[J^a, J^b]$, is the set of all the judgments on all the shortest paths in $G_{\mathcal{A},\Gamma}$ from J^a to J^b .

A subset $S \subseteq V$ is convex if it is closed under $I_{\mathcal{A},\Gamma}$, namely when it includes all shortest paths between two vertices in S .

The convex hull of P on $G_{\mathcal{A},\Gamma}$, denoted $\text{CH}(P)$, is the smallest convex subset of V from $G_{\mathcal{A},\Gamma}$ that contains $\{P\}$.

The eccentricity of a judgment $J^a \in S \subseteq V$ is $e_S(J^a) = \max\{d(J^a, J^b) \mid J^b \in S\}$, i.e., the farthest that J^a gets from any other judgment in S .

A diameter of a set $S \subseteq V$ is $\text{max}_d(S) = \max\{e_S(J) \mid J \in S\}$, namely the maximal eccentricity of a vertex in S . All judgments for which $\text{max}_d(S) = e_S(J)$ are called peripheral judgments for S . For $S = \{P\}$ we call these judgments peripheral judgments of a profile P . If for two judgments J^a, J^b it holds that $d(J^a, J^b) = \text{max}_d(P)$, then these are called antipodal judgments of a profile P .

Example 2 Consider the Hamming graph $G_{\mathcal{A}}^h$ for the doctrinal paradox, presented in Figure 2a, and take the profile P as defined in Example 1. The $\text{CH}(P)$ -induced subgraph of $G_{\mathcal{A}}^h$ is given in Figure 3; this graph contains all the shortest paths, and nodes from $\mathcal{J}_{\mathcal{A},\Gamma}$, between $(0, 1, 0)$ and $(1, 0, 0)$, between $(0, 1, 0)$ and $(1, 1, 1)$, and between $(1, 0, 0)$ and $(1, 1, 1)$. Node $(0, 0, 1)$ is not in this $\text{CH}(P)$ -induced subgraph of $G_{\mathcal{A}}^h$ because this node is not on any of the shortest paths between the profile judgments.




Figure 3. The $\text{CH}(P)$ -induced subgraph of $G_{\mathcal{A}}^h$ for P in Example 1.

We have defined the convex hull of P because we will build our iteration algorithm on the principle of only allowing the agents to move from their current judgment to an adjacent judgment in the hull of P . By doing this we ensure that the agents do not disperse away from each other. The Proposition 1 shows that if agents in profile P each move to a judgment in $\text{CH}(P)$, thus creating profile P' , the diameter of the P' profile cannot be bigger than that of the P .

Proposition 1 For an agenda graph $G_{\mathcal{A},\Gamma}$ and profile $P \in \mathcal{J}_{\mathcal{A},\Gamma}^n$, if $S \subseteq \text{CH}(P)$, then $\text{CH}(S) \subseteq \text{CH}(P)$ and $\text{max}_d(\text{CH}(S)) \leq \text{max}_d(\text{CH}(P))$.

Proof This property follows from the fact that CH is a finitary closure operator [36, p. 6, Theorem 1.3]. Thus, for the convex hull $\text{CH}(P)$, it holds that if $S \subseteq \text{CH}(P)$, then $\text{CH}(S) \subseteq \text{CH}(\text{CH}(P))$, and $\text{CH}(\text{CH}(P)) = \text{CH}(P)$. ■

Definition 1 We say that the profile P has a k -cycle in $G_{\mathcal{A},\Gamma}$ if and only the $\text{CH}(P)$ -induced subgraph of $G_{\mathcal{A},\Gamma}$ has a simple cycle of length k . We say that P is a k -cycle in $G_{\mathcal{A},\Gamma}$ if and only if

the $\text{CH}(P)$ -induced subgraph of $G_{\mathcal{A},\Gamma}$ is a simple cycle in $G_{\mathcal{A},\Gamma}$ of length k .

The doctrinal paradox profile from Example 1 is a 4-cycle in $G_{\mathcal{A},\Gamma}^m$, and it has a 6-cycle in $G_{\mathcal{A}}^h$, as can be inferred from Figure 2.

We make the following observation.

Observation 2 No profile $P \in \mathcal{J}_{\mathcal{A},\Gamma}^n$ has a 3-cycle in $G_{\mathcal{A}}^h$ or in $G_{\mathcal{A},\Gamma}^m$, regardless of \mathcal{A} and Γ .

This exclusion of 3-cycles is due to the betweenness property of path distances: if J^b is between J^a and J^c in the graph, then $d(J^a, J^c) = d(J^a, J^b) + d(J^b, J^c)$.

Definition 2 An interval $I_{\mathcal{A},\Gamma}[J^a, J^b]$ is ϵ -connected in $G_{\mathcal{A},\Gamma}$ if and only if the maximal path distance between two adjacent rational judgments in the $I_{\mathcal{A},\Gamma}[J^a, J^b]$ -induced subgraph of $G_{\mathcal{A},\Gamma}$ is at most $\epsilon \in \mathbb{N}$. A profile P is ϵ -connected in $G_{\mathcal{A},\Gamma}$ if and only if every interval $I_{\mathcal{A},\Gamma}[P[i], P[j]]$ for $i, j \in N$ is ϵ -connected in $G_{\mathcal{A},\Gamma}$.

While every interval in $G_{\mathcal{A},\Gamma}^c$ and $G_{\mathcal{A},\Gamma}^m$ is ϵ -connected for every $\epsilon \geq 1$, this may not be the case for intervals in $G_{\mathcal{A}}^h$. As an example, consider the doctrinal paradox profile P from Example 1 and graph $G_{\mathcal{A}}^h$ in Figure 2a. Here, the interval $I_{\mathcal{A},\Gamma}[(1, 1, 1), (0, 1, 0)]$ is 2-connected because every shortest path between $(1, 1, 1)$ and $(0, 1, 0)$ passes only through judgments that are not rational.

4 Iteration Algorithm

Collective opinions in human societies are often formed not in one step, but rather in an intricate process that involves mutual information exchange, argumentation, persuasion, and opinion revision. Typically, social agents are motivated by two somewhat conflicting needs: on one hand, they want to form a unified stance with a significant enough part of the community; on the other hand, they do not want to concede too much of their own opinion. Here, we try to mimic this kind of behaviour – obviously, in a very simplified way. To this end, we design an iteration algorithm, Algorithm 1, based on an agenda graph $G_{\mathcal{A},\Gamma}$. As it is standard in judgment aggregation, we assume that the agents can only chose rational judgments at each iteration step. For an agenda \mathcal{A} and constraints Γ , each agent’s judgment is a node in the graph $G_{\mathcal{A},\Gamma}$. In the first step of the iterative procedure each agent announces which node she has chosen. Two or more agents may choose the same node. The agents do not have a constructed $G_{\mathcal{A},\Gamma}$ available. At each subsequent step, the agents (simultaneously) compute their adjacent nodes and try to “move” to one of these adjacent nodes along some shortest path towards the other agents. A move is possible if and only if the adjacent judgment is rational and it brings the agent closer to the rest of the profile, i.e., it decreases its aggregated path distance to the other judgments. More precisely, an agent $i \in N$ will move from $P[i]$ to a J iff there exist a rational $J \in \text{CH}(P)$ s.t. $d(J, P[i]) = 1$ and $\sum_{j \in N, j \neq i} d(P[i], P[j]) < \sum_{j \in N, j \neq i} d(J, P[j])$, where d is a path distance for a given $G_{\mathcal{A},\Gamma}$. Given a choice between two moves the agent chooses the one which better reduces the distance to the rest of the profile. If more than one move reduces the distance to the same extent, the agent chooses using some internal decision-making procedure which we do not model. We take it that in this case the agent chooses non-deterministically, with all move options being probable. The agents continue moving along the nodes of $G_{\mathcal{A},\Gamma}$ until no agent in the profile can move, or all of the agents “sit” on the same node, namely until a consensus is reached.

