




Opinion dynamics with median aggregation

Downloaded from: <https://research.chalmers.se>, 2026-06-16 00:30 UTC

Citation for the original published paper (version of record):

Berenbrink, P., Hofer, M., Kaaser, D. et al (2026). Opinion dynamics with median aggregation. *Artificial Intelligence*, 355. <http://dx.doi.org/10.1016/j.artint.2026.104527>

N.B. When citing this work, cite the original published paper.



Opinion dynamics with median aggregation

Petra Berenbrink ^{id a}, Martin Hofer ^{id b}, Dominik Kaaser ^c, Marten Maack ^a,
Malin Rau ^{id d}, Lisa Wilhelmi ^{id b,*}

^a Dept. of Computer Science, University of Hamburg, Vogt-Kölln-Str. 30, Hamburg, 22527, Germany

^b Dept. of Computer Science, RWTH Aachen University, Ahornstrasse 55, Aachen, 52074, Germany

^c Inst. for Data Engineering, TU Hamburg, Blohmstraße 15, Hamburg, 21079, Germany

^d Division of Applied Mathematics and Statistics, Chalmers University of Technology, Gothenburg, 412 96, Sweden

ARTICLE INFO

Keywords:

Opinion formation
Median voting
Nash equilibrium

ABSTRACT

Understanding the formation and evolution of opinions is of broad interdisciplinary interest. Many classical models for opinion formation focus on the impact of different notions of *locality*, e.g., locality due to network effects among agents or the role of the proximity of opinions. In practice, however, opinion formation is often governed by the interplay of *local* and *global* influences.

In this paper, we study these influences with a model for opinion formation of agents embedded in a social network. Each agent has a static intrinsic opinion as well as a public opinion that is updated asynchronously over time. Moreover, agents have access to a global aggregate (e.g., the outcome of a vote) of all public opinions. We focus on the popular median voting rule and show that pure Nash equilibria always exist. For every initial state of the dynamics, a pure equilibrium can be reached. The set of reachable equilibria forms a complete lattice, and extremal equilibria can be computed in polynomial time.

We show that by uniformly increasing the influence of the global median we can enforce that the median opinion is the same in every reachable equilibrium. We can compute the increase scheme that achieves this property in polynomial time. In contrast, when we can increase the influence of the global median for a set of at most k agents, finding the set that leads to a unique median opinion in every reachable equilibrium is NP-complete.

1. Introduction

Opinion formation processes [1–3] model how individuals develop, modify and express opinions in a social context. Opinion formation is highly influenced by various factors, including the social environment, domain experts, or professionals for PR and advertising. Such processes play a crucial role in public discourse, decision-making, and collective behavior. They decide the outcome of elections and the fate of political parties, the success of new products, companies, and entire economies; they influence whether political movements are successful, and they lead to new trends and directions in research and development. Understanding how opinions are formed is essential for promoting constructive dialogue, managing conflicts, and designing effective communication strategies. Research on opinion formation processes sheds light on the emergence of collective phenomena, such as public opinion shifts [4], social movements [5], and the spread of misinformation [6]. Understanding the formation and evolution of opinions is of broad interdisciplinary interest, including artificial intelligence, sociology, economics, mathematics, and physics.

* Corresponding author.

E-mail addresses: petra.berenbrink@uni-hamburg.de (P. Berenbrink), mhofer@cs.rwth-aachen.de (M. Hofer), dominik.kaaser@tuhh.de (D. Kaaser), marten.maack@uni-hamburg.de (M. Maack), malin.rau@chalmers.se (M. Rau), wilhelmi@algo.rwth-aachen.de (L. Wilhelmi).

<https://doi.org/10.1016/j.artint.2026.104527>

Received 8 May 2025; Received in revised form 3 March 2026; Accepted 26 March 2026

Available online 30 March 2026

0004-3702/© 2026 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Published by Elsevier B.V. This is an open access article under the CC BY license

Many classical models for opinion formation focus on the impact of social network effects (such as, e.g., Friedkin-Johnsen [7] or voter models [8]) or the role of the proximity of opinions (such as, e.g., Deffuant-Weisbuch [9] or Hegselmann-Krause [10]). These models introduce notions of *locality* between agents and/or *distance* between opinions. These are key aspects of problematic developments such as filter bubbles and polarization, which are known to occur in real-world social networks and describe the fact that algorithmic curation in social networks disconnects users from information that disagrees with their viewpoints. While locality aspects are important, opinion formation is usually also influenced by *global* information and, more precisely, the interplay of *local* and *global* influences. Agents are exposed to the opinions of their friends, but they also can have access to global information such as public opinion polls, media reports, research studies, or other forms of aggregated information about the global opinion landscape. In turn, opinion formation also heavily influences results of voting procedures and elections. These issues are of fundamental importance in the design and analysis of multi-agent and AI systems.

In this paper, our goal is to shed light on this interplay between global and local aspects of opinion formation. We study an opinion formation process within a game-theoretic model which we term *median opinion game*. In a static social network $G = (V, E)$ each node corresponds to an agent. Each agent $i \in V$ has a (static) *intrinsic opinion* $s_i \in \mathbb{R}$ that they keep to themselves and a *public opinion* $z_i \in \mathbb{R}$ that they disclose to their neighbors in G . In addition to these local opinions, agents have access to a global aggregate opinion $f(\mathbf{z})$ of the *strategy profile* \mathbf{z} , i.e., the vector \mathbf{z} of all agents' public opinions. We assume that agents update their public opinion in a sequential fashion. For the choice of the public opinion z_i , agent i strives to strike a balance between (1) their intrinsic opinion s_i , (2) the public opinions z_j in their local neighborhood $N(i) = \{j \mid (i, j) \in E\}$, and (3) the global aggregate opinion $f(\mathbf{z})$. Balance here means to minimize an individual *cost function* incorporating distances between the public opinions, the intrinsic opinion, and the global aggregate opinion $f(\mathbf{z})$. Formally, the cost function of agent i is given by

$$\text{cost}_i(\mathbf{z}) = \alpha_i |s_i - z_i| + \sum_{j \in N(i)} \beta_{ij} |z_j - z_i| + \gamma_i |f(\mathbf{z}) - z_i|$$

with weights α_i , β_{ij} , and γ_i for each $j \in N(i)$. As a consequence, a best response for agent i is given by a weighted median of the public opinions, the intrinsic opinion and the global aggregate opinion (see [Proposition 1](#) below). Optimization using the (weighted) median represents a plausible behavior of the agents in opinion formation games, see the discussion in [11]. For global aggregation, we focus on the popular *median voting rule* $f(\mathbf{z}) = \text{med}(\mathbf{z})$. Median voting has many favorable properties, e.g., optimality properties, as well as incentive compatibility for all agents [12,13]. Median opinion games combine ideas from both voting and opinion formation in multi-agent systems. Both domains are important areas of AI that have received substantial interest in the past. Median opinion games are a first step to broaden our understanding of the *interplay* of these aspects within a joint opinion-based, networked voting model.

Overview of Results. Our first set of results concern structural properties of improvement dynamics and equilibria in median opinion games. We show that for any agent i the cost function is minimized by choosing their public opinion as a weighted median of the intrinsic opinion, the public opinions of the neighbors, and the aggregate opinion. We show that for any instance (defined by the network and the intrinsic opinions) there exists a (*pure Nash*) *equilibrium* where no agent can further reduce their costs. Moreover, for any instance and an initial vector of public opinions, an equilibrium can be reached by a sequence of best-response dynamics. The set of *all* equilibria reachable by sequential improvement dynamics forms a complete lattice, i.e., there is a pointwise maximal and a pointwise minimal reachable equilibrium. These equilibria can be computed in polynomial time.

Beyond these structural properties, our results highlight how a key property of equilibria – the voting result of the global median – are affected by the *order* in which the agents update their public opinion. We show that there exists a family of instances in which any of the intrinsic opinions can become the final median, depending on the order in which the agents are updated. Furthermore, in this family, the number of reachable equilibria is exponentially large in the number of opinions and agents. The update sequence of the agents is crucial to steer the process towards one or another equilibrium. As such, the dynamics can potentially terminate in many different equilibria with fundamentally different aggregation outcomes. With a predominant influence of locality, small changes in the formation process can lead to very heterogeneous outcomes in agent behavior.

For our second set of results we shed light on the impact of *global influence* of median voting on the equilibria. We observe that the range of reachable equilibria is decreasing as soon as the *impact of the median opinion* increases. Here, increasing the impact of the median opinion means that each component of the vector γ is increased by an additive value δ . We show that for every instance such that the global median is not unique in all reachable equilibria, there is a threshold δ such that the following observation holds: if the median weight of all agents is increased by at least δ , then the global median is the same in all reachable equilibria of the modified instance. Additionally, we show that the threshold δ can be computed efficiently. As a consequence of this result, suppose a designer can uniformly increase the awareness of the result of median voting. Then reachable equilibria can become much more robust to changes in the formation process and more predictable with respect to the outcome of the voting.

We contrast this property with a different perspective on the impact of median voting. Instead of increasing γ for all agents by an additive term δ , suppose we select k agents to increase their median weights. Similar approaches have been popular in opinion models for information diffusion. Here one selects k “seed agents” to maximize the eventual spread of a new technology in a network via viral marketing (see [14] and the enormous amount of subsequent literature). In our scenario, the decision problem STABLEMEDIAN asks if for a given k there exists a set of k agents such that each reachable equilibrium has the same global median. We show that this problem is NP-complete. Thus, selecting the right “seeds” for the influence of median voting to increase consistency in reachable equilibria is a hard problem. On a technical level, it is easy to see that our seeding problem is also lacking submodularity properties

(which were a key ingredient for the constant-factor approximation results in [14] and many other problems in AI [15]). We leave an approximation result in polynomial time as an interesting open problem.

2. Related work

The literature on opinion formation is vast. A comprehensive survey is beyond the scope of this paper. We give a brief overview of models that are commonly analyzed in AI and multi-agent systems. For a large number of additional references we refer to [16].

Early approaches on modeling opinion formation include influential work by Downs [17], Abelson [18], and French [19] as well as DeGroot [20]. In their seminal paper, Friedkin and Johnsen [7] introduce a model of opinion formation where each agent has an intrinsic opinion and a public opinion. Agents update their public opinion based on their own intrinsic opinion and the public opinions of their neighbors in a social network. Co-evolutionary and game-theoretic variants of this process are studied in a large number of works [21–26], with a focus on the existence of equilibria and their *social quality*, measured by the price of anarchy.

Opinion formation processes are studied intensively in AI and multi-agent systems. A substantial body of work considers opinion diffusion, the spread of information and opinions among agents embedded in social networks [27–33]. Closer related to our focus is work that (1) studies median aggregation of agents or (2) connects opinions dynamics with voting.

Towards the first direction, Mei et al. [11,34] study median aggregation as a method to choose opinions in the French-DeGroot model. They provide structural and complexity results for consensus properties in equilibrium. Zhang et al. [35] analyze the convergence of an opinion model related to the FJ model. They also assume agents rely on median aggregation. Besides the absence of global voting aspects, a main difference is that they assume a heterogeneous aggregation method for each agent, in which the public opinion is chosen as a linear combination of the intrinsic opinion and the weighted median of the neighbors' public opinions (rather than a weighted median of intrinsic *and* neighbor opinions).

