Group-Decision Making With Balanced Procedures

Thomas Boyer-Kassem and Cyrille Imbert
working paper – November 2024

Abstract

Improving how epistemic groups work is one of the goals of social epistemology. Here, we focus on oral procedures for group decision-making. Social groups, such as expert committees, juries, advisory boards, hiring committees, or local councils, routinely make decisions through oral sequential procedures. Unfortunately, the quality of these oral decisions may be spoiled when agents, because of perceived social pressure, publicly conform to what they take to be the collective view of others while not endorsing this view privately. This paper presents a novel oral voting procedure that almost eradicates this undesirable effect by keeping the voting count balanced as long as possible. Oral votes in favor of the competing views are collected alternately through a procedure that preserves the sincerity of votes till nobody volunteers to support one of these views. Using a simple agent-based model reflecting central empirical findings concerning how people actually behave within groups, we normatively assess the merits of this procedure and show that it makes conformity effects disappear in one single voting round. We also discuss the contextual problem of how to initiate the procedure best by choosing the first voter or the first supported option. Finally, we explain why this procedure seems proof against unfortunate misapplications or deliberate spoiling.

Key words. Social Epistemology. Epistemic Groups. Group Decision-Making. Oral voting. Voting Procedures. Voting order. Social psychology. Social Conformity. Conformity cascades. Balanced Procedures. Alternate vote. Agent-based Models.

1. Introduction

Oral group decisions are often taken across various contexts involving formal and informal procedures, experts and laypeople, and private or public situations. They pertain to political, economic, and moral, but also factual and scientific issues. For instance, expert groups, boards, and local councils often use oral roundtables without secret ballots to make decisions. Various types of mechanisms and biases may alter the quality of oral decisional processes. Public belief conformity under perceived social pressure to what seems to be the majority's or opinion leaders' view is one such major mechanism, much investigated by psychologists, where the internal and external views of individuals may differ, and group decisions may consequently differ from what the majority privately believes[1]. Therefore, trying to improve oral voting is a valuable goal for epistemologists studying groups. Improvements may, for example, pertain to voting rules, the number of voting rounds, the types of publicity, or the choice of a voting order.

Here, we explore further how the detrimental effects of such external conformity on group decisions may be significantly reduced by specific voting orders, be it on preference-related

[1] The recent literature has investigated cases of conformity in complex social networks without psychological dissonance (Mohseni and Williams 2021; Weatherall and O'Connor 2021). (Fazelpour and Steel 2022) is sensitive to the various aspects and types of conformity.

questions or factual ones. This hypothesis looks promising in the light of results already showing that deliberative order may determine which consensus emerge in groups (Hartmann and Rafiee Rad 2020).

Formal inquiries, in particular model-based ones, though hard to tame, can be stringent and insightful (Martini and Fernández Pinto 2017; Šešelja 2023), for example, concerning small groups (Bright, Dang, and Heesen 2018). They provide ways of integrating both theoretical knowledge and empirical evidence, investigating the effects of specific mechanisms when data are hard to collect, and finding out how much results and claims are context-sensitive. Using a model reflecting empirical findings concerning how people *actually* behave within groups, we investigate the potential of novel "balanced" procedures, in which voters are alternately asked whether someone is ready to endorse one of the competing views and show how these procedures almost eradicate undesired conformity effects.

The paper is structured as follows. First, we present standard oral voting procedures and their vulnerabilities to conformity effects (section 2) and highlight that they can hardly be replaced by globally superior procedures (section 3). We then investigate balanced oral procedures (section 4): after justifying the model and its empirical credentials (4.1), we introduce alternate procedures and show that they can eradicate the effects of public conformity (4.2-4.3). Finally, we discuss the validity, limits, scope, and contextual adequacy of the results (section 5). Technical details are presented in an Appendix.

2. Oral voting and its problems

Groups routinely make decisions in which votes are public or semi-public, for instance, in governments, boards, juries, expert committees, hiring committees, local councils, or families. Prominent cases involve the Federal Reserve's Federal Open Market Committee (Meade and Stasavage 2008), the Food and Drug Administration advisory committees (Urfalino and Costa 2015), or major political groups (Elster 2015, passim) Oral decision-making is a well-entrenched practice within human communities, which has survived the development of other forms of decision-making and is unlikely to disappear.

2.1 Oral procedures: hardly dispensable but fraught by psycho-socio mechanisms

Oral decision-making can be quick, practical, and informal and lets others know who endorses which view. Publicity may be merely internal to the group, partly external (if observers are allowed and reporting is permitted), or total if the votes are published with or without delays. Whenever their views are known, voters may have to answer for them and be held responsible, which is why public vote is central across democratic countries. This may drive voters towards views that agree with objective reasons and the general interest. However, voters may suffer from retaliation from those displeased by their views (e.g., economic lobbies, political parties, employers, influential colleagues, or community members) and may engage in window dressing to avoid displeasing third parties.

Much remains to be done about such issues. First, this literature and its arguments often rely on qualitative observations. More systematic empirical evidence, both from sociology and social psychology, should be relied upon. Second, aspects specific to voting and those pertaining to deliberations should be more sharply distinguished (Gosseries and Parr 2022, introduction, §2.4), since group performance and collective behavior vary with the details of tasks and activities (Kerr and Tindale 2004). Finally, the analysis should be sensitive to the complexity of contexts and when precisely transparency and secrecy are appropriate (Chambers 2004).

The effects of conformity on oral voting call for specific scrutiny. Psychologists distinguish between various forms of conformity. It may be internal (e.g., belief or feeling adjustment) or external (e.g., behavior adjustment). Here, we investigate cases where, under perceived pressure, voters publicly endorse the view already expressed by previous voters, even if it differs from what they privately believe. Psychologists have intensively studied this conformist behavior for decades and documented its existence on both epistemic and normative issues, even without social pressure and risks of retaliation, both for laypeople and experts. Its effects are often large. In Asch's initial experiment (that is, without social hierarchy and explicit pressure), a wrong answer to a trivial question is given in up to 32% of cases (Asch 1951). Different mechanisms may trigger such social conformity: agents may duly or not *perceive* some social pressure, they may implicitly take groups' beliefs as truth-conducive, they may dislike to disagree publicly with other people, or try to maintain favorable self-concepts (Cialdini and Griskevicius 2010). Social psychologists still discuss the details of the functions describing quantitatively social influence across contexts, typically concerning the under or over-influence of minorities and majorities, the effects of majority defectors (Kerr and Tindale 2004), or the importance of group size and threshold effects (Bond 2005). However, they agree that, be it rational or not in its goals and forms[2], conformity under perceived pressure is clearly widespread, usually important, and may significantly influence oral votes.