Data: $\epsilon > 0$, $\mathcal{J}_{\mathcal{A},\Gamma}$, N , own identifier $i \in N$, initial judgment

J_i^0

Result: $P \in \mathcal{J}_{\mathcal{A},\Gamma}^{|N|}$

$t := 0$; MOVES := \emptyset ; $J_i := J_i^0$; $P :=$ empty list;

repeat

$P' := P$;

$P[i] := J_i$;

for $j \in N, j \neq i$ **do**

 | $send(J_i, j), receive(J_j, i), P[j] := J_j$;

end

 MOVES := $\underset{J \in \text{CH}(P) \cap \mathcal{J}_{\mathcal{A},\Gamma}}{\text{argmin}} D(i, J, P) \cap$

$\{J \mid 0 < d(J, J_i) \leq 1 \text{ and } D(i, J, P) < D(i, P[i], P)\}$;

if MOVES $\neq \emptyset$ **then**

 | select $J \in \text{MOVES}, J_i := J$;

end

$t := t + 1$;

until P is unanimous or $P' = P$;

return P ;

Algorithm 1: Iteration algorithm

In Algorithm 1, $send(J_i, j)$ informs agent $j \in N$ that agent $i \in N$ has chosen to move to node $J_j \in V$, while $receive(J_j, i)$ denotes that the agent $i \in N$ has been informed that agent $j \in N$ has chosen to move to node $J_j \in V$. To ease readability we use $D(i, J, P) = \sum_{j \in N, j \neq i} d(J, P[j])$. We call $D(i, J, P)$ the distance of

J to the profile P_{-i} . In Algorithm 1, at each iteration t , MOVES is the set of judgments that are strictly closer to P than the current judgment J_i . Note that the algorithm is fully decentralised, in the sense that there is no need for any central authority to take over the iteration at any point of the process.

The starting profile P^0 collects the initial individual opinions of the agents. That is, it is the profile that would be aggregated under classical one-step social choice. We say that the algorithm *reaches consensus* J for P^0 , if it terminates starting from P^0 by returning the unanimous profile $\{P\} = \{J\}$. We first observe a necessary condition for reaching consensus.

Proposition 3 *If Algorithm 1 reaches consensus, then P^0 is ϵ -connected for $\epsilon = 1$.*

Proof Assume that the algorithm terminates with a J^* -unanimous P at some t . In each $t' < t$, every agent either keeps her own judgment $P[i]$, or moves to a new $J \in \mathcal{J}_{\mathcal{A},\Gamma}$ with $d(P[i], J) = 1$. Since a J^* is reached by every agent, there must exist a 1-connected path between any two judgments in P^0 . Thus P^0 must be 1-connected. ■

Note also that if P^0 is ϵ -connected, so is P at any step $t > 0$. The interesting question is: what are the sufficient conditions for reaching consensus by Algorithm 1? We address the question in Section 5.

5 Reaching Consensus

In this section, we examine the sufficient conditions for reaching consensus by Algorithm 1. We begin by looking at the iteration over the complete agenda graph $G_{\mathcal{A},\Gamma}^c$, and then we move on to the more interesting cases of $G_{\mathcal{A}}^h$ and $G_{\mathcal{A},\Gamma}^m$.

5.1 Iteration with $G_{\mathcal{A},\Gamma}^c$

Theorem 4 *If P contains a unique plurality judgment J , then Algorithm 1 always reaches consensus in one step.*

Proof On $G_{\mathcal{A},\Gamma}^c$, the path distance d_e between any two judgments that are different is 1. Let J be the unique plurality judgment in




Figure 4. An agenda graph of a non-equidistant profile $P = (J_1, J_2, J_3)$

$P \in \mathcal{J}_{\mathcal{A},\Gamma}^n$, selected by k agents. For every $J_i = J \neq J_j$, we have $D(i, J, P) = n - k \leq D(i, P[j], P)$, so the agents selecting J can not change their judgments. Moreover, switching from J_j to J decreases the distance to P_{-j} most, so all the other agents will switch to J in the first iteration. ■

What about starting profiles with several plurality judgments? They converge towards consensus under reasonable conditions.

Theorem 5 *If N consists of an odd number of agents, then Algorithm 1 probabilistically reaches consensus, i.e., it reaches consensus with probability 1.*

Proof If there is a single plurality judgment in P , then the algorithm converges in one step. If there are two or more plurality judgments $J^1, \dots, J^k \in \mathcal{J}_{\mathcal{A},\Gamma}$, then those agents swap non-deterministically between J^1, \dots, J^k , and the other agents move to one of J^1, \dots, J^k . In the next round, the same argument applies. Eventually, P converges either to the unanimous profile P' such that $\{P'\} = \{J\}$ for some $J \in \{J^1, \dots, J^k\}$, or to a profile P'' such that $\{P''\} = \{J^1, \dots, J^m\}$ for an odd m , each favoured by the same amount of agents. From then on, the agents keep swapping judgments until one judgment gets plurality in the profile, and wins in the next round.

Formally, let $\text{MOVES}_{i,t}$ be the set of moves available to agent i at the step t of Algorithm 1. We assume that there is some $\delta > 0$ such that, for each step t , agent i selects judgment $J \in \text{MOVES}_{i,t}$ with probability $p_i(J) \geq \delta$. Then, there exists $\delta' > 0$ such that the probability of all the agents “hitting” a profile with no plurality in the next round is at most $1 - \delta'$. Hence, the probability that the profile stays with no plurality in m steps is at most $(1 - \delta')^m$, which converges to 0 as m increases. ■

The Algorithm 1 has good convergence properties on $G_{\mathcal{A},\Gamma}^c$ but the consensus it reaches is limited to the judgments that are already in the starting profile. On the Hamming and Model agenda graphs Algorithm 1 surpasses this limitation. However, its convergence becomes a subtler issue.

5.2 Iteration with $G_{\mathcal{A}}^h$ and $G_{\mathcal{A},\Gamma}^m$

In this section, we will use $G_{\mathcal{A},\Gamma}$ to refer to one of $G_{\mathcal{A},\Gamma}^m, G_{\mathcal{A}}^h$ in order to avoid stating and proving the same properties for $G_{\mathcal{A},\Gamma}^m$ and $G_{\mathcal{A}}^h$ separately when the same proof can be applied.

We start with a negative result. Let us call *equidistant* those profiles P such that, for any $i, j, r \in N$ with $i \neq j, i \neq r, j \neq r$, it holds that $d(P[i], P[j]) = d(P[i], P[r])$.