Towards the second direction, Brederick et al. [36] study a process where agents in a network have preference rankings over a set of alternatives. Agents update their ranking sequentially by applying an aggregation method to the rankings in their neighborhood. The authors focus on single-peaked rankings and examine which aggregation rules preserve this property. Epitropou et al. [24] build upon the standard FJ model and study aggregation using the average public opinion. They identify conditions in which the opinion formation process converges quickly towards a unique equilibrium, even with outdated information on the average. For (two) discrete opinions, Auletta et al. [37] study in which way asynchronous majority dynamics can result in an initial minority opinion to become a majority. Doucette et al. [38] study a stochastic inference model with two opinions where agents public opinions are noisy and correlated based on their intrinsic opinion and the public opinions of their neighbors. Beyond these rather sporadic additions to the literature there is – up to our knowledge – no previous work that directly analyzes the outcome of non-trivial global aggregation rules (such as voting mechanisms) in connection with opinion formation dynamics.

Instead of voting and aggregation, a common objective in the analysis of opinion formation is to detect conditions when the agents achieve *consensus*. This is a very desirable property in models without intrinsic opinions with continuous (e.g., French-DeGroot [19,20] or bounded confidence models [39]) or discrete opinions (such as the voter model [40,41] or k -majority [42]). Note that consensus makes a subsequent voting step trivial. In contrast, the FJ model does not lead to consensus due to distortion by heterogeneous intrinsic opinions. Some model aspects have been studied to achieve it, e.g., time pressure [43]. This variant represents a co-evolutionary process, in which the cost of the distance to neighbor opinions increases over time. The cost for the intrinsic opinion remains constant and thereby has diminishing impact. The authors study the convergence time of deterministic and randomized dynamics and characterize network structures that lead to consensus.

In addition, there is an extremely large body of literature on *influencing* agents in opinion formation models, which is too vast to be reviewed here in detail. Generally, the goal in this literature is to influence the intrinsic opinions or the network structure (using seeding, prices, biases, information design, contracts, etc.) in order to ensure favorable equilibrium properties (such as, e.g., consensus or maximizing the spread of a certain opinion). While our approach of changing the susceptibility to the voting outcome falls into this domain, we are not aware of any previous work studying this particular aspect.

An extended abstract of this paper appears in the proceedings of AAMAS 2025 [44].

3. Model and preliminaries

A *median opinion game* is a tuple $\mathcal{G} = (G, s, \alpha, \beta, \gamma)$. We are given a set V of n agents. They constitute the node set of a directed graph $G = (V, E)$. Each agent $i \in V$ has an intrinsic opinion $s_i \in \mathbb{R}$. Moreover, each agent can choose a public opinion $z_i \in \mathbb{R}$ as a strategy. The individual cost for agent i is given by the weighted sum of distances to their intrinsic opinion, the public opinions of their out-neighbors $N(i)$, and the median of public opinions $\text{med}(\mathbf{z})$:

$$\text{cost}_i(\mathbf{z}) = \alpha_i |s_i - z_i| + \sum_{j \in N(i)} \beta_{ij} |z_j - z_i| + \gamma_i |\text{med}(\mathbf{z}) - z_i|.$$

All $\alpha_i, \gamma_i, \beta_{ij} \in \mathbb{N}_0$ are constant and non-negative integers. β_{ij} is defined as the weight of the edge (i, j) . In case that all $\alpha_1 = \alpha_2 = \dots = \alpha_n$, $\gamma_1 = \gamma_2 = \dots = \gamma_n$, and $\beta_{i,j} = \beta_{i',j'}$ for all $(i, j), (i', j') \in E$ we call the opinion game *uniform*. We assume that the median function $\text{med}(\cdot)$ always refers to the lower median, i.e., the $\lceil n/2 \rceil$ th position in the ranking when ordered from smallest to largest number.

We consider opinion dynamics starting from any initial strategy profile \mathbf{z} . In each step, one agent i is chosen, either randomly or by some deterministic selection rule. Agent i updates their public opinion to a best response $z'_i \in \arg \min_{y \in \mathbb{R}} \text{cost}_i(y, \mathbf{z}_{-i})$, i.e., an opinion that minimizes their cost function. Note that this choice depends on $\text{med}(y, \mathbf{z}_{-i})$ rather than $\text{med}(\mathbf{z})$. Clearly, there might be

several best responses (c.f. an example directly after Definition 1). We use $B_i(\mathbf{z}) = \arg \min_{y \in \mathbb{R}} \text{cost}_i(y, z_{-i})$ to denote the set of all best responses for i in \mathbf{z} . We define $\text{best}_i^-(\mathbf{z}) = \min B_i(\mathbf{z})$ and $\text{best}_i^+(\mathbf{z}) = \max B_i(\mathbf{z})$ as the smallest and largest best-response for agent i in profile \mathbf{z} , respectively. For some of our results, we assume that each agent chooses a best response that minimizes the distance to their latest public opinion. We call this the *lazy* best-response and denote it by $\text{best}_i^l(\mathbf{z})$. Note that the intrinsic opinion of an agent remains static.

In the following we characterize best-response strategies. We use the following definition.

Definition 1 (Weighted Median). Given a vector $\mathbf{x} = (x_i)_{i \in [n]}$ with n values x_i where $x_i \leq x_{i+1}$, and a vector $\mathbf{w} = (w_i)_{i \in [n]}$ of weights $w_i \geq 0$, a value a is called *weighted median* of $(\mathbf{x}; \mathbf{w})$ if

$$\sum_{\substack{j \in [n] \\ x_j < a}} w_j \leq \frac{\|\mathbf{w}\|_1}{2} \quad \text{and} \quad \sum_{\substack{j \in [n] \\ x_j > a}} w_j \leq \frac{\|\mathbf{w}\|_1}{2}.$$

Note that for each $(\mathbf{x}; \mathbf{w})$ there is at least one $j \in [n]$ such that x_j is a weighted median. The weighted median is unique, unless we have $\sum_{j=1}^{i^*} x_j = \|\mathbf{w}\|_1/2$ for some value $i^* \in [n-1]$, because then each value $a \in [x_{i^*}, x_{i^*+1}]$ is a weighted median.

Proposition 1. A value z_i^* is a best response for agent i against z_{-i} if and only if it is a weighted median of

$$(s_i, z^{N(i)}, \text{med}(z_i^*, z_{-i}); \alpha_i, \beta^{N(i)}, \gamma_i),$$

where $z^{N(i)}$ is the vector of public opinions z_j of all neighbors $j \in N(i)$ of agent i and $\beta^{N(i)}$ the corresponding vector of weights β_{ij} for all $j \in N(i)$.

Proof. Consider $\text{cost}_i(z_i, z_{-i})$ as a function of z_i . We show that the function is piecewise linear and decreases until the first weighted median is reached, stays constant until the last weighted median is exceeded, and increases afterwards.

Let $V' = \{i\} \cup N(i)$, $x_j = z_j$, $x_i = s_i$, $w_i = \alpha_i$, and $w_j = \beta_{ij}$ for each $j \in N(i)$.

First, assume for simplicity that the weight γ_i of the median of public opinions is zero and can therefore be ignored. If z_i is strictly smaller than the smallest weighted median, increasing z_i results in a linear cost reduction of the function, until the next value from $\{z_j \mid j \in N(i)\} \cup \{s_i\}$ is reached. The gradient of this linear function is $\sum_{j \in V' : x_j \leq z_i} w_j - \sum_{j \in V' : x_j > z_i} w_j$. Each time a value from $\{z_j \mid j \in N(i)\} \cup \{s_i\}$ is reached, the slope changes but remains negative until the first weighted median is reached and turns positive only after the last weighted median is passed.

Now if the median of public opinions has a weight strictly larger than zero, the argument is slightly more complicated. The contribution of the median of public opinions remains the same as long as the public opinion of agent i is strictly larger or strictly smaller than the median. If z_i is changed in the direction of the median, then the cost due to the median decreases linearly until the median is reached. At this point the public opinion of agent i itself is the median of public opinions. If z_i is changed further in the same direction, the cost contribution of the median remains zero until the value of the next opinion is reached and that value takes on the role of the median. Afterwards, the cost contribution starts to grow linearly again. Hence, the cost of agent i depending on z_i remains a convex and piecewise linear function. The minima of this function are exactly the weighted medians. \square

Proposition 1 has many consequences. Best responses represent a closed interval in \mathbb{R} , i.e., $B_i(\mathbf{z}) = [\text{best}_i^-(\mathbf{z}), \text{best}_i^+(\mathbf{z})]$. Moreover, the smallest, largest, and lazy best responses of every agent i are monotone in the strategy profile \mathbf{z} .

Corollary 1. Suppose $\mathbf{z} \leq \mathbf{z}'$ pointwise. Then $\text{best}_i^-(\mathbf{z}) \leq \text{best}_i^-(\mathbf{z}')$, $\text{best}_i^+(\mathbf{z}) \leq \text{best}_i^+(\mathbf{z}')$, and $\text{best}_i^l(\mathbf{z}) \leq \text{best}_i^l(\mathbf{z}')$.

Yet another consequence of Proposition 1 is a restriction of the strategy space. If agents always choose their best responses from $\text{best}_i^+(\mathbf{z})$, $\text{best}_i^-(\mathbf{z})$ or $\text{best}_i^l(\mathbf{z})$, then at any point of time the opinions in the strategy profile remain a subset of the initial public and intrinsic opinions.

Observation 1. In a state \mathbf{z} , suppose an agent i chooses one of the three best responses $\text{best}_i^-(\mathbf{z})$, $\text{best}_i^+(\mathbf{z})$, and $\text{best}_i^l(\mathbf{z})$. Then the chosen strategies come from $\{s_i\} \cup \{z_j \mid j \in V\}$.

Proof. Proposition 1 guarantees that all three best responses are a weighted median of

$$(s_i, z^{N(i)}, \text{med}(\text{best}_i(\mathbf{z}), z_{-i}); \alpha_i, \beta^{N(i)}, \gamma_i). \tag{1}$$

If the weighted median is unique, then $\text{best}_i^-(\mathbf{z}) = \text{best}_i^+(\mathbf{z}) = \text{best}_i^l(\mathbf{z})$ and the median has to be one of the values from the vector in (1), which are always included in $\{s_i\} \cup \{z_j \mid j \in V\}$. If $B_i(\mathbf{z})$ is an interval, then each of the three best responses is one of the borders, which are again included in $\{s_i\} \cup \{z_j \mid j \in V\}$. \square

Observe that no new opinions are introduced during such dynamics. For a given initial profile, let $I = \{s_j, z_j \mid j \in V\}$ be the set of *initial opinions* in \mathbf{z} and s . When considering such dynamics, we can assume w.l.o.g. that the set of possible opinions is restricted to I with at most $2n$ opinions. Moreover, the lower weighted median is independent of the numeric opinion values and based only on (weights and) the total ordering of opinions. As such, we can assume that the available opinions are the integers $1, 2, \dots, |I| \leq 2n$.

4. Analysis of equilibria

We start by analyzing the properties of equilibria and consider convergence of sequential best-response dynamics. Towards this end, we study an efficient method that schedules best-response dynamics to enable convergence to equilibrium, the 2-PHASE algorithm. We show that the set of equilibria has substantial structure. For any initial profile \mathbf{z} , let $\Xi = \{\mathbf{z}' \mid \min I \leq z'_i \leq \max I \text{ for all } i \in V\}$ be the subset of strategy profiles bounded by initial opinions.