Also, deliberation and decision-making are sometimes entangled, e.g., when people present their views and voting is organized only when this is perceived as necessary. There, additional mechanisms, generating conformity or not, may be involved, such as anchoring effects (Kahneman, Slovic, and Tversky 1982), shared information bias (Stasser and Titus 1985), groupthink (Janis 1983) or polarization (see, e.g., (Tindale and Winget 2019; Forsyth 2019) for recent surveys by psychologists about group decision-making). We shall not consider such cases here.

2.2 Orientation and methodology

Here, we consider group decisions about preference-related or factual issues in which deliberations and decisions proper are temporally separated, and we focus on the final voting stage. Agents' private views are supposed to have already converged and remain constant; however, the views that agents actually endorse publicly may still evolve due to social conformity. Then, it is legitimate to investigate how to avoid these effects to "lessen the efficacy of undue influence" (Bentham 1999, 106). Accordingly, we consider groups that have to choose between *two* options (e.g., A or B, accept or reject) and decide with the majority rule (i.e., agents side for A or B, and the option with the most votes wins)[3]. When individuals vote differently from what they privately believe under social pressure, this may influence additional voters, create cascades, and flip unduly the final group decision. To assess this phenomenon and investigate how voting order may help, we compute the percentage of such "distorted" cases where group decisions differ with and without the effect of social conformity.[4]

Naturally, one might worry that the depicted situation is rarely met. Is in-group conformity always high and impactful when much is at stake, especially when voters are trained professionals used to resisting pressure or have incentives towards one option? This objection does not resist scrutiny.

[2] Social conformity may serve different purposes, such as truth-seeking or minimizing retaliation (Kuran 1997). Even if simple majority-based rules depart from strict rationality standards like Bayesian update rules, they may constitute efficient and rational heuristics (Mercier and Morin 2019).

[3] Multi-option cases are trickier since they provide opportunities for strategic voting (Gibbard 1973).

[4] Social influence is often beneficial and epistemic equality may be sometimes be counterproductive, e.g., in group problem-solving activities. Additionally, reducing tension and creating consensus are legitimate goals. Here, we focus on cases where equality and representativeness should prevail.

Various important cases do involve conformity effects. Conformity is present whenever reputational costs exist, or agents are sensitive to how they are perceived. This includes, e.g., situations where lay people are involved in group decisions, contexts involving hierarchies, and cases involving experts, like judges within appeal courts (Sunstein 2003). This latter case also illustrates how much can be at stake without the voters being directly impacted by the decisions. Finally, as we shall show, even moderate conformity can have detrimental effects – hence the need to investigate all degrees of conformity and find procedures beneficial for all such degrees.

3. Existing solutions: merits and limits

Before going further, we first emphasize that no simple procedure keeps the advantages of oral voting while avoiding the detrimental effects of social conformity. We highlight in this brief review features of existing procedures that, ideally, one may wish to retain.

Turning to secret voting to shield voters from external pressure brings different problems. Secret balloting is heavier and more formal than oral voting, which explains why the latter is often used. Also, in case of apparent consensus during deliberation, asking for a secret vote may look superfluous (and may reveal people's views). Finally, secrecy conflicts with accountability. Overall, secret voting is not a panacea. Thus, oral voting is not always dispensable, and finding ways of improving it would be welcome. How could this be done?

Making the atmosphere less abrasive (Imbert et al. 2020) may be considered in daily contexts but is unlikely to do much work when a lot is at stake. In practice, oral procedures are often sequential. This may amplify conformity effects by favoring cascading effects. Vote manipulation is also easier (e.g., if the boss votes first), which also contradicts the idea of political or decisional equality (all votes should count equally).

Can public synchronous voting be the answer to such problems? Unfortunately, not. Synchronicity also favors pluralistic ignorance since agents are influenced by their prior beliefs about other views. Further, with standard procedures like raising hands, group effects can be even more massive since agents are known to take their clues from others, typically influential voters (Urfalino and Costa 2015; Bentham 1999, 107).

Organizing several rounds of vote may help when the effects of mechanisms that spoil oral procedures dampen over time (Imbert et al. 2020, §4). Indeed, conformity is highest when individuals face a unanimous group and significantly decreases when unanimity is broken (Asch 1951). Nevertheless, long voting procedures are inconvenient, the sincerity of votes would be questioned if people switched their votes, and people are usually reluctant to revise their public views (Meade and Stasavage 2008; Gosseries and Parr 2022, §2.2.2). Thus, efficient one-round voting procedures should be preferred.

Novel voting rules may also improve group decision-making. Beyond the usual majority rule ("follow the larger number of voiced options"), other rules have been proposed: the Borda rule ("rank all options"), the Condorcet rule ("make pairwise comparisons"), and more recently Majority Judgment ("grade each option"). In subsequent Sections, we concentrate on the problem of choosing between two options, for which the Borda and Condorcet rules amount to the majority rule. Majority Judgment also coincides with it for the model we consider below[5]. So, even if it brings strong theoretical advantages over the majority rule and answers some inconveniences of

[5] Technically, they coincide for a polarized electorate (Balinski and Laraki 2020, Theorem 4 p. 453), which is a sensible hypothesis in the model introduced in Section 4.

voting procedures (Balinski and Laraki 2020),[6] we shall consider the majority rule in the remainder of the paper.

To make the best of sequentiality, one may impose well-designed orders on voters, such as an anti-seniority rule in which people speak in reverse age order (Ottaviani and Sørensen 2001) or anti-hierarchy rules more generally. Such orders are designed to help all views be heard, thus contributing to a better epistemic decision; the downplay is that some expert views may be somewhat ignored. Still, anti-hierarchical rules display other drawbacks: various influential hierarchies, public or not, may co-exist, and using discrimination based on specific public hierarchies may be difficult to justify.