Proposition 6 *Consider $N = \{1, 2, 3\}$ agents and a 1-connected P^0 . If Algorithm 1 reaches a P that is an equidistant k -cycle, then Algorithm 1 will not terminate with a consensus.*

Proof If P is an equidistant k -cycle then no agent can reduce the distance to one agent by 1 without increasing the distance to the other agent by 1. Thus no agent has a possible move. ■

Note that the same applies to some non-equidistant k -cycles. For example, the profile in Figure 4 is not equidistant, but it is easy to check that each agent has an empty set of moves.

3-agent profiles that form a simple cycle are problematic because an agent may not be able to get closer to one of the other agents without distancing itself from the third. For profiles of more than three agents that form a simple cycle, the judgments cannot be equidistant, and this is no longer a problem.

Lemma 7 *If P is a (1-connected) k -cycle for $n > 3$ agents at step t of Algorithm 1, then the set MOVES at t is nonempty for some $i \in N$.*

Proof Take any judgment $P[i]$ which is peripheral in $\text{CH}(P)$, and consider its antipodal judgment $P[j]$. Let $\mathbf{p}_1, \mathbf{p}_2$ be the two paths from $P[i]$ to $P[j]$ in $\text{CH}(P)$, and let $J_m \in \text{CH}(P)$ be the node adjacent to $P[i]$ on path \mathbf{p}_m , $m = 1, 2$. We have that $d(J_1, P[r]) = d(P[i], P[r]) - 1$ for every $P[r]$ on \mathbf{p}_1 , while $d(J_2, P[s]) = d(P[i], P[s]) + 1$ for every $P[s]$ on \mathbf{p}_2 . If there are more profile judgments on \mathbf{p}_1 than on \mathbf{p}_2 , then $J_1 \in \text{MOVES}$, otherwise $J_2 \in \text{MOVES}$. If there are exactly as many judgments on \mathbf{p}_1 as there are on \mathbf{p}_2 , then both $J_1 \in \text{MOVES}$ and $J_2 \in \text{MOVES}$ because in that case $d(J_1, P[j]) = d(J_2, P[j]) = d(P[i], P[j]) - 1$ and consequently $D(i, J, P) = D(i, J', P) < D(i, P[i], P)$. ■

Let us consider the case of a 1-connected P^0 for $|N| > 3$. Let P be the profile produced by Algorithm 1 at step $t \geq 0$, and let P' be the profile produced by Algorithm 1 at step $t + 1$. We begin by showing that for graphs in which no judgment has a degree higher than 2, it is never the case that $P = P'$, i.e., there exist at least one agent for which $\text{MOVES} = \emptyset$ for P .

Lemma 8 *Let P be a profile for $n > 3$ agents, produced by Algorithm 1 at step $t \geq 0$, and let P' be the profile produced by Algorithm 1 at step $t + 1$. Assume that P is 1-connected on $G_{A,\Gamma}$. If the $\text{CH}(P)$ -induced subgraph of $G_{A,\Gamma}$ is such that no vertex in it has a degree higher than two, and $\text{max}_d(P) > 1$, then $P \neq P'$.*

Proof We show that at least the agent $i \in N$ with a peripheral judgment for P has a possible move in P .

Case a. There exists a peripheral judgment in P with degree 1, belonging to $i \in N$.

Let $P[j]$ be an antipodal of $P[i]$. Since $d(P[i], P[j]) > 1$ and P is 1-connected, there must exist exactly one judgment $J \in \mathcal{J}_{A,\Gamma}$, such that $d(P[i], J) = 1$ and that is between judgments $P[i]$ and $P[j]$. For every $r \in N$, $r \neq i$ it holds that $d(J, P[r]) = d(P[i], P[r]) - 1$. Thus J is a move for $P[i]$.

Case b. All peripheral judgments in P are with degree 2. Consider the antipodal judgments $P[i]$ and $P[j]$. There are exactly two shortest paths connecting them: \mathbf{p}_1 and \mathbf{p}_2 . All other profile judgments $P[r]$ are: either on \mathbf{p}_1 , or on \mathbf{p}_2 , or have a shortest path to $P[j]$ that intersects either \mathbf{p}_1 or \mathbf{p}_2 , possibly both. We can apply the same reasoning as in the proof of Lemma 7.

Consider $J \in \text{CH}(P)$ adjacent to $P[i]$ on \mathbf{p}_1 and $J' \in \text{CH}(P)$ adjacent to $P[i]$ on \mathbf{p}_2 . We have that $d(J, P[r]) = d(P[i], P[r]) - 1$ for every $P[r]$ on \mathbf{p}_1 or whose shortest path to $P[j]$ intersects \mathbf{p}_1 , while $d(J, P[s]) = d(P[i], P[s]) + 1$ for every $P[s]$ on \mathbf{p}_2 or whose shortest path to $P[j]$ intersects \mathbf{p}_2 , but does not intersect \mathbf{p}_1 . If there are more agents r than agents s , then $J \in \text{MOVES}$, otherwise $J' \in \text{MOVES}$. If there are exactly as many agents r as agents s , then both $J \in \text{MOVES}$ and $J' \in \text{MOVES}$ because in that case $d(J, P[j]) = d(J', P[j]) = d(P[i], P[j]) - 1$ and consequently $D(i, J, P) = D(i, J', P) < D(i, P[i], P)$. ■

Observe that if the $\text{CH}(P^0)$ -induced subgraph on $G_{A,\Gamma}$ is such that every vertex in it is of degree at most two, then for every subsequently constructed P in Algorithm 1, it will hold that the $\text{CH}(P)$ -induced subgraph on $G_{A,\Gamma}$ is such that every vertex in it is of degree

at most two. This is due to the fact that, at each step of Algorithm 1, the agents can only chose judgments from the $\text{CH}(P^0)$.

From Observation 2 we know that a $\{P\}$ -induced subgraph of $G_{A,\Gamma}^h$ and $G_{A,\Gamma}^m$ does not have 3-cycles. If the $\text{CH}(P^0)$ -induced subgraph of $G_{A,\Gamma}^h$, respectively $G_{A,\Gamma}^m$ contains no k -cycles for $k > 3$, then this induced subgraph contains no cycles and it is by definition a tree. From the Case a. of the proof of Lemma 8, we immediately obtain the following corollary.

Corollary 9 *Let P be a profile produced in Algorithm 1 at step $t \geq 0$ and let P' be profile produced in Algorithm 1 at step $t + 1$. Assume that P is 1-connected on $G_{A,\Gamma}$. If the $\text{CH}(P)$ -induced subgraph of $G_{A,\Gamma}$ is a tree, and $\text{max}_d(P) > 1$, then $P \neq P'$.*

Proof The proof follows from the Case a. of the proof of Lemma 8 and the observation that: all the subgraphs of a tree are trees, and the peripheral vertices of a tree have a degree 1. ■

We now need to show that not only the profile changes in each iteration, it also changes towards a consensus.

From Proposition 1 we have that $\text{max}_d(\text{CH}(P))$ does not increase with each step of the Algorithm 1. It is possible that $\text{max}_d(\text{CH}(P)) = \text{max}_d(\text{CH}(P'))$ for P' being constructed immediately after P in Algorithm 1. From the proof of Lemma 8 we have the following corollary.