Proposition 2. *For any median opinion game \mathcal{G} with initial profile \mathbf{z} , the set $\{\mathbf{z}^* \mid \mathbf{z}^* \in \Xi \text{ is an equilibrium}\}$ forms a complete lattice with respect to \leq (componentwise).*

Proof. The set of strategy profiles Ξ forms a complete lattice with respect to \leq (componentwise). By [Corollary 1](#) the multi-function f that maps each component to $B_i(\mathbf{z})$ is monotone. By [Observation 1](#) and monotonicity we see that if $\mathbf{z} \in \Xi$, then $B_i(\mathbf{z}) \subseteq \Xi$ for all $i \in V$. Hence, $f : \Xi \rightarrow 2^\Xi$, and the fixed-points of f (in which $z_i \in B_i(\mathbf{z})$ for all $i \in V$) are exactly the equilibria from Ξ . An extension of the Knaster-Tarski theorem to order-preserving multi-functions [[45–47](#)] shows that the set of fixed-points of f forms a complete lattice. \square

Since every complete lattice consists of at least one element, [Proposition 2](#) implies that every game has at least one equilibrium. Moreover, when we restrict ourselves to equilibria from Ξ , the componentwise maximal and minimal equilibria are both unique.

Beyond the mere existence of equilibria, we are interested in equilibria that can be *reached* by best-response dynamics. We call such an equilibrium *reachable* (from initial profile \mathbf{z}). Indeed, for every game and initial strategy profile \mathbf{z} , there is at least one reachable equilibrium, and it can be computed efficiently using the following 2-PHASE algorithm (see [Algorithm 1](#)).

Algorithm 1: 2-PHASE algorithm (ascending).

Input: arbitrary initial configuration \mathbf{z}

Phase 1

$U(\mathbf{z}) \leftarrow \{i \mid z_i \notin B_i(\mathbf{z}), \text{best}_i^+(\mathbf{z}) > z_i\}$

while $U(\mathbf{z}) \neq \emptyset$ **do**

let $i \in \arg \max\{\text{best}_i^+(\mathbf{z}) \mid i \in U(\mathbf{z})\}$
 $z_i \leftarrow \text{best}_i^+(\mathbf{z})$
 $U(\mathbf{z}) \leftarrow \{i \mid z_i \notin B_i(\mathbf{z}), \text{best}_i^+(\mathbf{z}) > z_i\}$

Phase 2

$D(\mathbf{z}) \leftarrow \{i \mid \text{best}_i^+(\mathbf{z}) < z_i\}$

while $D(\mathbf{z}) \neq \emptyset$ **do**

let $i \in \arg \min\{\text{best}_i^+(\mathbf{z}) \mid i \in D(\mathbf{z})\}$
 $z_i \leftarrow \text{best}_i^+(\mathbf{z})$
 $D(\mathbf{z}) \leftarrow \{i \mid \text{best}_i^+(\mathbf{z}) < z_i\}$

In the beginning of each round $k \geq 1$ in phase 1, consider the current state $\mathbf{z}^{(k)}$ (where $\mathbf{z}^{(1)} = \mathbf{z}$). We consider the set of agents that want to deviate to a larger opinion. We pick an agent i that wants to deviate to the largest best response and allow them to deviate. We continue phase 1 until no agent wants to deviate to a higher opinion. Let $\hat{\mathbf{z}}$ be the state that emerges after phase 1.

In the beginning of each round $k \geq 1$ in phase 2, consider the current state $\mathbf{z}^{(k)}$ (where $\mathbf{z}^{(1)} = \hat{\mathbf{z}}$). We consider the set of agents that want to deviate to a lower opinion. We pick an agent i whose largest best response is smallest and allow them to deviate. We continue phase 2 until no agent wants to deviate to a lower opinion. Let $\hat{\mathbf{z}}$ be the state that emerges after phase 2.

There are two variants of the algorithm. We call the one described above the *ascending* 2-PHASE algorithm. In the *descending* 2-PHASE algorithm, the phases are reversed, we consider smallest best responses best_i^- in both phases, and among the agents that want to deviate we choose one for which best_i^- is smallest in phase 1 and largest in phase 2. If not stated otherwise, we only consider the ascending 2-PHASE algorithm in the following proofs, as the arguments for the two variants are very analogous. We first observe that our algorithms indeed converge to an equilibrium.

Lemma 1. *For any median opinion game \mathcal{G} and initial profile \mathbf{z} , both variants of the 2-PHASE algorithm compute an equilibrium in $O(n)$ steps. The resulting equilibrium is unique over possible tie-breaking in choosing the deviating agent i in each step.*

Proof. We only consider the ascending 2-PHASE algorithm. Consider any round k of phase 1 in which some agent i is chosen to deviate. Profile $\mathbf{z}^{(k+1)}$ evolves after i deviated to $z_i^{(k+1)} = \text{best}_i^+(\mathbf{z}^{(k)})$. Consider any other agent $j \neq i$. Due to monotonicity ([Corollary 1](#)), $\text{best}_j^+(\mathbf{z}^{(k+1)}) \geq \text{best}_j^+(\mathbf{z}^{(k)})$. The set of agents that want to deviate to smaller opinions can only shrink. Since best responses are weighted medians, $\text{best}_j^+(\mathbf{z}^{(k+1)})$ can grow to at most $z_i^{(k+1)}$. Thus, if $\text{best}_j^+(\mathbf{z}^{(k)}) > z_i^{(k+1)}$ previously, then after the deviation $\text{best}_j^+(\mathbf{z}^{(k+1)}) = \text{best}_j^+(\mathbf{z}^{(k)})$ still holds. Overall, for each agent $j \in V$ we obtain $\text{best}_j^+(\mathbf{z}^{(k+1)}) \leq \max\{\text{best}_j^+(\mathbf{z}^{(k)}), z_i^{(k+1)}\}$. We choose agent i to deviate such that $\text{best}_i^+(\mathbf{z}^{(k)})$ is maximal. Hence, all agents in $U(\mathbf{z}^{(k+1)})$ want to deviate to at most $z_i^{(k+1)}$. Thus, $z_i^{(k+1)} = \text{best}_i^+(\mathbf{z}^{(k+1)})$ remains the largest best-response throughout the rest of phase 1. Each agent i deviates at most once in phase 1.

If there are several agents $i \in U(\mathbf{z}^{(k)})$ with maximal $z'_i = \text{best}_i^+(\mathbf{z}^{(k)}) = \max\{\text{best}_j^+(\mathbf{z}^{(k)}) \mid j \in U(\mathbf{z}^{(k)})\}$, by the above arguments they will be chosen sequentially in the algorithm and all deviate to z' independent of their ordering. As such, the state $\hat{\mathbf{z}}$ emerging at the end of phase 1 is unique.

The analysis of phase 2 follows symmetrically. By the same argument, we choose the agent from $D(\mathbf{z}^{(k)})$ with smallest $\text{best}_i^+(\mathbf{z}^{(k)})$ in phase 2. In the emerging profile $\mathbf{z}^{(k+1)}$, by monotonicity, we only lower all $\text{best}_j^+(\mathbf{z}^{(k+1)}) \leq \text{best}_j^+(\mathbf{z}^{(k)})$. This maintains the invariant that no player wants to deviate to a higher opinion. Moreover, by the properties of weighted medians, $\text{best}_j^+(\mathbf{z}^{(k+1)}) \geq \min\{\text{best}_j^+(\mathbf{z}^{(k)}), z_j^{(k+1)}\}$. Thus, each agent deviates at most once, and if there are several agents $i \in D(\mathbf{z}^{(k)})$ with smallest $\text{best}_i^+(\mathbf{z}^{(k)})$, the order in which they are allowed to deviate does not affect their chosen opinion. \square

For each initial profile \mathbf{z} , the two versions of the 2-PHASE algorithm yield two unique reachable equilibria. These equilibria bound all of the reachable equilibria from \mathbf{z} .

Theorem 1. *For any median opinion game \mathcal{G} and initial profile \mathbf{z} , the ascending (descending) 2-PHASE algorithm computes the unique maximum (minimum) reachable equilibrium in polynomial time. The maximum (minimum) median opinion among all reachable equilibria occurs at the maximum (minimum) equilibrium.*

Proof. Consider the state $\hat{\mathbf{z}}$ after phase 1. We will first show inductively that $\hat{\mathbf{z}} \geq \bar{\mathbf{z}}$ (componentwise) for any state $\bar{\mathbf{z}}$ reachable by best-response dynamics from \mathbf{z} . We number the agents based on the round they deviated in phase 1. The statement clearly holds for agent 1. $\text{best}_1^+(\mathbf{z})$ is the largest opinion any agent wants to deviate to in \mathbf{z} . No deviation of any agent j to $x \leq \text{best}_1^+(\mathbf{z})$ can make $\text{best}_k^+(\mathbf{z})$ grow to beyond x , for any $k \in V$. Hence, in any sequence of best responses starting from \mathbf{z} , we maintain the invariant that no agent ever wants to deviate to an opinion $x > \text{best}_1^+(\mathbf{z})$. This proves the statement for agent 1. Now, inductively, given that the statement holds for agents $1, \dots, k-1$, consider agent k . We again use \mathbf{z}_k to denote the state in the beginning of round k of phase 1. Agents $j = 1, \dots, k-1$ are at their maximal reachable opinions $z_j \geq \text{best}_k^+(\mathbf{z}_k)$. By monotonicity, this only increases $\text{best}_j^+(\mathbf{z}_j)$ for $j = k, k+1, \dots$. By the choice of k , all other agents $j \geq k$ want to deviate to opinions at most $\text{best}_k^+(\mathbf{z}_k)$, even after k deviates to $\text{best}_k^+(\mathbf{z}_k)$. Thus, for any sequence of best responses from \mathbf{z} , given that agents $1, \dots, k-1$ never play higher opinions, agents $k, k+1, \dots$ never want to deviate to any opinion $x > \text{best}_k^+(\mathbf{z}_k)$. This proves the statement for agent k , and thus all agents that deviate during phase 1 of the 2-PHASE algorithm. Now finally, consider all agents i with $z_i = z_i$ who did not deviate during phase 1. By monotonicity, these agents must also have $\hat{z}_i \leq z_i$ in any reachable state.

Consider any reachable equilibrium $\bar{\mathbf{z}}$. Since $\hat{\mathbf{z}} \geq \bar{\mathbf{z}}$ for any reachable state $\bar{\mathbf{z}}$, phase 2 maintains the invariant that $\text{best}_i^+(\mathbf{z}) \geq \text{best}_i^+(\bar{\mathbf{z}})$ for all $i \in V$, by monotonicity. Thus, the equilibrium computed by 2-PHASE fulfills $\hat{\mathbf{z}} \geq \bar{\mathbf{z}}$.

An analogous argument using the descending 2-PHASE algorithm proves the statement regarding the minimum equilibrium. By monotonicity, the maximum and minimum median opinions among the reachable equilibria have to occur at the (componentwise) maximum and minimum equilibria, respectively. \square

Lemma 1 shows that for every game \mathcal{G} with initial profile \mathbf{z} there exists a sequence of deviations that converges in an equilibrium. In other words, median opinion games are *weakly acyclic*. The 2-PHASE algorithm constructs such a sequence using a mild form of coordination in which the agent that wants to deviate to the highest or lowest opinion updates in each step. Notably, the sequence requires only a number of deviations linear in the number of players.

Interestingly, if the agents deviate in an *arbitrary order* without coordination, the sequence may not even converge to an equilibrium and run into cycles. As such, some mild form of coordination is necessary to steer agents towards an equilibrium.

Example 1. Consider a game with three agents $i, j, k \in V$ where the edges (i, j) , (j, k) and (k, i) form a directed cycle. All edge weights are 1 and all other weights are sufficiently small. Furthermore, the intrinsic opinion of all agents equals 0 whereas the initial public opinions are given by $z_i = 0, z_j = 1, z_k = 1$.