Against influence cascades, another option is to draw the order of voters randomly (Imbert et al. 2020). This requires no information about the voter's features, is procedurally fair, can be integrated into a controllable procedure, and significantly decreases the number of distorted decisions. The random procedure has its own theoretical and practical drawbacks. It makes the procedure more formal and heavier. It necessitates organizing a reliable random draw, given that usual tricks like starting somewhere in a list work imperfectly and using urns or computers requires material and time. Also, random draws still yield voting orders for which influence cascades may develop. In practice, actual voting orders often involve hidden correlations (e.g., when sitting locations determine voting order), and their average distortion lies somewhere between that of orders that favor cascades most (by increasing order of preference for some view) and of random orders.

Overall, no simple solution presently eradicates the defects of sequential oral voting, hence the need to improve existing procedures. Here, we investigate how specific voting orders may avoid the unwanted effects of social conformity. More precisely, can some simple, informal, one-round procedures generate orders that produce less distortion than random draws? We answer affirmatively by investigating the power of balanced procedures in which group members are requested to volunteer to vote alternately for the competing views.

4 Balanced procedures – participative and alternate

Here, our aim is to investigate whether alternate oral voting procedures may counter the cascading effects of conformist behavior. Thus, we test these procedures in a model that represents this behavior and investigate it by simulation and mathematical means. Formal details about this model and quantitative results about distortion rates for the different voting orders and degrees of social influence are presented in the Appendix, together with relevant theorems.

4.1 The model

To play its role correctly, the model should be descriptive in purpose and reflect the conformist behavior highlighted by social psychologists. What produces conformist behavior, whether it is irrational or not, and whether Bayesian rules[7] and rational choice theory may account for conformist behavior are independent issues debated both by psychologists and economists and not discussed here.[8] For the present purpose, we use a simple behavioral model that plausibly integrates empirical findings concerning social conformity. Importantly, the model is not aimed at making precise

[6] Majority Judgment may be adapted for oral voting when used with few grades (see, e.g., (Balinski and Laraki 2022, §3.4) for three grades), but it is not practical with many grades.

[7] By contrast, Dekel and Piccione (2000) explore sequential voting as a game in which voters gain information through others' votes and may behave rationally.

[8] Kuran shows that conformity can be rational since it integrates individual reputation and retaliation costs in addition to the preferences of agents for the different options (Kuran 1997).

quantitative predictions for specific votes (which would be hazardous since the degree of social conformity varies contextually) but at analyzing robust effects of social conformity and investigating mechanisms that eradicate them across contexts.

The model represents cases where the agents' views are gradual and can be represented as numbers between 0 and 1. The group is to decide between options A or B, which, for symmetry reasons, correspond to 0.25 and 0.75.[9] Agents are supposed to have already converged towards stable private views and reached the voting stage, where a collective decision is cooked up. When they vote orally, agents only state that they prefer A or B without describing the strength of their support nor their motivations, and other agents do not exactly know where their private views lie between 0 and 1. Without social influence, agents would publicly choose the option closest to their private view, for example, 0.25 for a private view at 0.40. Because of social influence, agents' oral votes are impacted by the public votes of previous voters (if any)[10]. Importantly, the focus is on cases where direct social influence only comes from group members and external influence (e.g., by lobbies), if present, goes through in-group members. Hence, only internal group influence needs to be explicitly modeled.

While agents are influenced by majority sizes (Mercier and Morin 2019, §2.1), how they aggregate others' views is a complex matter. Here, we follow the widespread choice by psychologists to consider the simple average as a satisfactory picture of what the group is taken to believe – especially for discrete decision alternatives (Tindale and Winget 2019, 4; Kerr and Tindale 2004, 634). In other words, if two voters have already sided with 0.25 and one voter with 0.75, the 4th agent takes the group's view to be (0.25·2+0.75)/3, that is, 0.416.

The precise shape of the relationship between the size of groups and majorities and how much they are followed (e.g., linear, logarithmic, *s*-shaped, inverted-*s*-shaped) is a debated issue (see (Mercier and Morin 2019) for a review). A natural expectation is that social influence depends on the size of the group that agents are faced with. A crucial finding is that for unanimous groups of size k, maximum conformity is reached for $k = 3$ and does not increase beyond (Asch 1951; Bond 2005), and social influence remains bounded. Accordingly, we represent the views that an agent is willing to endorse publicly as a weighted average of what they take to be the group's view and of their own view. For example, if a voter believes A and takes B to be the group's view, she will be willing to defend publicly the view $\alpha \cdot B + (1 - \alpha) \cdot A$. In this way, the group's influence is bounded by α. As described above, agents do not publicly express this weighted view but vote for the option (0.25 or 0.75) that is closest to it. Since the importance of social influence remains a contextual issue, the model is investigated for all possible values of α, and we focus on patterns that are robust across these values.

A seemingly inconvenient aspect of this simple α bounded model is that the group influence does not depend on the number of voters for $1 \leq k \leq 3$. But this is actually not a problem for our inquiry. We have systematically checked that the adoption of a linear model up to $k = 3$ (that is, with influence $\alpha/3$; $2\alpha/3$; $3\alpha/3$) dampens influence but does not change the global distortion patterns documented below (as already described in (Imbert et al. 2020, p. 415). Thus, assuming a constant α amounts to a simpler, worst-case hypothesis: if a voting procedure lowers distortion rates in that case, then it *a fortiori* does so if α is smaller for the first voters.

Overall, agents tend to conform to already expressed views when voting, which may lead them to vote for the view they do not favor. For example, an agent giving equal weight to her private view 0.58 and to the average view 0.416 votes for 0.25, and such a flipped vote may distort the vote result. Importantly, individual votes and the final results are sensitive to voting order. For instance, suppose agents P, Q, R, S, and T privately believe 0.18, 0.48, 0.58, 0.59, and 0.76, respectively (so

[9] The values are not set to the extremes 0 and 1 to enable the representation of more extreme views than the options between which agents choose.

[10] This model was already used for larger deliberative purposes in (Imbert et al. 2020). It is used here more stringently.