Corollary 10 *Let $P \in \mathcal{J}_{A,\Gamma}^n$ be a profile produced in Algorithm 1 at step $t \geq 0$ and let $P' \in \mathcal{J}_{A,\Gamma}^n$ be profile produced in Algorithm 1 at step $t + 1$. If $\{P\} = \{P'\}$, then the $\{P\}$ -induced graph of $G_{A,\Gamma}$ has at least one k -cycle, where $2m + 2 \geq k \geq 2m$.*

Clearly if the agents whose judgments are antipodal in P can choose to move towards each other via two different shortest paths between their judgments causing $\{P\} = \{P'\}$. These agents however, also have the possibility to chose to move towards each other on the same shortest path between their judgments. As soon as two agents use the same shortest path, the k -cycle will be broken in the next step of the algorithm and $\{P\} \neq \{P'\}$.

Let us consider the case when $\text{max}_d(P) = 1$.

Lemma 11 *Let P be a 1-connected profile for $n > 3$ agents at step t with $\text{max}_d(P) = 1$ and let P' be a profile obtain from it by Algorithm 1 at step $t + 1$. If n is odd then $\{P\} \neq \{P'\}$.*

Proof In this case the Algorithm 1 behaves as on the $G_{A,\Gamma}^c$ graph, see Theorems 4 and 5, except the P -induced subgraph of $G_{A,\Gamma}$ will have no 3-cycles (or any size cycles since $\text{max}_d(P) = 1$). Namely, if there is one plurality judgment J in P , all the agents can reach it, because $\text{max}_d(P) = 1$ and P is 1-connected. Consequently $\{P'\} = 1$. If more than one plurality judgment exists, the agents whose judgment is this plurality judgment will not have a move, while and all the other agents will move to their choice of a plurality judgment. If n is odd P' will have exactly one plurality judgment and the profile P'' constructed by Algorithm 1 in step $t + 1$ is a consensus. If however n is even, as with $G_{A,\Gamma}^c$, P can be such that half of the agents have a judgment J , while the other half have an adjacent judgments J' . Namely $\{P\} = \{J, J'\}$ and $d(J, J') = 1$. If such P is reached the Algorithm 1 forces the agents to infinitely “swap” between J and J' . ■

Lemma 12 *Assume an odd number of agents $n > 3$, that in the P -induced subgraph on $G_{A,\Gamma}$ each vertex has a degree at most 2. Let $P \in \mathcal{J}_{A,\Gamma}^n$ be s.t. $\text{CH}(P)$ has at least one k -cycle for $k > 3$. Let $p_i(J) > 0$ be the probability that an agent i will choose a possible move J from the set MOVES at a step t_1 in the Algorithm 1. Then the algorithm will reach a point $t_2 > t_1$ where $P' \in \mathcal{J}_{A,\Gamma}^n$ is obtained s.t. $\text{CH}(P) \subset \text{CH}(P')$ with probability 1.*

Proof If a profile P'' is reached such that all antipodal judgments have degree two, it is sufficient that only one antipodal pair “breaks” the cycle for a profile P' to be reached. To do so, two agents with antipodal judgments have to choose to move along the same path towards each other. Consider a pair of antipodal judgments in P , $P[i]$ and $P[j]$. Assume that at the non-deterministic step of the algorithm there exists a probability $1 > p_i(J) > 0$ that the agent i selects $J \in \text{MOVES}$ that is on a shortest path \mathbf{p} between $P[i]$ and $P[j]$ and probability $p_i(J') = 1 - p_i(J)$ that she selects $J' \in \text{MOVES}$ that is on a different path \mathbf{q} between $P[i]$ and $P[j]$. Similarly, let those probabilities be $1 > p_j(J'') > 0$ that agent j selects to move to J'' on path \mathbf{p} and $p_j(J''') = 1 - p_j(J'')$ for the probability that j moves to J''' on some other path \mathbf{q}' (\mathbf{q} and \mathbf{q}' may not be the same). Since the agents decide on their moves independently, the probability that agent i will choose the same path as j is $p_i(J) \cdot p_j(J'') > 0$. Since the two peripheral judgments $J[i]$ and $J[j]$ are no longer part of the new profile P' , $\{P'\} \subset \text{CH}(P)$ and from Corollary 10 we get that $\text{CH}(P) \subset \text{CH}(P')$ is reached after a finite time with probability 1. ■

Let us call **Class A** for $G_{\mathcal{A},\Gamma}$ the set of all $\text{CH}(P)$ -induced subgraphs of $G_{\mathcal{A},\Gamma}$ that are tree graphs. Let us call **Class B** for $G_{\mathcal{A},\Gamma}$ the set of all $\text{CH}(P)$ -induced subgraphs of $G_{\mathcal{A},\Gamma}$ whose vertices have a degree of at most 2. For instance, the doctrinal paradox profile from Example 1 is in Class B for $G_{\mathcal{A},\Gamma}^m$, see Graph b in Figure 2. On the other hand, it is neither in Class B nor in Class A for $G_{\mathcal{A}}^h$, see Example 2 and Figure 3.

We can now state the following theorem whose proof follows from Lemma 8, Corollary 10, Corollary 9, Lemma 11 and Lemma 12.

Theorem 13 *Let $P^0 \in \mathcal{J}_{\mathcal{A},\Gamma}^n$ be a 1-connected profile belonging to Class A or to Class B for $G_{\mathcal{A}}^h$ or $G_{\mathcal{A},\Gamma}^m$. If $n > 3$ is odd, and each element of MOVES has a non-null probability of being selected in the non-deterministic choice step, then the Algorithm 1 reaches consensus with probability 1 on $G_{\mathcal{A}}^h$, respectively $G_{\mathcal{A},\Gamma}^m$.*

6 Properties of Consensus

In this section, we compare the output and performance of our iteration procedure to those of standard distance-based judgment aggregation rules. We first discuss the “quality” of the consensual decision. Then, we look at the computational complexity of the procedure.

6.1 Consensus Quality

Distance-based judgment aggregation [28, 39, 21, 19] combines an algebraic aggregation function \star with a distance function d (not necessarily a path distance in some agenda graph) in order to select the collective opinion that is closest to the given profile. Given $P \in \mathcal{J}_{\mathcal{A},\Gamma}^n$, the distance-based aggregation function $F^{d,\star} : \mathcal{J}_{\mathcal{A},\Gamma}^n \rightarrow 2^{\mathcal{J}_{\mathcal{A},\Gamma}} \setminus \emptyset$ is defined as

$$F^{d,\star}(P) = \underset{J \in \mathcal{J}_{\mathcal{A},\Gamma}}{\text{argmin}} \star (d(P[1], J), \dots, d(P[n], J)).$$

Natural questions to ask are:

- How does Algorithm 1 perform in comparison to $F^{d,\Sigma}$ when d is a path distance in an agenda graph?
- How do the collective judgments $F^{d,\Sigma}(P^0)$ compare to the consensus judgment reached by Algorithm 1 for a given starting profile P^0 ?

The questions cannot be fully explored within the scope of this paper. However, we establish some initial properties below.

A property generally deemed desirable in judgment aggregation is that of *propositional unanimity* [21, 39, 16]. Propositional unanimity requires that, if every agent in profile P has the same value for some

issue $\varphi \in \mathcal{A}$, then the same value for $\varphi \in \mathcal{A}$ shows up either in at least one of the judgments in $F^{d,\star}$ (weak unanimity) or in all of the judgments in $F^{d,\star}$ (strong unanimity). It is interesting to note that the most popular distance based judgment aggregation rule $F^{d_h,\Sigma}$ does not satisfy even the weak version of the property [38] and the same applies to $F^{d_m,\Sigma}$ and $F^{d_c,\Sigma}$ [22]. In this respect, iterative consensus building behaves better.