Clearly, the best response of agent k is to deviate to opinion 0 while the best response of i is 1. After an update of k first and then i , we arrive at a symmetric state. It is easy to see that the deviation sequence of k, i, j, k, i, j recovers the initial strategy profile.

Let us also briefly discuss a special class of games with more symmetry. We consider games in which G is an undirected graph with symmetric edge weights (i.e., $\beta_{ij} = \beta_{ji}$) and we have uniform median weights (i.e., $\gamma_i = \gamma$). In these games, we can adapt arguments in [48, Claim 2.4] to observe that they admit an *exact potential function* $\Phi(\mathbf{z})$. Whenever an agent performs a unilateral deviation and changes their cost by a certain value $\Delta > 0$, the value of the potential function changes by exactly Δ .

Proposition 3. *For any median opinion game $\mathcal{G} = (G, s, \alpha, \beta, \gamma)$ where $\beta_{ij} = \beta_{ji}$ and $\gamma_i = \gamma$, for all $i, j \in V$,*

$$\Phi(\mathbf{z}) = \sum_{i \in V} \alpha_i |s_i - z_i| + \frac{1}{2} \sum_{i \in V} \sum_{j \in N(i)} \beta_{ij} |z_i - z_j| + \sum_{i \in V} \gamma |\text{med}(\mathbf{z}) - z_i|$$

is an exact potential function.

Proof. Suppose an agent i deviates from opinion z_i to z'_i . Let \mathbf{z} be the initial state and $\mathbf{z}' = (z_{-i}, z'_i)$ be the state after deviation of i .

Case 1: Assume that $\text{med}(\mathbf{z}) = \text{med}(\mathbf{z}')$. In such a deviation, we can express the distance to the median (in cost cost_i and potential Φ) equivalently as part of the neighborhood distances: Suppose that, instead of measuring the distance to the median, there is an auxiliary agent a that plays strategy $z_a = \text{med}(\mathbf{z})$, and that $\beta_{ia} = \beta_{ai} = \gamma$ for all $i \in V$. In this formulation, it is straightforward to apply symmetry arguments exactly as in [48, Claim 2.4] to verify that $\Phi(\mathbf{z}') - \Phi(\mathbf{z}) = \text{cost}_i(\mathbf{z}') - \text{cost}_i(\mathbf{z})$.

Case 2: Now suppose $\text{med}(z) < \text{med}(z')$ (the case $\text{med}(z) > \text{med}(z')$ can be argued symmetrically). This means that $z_i \leq \text{med}(z)$ and $z'_i \geq \text{med}(z')$. We split the deviation into (at most) three parts with the values $z_i \leq z_i^{(1)} \leq z_i^{(2)} \leq z'_i$, where we denote $z^{(\ell)} = (z_{-i}, z_i^{(\ell)})$ for each $\ell = 1, 2$. Let $z_i^{(1)} = \min\{x \in [z_i, z'_i] \mid \text{med}(z_{-i}, x) = x\}$ and $z_i^{(2)} = \max\{x \in [z_i, z'_i] \mid \text{med}(z_{-i}, x) = x\}$. Now assume that i first deviates from z_i to $z_i^{(1)}$, then from $z_i^{(1)}$ to $z_i^{(2)}$, then from $z_i^{(2)}$ to z'_i . In the first part, i raises their opinion to meet $\text{med}(z) = \text{med}(z^{(1)})$. Thus, by Case 1, Φ changes exactly as cost_i in this first part. Similarly, for the third part we see that $\text{med}(z^{(2)}) = \text{med}(z')$, so Case 1 applies. For the rest of the proof we consider the second part.

First suppose the number of agents n is even. In this case, we introduce an auxiliary agent b with opinion $-B$ for some prohibitively small value $-B$. This agent has $\alpha_b = 0, \beta_{ib} = \beta_{bi} = 0$ for all $i \in V$, and they are never selected for an improvement step. The agent ensures that n becomes odd. Since z_b is an extremely small number, the existence of b only ensures that the lower median becomes the unique median. As such, b does not influence the improvement dynamics in the game. It shifts the value of the potential function by $|\text{med}(z) + B|$. Observe that all previous arguments remain unaffected by this shift.

Hence, we can assume for simplicity that the number of agents n is odd. The deviation in the second part raises the median by $\delta = z_i^{(2)} - z_i^{(1)}$. We can assume i is the *unique* median agent before and after this deviation. Therefore, it is possible to partition the set of agents into exactly $(n - 1)/2$ agents, which are at or above the median in both $z^{(1)}$ and $z^{(2)}$, and $(n - 1)/2$ agents, which at or below the median in both $z^{(1)}$ and $z^{(2)}$. For the ones above, Φ decreases by δ for each of them; for the ones below, Φ increases by δ for each of them. Since i has the median opinion, their contribution to Φ is 0. Thus, overall there is no change in Φ due to the change in the median – all changes in Φ are due to the distances of i in their neighborhood and to their intrinsic opinion. Similarly, in the cost cost_i the distance to the median remains 0 since i is playing exactly the median opinion before and after the deviation. Changes in cost_i also result only from distances of i in their neighborhood and to their intrinsic opinion. These changes are exactly the ones in Φ by the same calculations as in [48, Claim 2.4] and for Case 1. \square

The proposition implies that in every game of this class *every sequence* of best-response deviations is cycle-free (in contrast to general games as in Example 1). Since we restrict attention to strategies from the set of initial opinions $I = \{s_j, z_j \mid j \in V\}$ (c.f. Observation 1), we see that every sequence of best-response deviations is finite and must converge to an equilibrium. However, the length of such a sequence might still be very long. Indeed, even in this special class of games, there exist games and initial states from which there is an exponentially long sequence to an equilibrium – the construction described in [49, Theorem 3.8] can be easily adapted to median opinion games (with $\gamma \approx 0$).

5. Number of equilibria

By Proposition 1, best responses form a closed interval. Hence, for some initial state z , there are trivial examples with infinitely many reachable equilibria. In this section we will show that even for *lazy* best-response dynamics (in which all best responses remain within the set I of initial opinions) the number of reachable equilibria can be very large, and the median opinions can be very diverse. We can show these properties even for instances with *undirected* graphs.

Let $k = 2^q$ with $q \in \mathbb{N}, q \geq 2$. In the following we define a uniform family of instances T_k (games and initial states) with k opinions and $n = \text{poly}(k)$ agents with the following properties:

- For each opinion j ($1 \leq j \leq k$), there exists a reachable equilibrium where the median of the public opinions equals the public opinion j .
- The number of reachable equilibria is exponentially large in k .

The family $\mathcal{G}_k = (T_k, s, \alpha, \beta, \gamma)$ with corresponding initial profile z is defined as follows. For a fixed $k, T_k(\alpha, \beta, \gamma)$ consists of $\log k + 1$ layers with $s \subset \{0, \dots, k\}$. See Fig. 1 for an example. The first $\log k$ layers consist of k agents each, while the last layer has $2k \log k$ agents.

- The j th agent of layer 0 is denoted p_j (colored orange) and is initialized with intrinsic opinion $z_{p_j} = s_{p_j} = j$. We call these agents *player agents*.
- For $1 \leq \ell \leq \log k$ the j th agent of the ℓ th layer is denoted $t_{j,\ell}$. All agents $t_{j,\ell}$ with $1 \leq j \leq k, 1 \leq \ell \leq \log k$ (colored blue) are initialized with opinion $z_{t_{j,\ell}} = s_{t_{j,\ell}} = \lfloor \frac{j-1}{4\ell} \rfloor \cdot 4\ell + ((j + 2\ell) \bmod 4\ell)$. They are called *transfer agents*.
- The j th agent from the last layer $\log k$ is denoted m_j (colored yellow) and they are initialized with opinion $z_{m_j} = s_{m_j} = k/2$. These agents are called *median stabilizing agents*.

In addition to these layers, we have $\log k - 1$ *intermediate layers* (colored pink), layer $1 \leq i \leq \log k - 1$ containing $2^{\log k - 1 - i}$ many agents. The j th agent of the ℓ th layer is denoted $e_{j,\ell}$ and is initialized with opinion $z_{e_{j,\ell}} = s_{e_{j,\ell}} = 1$. We call these agents *elimination agents*. It remains to define the edges of T_k .

- **Player agents:** for each $i = 2j - 1, j \in \mathbb{N}$, we connect p_i and p_{i+1} with an edge of weight $w_1 = w_2 + \alpha + \beta + \gamma + 1$. Additionally, p_j will be connected to a transfer agent $t_{j,1}$ (blue) with an edge of weight w_2 .
- **Elimination agents:** agents at level $1 \leq \ell \leq (\log k) - 2$ are connected to transfer agents at level ℓ and level $(\ell + 1)$ with edges of weight

$$w_{2\ell+1} = 2^{\ell+1} \cdot w_{2\ell+2} + \alpha + \beta + \gamma + 1$$

and $w_{2\ell+2} = w_{2\ell+3} + \alpha + \beta + \gamma + 1,$




Fig. 1. Tournament graph T_k for $k = 8$ with 8 player agents (orange), 16 transfer agents (blue), 3 elimination agents (purple), and 48 stabilizing agents (yellow) from which only 8 are depicted. Public opinion is presented at the top, and intrinsic opinion is at the bottom of each node. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

respectively. This means that weights are propagated *bottom up*, with smaller values at larger levels. Elimination agent $e_{j,\ell}$ is connected to transfer agents $t_{(j-1)2^{\ell+1},\ell}$ to $t_{j,2^{\ell+1},\ell}$ and transfer agents $t_{(j-1)2^{\ell+1},\ell+1}$ to $t_{j,2^{\ell+1},\ell+1}$ if $\ell < (\log k) - 1$. For $\ell = (\log k) - 1$, the single elimination agent $e_{1,\ell}$ at intermediate layer $\ell = (\log k) - 1$ is connected to all transfer agents at layer $(\log k) - 1$ with edge weight $w_{2^{\ell+1}} = 2^{\ell+2} \cdot (\ell + 1) \cdot w_1 + \alpha + \beta + \gamma + 1$. Furthermore, the agent $e_{1,\ell}$ is connected to all median stabilizing agents with edge weight $w_{2^{\ell+2}} = \alpha + \beta + \gamma + 1$.

Lemma 2. Suppose that $k = 2^q$ with $q \in \mathbb{N}, q \geq 2$. Let $T_k(\alpha, \beta, \gamma)$ be a uniform game with $\alpha, \beta, \gamma \in \mathbb{N}_0$ with k opinions and $O(k \log k)$ agents. For each $1 \leq j \leq k$ there exists an equilibrium where the global median takes on the value j .

Proof. For each opinion $j \in [k]$, we will describe an update sequence, after which the opinion j will be the median opinion. The main idea is to propagate the opinion j through the elimination rounds of the tournament until it reaches the final elimination agent $e_{1,\log(k)-1}$. After opinion j reaches this agent, all median stabilizing agents are updated. They adopt opinion j , after which the global median becomes opinion j . Finally, all other nodes are updated, starting from round one and iterating down through the elimination rounds to the last elimination agent. In this last step, agents who have been previously updated cannot change their opinion anymore due to the design of the weights w_j , resulting in an equilibrium where opinion j is the global median.