0.75 is the majority view). Order Q-T-P-R-S yields votes (0.25; 0.75; 0.25; 0.25; 0.25), and the minority view wins, which also happens for increasing order P-Q-R-S-T. By contrast, order Q-T-P-S-R makes the majority view win. Thus, procedures generating orders that reduce this distortion are preferable. In this perspective, let us note that if the second voter expresses early dissent, she neutralizes influence; also, it is preferable that the majority view be supported immediately by the first voter.

Finally, what if the first agents have some clues about what the group thinks? To keep things simple, we start with cases where such beliefs are neutral and can be discounted. We explain in Section 5 that, for the specific purpose of this inquiry (viz, discussing the merits of balanced procedures), reintegrating this aspect does not change the conclusions.

The basic result is that distortion rates are globally high even when one adopts a random voting order, which disrupts psychological or sociological biases in the voting order (see the discussion of figure 1 in the Appendix). Accordingly, the question is how to find a voting order that minimizes distortion without supposing that the chair is like some virtuous God that can read the minds of voters.

4.2 Adopting balanced procedures.

Social conformity exerts major influence when the first votes unanimously support a view, and the subsequent voters feel pressured to endorse it. In order to avoid this, each view should be supported as early as possible. As social pressure varies with the ratio of votes supporting each view, whatever the exact form of this dependence, pressure will be minimal if this ratio remains close to 1, which suggests trying to alternate votes for each option. But how can this be achieved since private views are… private and clues about these views are uncertain and hardly integrable into formal procedures?

Our idea is to involve those for whom these views are known – voters themselves – and generate the voting order from within the group. The chair starts by asking, “Is there one person willing to vote for A?” (say), and waits till one person, if any, says “Me.” After this first person has voted uninfluenced by previous voters, the chair immediately asks whether one person is willing to vote for B, and so on, alternately for A and B. In this procedure, an agent ready to say “Me” is someone for whom the weighted average between her private view and the average of already expressed views tips on the side of the view called by the chair. Since social pressure for a view evolves with the votes, the pool of agents willing to publicly endorse a given view may change over time. For example, someone who weakly prefers B may resent voting for B when the count is 1-0 against B, but may find it acceptable when it is 3-2 and social pressure is weaker. If nobody wants to vote for A (resp. B) at some point, all remaining votes count for B (resp. A).

This procedure keeps the public support for each view roughly equal as long as this is possible. Social pressure is washed out progressively: for even ranks, it remains neutral, and for odd ranks, the voting ratio tends to 1, the pressure faced by voters keeps decreasing, and voters are increasingly in a position to vote sincerely. Further, all voters who say “me” along the alternate voting scheme are sincere (cf. Appendix, Proposition 1), either because they vote in odd positions where the count is even and there is no pressure or because they vote for B in even positions despite a count that favors A and only B supporters who can endure this social pressure for A volunteer to vote for B. The self-generated alternate procedure promotes sincere votes by detecting people willing to endorse their private views publicly despite unfavorable social pressure. By accepting to speak up, these voters immediately re-balance the vote count, prevent an early onset of cascades, decrease social pressure, and contribute to unlocking progressively faint-hearted or lukewarm voters. Thereby, vote sincerity is preserved until nobody answers the call for some view.

To apply a balanced procedure, one must decide which view is called first. We investigate two promising possibilities, named BOD and BAD.

4.3 BOD – Balanced, Option Drawn

The even and odd positions in the alternate scheme are not equivalent. The first view (say A) that receives the first vote gets an influence privilege since, to support the other view, the second voter needs to express public disagreement without any social support, which is the most uncomfortable position, and even voters who need to support B always face an unfavorable voting count. Further, the first view gets a result privilege: it cannot be defeated unduly and may win the vote unduly (Proposition 2 in the Appendix). Indeed, supporters of the first view always face balanced social pressure and may always sincerely vote for the view they privately support. Therefore, distorted decisions correspond to cases where most voters privately support the second view but lose the vote. Conversely, the second view cannot win the vote unduly but may be defeated unduly. Hence, the choice of the first view called by the chairperson is a crucial issue and should be treated advisedly. An intuitive suggestion is to draw randomly to treat both views equally, which we call the BOD procedure ("Balanced, Option Drawn").

BOD brings two major improvements over the Random Order procedure (see the Appendix for details). First, the number of distorted decisions is significantly lower, provided conformity is not too high, which is often the case. Second, this improvement is obtained in one single round (see Propositions 3 and 4 in the Appendix), whereas the random procedure needs several rounds to reach its minimal distortion rate. Indeed, when people vote in a predetermined order, they may misrepresent their private views when the voting count is unfavorable; since the voting count often gets more balanced at later ranks, these people would have sincerely voted if they had voted later. Thus, making several rounds may help (provided people accept to change their votes publicly). This is useless within an alternate series since all people who answer the call are sincere.

However, the BOD procedure displays an inconvenience: a random draw gives equal chances to both options, irrespective of the number of agents who privately support them. Thereby, agents in the minority are given proportionally more chances to defend their views first, benefit from the influence privilege, and unduly win the vote. Expectedly BOD yields more distortion for high social conformity than Random Order (see Fig. 1 in the Appendix). Accordingly, we present another alternate procedure to circumvent this problem.

4.4 BAD – Balanced, Agent Drawn

To limit the disproportionate influence of minorities, the first speaker should be representative of what voters privately prefer. But how to achieve this, since neither private views nor the majority are known (which would make voting unnecessary)? A solution is to randomly pick the first voter, with equal chances for all voters. That is, the chairperson first calls a person at random and asks, "Which option do you support?", and her answer initiates the alternate, ABABA or BABAB, sequence. With this BAD procedure ("Balanced, Agent Drawn"), the chance that the first speaker belongs to the minority is statistically lowered to the proportion of that minority.

Expectedly, BAD also reaches full efficiency in one voting round; compared to BOD, it slightly lowers the rate of distorted decisions when conformity is low to moderate and lowers it more when conformity is higher (see the Appendix). Indeed, when conformity is important, so is the influence privilege for the first view and the need for a representative first voter. Unsurprisingly, when conformity is extreme, and people follow the first voter, BAD lowers distortion to that of Random. Overall, BAD seems to combine the advantages of BOD and Random Order.