Proposition 14 *If Algorithm 1 terminates with a consensus on $G_{\mathcal{A}}^h$, then the consensus satisfies strong unanimity with respect to the initial profile P^0 .*

Proof Note that, for $G_{\mathcal{A}}^h$, judgment J' is between judgments J and J'' iff vertex J' is on the shortest path between vertices J and J'' in the graph. Consequently, if all the agents in P^0 give the same truth-value on an issue, then $\text{CH}(P^0)$ cannot contain judgments that assign different truth-value to this issue. ■

The same is not the case for $G_{\mathcal{A},\Gamma}^m$.

Proposition 15 *There is an initial profile P^0 such that Algorithm 1 terminates with a consensus on $G_{\mathcal{A},\Gamma}^m$, and the consensus does not satisfy weak unanimity with respect to P^0 .*

Proof As a counter-example consider Example 1 and Graph b in Figure 2. The vertex $(1, 1, 1)$ is between vertices $(0, 1, 0)$ and $(1, 0, 0)$, but the judgment $(1, 1, 1)$ is not between judgments $(0, 1, 0)$ and $(1, 0, 0)$. Thus the agents can move from $(0, 1, 0)$ and $(1, 0, 0)$ to $(1, 1, 1)$ thus violating propositional unanimity on the last issue. ■

A big advantage of one-shot distance-based aggregation $F^{d,\Sigma}$ is that it produces output (a winner or winners) on any profile P^0 , while our Algorithm 1 is more restricted in this respect. As we have seen, a necessary condition for successful termination of Algorithm 1 is that P^0 is 1-connected. Sufficient conditions are even more restrictive. Still, Proposition 14 demonstrates that, when Algorithm 1 reaches a consensus, it is structurally “better behaved” than a distance-based judgment aggregation rule for the most popular approach based on the sum of Hamming distances. In the next subsection we show that Algorithm 1 is also “better behaved” in the sense of computational complexity.

6.2 Complexity of Reaching Consensus

An important drawback of distance-based judgment aggregation is the computational complexity of producing the output, i.e., the winning judgment or judgments. The winner determination problem for $F^{d_h,\Sigma}$ is known to be Θ_2^p -complete [14], and the result extends to most other distances d and aggregation functions \star [19]. The computational complexity of determining the collective judgment sets by $F^{d_m,\Sigma}$ is actually not known. How does it work for the iteration procedure formalized with Algorithm 1?

We have shown that the algorithm reaches consensus for an odd number $n > 3$ of agents on 1-connected, not equidistant profiles. How costly is it to reach the consensus? On $G_{\mathcal{A},\Gamma}^c$, it is evident that Algorithm 1 performs well, but the resulting consensus is not very exciting. For the $G_{\mathcal{A},\Gamma}^m$ graph, the consensus-friendly attitude may not earn much in terms of computational complexity, when compared to $F^{d_m,\Sigma}$. For each $P[i]$, we need to find every $J \in \mathcal{J}_{\mathcal{A},\Gamma}$ s.t. there is no rational judgment between $P[i]$ and J . It is not difficult to show, by a reduction to coSAT, that checking whether there is no rational judgment between two given rational judgments is in general coNP-complete. This has to be repeated for multiple candidate judgments to compute the set MOVES, and on top of that with every iteration of the algorithm. As a consequence, we get the following.

Theorem 16 *For $G_{\mathcal{A},\Gamma}^m$, determining MOVES of a single iteration of Algorithm 1 is coNP-hard.*

Note that the hardness result is not really due to the iteration procedure, but rather due to the inherent complexity of computing d_m , which requires to determine nonexistence of particular rational judgments, i.e., to solve the Boolean co-satisfiability problem.

In contrast, the Hamming distance d_h can be always computed efficiently. Consequently, when Algorithm 1 reaches a consensus on $G_{\mathcal{A}}^h$, it is also “better behaved” computationally than the distance-based judgment aggregation rule $F^{d_h, \Sigma}$. We demonstrate it formally.

Proposition 17 For $G_{\mathcal{A}}^h$, a single iteration of Algorithm 1 runs in deterministic polynomial time.

Proof Follows from the fact that the set MOVES can be constructed by checking at most $|\mathcal{A}|$ candidate judgments. ■

By Corollary 10, if the $\text{CH}(P^0)$ induced subgraph of $G_{\mathcal{A}}^h$ has no cycles, then the diameter of $\text{CH}(P)$ is strictly shrinking with each non-terminating step t . In consequence, if Algorithm 1 reaches consensus for such P^0 , then it does so in polynomially many steps. However, in case of cycles in the $\text{CH}(P^0)$ induced graph in $G_{\mathcal{A}}^h$, the algorithm may run into such a cycle and take some time until the agents “stray” from the loop. When it happens, any judgment occurring on the loop can be the consensus. Using this observation, we propose the following modification of Algorithm 1.

Algorithm 2: Same as Algorithm 1, only it stops the iteration when $\{P_t\} = \{P_{t'}\}$ for some $t > t'$, and non-deterministically chooses one $J \in \{P\}$ as the consensus, producing in the next step P_{t+1} with $\{P_{t+1}\} = \{J\}$.

Unlike Algorithm 1, Algorithm 2 avoids looping and waiting until two or more agents “move” in the same direction. It also avoids infinite loops in the case of profiles with evenly many agents. On the other hand, Algorithm 2 is no longer decentralised, which is a clear disadvantage. We suggest that it can be treated as a technical variant of Algorithm 1 that potentially reduces its running time by employing a trusted third party which simulates probabilistic convergence of the profile in Algorithm 1 by one-step non-deterministic choice in Algorithm 2. The following formal results, which are straightforward consequences of our analysis above, justify the suggestion.

Theorem 18 Consider $G_{\mathcal{A}}^h$ and N such that $|N|$ is odd and larger than 3. If Algorithm 1 can reach consensus with J then also Algorithm 2 can reach consensus with J .

Theorem 19 Consider $G_{\mathcal{A}}^h$ and N such that $|N|$ is odd and larger than 3. Moreover, let $P^0 \in \mathcal{J}_{\mathcal{A}, \Gamma}^n$ be 1-connected and not equidistant. Algorithm 2 reaches consensus from P^0 on $G_{\mathcal{A}}^h$ in deterministic polynomial time.

Lastly, let us observe that checking whether P^0 is equidistant can be done in linear time of the number of agents. For a graph G , determining if it has a simple cycle of size k , for k fixed, is a polynomial time problem over the size of G , see [1], however we do not generate the full $G_{\mathcal{A}, \Gamma}$ when (or before) we run the iteration algorithm.

7 Related Work

List [25] considered judgment transformation functions $\tau : \mathcal{J}_{\mathcal{A}, \Gamma}^n \rightarrow (\{0, 1\}^{\mathcal{A}})^n$ as means to building iteration procedures for judgment aggregation problems. He showed that for a set of desirable properties no transformation functions exists. Such impossibility results exist for judgment aggregation functions, however, by relaxing some of the properties, specific judgment aggregation operators have been constructed: quota-based rules [6], distance-based rules [37, 28, 14, 11], generalisations of Condorcet-consistent voting rules [30, 29, 21], and rules based on the maximisation of some scoring function [21, 5, 42]. To the best of our knowledge, specific iteration

procedures for judgment aggregation problems have not been proposed in the literature.