The first updated agent will be the player agent $p_{j+(-1)^{j+1}}$, which is connected to p_j by an edge with weight w_1 . When updating agent $p_{j+(-1)^{j+1}}$ they will adopt opinion j , as the weight $w_1 = 2w_2 + \alpha + \gamma + 1$ is larger than the sum of all other weights connected to this agent. Next, the connected transfer agents $t_{j,1}$ and $t_{j+(-1)^{j+1},1}$ are updated. Both transfer agents will adopt the opinion j as well, as $w_2 = w_3 + \alpha + \gamma + 1$. Now the elimination agent $e_{\lfloor j/4 \rfloor, 1}$, to which the transfer agents $t_{j,1}$ and $t_{j+(-1)^{j+1},1}$ are connected to, is connected to $3 = 2^1 + 1$ agents with opinion j . When updating this agent, the three transfer agents with the opinion j are enough to set the weighted local median to j , as we have chosen $w_3 = \gamma + \alpha + 4 \cdot w_4 + 1$. From here on, we proceed by alternatingly updating all transfer agents from the next level that are connected to the updated elimination agent and then updating the next level elimination agent connected to the transfer agents. Recall that we set $w_{2^{\ell+1}} = 2^{\ell+1} \cdot w_{2^{\ell+2}} + \alpha + \beta + \gamma + 1$. This choice for the weights ensures the following behavior: if $2^\ell + 1$ transfer agents from level ℓ with the same opinion j are connected to one elimination agent $e_{\cdot,\ell}$, the local median of $e_{\cdot,\ell}$ will be j . Therefore, our alternating update strategy propagates opinion j to the final elimination agent $e_{1,\log(k)-1}$.

After agent $e_{1,\log(k)-1}$ has adopted the opinion j , all the median stabilizing agents $m_1, \dots, m_{n'}$ are updated. As the edge weight $w_{2(\log(k)-1)+1} = \alpha + \gamma + 1$, these agents will adopt the opinion of the elimination agent $e_{1,\log(k)-1}$, which is j . As $n' = 2k \log k$, and there are $(\log(k) - 1) \cdot k$ transfer agents, k player agents, and $k/2 - 1$ elimination agents, more than half of all agents have the opinion j , implying that the median opinion is also j . See Fig. 2 for an example of the state at this step.

Finally, the residual agents have to be updated to find an equilibrium. To find the equilibrium, we update the agents layer by layer. First, all player agents are updated in an arbitrary order. Afterwards from $\ell = 1$ to $\ell = (\log(k) - 1)$, first all transfer agents $t_{\cdot,\ell}$




Fig. 2. Tournament graph T_k for $k = 8$ after opinion 7 has been propagated to the stabilizing nodes. Public opinion is presented at the top, and intrinsic opinion is at the bottom of each node.

are updated, followed by all elimination agents $e_{i,c}$. Note that all transfer agents and elimination agents who were updated before this final phase will not change their opinion when updated a second time. After this update phase, all agents are stable, as the player agents will not change their opinion, the transfer agents will only change their opinion if the previous level was changed, and the elimination agents will only change their opinion if the transfer agents from the previous level have changed before. \square

For each player pair p_i, p_{i+1} for $i = 2j + 1, j \in \mathbb{N}$, in each equilibrium both agents have either opinion i or opinion $i + 1$. The following corollary is a consequence.

Corollary 2. $T_k(\alpha, \beta, \gamma)$ has an exponential (in k) number of reachable equilibria.

Proof. We consider two neighboring player agents p_i and p_{i+1} that have not been updated and are connected by an edge with weight w_1 . Depending on which agent is updated first, both agents will adopt either opinion i or opinion $i + 1$. As there are k opinions, in the end, there are $k/2$ different opinions present in the player agents, as player agents are never influenced by their transfer agents. Since for each pair p_i and p_{i+1} both opinions i and $i + 1$ can win, there are at least $2^{k/2}$ reachable equilibria for the tournament graph T_k . \square

6. Stabilizing the median

Let $\mathcal{G} = (G, s, \alpha, \beta, \gamma)$ be a median opinion game with initial profile \mathbf{z} . We call the pair \mathcal{G}, \mathbf{z} stable if all equilibria reachable from \mathbf{z} share the same global median. In this section we consider the question of how to change the median weights γ_i of a given opinion game \mathcal{G} with a fixed initial profile \mathbf{z} such that the pair becomes stable. We say that these changes stabilize \mathcal{G} and \mathbf{z} .

Observe that stabilizing \mathcal{G} and \mathbf{z} is always possible when all the median weights γ_i are increased by some sufficiently large value δ . In particular, there is only one reachable equilibrium – and hence a unique global median – if γ_i is larger than the weight of the intrinsic opinion plus the total weight of all neighbors in $N(i)$, for each agent i . This is the case, e.g., if we set δ to $\xi = 1 + \max\{\alpha_i + (\sum_{j \in N(i)} \beta_{ij}) - \gamma_i \mid i \in V\}$. However, intuitively, in many cases much lower values for δ will suffice; the minimum value of δ should depend on \mathcal{G} and \mathbf{z} . The main result of this section is that we can efficiently compute the minimum value of δ with the help of the 2-PHASE algorithm described in Section 4.

Theorem 2. There exists a polynomial-time algorithm that, given a median opinion game $\mathcal{G} = (G, s, \alpha, \beta, \gamma)$ and an initial profile \mathbf{z} , finds the minimal value $\delta \geq 0$ such that the global median of \mathbf{z} is stabilized when all the median weights γ_i are increased by δ .

We prove Theorem 2 in several steps, taking a closer look at the two phases of (the two variants of) the 2-PHASE algorithm. Let us first introduce some notation. Let $\hat{\mathbf{z}}$ and $\tilde{\mathbf{z}}$ denote the strategy profiles that are obtained from \mathbf{z} after phase 1 and after phase 2 of the ascending 2-PHASE algorithm, respectively. Note that Lemma 1 and its proof imply that both $\hat{\mathbf{z}}$ and $\tilde{\mathbf{z}}$ are well-defined and independent of tie-breaking in the ordering of the updating agents in the two phases. We define $\hat{\mathbf{z}}$ and $\tilde{\mathbf{z}}$ accordingly with respect to

the descending 2-PHASE algorithm. Furthermore, we will consider these four profiles with respect to different values δ that are added to the median weight and write $\hat{z}(\delta)$, $\check{z}(\delta)$, $\hat{z}(\delta)$ and $\check{z}(\delta)$, respectively. The proof of [Theorem 2](#) relies upon the following two lemmas.

Lemma 3. For $\delta \in [0, \xi]$, the value $\text{med}(\hat{z}(\delta))$ is monotonically non-increasing in δ . The value $\text{med}(\check{z}(\delta))$ is monotonically non-decreasing in δ .

Lemma 4. Let $\lambda, \lambda', \rho, \rho' \in [0, \xi]$ with $\lambda < \rho$ and $\lambda' < \rho'$ such that $\text{med}(\hat{z}(\lambda)) = \text{med}(\check{z}(\rho))$ and $\text{med}(\hat{z}(\lambda')) = \text{med}(\check{z}(\rho'))$.

- For $\delta \in [\lambda, \rho]$ the median $\text{med}(\hat{z}(\delta))$ is monotonically non-decreasing in δ , and
- for $\delta \in [\lambda', \rho']$ the median $\text{med}(\check{z}(\delta))$ is monotonically non-increasing in δ .

With these lemmas we can show the main theorem of this section.

Proof of Theorem 2. Recall that the ascending 2-PHASE algorithm computes the unique maximum reachable equilibrium (see [Theorem 1](#)). In this algorithm each activated agent i updates to $\text{best}_i^+(z)$, which is either one of the intrinsic opinions from \mathcal{G} , or an opinion from the initial strategy profile z (see [Observation 1](#)). Hence, the median can adopt at most $2n$ different values during the whole run of the algorithm. This observation combined with [Lemma 3](#) implies that $\text{med}(\hat{z}(\delta))$ is a monotonically non-increasing step function that can have only $O(n)$ many steps.

In order to compute the step function we consider any given state z (with opinions from I). We examine when a deviation to some opinion $z'_i \in I$ of any agent i will become or stop being a best response, respectively, and how this property depends on δ . For this we compare agent costs of any strategy choice z'_i in a (remaining) state z_{-i} to the cost of any other choice z''_i . We denote $z' = (z_{-i}, z'_i)$ and $z'' = (z_{-i}, z''_i)$. Then, using the following notation

$$A_i(z) = \alpha_i |s_i - z_i| + \sum_{j \in N(i)} \beta_{ij} |z_j - z_i| + \gamma_i |\text{med}(z) - z_i| \quad \text{and}$$

$$B_i(z) = |\text{med}(z) - z_i|$$

we can express the cost function by

$$\text{cost}_i(z', \delta) = A_i(z') + \delta B_i(z') \quad \text{and}$$

$$\text{cost}_i(z'', \delta) = A_i(z'') + \delta B_i(z'').$$

Equality $\text{cost}(z') = \text{cost}(z'')$ emerges exactly when

$$\delta = \frac{A_i(z') - A_i(z'')}{B_i(z'') - B_i(z')}.$$

We call such a value a *critical value* for δ . At this value, the preference of i for z'_i vs. z''_i changes (when increasing δ from below to above this value). Note that we assume $B_i(z') \neq B_i(z'')$ – otherwise changing δ does not change the preference of i for z'_i and z''_i , so no critical value emerges.

A critical value for δ is a rational number. Since all opinions are integers from I and all parameters $\alpha_i, \beta_{ij}, \gamma_i$ are integers, the numbers $A(\cdot)$ and $B(\cdot)$ are integers as well. More precisely, the numerator of a critical value is an integer in (D', D') where $D' = (\sum_i (\alpha_i + \gamma_i) + \sum_{i,j} \beta_{ij}) \cdot \max\{z - z' \mid z, z' \in I\}$. Likewise, the denominator is some integer in $[-D, D]$ where $D = \max\{z - z' \mid z, z' \in I\}$. For any δ in between two critical values, the execution of the 2-Phase algorithm remains invariant, since every largest best response from I in each of the emerging states remains the same. Thus, we limit our attention to critical values for δ for the step borders.

We use a separate binary search to find each step border. After $\lceil \log(D^2) \rceil + 1 \leq 2 \lceil \log D \rceil + 1$ rounds, binary search has narrowed down the search to an interval of length at most $1/D^2$. The time required for this is polynomial in the representation of the instance. The result of binary search is an interval $[x/2^k, (x+1)/2^k]$ for $k = \lceil \log(D^2) \rceil$ such that there is a step $\text{med}(\hat{z}(x/2^k)) > \text{med}(\hat{z}((x+1)/2^k))$.

Since the denominator of each critical value is an integer of (absolute) value at most D , there can be at most a single critical value for δ inside the interval. In order to find this critical value, we examine the (polynomially) many states visited by the 2-Phase algorithm executed for the lower end $\delta = x/2^k$. We check for every visited state and every agent, every pair of opinions from I and calculate the resulting critical value for δ (if any). These are only polynomially many critical values in total. Exactly one of them must be the step border in the interval, which can be seen by contradiction: Suppose by going from $\delta = x/2^k$ to $(x+1)/2^k$ we do not pass by any of these critical values. Then the largest best response for every agent in every state during the execution of the 2-Phase algorithm would remain unchanged, and hence no step in $\text{med}(\hat{z}(\delta))$ could evolve.

The step borders partition $[0, \xi]$ into $O(n)$ many non-overlapping, discrete intervals such that $\text{med}(\hat{z}(\delta))$ is invariable for each such interval $[\lambda, \rho]$. Following this, we can employ the same approach for each such interval utilizing [Lemma 4](#) and the binary search approach outlined above to derive a partition of $[0, \xi]$ into $O(n^2)$ many non-overlapping, discrete intervals such that $\text{med}(\hat{z}(\delta))$ is invariable for each such interval.