5. Discussion

Some features of the BOD and BAD procedures are worth specifically highlighting. First, they offer an applicable way to build an alternate order faithful to private views, which is an achievement, given the non-publicity of these views. Second, they are a very efficient way to reduce distortion across contexts. Within the model, they are even optimal in the sense that, given the first voter, no other alternate procedure can yield more sincere votes (see Prop. 5 in the Appendix). Third, the participative mechanism they rely on provides efficient uses of contingent but influential features. When agents volunteer to defend a view despite an unfavorable voting count, they do so because of the strength of their convictions, psychological features, or hierarchical positions, which make them immune to social pressure. In other words, some hidden features of agents, which may often be epistemically detrimental, here contribute to the sincerity of the votes. Finally, since the voting order emerges endogenously from the group, it can hardly be manipulated by internal or external parties – contrarily to cases in which order is chosen discretionary by an agent. We now discuss the scope of the results and analyze potential objections.

5.1 Idealized features of the model

We now discuss two simplified aspects of the model and argue that they do not threaten the general claim that balanced voting orders globally improve on other voting orders.

First, what about cases where agents have initial non-neutral clues about what the majority thinks? If present, such clues are likely to influence voters in one direction, say A. This influence cannot be larger than α, which describes social conformity in the model. With balanced procedures, it should primarily impact the first voters since subsequent voters are likely to give weight to the actual votes showing that there is support for both views (remember that public dissent strongly dampens conformity effects). Now, remember that clues are fallible indicators: for example, clues stemming from previous interactions can be misleading since exchanges may be partial, views may evolve, and oral interventions do not unambiguously indicate how people will finally vote. Thus, two cases should be distinguished. When token clues unreliably push voters toward A, neutralizing their effects on actual votes is all the more important, especially at the beginning. This is precisely what balanced procedures do by selecting early voters still willing to endorse B publicly. Thus, such procedures are likely to keep their differential advantage. Now, suppose token clues happen to be reliable. Then, they will influence the first voters in a direction that reduces distortion for all voting orders. For low and intermediate social pressure, this beneficial influence may partially decrease the differential advantage of balanced procedures (compared to other orders), which is very large (cf. Fig. 1). However, it is unlikely to reverse the compared profitableness of procedures since the mechanisms that spoil other orders and make balanced orders useful remain unchanged. For contexts where social pressure gives the first voters decisive influences, all procedures are similarly tipped in the right direction, and BAD remains unchallenged because, statistically, it provides a representative first voter. Finally, when clues suggest a strong majority favoring some view, starting the alternate vote with the minority view may be a good policy to guard against potential situations of pluralistic ignorance. In brief, the existence of preexisting beliefs or clues, reliable or not, does not threaten the central point that balanced procedures are beneficial.

Let us finally consider cases where voters are directly sensitive to asymmetric influence from outside the group (e.g., public opinion or lobbies). This external influence may vary with the types of publicity (e.g., live broadcasting, situations with public observers like in trials, or semi-public procedures when publicity is delayed to shield voters from pressure) and what is made public (the

details of deliberations, the winning option only, the scores of each option, or minority reports). Measuring the existence and importance of such conformist effects without live interaction is a difficult empirical issue (Mallinson and Hatemi 2018). Be this as it may, if present, asymmetric effects of external influence will globally increase distorted decisions for all oral voting orders. Then, using balanced procedures to neutralize this spurious influence will be all the more crucial, and the compared benefits of balanced procedures should remain. This may be investigated more quantitatively in future works by changing the thresholds at which agents flip their public view from A to B. And, if external pressure clearly favors A (say), a virtuous chair may start with the underprivileged option B (see the BMC procedure below).

5.2 Procedures and their contextual adequacy

Social conformity comes in various forms, and different procedures may be adequate across contexts. For example, not acquainted experts deliberating on a subject and belonging to equally relevant fields may perhaps be equally influenced by each other. In other cases, social influence follows a public hierarchy, and a few popular or powerful agents may be particularly influential. For instance, tenure-track professors are more liable to conform to senior professors at department meetings; similarly, selective influence may be at work between experts engaged in deference relationships. Such contextual differences should be considered to assess when and why each procedure should be preferred. Accordingly, we now discuss the contextual advantages of the BAD and BOD procedures while keeping in mind that since agents may not be fully aware of which contexts they are in and simplicity is crucial for the success of social practices, procedures that do reasonably well across many contexts and are easily applicable should be globally privileged.

Self-generated voting orders are not perfect. Participative procedures provide opportunities for influential voters to speak immediately – typically, the boss may speak first, impress others, and make the minority view win the vote – even if they attenuate this inconvenience by detecting people that may oppose hierarchies and play the costly role of (first) dissenter (Asch 1951; Granberg and Bartels 2005).

As mentioned above, the BAD procedure, which draws the first *agent* at random, takes the best of Random Order (for high levels of conformity) and BOD (for low ones). Still, should it be preferred unconditionally across all real situations? Selecting a representative first voter is beneficial only if this person sides publicly for what she privately believes. This may not happen if she fears, rightly or not, that she might be in the minority or might oppose existing hierarchies. That is, a random pick may select someone too sensitive to social pressure, thereby ruining the advantage of representativeness. One faces a quandary here: a representative first speaker is specifically needed when social pressure is high to prevent an influential minority from unduly winning the vote; however, a higher social pressure may impact a randomly picked voter more significantly. By contrast, as BOD selects an option at random, the first voter is a volunteer, and she has every reason to be sincere – but this still gives the very influential boss a 50% chance of speaking first. Depending on the context, the glass may be seen as half-full or half-empty.

Other potential improvements may be worth exploring more systematically. One may start alternate procedures with the *apparent* minority view – call this the BMC procedure: “Balanced, apparent Minority Chosen.” The rationale for BMC is that views in the apparent minority may be in the majority, typically when their defenders are not influential, are more likely to be misrepresented orally, and crucially need support; views in the apparent majority are easy to endorse publicly and are unlikely to be overthrown if they correspond to the actual majority. For example, when discussing the issue of binge drinking with students, asking whether someone dislikes this practice may help unravel the hidden fact that most students do not support it. However, such a procedure faces difficulties. Which position is in the apparent or perceived minority cannot be ascertained

with certainty by formal means. Deferring this decision to a chair reintroduces risks of manipulation, and the procedure may be detrimental in the presence of influential hierarchy-rooted minorities.