List [25] argues that the desirable conditions for judgment transformation functions should satisfy the following properties: universal domain, rational co-domain, consensus preservation, minimal relevance, and independence. Universal domain is satisfied when the transformation function accepts as admissible input any possible profile of rational judgments. Rational co-domain is satisfied when the function always outputs a profile of rational judgments. Consensus preservation is satisfied when τ always maps unanimous profiles into themselves. Minimal relevance is a weak property. It is satisfied when for each $P[i]$ there exists a profile P' to which P can be transformed such that $P'[i] = P'[i]$. In other words, the transformation function should be such that it does not allow one agent to never change her judgment regardless of what the other profile judgments are. Lastly independence is satisfied when for each agenda issue, $J'_i(\varphi)$ depends only on $J(\varphi)$, and not on $J(\phi)$ for some other $\phi \in \mathcal{A}$; $J'_i = P'[i]$, $J_i = P[i]$, $P' = \tau(P)$.

Each step of Algorithm 1 can be seen as a (distributed) function that transforms an input profile into an output profile, namely as a List judgment transformation function. Given a profile $P \in \mathcal{J}_{\mathcal{A}, \Gamma}$, let $T_d(P) = \{P' \mid P' \in \mathcal{J}_{\mathcal{A}, \Gamma}^n, d(P[i], P'[i]) \leq 1, i \in [1, n]\}$. We can define the transformation function τ_d that maps a profile $P \in \mathcal{J}_{\mathcal{A}, \Gamma}^n$ to a profile $P' \in T_d(P)$. Although Algorithm 1 does not terminate for each profile, τ_d does satisfy universal domain in the case of $G_{\mathcal{A}, \Gamma}^c$ and $G_{\mathcal{A}, \Gamma}^m$, because each step of the algorithm transforms the profile (possibly into itself). Universal domain is not satisfied on $G_{\mathcal{A}}^h$ because profiles on this graph do not always satisfy the necessary conditions for termination with a consensus. The rational co-domain and the consensus preservation properties are also trivially satisfied. It is not difficult to show that the minimal relevance property is also satisfied. Independence is the desirable property that is violated, and in fact List [25] argues that relaxing independence is the most plausible path towards avoiding the impossibility result.

In voting, deliberation and iterative consensus have been studied, although perhaps not axiomatically. As most similar with our work we distinguish [18] and [15]. Voting problems can be represented as judgment aggregation problems, see e.g., [6, 23], therefore it is possible to compare these works with ours. First we show how voting problems are represented in judgment aggregation.

A voting problem is specified with a set of agents N and a set of candidate options $O = \{x_1, x_2, \dots, x_m\}$. Let \mathcal{O} be the set of all total, transitive, and antisymmetric orders over the elements of O . A vote \succ is an element of \mathcal{O} and a voting profile is a collection of votes, one for each agent in N . The preference agenda \mathcal{A}_o is constructed by representing each pair of options x_i and x_j , where $i < j$ with an issue $x_i P x_j$. The constraint Γ_{tr} is the transitivity constraint defined as $\Gamma_{tr} = \bigwedge_{x_i P x_j, x_j P x_k, x_i P x_k \in \mathcal{A}_o} ((x_i P x_j) \wedge (x_j P x_k) \rightarrow (x_i P x_k))$.

For each vote $\varphi \in \mathcal{O}$ we obtain a rational judgment J_{\succ} such that $J_{\succ}(x_i P x_j) = 1$ iff $x_i \succ x_j$ and $J_{\succ}(x_i P x_j) = 0$ iff $x_j \succ x_i$.

A Condorcet winner for a voting profile, when it exists, is the option that wins the majority of pairwise comparison for every other option in O , see e.g., [32]. The corresponding concept in the judgment aggregation representation of a voting problem is called *majority consistency*. A judgment profile is majority-consistent if the judgment obtained by taking the value for each issue assigned by a strict majority of agents in the profile is rational. The doctrinal paradox profile from Example 1 is not majority-consistent. It was shown [23, 31] that if a judgment profile on the preference agenda is majority-consistent, then the corresponding voting profile has a Condorcet winner.

Hassanzadeh *et al.* [18] consider an iterative consensus algorithm for voting profiles. In their algorithm, each agent is allowed to (simultaneously with other agents) move from vote \succ_i to vote \succ if she can flip the order of two adjacent options without violating transitivity. This corresponds to the agents moving to an adjacent judgment

in the agenda graph $G_{\mathcal{A}_o, \Gamma_{tr}}^h$. Hassanzadeh *et al* consider the majority graph for a voting profile (for an odd number of agents): the vertices in this directed graph are the elements of O and there is an edge from x_i to x_j if there are more agents in the profile who prefer x_i to x_j than agents who prefer x_j to x_i . The majority graph corresponds to a judgment $J \in \mathcal{J}_{\mathcal{A}_o, \Gamma_{tr}}$ for which $J(x_i P x_j) = v$, $v \in \{0, 1\}$ if there is a strict majority of agents $r \in N$ for which $J_r(x_i P x_j) = v$, $J_r = P[r]$. Hassanzadeh *et al* show that their algorithm terminates with a consensus on the Condorcet winner when the majoritarian graph has no cycles. If the majority graph of a voting profile has no cycles, then the voting profile has a Condorcet winner.

Goel and Lee [15] consider an iteration procedure in which the agents “move” along adjacent vertices along (what corresponds to) the graph $G_{\mathcal{A}, \Gamma}^m$. They do not commit to the nature of their vertices, so they are not exactly judgments or alternatives, just allowed options for iteration. In their algorithm not all agents move individually, but three agents at a time first reach a consensus and then all three move to the consensus option in the graph. Goel and Lee consider line graphs, graphs in which two vertices have degree 1 and all other vertices have degree 2, and show that the consensus produced by their algorithm is the generalised median. Namely, if the options in their algorithms were judgments from $\mathcal{J}_{\mathcal{A}, \Gamma}$ the consensus their algorithm reaches for these graphs is an approximation of $F^{d_m, \Sigma}$.

Both [18] and [15] offer interesting directions for future study in context of our algorithm: to consider the profiles that have a Condorcet winner (see e.g., [23] for the concept of Condorcet winner in judgment aggregation) and to consider triadic iteration, allowing three agents to coordinate their moves with respect to each other and then see when a consensus emerges.

It is an open question of how our algorithm would perform on the special case of voting problems represented in judgment aggregation. The $G_{\mathcal{A}_o, \Gamma_{tr}}^h$ graph on the preference agenda has a more regular topology in comparison to general judgment aggregation problems, it is a permutahedron. For example, for an agenda of three options, the graph $G_{\mathcal{A}_o, \Gamma_{tr}}^h$ is a cycle of length 6. For every $J, J' \in \mathcal{J}_{\mathcal{A}_o, \Gamma_{tr}}$, $d_h(J, J') = d_m(J, J')$, thus the necessary conditions for reaching consensus for Algorithm 1 would be satisfied even on $G_{\mathcal{A}}^h$ because every profile on the preference agenda and judgments rational for the transitivity constraint is 1-connected in $G_{\mathcal{A}_o, \Gamma_{tr}}^h$. The graph $G_{\mathcal{A}_o, \Gamma_{tr}}^m$ always has $\frac{|O| \cdot (|O| - 1)}{2}$ vertices and each of these vertices has a degree $|O| - 1$. We leave for future work the study of whether our algorithm terminates for voting profiles. In particular, we conjecture that the Algorithm 1 for an odd number of agents will converge on the Condorcet winner in the case of voting profiles from the single crossing domain [4], also studied in judgment aggregation [8]. This is because profiles in this domain would have a hull whose induced graph is a line on $G_{\mathcal{A}_o, \Gamma_{tr}}^h$.