An analogous approach can be used for the descending 2-PHASE algorithm. This shows that $[0, \xi]$ can be partitioned into $O(n^2)$ many non-overlapping, discrete intervals such that $\text{med}(\check{z}(\delta))$ remains constant over each such interval.

At this point we have $O(n^2)$ many interval border values (from the ascending and descending 2-PHASE algorithm). We can consider these values in increasing order. For each border value δ , at least one of the values $\text{med}(\hat{z}(\delta))$ or $\text{med}(\check{z}(\delta))$ changes. We are interested in the minimum value $\bar{\delta}$ with $\text{med}(\hat{z}(\bar{\delta})) = \text{med}(\check{z}(\bar{\delta}))$. This value $\bar{\delta}$ is indeed the minimum non-negative value such that each reachable equilibrium has the same median if all the median weights γ_i are increased by $\bar{\delta}$ (see [Corollary 1](#)). \square

It remains to prove [Lemmas 3](#) and [4](#). For both lemmas, we will only prove the statements for the ascending 2-PHASE algorithm, i.e., regarding \hat{z} and \hat{z} . The statements for the descending 2-PHASE algorithm work analogously. To show the lemmas, we consider two values $\delta_1, \delta_2 \in [0, \xi] \cap \mathbb{Z}$ with $\delta_1 < \delta_2$. To distinguish the runs of the 2-PHASE algorithm w.r.t. δ_1 and δ_2 , we will talk about process 1 and process 2, respectively. We abbreviate $\mathbf{z}^{(\ell)} = \mathbf{z}(\delta_\ell)$ and $\mu^{(\ell)} = \text{med}(\mathbf{z}(\delta_\ell))$ for $\ell = 1, 2$. Moreover, for any opinion x and comparative operator $\succ \in \{<, >, \leq, \geq, =\}$, we set $V_{\geq x}^{(\ell)} = \{i \in V_{z_i^{(\ell)}} \mid z_i^{(\ell)} \geq x\}$. This notation is further extended in an obvious fashion, e.g., $\hat{\mu}^{(\ell)} = \text{med}(\hat{\mathbf{z}}^{(\ell)})$. Finally, we will examine combined states in the progress of both processes in more detail. Formally, such a (combined) state s is a pair $(s^{(1)}, s^{(2)})$, where $s^{(\ell)}$ denotes the state of process $\ell \in \{1, 2\}$. We will write $\mathbf{z}^{(\ell, s)}$ to denote the strategy profile of process ℓ in state s and extend this notation in an obvious fashion as above.

Proof of Lemma 3. Recall that $\hat{\mu}^{(2)}$ is defined as the median that is reached after the first phase of process 2. For slightly simpler notation in the following, we set $\bar{\mu} = \hat{\mu}^{(2)}$.

We define the combined state $s = (s^{(1)}, s^{(2)})$ as follows. State $s^{(1)}$ denotes the first time in phase 1 of process 1 in which no agent remains that can deviate to an opinion greater than or equal to $\bar{\mu}$. State $s^{(2)}$, on the other hand, denotes the first time in phase 2 of process 2 that the median takes the value $\bar{\mu}$. We claim that $V_{\geq \bar{\mu}}^{(1, s)} \supseteq V_{\geq \bar{\mu}}^{(2, s)}$. Note that by definition, we have $\mu^{(2, s)} = \bar{\mu}$. Hence, the claim implies $\mu^{(1, s)} \geq \bar{\mu}$ and therefore shows the lemma.

Assume that the claim is not true. Since the two processes start with the same strategy profile, this implies that there has to be a first agent i^* in process 2 with $i^* \in V_{\geq \bar{\mu}}^{(2, s)} \setminus V_{\geq \bar{\mu}}^{(1, s)}$. We consider all the influences supporting the deviation of i^* in process 2. The choice of i^* implies that any neighbor i' of i^* that supports the deviation of i^* to at least $\bar{\mu}$ in process 2, also does so in process 1 in state s . Furthermore, the definition of state $s^{(2)}$ implies that the median works against the deviation of i^* in process 2. This may or may not also be the case regarding process 1 and state s but in any case the median has a smaller influence in process 1 since $\delta_1 < \delta_2$. Lastly, the intrinsic opinion of i^* obviously has the same influence in both processes. Hence, i^* has an incentive to deviate to at least $\bar{\mu}$ in process 1 in state s , but no agent in process 1 has such an incentive in state s – a contradiction. \square

In the proof of [Lemma 4](#) we are interested in the case that the two processes reach the same median at the end of phase 1. Hence, we assume $\hat{\mu}^{(1)} = \hat{\mu}^{(2)}$ in the following analysis and again denote this value as $\bar{\mu}$. We can use a similar argument as we did above, but first have to take a closer look at phase 1 of the two processes. Intuitively, the fact that $\delta_1 < \delta_2$ should lead to the opinions being clustered more closely around $\bar{\mu}$ for $\hat{\mathbf{z}}^{(2)}$ than for $\hat{\mathbf{z}}^{(1)}$. This is indeed utilized in the following.

Proof of Lemma 4. Recall that $\hat{\mu}^{(2)}$ is defined as the median that is reached after phase 2 of process 2. Both processes have the median value $\bar{\mu}$ at the end of phase 1 and the median can only decrease in phase 2. Hence, there is nothing further to show if $\hat{\mu}^{(2)} = \bar{\mu}$ and we assume $\hat{\mu}^{(2)} < \bar{\mu}$ in the following.

We consider the combined state a , where state $a^{(\ell)}$ of process $\ell \in \{1, 2\}$ denotes the first time in phase 1 in which no agent remains that can deviate to an opinion larger or equal to $\bar{\mu}$. We claim that $V_{\geq \bar{\mu}}^{(1, a)} \subseteq V_{\geq \bar{\mu}}^{(2, a)}$.

Assume for the sake of contradiction that this is not true. Let $i^* \in V_{\geq \bar{\mu}}^{(1, a)} \setminus V_{\geq \bar{\mu}}^{(2, a)}$ and let i^* be the first such agent in process 1. The choice of i^* implies that any neighbor i' of i^* that supports the deviation of i^* to at least $\bar{\mu}$ in process 1 also does so in process 2 in state a . Moreover, the median may or may not support the deviation of i^* in process 1. But in process 2 the median $\bar{\mu}$ is reached in state a , and hence the median does support a deviation to at least $\bar{\mu}$ and with a stronger influence since $\delta_2 > \delta_1$. Hence, i^* has an incentive to deviate to at least $\bar{\mu}$ in process 2 in state a , but no agent has such an incentive in state a – a contradiction.

We continue with the combined state b in which both processes finished phase 1 of the 2-PHASE algorithm. We claim $V_{< \bar{\mu}}^{(1, b)} \supseteq V_{< \bar{\mu}}^{(2, b)}$ and $z_i^{(1, b)} \leq z_i^{(2, b)}$ for each $i \in V_{< \bar{\mu}}^{(1, b)}$. Now, the first statement of the claim is directly implied by the claim for state a , and we assume, for the sake of contradiction, that the second statement is not true. For $i \in V_{< \bar{\mu}}^{(1, b)} \setminus V_{< \bar{\mu}}^{(2, b)}$, we already know that i deviates to a value smaller than $\bar{\mu}$ in phase 1 of process 1 but at least $\bar{\mu}$ in phase 1 of process 2. Hence, there has to be a first agent i^* in process 1 with $i^* \in V_{< \bar{\mu}}^{(1, b)} \cap V_{< \bar{\mu}}^{(2, b)}$ and $z_{i^*}^{(1, b)} > z_{i^*}^{(2, b)}$. Since the final median $\bar{\mu}$ for both processes in phase 1 is already reached in state a , the median supports a deviation of i^* to at least $z_{i^*}^{(1, b)}$ in process 2 in state b . Furthermore, this support is stronger than the one for the deviation of i^* in process 1 because $\delta_1 < \delta_2$. Next, we consider the neighbors of i^* that support their deviation in process 1. Any such neighbor i' also supports a deviation of i^* to $z_{i^*}^{(1, b)}$ in process 2 in state b . To see this, note that i' either did not deviate in phase 1 of process 1 before i^* or it did to either an opinion of at least $\bar{\mu}$ or at most $\bar{\mu}$. The support for a deviation of i^* to $z_{i^*}^{(1, b)}$ in process 2 in state b is guaranteed in the first case, since the two processes have the same initial strategy profile; in the second, due to the claim for state a ; and in the third case, because of the choice of i^* . Hence, i^* has an incentive to deviate to a higher opinion in state b in process 2, but no agent has an incentive to do so in state b – a contradiction.

Finally, we define the combined state c as follows. State $c^{(1)}$ of process 1 denotes the first time in phase 2 in which no agent remains that can deviate to an opinion smaller or equal to $\hat{\mu}^{(2)}$. State $c^{(2)}$ of process 2, on the other hand, denotes the first time in phase 2 that the median takes the value $\hat{\mu}^{(2)}$. For clarity of notation, we set $\hat{\mu} = \hat{\mu}^{(2)}$. We claim that $V_{\leq \hat{\mu}}^{(1, c)} \supseteq V_{\leq \hat{\mu}}^{(2, c)}$. Note that by definition, we have $\mu^{(2, c)} = \hat{\mu}$. Hence, the claim implies $\mu^{(1, c)} \leq \hat{\mu}$ and therefore suffices to show the lemma.

Assume that the claim is not true. Remember that we have $\hat{\mu} < \bar{\mu}$, as discussed at the start of the proof. Moreover, regarding state b , we did show $V_{< \bar{\mu}}^{(1, b)} \supseteq V_{< \bar{\mu}}^{(2, b)}$ and $z_i^{(1, b)} \leq z_i^{(2, b)}$ for each $i \in V_{< \bar{\mu}}^{(1, b)}$. This implies that the corresponding claim is true at the end of phase 1, i.e., for the combined state b . Hence, there has to be a first agent i^* in process 2 with $i^* \in V_{\leq \hat{\mu}}^{(2, c)} \setminus V_{\leq \hat{\mu}}^{(1, c)}$. The definition of state $c^{(2)}$ implies that the median does not support the deviation of i^* in process 2. Moreover, a neighbor i' supporting the deviation of i^* to a value of at most $\hat{\mu}$ in process 2 also does so in process 1 in state c . To see this, note that i' either did deviate in phase 2 of process

2 prior to i^* or it did not. In the first case, the choice of i^* implies $i' \in V_{\leq \mu}^{(1,c)}$ and in the second i' has opinion $z_{i'}^{(2,b)} \geq z_{i'}^{(1,b)}$ when i^* deviates in process 2. Hence, we have a contradiction. \square

The proof technique of [Theorem 2](#) can be used to show further results. For example, [Lemmas 3](#) and [4](#) work as well if the median weight is only increased for a given subset V' of agents. However, in this case, it can be impossible to stabilize the median. Rather, we aim to obtain the smallest possible difference between the minimum and maximum median of reachable equilibria. Let this value be $\tau_{V'}$. For a collection $\mathcal{V} \subseteq 2^{V'}$ of subsets of agents, we may be interested in the smallest $\tau_{\min} = \min\{\tau_{V'} \mid V' \in \mathcal{V}\}$. Moreover, from the subsets that guarantee τ_{\min} , we want to choose one that achieves this optimal difference using the smallest (integer) increase in median weight, i.e., $\delta_{\min} = \min\{\delta_{V'} \mid \tau_{V'} = \tau_{\min}, V' \in \mathcal{V}\}$. We say that a pair $(V', \delta_{V'})$ with $\delta_{V'} = \delta_{\min}$ has an *optimal median gap with smallest increase* in \mathcal{V} . It is easy to show the following result:

Corollary 3. *Let $k \in \mathbb{Z}_{\geq 0}$ be a constant and $\mathcal{V} = \{S \subseteq V \mid |S| \leq k\}$. There is a polynomial-time algorithm to find a pair $(V', \delta_{V'})$ that has an optimal median gap with smallest increase in \mathcal{V} .*

Proof. First note that increasing any median weight by more than ξ will lead to the same result as increasing it by ξ . Hence, we can simply enumerate all $O(|V|^k)$ many possibilities for V' . By increasing the weight by ξ for each, we determine τ_{\min} . Then, by restricting attention to sets V' with τ_{\min} , we can use the approach from the proof of [Theorem 2](#) to find $\delta_{V'}$ for each of them. In this way, we obtain a pair $(V', \delta_{V'})$ that has optimal median gap with smallest increase. \square

This polynomial-time enumeration can be applied even when we have further restrictions on \mathcal{V} . One might wonder about the complexity of the problem if k is not part of the input (i.e., not necessarily constant). A special case of the problem with arbitrary k is studied in more detail in the next section.