Overall, the following typology of cases may be sketched. When social influence is low to moderate, the BAD procedure should be used. When social influence is high but equally distributed (that is, without any obvious social hierarchy) and no apparent consensus may put pressure on the first voter, the BAD procedure may still be adequate. In cases of apparent consensus or of asymmetric outside pressure, one may adopt the BMC procedure to avoid potential situations of pluralistic ignorance. When conformity is high and rooted in social hierarchies, different cases should be considered. When disagreeing influential voters can be expected to support both views at the first two voting ranks, the BOD procedure may be a cautious choice to avoid selecting a vulnerable (although representative) first voter. When social pressure is extreme or an influential social hierarchy structure is likely to determine votes, oral procedures may be inadequate, and a secret vote should be considered. Or, one could switch to the following heavier, provisionally secret voting procedure, which still preserves accountability and transparency. First, voters are asked to cast a secret ballot so all voters may safely express their private views and learn where the majority lies. Then, to validate the vote for public records if transparency is mandatory or to make the real vote, the chair asks who supports the majoritarian option or uses the BOD procedure (starting with the majoritarian view), and voters supporting this view may say, "Me" safely with social support and full confidence that the decision will not be stolen.[11]

5.3 Application issues

Alternate oral procedures are easily applicable and should be quicker than a secret ballot in small to moderate-size groups. Also, they remain rather informal and encourage a respectful and open stance toward all views, which matters for decisions on sensitive issues. But do not they also have practical drawbacks?

First, both BOD and BAD rely on a random draw: of the option for the former and of an agent for the latter. What about the practical problems highlighted above concerning the production of random voting orders? Actually, these inconveniences do not apply to BOD because a single coin toss may be used to select an option. They do apply to BAD in a milder way since one single random draw is needed to select an agent. Whenever this attenuated inconvenience is deemed too important, BOD may be preferred.

Second, the reliance of alternate procedures on a self-designation mechanism raises the issue of misapplications or the potential sabotage of the procedure. Voting procedures require good rules of demeanor but must also be proof against such situations. Some people may declare their oral vote by saying "Me" simultaneously. Or ill-intentioned voters may do so deliberately to spoil the procedure by trying to launch a cascade while pretending it was not deliberate. Nevertheless, if people are well-intentioned, the risk of misapplications is low for small to medium-sized groups since verbal voting signals can be emitted and perceived in a fraction of a second (hence the suggestion that people say "Me" and not raise their hands, which is lengthier and may be used for coordination purposes). Still, what should be done whenever two voters almost simultaneously say "Me"? In an alternate voting order (say, ABAB…), the voting ratio is tipped towards A every two votes, even if decreasingly. If two people vote for B simultaneously, the chair may pursue the alternate order, which now becomes ABABBABA…, since, in this way, the periodic decreasing advantage has merely been switched from A to B. However, if two people simultaneously vote for A, the count now favors A by a two-vote margin, which is an important advantage. The chair may try to compensate and call for B voters twice in a row, or she may cancel the oral vote and switch back to the heavier solution presented earlier: a secret ballot followed by an oral participative

[11] An anonymous secret straw vote is sometimes used in juries before or after deliberations (Elster 2015, 5).

alternate vote. In this way, the mere possibility of defeating sabotage by returning to this heavier procedure may be sufficient to discourage ill-intentioned voters. At the end of the day, how frequently such application problems indeed occur and whether the proposed solutions work should be assessed empirically.

5.4 Directions for future research

The present proposals may be extended or refined in various theoretical or empirical directions. On the modeling side, one could investigate cases where agents have different degrees of social influence and pressure sensitivities. Comparative investigations could also assess at a finer-grained level the effects of uncertain priors, higher-order beliefs, or of different types of conformist behavior (see Section 2.1). Other procedures, such as anti-hierarchical orders, may also be worth investigating further. For example, when social hierarchies are public, the group may vote alternately within the different categories, starting with the lowest one. However, mixing anti-hierarchical orders with balanced procedures (BAD or BOD) is not obvious: when nobody wants to support view A in a category, should the remaining votes of that category count for the other view, or should one look for an A-supporter in other categories? Additionally, investigations about balanced alternate procedures for votes with more than two options, for cases where deliberations and votes are intermingled (that is, when people give their views and justify them), or for deliberations proper, would be valuable. In any case, empirically analyzing alternate procedures through lab or field experiments would prove useful. The degree of satisfaction of group members after the vote (to assess efficiency) would be interesting to measure; similarly, the issue of misapplication and sabotage could be monitored with special care.

6. Conclusion

Conformity cascades often spoil the results of sequential oral votes. Anti-hierarchical and random voting orders may be used to dampen their effects, but also have their defects. We have explored the potential of some balanced participative procedures in which group members alternately volunteer to vote for the competing views (“Is there one person willing to vote for A?”, “Me.”). Our results highlight that these alternate rules can effectively eradicate cascades in one voting round, thus outperforming other oral procedures. They do this by letting agents volunteer to vote alternately for the two options. With such endogenously generated voting orders, agents willing to speak their minds despite contrary pressure may vote early; both views can be equally supported so long as the private views of voters allow it; undue pressure and influence keep decreasing; influenceable voters may vote late (when the voting count is balanced enough); voting sincerity is more preserved; and, overall, voting results are less distorted. With these procedures, the first position to be supported may be chosen differently. The chair may randomly draw the first agent or option (BAD and BOD procedures, respectively) or choose the view in the apparent minority (BMC procedure). While the BAD procedure may be more efficient at picking a representative first voter and eradicating distorted decisions, the BOD order may be preferred when extreme social pressure may overinfluence a randomly picked first voter.

Overall, these results suggest applying these procedures in contexts where oral voting is already used, such as clubs, department meetings, ethical review boards, or boards of directors, and, in a second step, in situations with explicit formal voting procedures, such as FDA advisory committees and expert agencies, or political councils. Tests in political participative contexts, typically at the end of deliberative polls, would be specifically welcome.