Lastly, we must mention [33]. Obraztsova *et al* [33] consider a graph similar to our Hamming agenda graph. They work with preferences, not judgments, but most importantly, the vertices of their graph are elements of (what would correspond to) $\mathcal{J}_{\mathcal{A}, \Gamma}^r$, i.e., the vertices are profiles of votes. There exists a connection between two profiles if one profile can be obtained from the other by making exactly one swap between adjacent options in one vote. Obraztsova *et al* [33] study the properties of voting rules with respect to the “geometry” of the profiles in their graph.

8 Conclusions

In this paper, we propose a decentralised algorithm for iterative consensus building in judgment aggregation problems. We study the termination conditions for this algorithm, some of its structural properties, and its computational complexity.

In order to reach a consensus, our algorithm exploits the topology of a graph. All the available judgments that the agents can chose from are vertices in the graph. The algorithm models an agent’s change of

mind as a move between adjacent judgments in the graph. We define three natural graphs that can be constructed for a set of rational judgments $\mathcal{J}_{\mathcal{A}, \Gamma}$: the complete graph $G_{\mathcal{A}, \Gamma}^c$, the Hamming graph $G_{\mathcal{A}}^h$, and the model agenda graph $G_{\mathcal{A}, \Gamma}^m$. We prove that our algorithm always terminates for an odd number of agents on the graph $G_{\mathcal{A}, \Gamma}^c$, but it necessarily selects one of the judgments proposed in the first round of iterations. For the graphs $G_{\mathcal{A}}^h$ and $G_{\mathcal{A}, \Gamma}^m$ we show a class of profiles for which the algorithm terminates with a consensus and a class of profiles for which it does not terminate with a consensus.

If the agents initially chose judgments such that the convex hull of the profile of these judgments induces a subgraph of $G_{\mathcal{A}}^h$, or $G_{\mathcal{A}, \Gamma}^m$ in which each vertex has a degree of at most 2, then our algorithm probabilistically terminates with a consensus for an odd number of (more than 3) agents.

The list of profiles we give here, for which Algorithm 1 terminates with a consensus, is clearly not exhaustive. For example, it is easy to show that, for an odd number of agents, Algorithm 1 terminates with a consensus if the $\text{CH}(P^0)$ induced subgraph of $G_{\mathcal{A}}^h$, or $G_{\mathcal{A}, \Gamma}^m$, is such that it contains only k -cycles, where $k = 2 \cdot \text{max}_d(P^0) + 1$. This is because for such profiles there exists at least one pair of antipodal judgments with degree no more than 2 who will have a nonempty set MOVES. An immediate direction for future work is to strengthen our results with other classes of consensus terminating agenda graph topologies, particularly those corresponding to profiles on the preference agenda (and transitivity constraints).

A step of our algorithm implements a judgment profile transformation function of the type defined in [25]. List [25] gives an impossibility characterisation of such functions. Our function “escapes” this impossibility result by not satisfying the independence property on all agenda graphs and the universal domain on $G_{\mathcal{A}}^h$.

While $G_{\mathcal{A}, \Gamma}^c$, and $G_{\mathcal{A}, \Gamma}^m$ satisfy the necessary conditions for termination of Algorithm 1 for any \mathcal{A} and Γ , this is not the case with $G_{\mathcal{A}}^h$, which is why the transformation function fails to satisfy universal domain on $G_{\mathcal{A}}^h$. On $G_{\mathcal{A}}^h$, sometimes all the adjacent judgments to a rational judgment J are not rational and thus not allowed to move to. In our future work we aim to explore modifications of the algorithm allowing the agents to make “longer” moves, i.e., to “jump over” a vertex that is not a rational judgment.

In Section 6 we gave two results with respect to the quality of the consensus reached by Algorithm 1 with respect to the widely used distance-based aggregation function $F^{d_h, \Sigma}$. This function $F^{d_h, \Sigma}$ is also known as the median aggregation rule and it is widely used in many domains, e.g., generalises the Kemeny voting rule, see [23], and for measuring dissimilarity between concepts in ontologies [10]. We merely scratched the surface of this consensus quality analysis and this line of research merits further attention.

Lastly, a more long-term goal for our future work is to explore versions of iteration on an agenda graph where the agents do not try to move to reduce the path distance to all of the other agents, but only to their neighbours in a given social network, or as in [15], to two randomly selected two agents.

Acknowledgements. We are grateful to Edith Elkind and the anonymous reviewers for their help in improving this work. Wojciech Jamroga acknowledges the support of the 7th Framework Programme of the EU under the Marie Curie IEF project ReVINK (PIEF-GA-2012-626398). We also acknowledge the support of ICT COST Action IC1205.

REFERENCES

- [1] N. Alon, R. Yuster, and U. Zwick, ‘Finding and counting given length cycles’, *Algorithmica*, **17**(3), 209–223, (1997).
- [2] G. Boella, G. Pigozzi, M. Slavkovik, and L. van der Torre, ‘Group intention is social choice with commitment’, in *Coordination, Organizations, Institutions, and Norms in Agent Systems VI*, eds., Marina De Vos, Nicoletta Fornara, Jeremy Pitt, and George Vouros, volume 6541 of