7. Hardness results

In this section, we will first prove the following theorem, and then we discuss a variant with unit weights. Remember that we call a median opinion game \mathcal{G} and an initial profile z *stable* (with respect to the median of reachable equilibria) if all equilibria reachable from z share the same global median. We define the decision problem $\text{STABLEMEDIAN}(\mathcal{G}, z)$ as the problem to decide if the global median of the initial profile z in the median opinion game \mathcal{G} can be stabilized by increasing the influence of the weight of the aggregation function γ_i for k agents $i \in V$ by any value $\delta \geq 0$.

Theorem 3. *STABLEMEDIAN is NP-complete.*

Proof. Considering [Theorem 1](#), we observe that STABLEMEDIAN is indeed in NP: for a given guess of k agents, the problem can be decided by increasing their median weights by ξ and running the two variants of the 2-PHASE algorithm (as was done in the proof of [Corollary 3](#)).

We show NP-hardness starting from the classical vertex cover problem. In this problem, a graph $G' = (V', E')$ and an integer $k \in \{1, \dots, |V'|\}$ are given, and the goal is to decide whether a selection $C \subseteq V'$ of (at most) k vertices exists such that each edge is connected to at least one vertex in C . We will assume w.l.o.g. that $|E'| > k$.

The reduction uses a construction that can be divided into three gadgets. There is a vertex gadget, an edge gadget, and a pivot gadget. There are only two opinions: 0 and 1, and 1 is the initial median opinion. In the vertex gadget, there are vertex agents with public opinion 1 who may change their opinion. The edge gadget contains edge agents again with initial opinion 1 who are connected to their respective vertex agents and may change their opinion if both neighboring vertex agents do so as well. Finally, the pivot gadget contains many agents with public opinion 1 who may deviate to 0 as soon as one of the edge agents does. This can be realized in a way that the existence of a suitable vertex cover guarantees a suitable selection of (vertex) agents and vice versa. We proceed with a detailed description of the gadgets. The median weight of each agent is 1, and the weights of the intrinsic opinions are 0. The construction uses parameters a, b, c, d, g, h with $a = \Delta b + 2$ with Δ the maximum degree in G' , $b = (g \cdot c + 2)/2$, $c = 1$, $d = |E'|c$, $g = 2|V'| + 2|E'|$, and $h = 2|V'|$. See [Fig. 3](#) for an overview.

- *Vertex gadget:* For each vertex $v \in V'$, there is a *vertex agent* $\text{VAgt}(v)$ with initial opinion 1 and a *vertex flip agent* $\text{VFAgt}(v)$ with initial opinion 0. The agents $\text{VAgt}(v)$ and $\text{VFAgt}(v)$ are connected with weight a .
- *Edge gadget:* For each edge $e \in E'$, there is an *edge agent* $\text{EAgt}(e)$ with initial opinion 1. $\text{EAgt}(\{u, v\})$ is connected to the vertex agents $\text{VAgt}(u)$ and $\text{VAgt}(v)$ with edges of weight b .
- *Pivot gadget:* For each $i \in [g]$ there is one *pivot agent* $\text{PAgt}(i)$ with initial opinion 1 and one *pivot balance agent* $\text{PBAgt}(i)$ with initial opinion 0. Each pivot agent is connected to each edge agent with weight c . Moreover, $\text{PAgt}(i)$ is connected to $\text{PBAgt}(i)$ with weight d .
- *Dummy agents:* For each $i \in [h]$ there is one *dummy agent* $\text{DAgt}(i)$ with initial opinion 1. Each dummy agent is connected to each other dummy agent with weight 1, i.e., the dummy agents form a clique, and the dummy agents are not connected to any other agent.

To show correctness, we first consider the possible equilibria without change in the median weights. Initially, there are $|E'| + 2|V'|$ many more agents with opinion 1 than there are agents with opinion 0, and hence the initial median is 1. Running the descending 2-PHASE algorithm will cause all the vertex agents to deviate to 0 since $a > \Delta b + 1$ in phase 1. Afterwards, all edge agents will deviate to 0 because $2b > gc + 1$, and lastly, all of the pivot agents have an incentive to deviate to 0 as soon as any edge agent deviates because $d + c > (|E'| - 1)c + 1$. Hence, an equilibrium is reached in which the opinion of each agent except for the dummy agents

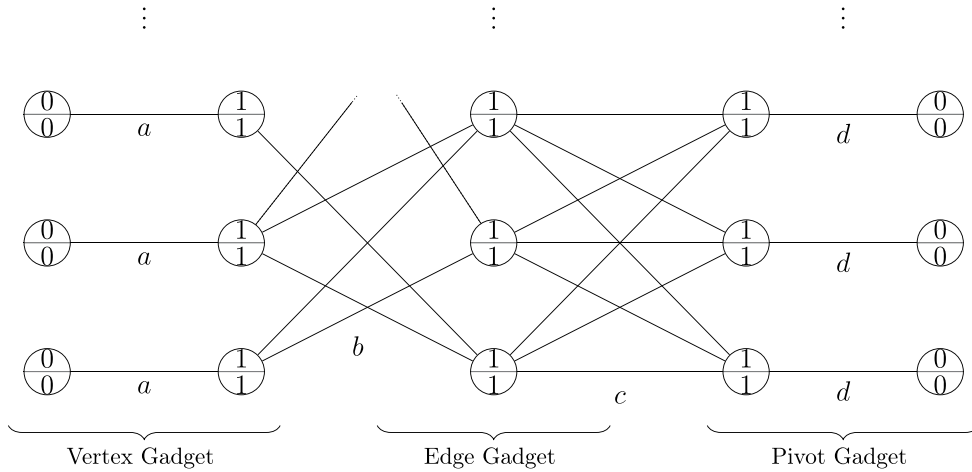


Fig. 3. The gadgets of the reduction used in the proof of [Theorem 3](#). The labels of the agents are the initial and intrinsic opinions. The vertex gadget includes $|V'|$ many pairs of agents, the edge gadget $|E'|$ many agents, and the pivot gadget g many pairs of agents. The h many dummy agents are not shown.

is 0. Therefore, also the median is 0. Running the ascending 2-PHASE algorithm, on the other hand, will cause all of the vertex flip agents and pivot balance agents to deviate to opinion 1 in phase 1, reaching an equilibrium with median opinion 1. In fact, changing the median weight for any agent will not change this behavior in this case. Therefore, we focus on the possibility of reaching an equilibrium with median 0 in the following analysis.

Now, assume that there exists a vertex cover C for G' with $|C| \leq k$. In this case, increasing the median weight of $\text{VAgT}(v)$ for each $v \in C$ to some suitable value, say $3a$, will stop these agents from deviating in phase 1 of the (descending) 2-PHASE algorithm. Since C is a vertex cover and each edge agent will only deviate after both of the respective vertex agents deviated, no edge agent will deviate. Hence, no further agent will deviate in phase 1. This implies that there are at least $g + 2|V'| + |E'|$ many agents with opinion 1 and at most $g + 2|V'|$ agents with opinion 0 and the median of public opinions is still 1 after phase 1 and therefore also after the second.

For the other direction, we consider the case that there exists a selection S of at most k agents for which increasing the median weight guarantees that no equilibrium with median 0 can be reached. As before, we can focus on a run of the descending 2-PHASE algorithm. During phase 1, each vertex agent outside of S will deviate. Afterward, any edge agent for which both respective vertex agents deviated will deviate unless they are included in S . If, at this point, no edge agent did deviate, there has to be a vertex cover C of size at most k in G' – we may construct C by adding each vertex v with $\text{VAgT}(v) \in S$ and adding either u or w if $\text{EAgt}(\{u, w\}) \in S$. Hence, for the remainder of the proof we assume for the sake of contradiction that at least one edge agent did deviate in phase 1. In this case, each pivot agent outside of S will deviate in phase 1. Hence, at least $2|V'| + 2g + 1 - k$ many agents have opinion 0 at some point in phase 1 and the median changes to 0. Therefore all remaining vertex, edge, and pivot agents will deviate as well. Furthermore, some dummy agents may deviate at this point if they are included in S . In the phase 2, no vertex, edge, or pivot agent will deviate to 1. An equilibrium with median 0 is reached – a contradiction. \square

Our final corollary shows that the large edge weights used in the reduction of [Theorem 3](#) are not necessary.

Corollary 4. *STABLEMEDIAN is NP-complete even if $\beta = 1$.*

Proof. We briefly discuss the changes to the reduction and the proof. First of all, all of the weights in the following analysis are 1 including the weight of the median and the intrinsic opinions, which have been irrelevant in the above reduction. The intrinsic opinion of an agent will always remain their initial opinion. We use parameters $a = \Delta + 3$, $b = g + 1$, $c = |E'| + 1$, $g = 2|V'| + 2|E'|$, and $h = (2 + (a - 1))|V'| + |E'|b + g(c - 1)$. The gadgets are changed as follows:

- **Vertex gadget:**
There is now one vertex flip agent $\text{VFAgt}(v, i)$ for each $v \in V'$ and $i \in [a]$. $\text{VFAgt}(v, i)$ is connected to $\text{VAgT}(v)$ for each $i \in [a]$.
- **Edge gadget:**
We introduce one edge balance agent $\text{EBAgt}(e, i)$ for each $e \in E'$ and $i \in [b]$ with initial (and intrinsic) opinion 0. $\text{EBAgt}(e, i)$ is connected to $\text{EAgt}(e)$ for each $i \in [g]$.
- **Pivot gadget:**
There is now one pivot balance agent $\text{PBAgt}(i, j)$ for each $i \in [g]$ and $j \in [c]$. $\text{PBAgt}(i, j)$ is connected to $\text{PAgt}(i)$ for each $j \in [c]$.
- **Dummy agents:**
The dummy agents stay the same (with an increased quantity h).

See [Fig. 4](#) for an overview.