7. Appendix

We present here technical details and quantitative results. The model pictures n agents with fixed private views *pi* drawn from a flat distribution over [0, 1]. These agents must choose between two options, A and B, represented respectively by 0.25 and 0.75. Several oral voting orders may be applied. For increasing (resp. decreasing) orders, the views p_i are ranked in increasing (resp. decreasing) order, and agents vote according to their rank in the list. For the random order, voting ranks are drawn randomly with equiprobable chances. Let us define the public vote function $F(p)$ that selects the closest option: $F(p)$ maps [0, 0.5] onto 0.25 and]0.5, 1] onto 0.75. Without social influence, agent i, who has private view p_i, would vote for $F(p_i)$.

Let us now define the group's expressed view $G(j)$ as the linear average of the last $j - 1$ expressed views for voting round 1 and voting rank j, and of the last $n - 1$ expressed views for the other voting rounds. For example, in round 1, for the second voter, the group's expressed view $G(2)$ is just the public view of the first voter. Parameter α in [0, 1] describes how strongly agents feel forced to misrepresent their views. Because of social influence, the view that agent i would accept to endorse publicly at rank j is $s_i(j) = \alpha\, G(j) + (1 - \alpha)\, p_i$, that is, a weighted linear average of her private view and what is for her the group's expressed view (for the first voter, $s_i(1) = p_i$). Then, when they vote publicly, agents side with the voting option closest to $s_i(j)$. Overall, if agent i votes at rank j, her vote $V_i(j)$ is $F(s_i(j))$. After a voting round, the group's decision is the option that has gathered the majority of the n oral votes. To avoid complications, n is assumed to be odd.

In alternate orders, each view is alternately supported in an ABAB… or BABA… sequence by agents who are, at each step, drawn randomly among the remaining voters who may accept to endorse the view publicly at this step (as described just above). With the BAD procedure, the first voter is drawn randomly; with BOD, the first option is drawn randomly; with BMC, the chair chooses the first option based on what seems to be the minority view. The alternate vote is pursued till rank m, where no public support for one view can be found, and the remaining votes count for the last supported view. Mathematically, when, at rank j, a vote for X_j is expected, the chair orally calls a vote for X_j. Let $K_{X,j}$ be the set of voters who have not yet voted and for whom $F(s_k(t)) = X_j$ (i.e., who would accept to vote publicly for X_j). An agent, drawn randomly from $K_{X,j}$, votes publicly for X_j. If $K_{X,j}$ is empty, rank m has been reached: nobody volunteers to vote for X_j, and the other view wins the vote.

We now prove some general propositions about alternate procedures.

Proposition 1. With alternate procedures like BAD, BOD, and BMC, with $\alpha \neq 1$, for any rank $j < m$, vote j is sincere (i.e., the agent publicly votes for the option she privately favors).

Proof. At odd rank, $G(j) = 0.5$, so voters do not feel pressured toward either view and vote sincerely. At even rank, $G(j) \neq 0.5$ and pushes against the view X that is called; thus, people within $K_{X,j}$ would accept to vote for X despite contrary pressure. Because $s_i(j)$ is a linear average with positive coefficients, this implies that they privately prefer X and vote sincerely. ■

Proposition 2. For alternate procedures BAD, BOD, and BMC, the first view to receive a vote cannot lose the vote unduly (i.e., if there is a majority of private views for the first view, the group's decision will be for this view) but may win the vote unduly. Conversely, the second view cannot win the vote unduly but may lose it unduly.

Proof. All votes till m are sincere by Prop. 1. At all odd ranks j till m, the vote count is even ($G(j) = 0.5$), so all remaining voters who privately support the called view X_j are within $K_{X,j}$ and may

vote sincerely without pressure. If m is odd, this means that all private supporters of X_j have already voted sincerely for X_j and, therefore, are in the minority. ■

Proposition 3. For alternate procedures BAD, BOD, and BMC, adding additional rounds of votes with the same voting order as in the first round does not decrease distortion.

Proof. Without loss of generality, assume the first called view is A (at 0.25). We now add subscripts to $G(j)$ to denote voting rounds. Because of Prop. 2, the only cases of distortion are those in which B unduly loses the vote at an even rank where $G_1(m) < 0.5$ and remaining B supporters dare not vote publicly for B. Therefore, only cases where A has won the vote need to be considered. Since the remaining votes count for A, at the final rank n of round 1, $G_1(n) < G_1(m)$. During the second round, A supporters will keep voting for A till rank m; in the most favorable case for B, B supporters will also keep voting for B till rank m, at which $G_2(m) = G_1(n) < G_1(m)$. Then, the remaining B supporters can even less vote publicly for B, and the vote result remains identical. ■

Proposition 4. For alternate procedures BAD, BOD, and BMC, adding additional rounds of votes in which the first view to be supported remains the same and the procedure is re-applied does not decrease distortion.

Proof. Without loss of generality, assume the first called view is A. As in the proof of Prop.3, only cases where A has won the vote need to be considered, and then, at the final rank n of round 1, $G_1(n) < G_1(m)$. Permuting the voting order of private A supporters does not make any difference, and such changes can be ignored. When the procedure is re-applied, the most favorable case for B is that some private B supporters vote for B and $G_2(j) = G_1(n)$ till m. Because $G_2(j) < G_1(m)$, this new voting order before m may merely involve permutations between B supporters who had voted for B before m in round 1. Then, at rank m, remaining private B supporters now face a worse situation than at round 1; thus, more votes for B cannot be made, and distortion cannot decrease. ■

Proposition 5. Given the first voter (and, consequently, the option which is to receive the first vote), no procedure that alternates votes between A and B can yield more sincere votes than participative procedures BAD, BOD, or BMC, and even an omniscient agent knowing all private views cannot generate a better alternate order. In this sense, asking people to volunteer to vote is optimal for alternate voting schemes.