- Lecture Notes in Computer Science*, 152–171, Springer Berlin / Heidelberg, (2011).
- [3] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, *Handbook of Computational Social Choice*, Cambridge University Press, March 2016.
- [4] R. Bredereck, J. Chen, and G. J. Woeginger, ‘A characterization of the single-crossing domain’, *Social Choice and Welfare*, **41**(4), 989–998, (2012).
- [5] F. Dietrich, ‘Scoring rules for judgment aggregation’, *Social Choice and Welfare*, 1–39, (2013).
- [6] F. Dietrich and C. List, ‘Arrow’s theorem in judgment aggregation’, *Social Choice and Welfare*, **29**(1), 19–33, (July 2007).
- [7] F. Dietrich and C. List, ‘Judgment aggregation by quota rules: Majority voting generalized’, *Journal of Theoretical Politics*, **19**(4), 391–424, (2007).
- [8] F. Dietrich and C. List, ‘Majority voting on restricted domains’, *Journal of Economic Theory*, **145**(2), 512–543, (2010).
- [9] F. Dietrich and P. Mongin, ‘The premiss-based approach to judgment aggregation’, *Journal of Economic Theory*, **145**(2), 562–582, (2010).
- [10] F. Distel, J. Atif, and I. Bloch, ‘Concept dissimilarity with triangle inequality’, in *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20–24, 2014*, pp. 614–617, (2014).
- [11] C. Duddy and A. Piggins, ‘A measure of distance between judgment sets’, *Social Choice and Welfare*, **39**, 855–867, (2012).
- [12] U. Endriss and R. de Haan, ‘Complexity of the winner determination problem in judgment aggregation: Kemeny, Slater, Tideman, Young’, in *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems, AAMAS ’15*, pp. 117–125, Richland, SC, (2015). International Foundation for Autonomous Agents and Multiagent Systems.
- [13] U. Endriss and U. Grandi, ‘Binary aggregation by selection of the most representative voters’, in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence*, pp. 668–674, (2014).
- [14] U. Endriss, U. Grandi, and D. Porello, ‘Complexity of judgment aggregation’, *Journal Artificial Intelligence Research (JAIR)*, **45**, 481–514, (2012).
- [15] A. Goel and D. Lee, ‘Triadic consensus’, in *Internet and Network Economics: 8th International Workshop, WINE 2012, Liverpool, UK, December 10–12, 2012. Proceedings*, ed., P. W. Goldberg, volume 7695 of *Lecture Notes in Computer Science*, 434–447, Springer Berlin Heidelberg, Berlin, Heidelberg, (2012).
- [16] U. Grandi and U. Endriss, ‘Lifting integrity constraints in binary aggregation’, *Artificial Intelligence*, **199–200**, 45–66, (2013).
- [17] U. Grandi, A. Loreggia, F. Rossi, K. Brent Venable, and T. Walsh, ‘Restricted manipulation in iterative voting: Condorcet efficiency and borda score’, in *Proceeding of the 3rd International Conference on Algorithmic Decision Theory (ADT-2013)*, pp. 181–192, (2013).
- [18] F.F. Hassanzadeh, E. Yaakobi, B. Touri, O. Milenkovic, and J. Bruck, ‘Building consensus via iterative voting’, in *2013 IEEE International Symposium on Information Theory Proceedings (ISIT)*, pp. 1082–1086, (July 2013).
- [19] W. Jamroga and M. Slavkovik, ‘Some complexity results for distance-based judgment aggregation’, in *AI 2013: Advances in Artificial Intelligence - 26th Australasian Joint Conference, Dunedin, New Zealand, December 1–6, 2013. Proceedings*, pp. 313–325, (2013).
- [20] E. M. Kok, J. J.Ch. Meyer, H. Prakken, and G. A. W. Vreeswijk, ‘A formal argumentation framework for deliberation dialogues’, in *Argumentation in Multi-Agent Systems: 7th International Workshop, ArgMAS 2010 Toronto, ON, Canada, May 10, 2010 Revised, Selected and Invited Papers*, eds., P. McBurney, I. Rahwan, and S. Parsons, volume 6614 of *Lecture Notes in Computer Science*, chapter Argumentation in Multi-Agent Systems, 31–48, Springer Berlin Heidelberg, Berlin, Heidelberg, (2011).
- [21] J. Lang, G. Pigozzi, M. Slavkovik, and L. van der Torre, ‘Judgment aggregation rules based on minimization’, in *TARK*, pp. 238–246, (2011).
- [22] J. Lang, G. Pigozzi, M. Slavkovik, L. van der Torre, and S. Vesic, ‘Majority-preserving judgment aggregation rules’, *CoRR*, **abs/1502.05888**, (2015).
- [23] J. Lang and M. Slavkovik, ‘Judgment aggregation rules and voting rules’, in *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, volume 8176 of *Lecture Notes in Artificial Intelligence*, pp. 230–244. Springer-Verlag, (2013).
- [24] O. Lev and J. Rosenschein, ‘Convergence of iterative voting’, in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2, AAMAS ’12*, pp. 611–618, Richland, SC, (2012). International Foundation for Autonomous Agents and Multiagent Systems.
- [25] C. List, ‘Group communication and the transformation of judgments: An impossibility result’, *Journal of Political Philosophy*, **19**(1), 1–27, (2011).
- [26] L. List and B. Polak, ‘Introduction to judgment aggregation’, *Journal of Economic Theory*, **145**(2), 441 – 466, (2010).
- [27] R. Meir, M. Polukarov, J. Rosenschein, and N. Jennings, ‘Convergence to equilibria in plurality voting’, in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11–15, 2010*, pp. 823–828, (2010).
- [28] M.K. Miller and D. Osherson, ‘Methods for distance-based judgment aggregation’, *Social Choice and Welfare*, **32**(4), 575 – 601, (2009).
- [29] K. Nehring and M. Pivato, ‘Majority rule in the absence of a majority’. Presented at the New Developments in Judgement Aggregation and Voting Theory Workshop, September 2011.
- [30] K. Nehring, M. Pivato, and C. Puppe, ‘Condorcet admissibility: Indeterminacy and path-dependence under majority voting on interconnected decisions’. <http://mpr.ub.uni-muenchen.de/32434/>, July 2011.
- [31] K. Nehring, M. Pivato, and C. Puppe, ‘The condorcet set: Majority voting over interconnected propositions’, *J. Economic Theory*, **151**, 268–303, (2014).
- [32] H. Nurmi, ‘Voting theory’, in *e-Democracy*, eds., D. Rios Insua and S. French, volume 5 of *Advances in Group Decision and Negotiation*, 101–123, Springer Netherlands, (2010).
- [33] S. Obraztsova, E. Elkind, P. Faliszewski, and A. Slinko, ‘On swap-distance geometry of voting rules’, in *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS ’13*, pp. 383–390, Richland, SC, (2013). International Foundation for Autonomous Agents and Multiagent Systems.
- [34] S. Obraztsova, E. Markakis, M. Polukarov, Z. Rabinovich, and N. Jennings, ‘On the convergence of iterative voting: How restrictive should restricted dynamics be?’, in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25–30, 2015, Austin, Texas, USA.*, pp. 993–999, (2015).
- [35] S. Ontañón and E. Plaza, ‘An argumentation-based framework for deliberation in multi-agent systems’, in *Argumentation in Multi-Agent Systems: 4th International Workshop, ArgMAS 2007, Honolulu, HI, USA, May 15, 2007, Revised Selected and Invited Papers*, eds., I. Rahwan, S. Parsons, and C. Reed, volume 4946 of *Lecture Notes in Computer Science*, chapter Argumentation in Multi-Agent Systems, 178–196, Springer Berlin Heidelberg, Berlin, Heidelberg, (2008).
- [36] I. M. Pelayo, *Geodesic Convexity in Graphs*, SpringerBriefs in Mathematics, 2013.
- [37] G. Pigozzi, ‘Belief merging and the discursive dilemma: an argument-based account to paradoxes of judgment aggregation’, *Synthese*, **152**(2), 285–298, (2006).
- [38] G. Pigozzi, M. Slavkovik, and L. van der Torre, ‘A complete conclusion-based procedure for judgment aggregation’, in *1st International Conference on Algorithmic Decision Theory, Venice, Italy, October 20–23, 2009. Proceedings*, eds., F. Rossi and A. Tsoukiàs, volume 5783 of *Lecture Notes in Computer Science*, pp. 1–13. Springer, (2009).
- [39] M. Slavkovik, *Judgment aggregation for multiagent systems*, Ph.D. dissertation, University of Luxembourg, 2012.
- [40] M. Slavkovik and T. Ågotnes, ‘Measuring dissimilarity between judgment sets’, in *Logics in Artificial Intelligence - 14th European Conference, JELIA*, pp. 609–617, (2014).
- [41] M. Slavkovik and G. Boella, ‘Recognition-primed group decisions via judgment aggregation’, *Synthese*, (2012).
- [42] W. Zwicker, ‘Towards a Borda count for judgment aggregation’. Working Paper, 2011.