First, note that the increased number of dummy agents counteracts the added agents with initial opinion 0 and hence there are still $|E'| + 2|V'|$ more agents with opinion 1 initially. Running the ascending 2-PHASE algorithm leads to an equilibrium in which




Fig. 4. The gadgets of the reduction used in the proof of Corollary 4. The labels of the agents are the initial and intrinsic opinions. The h many dummy agents are not included.

all agents have opinion 1 regardless of whether some median weights are increased. Moreover, if no median weights are increased, running the descending 2-PHASE algorithm, leads to an equilibrium in which all but the dummy agents have opinion 0 and the median is therefore 0 as well.

If a vertex cover C of size at most k exists, we choose the corresponding vertex agents and increase their median weights to at least $\Delta + 2$. This guarantees that the descending 2-PHASE algorithm leads into an equilibrium with median 1.

To see this, note that pivot agents will still only deviate to 0 after at least one edge agent deviated. Moreover, an edge agent will still only deviate if both of their neighboring vertex agents deviated first (as long as the median is 1).

The other direction follows by a similar set of arguments as in proof of Theorem 3. \square

8. Conclusion

In this paper we shed light on the interplay of opinion formation and voting. We analyze best-response dynamics of a novel opinion formation game that integrates local interactions with a global voting mechanism. In our proposed framework, a best response of an agent corresponds to a weighted median of their intrinsic belief, the opinions of their neighbors, and a global aggregate of the collective opinion landscape. Our results shed light on the structure and existence of equilibria, the number of equilibria reachable from a given median opinion game with initial profile as well as computational complexity regarding stabilization of the voting result in reachable equilibria.

Our research opens up many interesting directions for future research. It is an interesting and challenging open problem to design polynomial-time algorithms for the STABLEMEDIAN problem with provable approximation guarantees. More generally, it could be interesting to contrast our study with an approach, where agents measure the distance between the global aggregation and their *intrinsic belief* rather than their *public opinion*. In a different direction, is it possible to characterize how axiomatic properties of voting rules impact the structure and computational complexity of equilibria?

CRedit authorship contribution statement

Petra Berenbrink: Writing – original draft, Formal analysis; **Martin Hoefer:** Writing – original draft, Formal analysis; **Dominik Kaaser:** Writing – original draft, Formal analysis; **Marten Maack:** Writing – original draft, Formal analysis; **Malin Rau:** Writing – original draft, Formal analysis; **Lisa Wilhelmi:** Writing – original draft, Formal analysis.

Data availability

No data was used for the research described in the article.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests. Petra Berenbrink, Martin Hoefler, Malin Rau, Lisa Wilhelmi report financial support was provided by Deutsche Forschungsgemeinschaft (DFG). If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

Petra Berenbrink, Martin Hoefler, Malin Rau, and Lisa Wilhelmi gratefully acknowledge support by DFG Research Unit ADYN under grant DFG 411362735.

References

- [1] J.M. Epstein, R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, Brookings Institution Press, 1996.
- [2] J.M. Epstein, *Generative Social Science: Studies in Agent-Based Computational Modeling*, Princeton University Press, 2012. <https://doi.org/10.1515/9781400842872>
- [3] J.H. Miller, S.E. Page, *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*, Princeton University Press, 2009. <https://doi.org/10.1515/9781400835522>
- [4] J. Van Ginneken, *Collective Behavior and Public Opinion: Rapid Shifts in Opinion and Communication*, Routledge, 2003.
- [5] H. Johnston (Ed.), *Social Movements and Culture*, Routledge, 1995.
- [6] M. Del Vicario, A. Bessi, F. Zollo, F. Petroni, A. Scala, G. Caldarelli, H.E. Stanley, W. Quattrociocchi, The spreading of misinformation online, *Proc. Natl. Acad. Sci.* 113 (3) (2016) 554–559. <https://doi.org/10.1073/pnas.1517441113>
- [7] N.E. Friedkin, E.C. Johnsen, Social influence and opinions, *J. Math. Soc.* 15 (3–4) (1990) 193–206.
- [8] R.A. Holley, T.M. Liggett, Ergodic theorems for weakly interacting infinite systems and the voter model, *Ann. Probab.* 3 (4) (1975) 643–663.
- [9] G. Deffuant, D. Neau, F. Amblard, G. Weisbuch, Mixing beliefs among interacting agents, *Adv. Complex Syst.* 3 (1–4) (2000) 87–98.
- [10] R. Hegselmann, U. Krause, Opinion dynamics and bounded confidence models, analysis, and simulation, *J. Artif. Societ. Soc. Simul.* 5 (3) (2002) 2.
- [11] W. Mei, F. Bullo, G. Chen, J. Hendrickx, F. Dörfler, Micro-foundation of opinion dynamics: rich consequences of an inconspicuous change, *Phys. Rev. Res.* 4 (2) (2022) 023213.
- [12] H. Moulin, On strategy-proofness and single peakedness, *Public Choice* 35 (4) (1980) 437–455.
- [13] S. Ching, Strategy-proofness and “median voters”, *Int. J. Game Theory* 26 (1997) 473–490.
- [14] D. Kempe, J.M. Kleinberg, É. Tardos, Maximizing the spread of influence through a social network, in: *Proc. 9th Conf. Knowl. Discov. Data Min. (KDD)*, 2003, pp. 137–146.
- [15] D. Golovin, A. Krause, Adaptive submodularity: theory and applications in active learning and stochastic optimization, *J. Artif. Intell. Res.* 42 (2011) 427–486.
- [16] C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics, *Rev. Mod. Phys.* 81 (2) (2009) 591–646.
- [17] A. Downs, *An Economic Theory of Democracy*, Harper & Row New York, 1957.
- [18] R. Abelson, A. Bernstein, A computer simulation model of community referendum controversies, *Public Opin. Q.* 27 (1) (1963) 93–122.
- [19] J. French, Jr., A formal theory of social power, *Psychol. Rev.* 63 (3) (1956) 181–194.
- [20] M.H. DeGroot, Reaching a consensus, *J. Amer. Stat. Assoc.* 69 (345) (1974) 118–121.
- [21] D. Bindel, J.M. Kleinberg, S. Oren, How bad is forming your own opinion?, *Games Econ. Behav.* 92 (2015) 248–265.
- [22] K. Bhawalkar, S. Gollapudi, K. Munagala, Coevolutionary opinion formation games, in: *Proc. Symp. Theory Comput. (STOC)*, 2013, pp. 41–50.
- [23] V. Bilò, A. Fanelli, L. Moscardelli, Opinion formation games with dynamic social influences, *Theor. Comput. Sci.* 746 (2018) 73–87.
- [24] M. Epiropou, D. Fotakis, M. Hoefler, S. Skoulakis, Opinion formation games with aggregation and negative influence, *Theory Comput. Syst.* 63 (7) (2019) 1531–1553.
- [25] D. Fotakis, D. Palyvos-Giannas, S. Skoulakis, Opinion dynamics with local interactions, in: *Proc. 25th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2016, pp. 279–285.
- [26] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Robustness in discrete preference games, in: *Proc. 16th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*, 2017, pp. 1314–1322.
- [27] R. Brederbeck, E. Elkind, Manipulating opinion diffusion in social networks, in: *Proc. 26th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2017, pp. 894–900.
- [28] R. Brederbeck, L. Jacobs, L. Kellerhals, Maximizing the spread of an opinion in few steps: opinion diffusion in non-binary networks, in: *Proc. 29th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2020, pp. 1622–1628.
- [29] S. Botan, U. Grandi, L. Perrussel, Multi-issue opinion diffusion under constraints, in: *Proc. 18th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*, 2019, pp. 828–836.
- [30] A. Anagnostopoulos, L. Becchetti, E. Cruciani, F. Pasquale, S. Rizzo, Biased opinion dynamics: when the devil is in the details, in: *Proc. 29th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2020, pp. 53–59.
- [31] P. Faliszewski, R. Gonen, M. Koutecký, N. Talmon, Opinion diffusion and campaigning on society graphs, in: *Proc. 27th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2018, pp. 219–225.
- [32] A. De, S. Bhattacharya, N. Ganguly, Shaping opinion dynamics in social networks, in: *Proc. 17th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*, 2018, pp. 1336–1344.
- [33] V. Auletta, D. Ferraioli, V. Fionda, G. Greco, Maximizing the spread of an opinion when tertium datur est, in: *Proc. 18th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS)*, 2019, pp. 1207–1215.
- [34] W. Mei, J.M. Hendrickx, G. Chen, F. Bullo, F. Dörfler, Convergence, consensus, and dissensus in the weighted-median opinion dynamics, *IEEE Trans. Automatic Control* 69 (10) (2024) 6700–6714.
- [35] R. Zhang, Z. Liu, G. Chen, W. Mei, Convergence Analysis of Weighted-Median Opinion Dynamics with Prejudice, *IEEE Trans. Autom. Control* to appear 70 (6) 4155–4162.
- [36] R. Brederbeck, A. George, J. Israel, L. Kellerhals, Single-peaked opinion updates, in: *Proc. 31st Int. Joint Conf. Artif. Intell. (IJCAI)*, 2022, pp. 137–143.
- [37] V. Auletta, I. Caragiannis, D. Ferraioli, C. Galdi, G. Persiano, Minority becomes majority in social networks, in: *Proc. 11th Conf. Web and Internet Econ. (WINE)*, 2015, pp. 74–88.
- [38] J. Doucette, A. Tsang, H. Hosseini, K. Larson, R. Cohen, Inferring true voting outcomes in homophilic social networks, *Auton. Agents Multi Agent Syst.* 33 (3) (2019) 298–329.
- [39] C. Bernardo, C. Altafini, A. Proskurnikov, F. Vasca, Bounded confidence opinion dynamics: a survey, *Automatica* 159 (2024) 111302.
- [40] P. Berenbrink, G. Giakkoupis, A. Keramarrec, F. Mallmann-Trenn, Bounds on the voter model in dynamic networks, in: *Proc. 43rd Int. Colloq. Autom. Lang. Programming (ICALP)*, 2016, pp. 146:1–146:15.
- [41] C. Cooper, N. Rivera, The linear voting model, in: *Proc. 43rd Int. Colloq. Autom. Lang. Programming (ICALP)*, 2016, pp. 144:1–144:12.
- [42] C. Cooper, F. Mallmann-Trenn, T. Radzik, N. Shimizu, T. Shiraga, Asynchronous 3-majority dynamics with many opinions, in: *Proc. 36th Symp. Discret. Algorithms (SODA)*, 2025, pp. 4095–4131.
- [43] D. Ferraioli, C. Ventre, Social pressure in opinion games, in: *Proc. 26th Int. Joint Conf. Artif. Intell. (IJCAI)*, 2017, pp. 3661–3667.

- [44] P. Berenbrink, M. Hoefer, D. Kaaser, M. Maack, M. Rau, L. Wilhelmi, Opinion dynamics with median aggregation, in: Proc. 24th Conf. Auton. Agents and Multi-Agent Syst. (AAMAS), 2025, pp. 271–279.
- [45] R.E. Smithson, Fixed points of order preserving multifunctions, Proc. Amer. Math. Soc. 28 (1971) 304–310.
- [46] L. Zhou, The set of nash equilibria of a supermodular game is a complete lattice, Games Econ. Behav. 7 (2) (1994) 295–300.
- [47] F. Echenique, A short and constructive proof of Tarski's fixed-point theorem, Int. J. Game Theory 33 (2) (2005) 215–218.
- [48] F. Chierichetti, J. Kleinberg, S. Oren, On discrete preferences and coordination, J. Comput. Syst. Sci. 93 (2018) 11–29.
- [49] D. Ferraioli, P.W. Goldberg, C. Ventre, Decentralized dynamics for finite opinion games, Theor. Comput. Sci. 648 (2016) 96–115.