Proof. If all voters privately support the same view, all votes are sincere with any procedure; therefore, let us now consider non-consensual cases. By Prop. 2, with an alternate order, the first view to receive a vote cannot lose the vote unduly. An omniscient God may select the first voter so as to make the majority always win, hence the condition concerning the first voter. Without loss of generality, let us now assume that the first voter votes for A. As in the proof of Prop. 4, only cases where B loses the vote unduly need to be considered. With an alternate scheme, at rank j, $G_j = [j/2 \cdot 0.25 + 0.75 \cdot (j/2 - 1)]/(j - 1) = 0.5 - 1/[4(j - 1)]$. Since for odd j, $G_j = 0.5$, all permutations between A supporters are possible but do not make any difference. Define $L_{B,j}$ as the set of B supporters who would be willing to vote publicly for B at some even $j < m$, irrespective of whether they have already voted. Since for all even j inferior to m, G_j increases with j, any agent in $L_{B,j}$ would also be willing to vote for B at some larger j, that is, $L_{B,2} \subseteq L_{B,4} \subseteq \ldots \subseteq L_{B,m}$. Further, the values of the G_{2k} and, therefore, the $L_{B,2k}$ sets till m, do not depend on the specific order of voters within the alternate scheme. Thus, the set $\mathcal{P}$ of all possible alternate voting orders is obtained by sequentially picking, in all possible picking orders, one voter from each $L_{B,2k}$ after removing from this $L_{B,2k}$ all previously picked voters. That is, one picks a voter v from $L_{B,2}$ and removes her from all the $L_{B,2k}$, then one

picks another voter w from $L_{B,4} \setminus v$ and removes her from all the $L_{B,2k} \setminus v$, etc. For all such orders, all the $m/2$ members of $L_{B,m}$ vote for B before m and, by Prop. 1, their votes are sincere. Now, consider the set I of B supporters who have not voted at m yet. This set is the same for all alternate voting orders within $\mathcal{P}$ because these agents could not have voted earlier for B since, for all such orders, $G_j < G_m$ for all even j such that $j < m$. Thus, given a first voter, no alternate voting order can create a setting in which some members of I vote sincerely, all alternate orders yield the same number of sincere voters for B and, overall, the same number of sincere votes. Thus, all voting orders within $\mathcal{P}$, in particular, those produced by a participative alternate procedure, are equivalent and optimal. ■

Propositions 3, 4, and 5 delineate a sense in which procedures BAD, BOD, and BMC are optimal within our model. Their common point is to call for voters alternately for both views.

Finally, figure 1 plots an overview of the percentage of distorted decisions (i.e., group decisions that differ with and without social influence) for the BAD and BOD procedures, compared with those for the random and increasing orders (BMC cannot be simulated as it starts from a singular choice made by the chair). For each set of parameters, distortion results were computed by drawing individual private views from a flat distribution and then, for each draw and voting procedure, simulating sequential oral votes with a Java script and averaging results over tens to hundreds of thousands of replicas to warrant sufficient convergence. The complete Java script may be found here (provisional link):
https://sharedocs.huma-num.fr/wl/?id=c2rkqjhsjshRkreafn47QB0nYP3QNtOD.

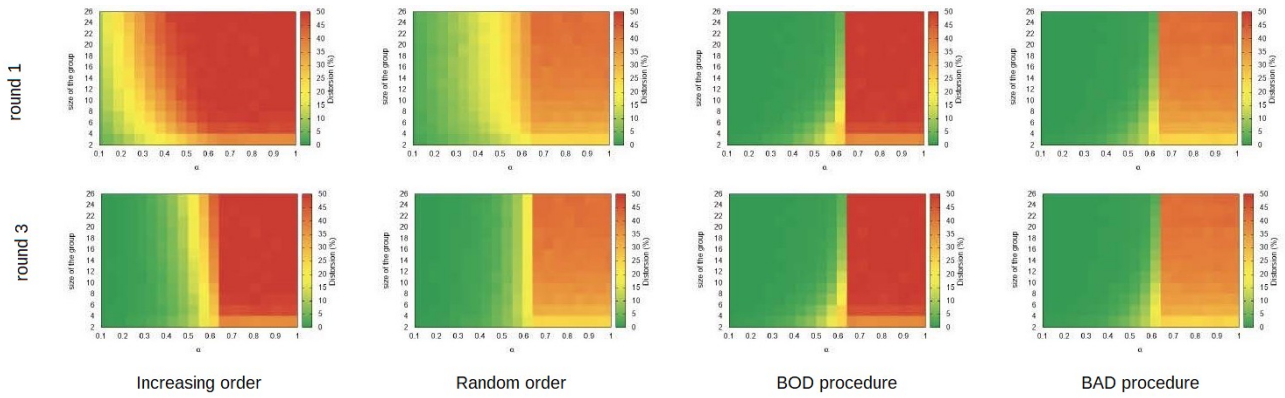

Fig. 1: Percentage of distorted decisions as a function of the degree of conformity α (x-axis) and of the size of the group (y-axis), for various voting orders and 1 or 3 rounds. Green means a minimal rate of distortion, and red means a maximal one.

To interpret the significance of this percentage, note that if all decisions were taken randomly or arbitrarily in favor of one view, on average, 50% of decisions would be distorted. Thus, even 15% of distorted decisions constitutes a poor performance for a procedure. This overview shows that the increasing order requires several rounds to reach acceptable distortion rates for low α but is still bad for high α. The random order is better for high α but still requires several rounds. The BOD procedure gives good results for moderate α in one round but bad ones for high α. The BAD procedure somehow takes the best of both worlds; it combines the advantage of the BOD procedure, being very good for small and moderate α (left green part), and of the random procedure, being not too bad for high α (right orange part for large α). Further, contrarily to the increasing and random orders, the BOD and BAD procedures yield their good results in one single round.

References

Asch, S. E. 1951. “Effects of Group Pressure upon the Modification and Distortion of Judgments.” *APA PsycNET*.

Balinski, Michel, and Rida Laraki. 2020. “Majority Judgment vs. Majority Rule.” *Social Choice and Welfare* 54 (2): 429–61.

———. 2022. “Majority Judgment vs. Approval Voting.” *Operations Research* 70 (3): 1296–1316.

Bentham, Jeremy. 1999. *Political Tactics*. The Collected Works of Jeremy Bentham, Ed. by James, Michael, Blamires, Cyprian, Pease-Watkin, Catherine. New York: Oxford University Press.

Bond, Rod. 2005. “Group Size and Conformity.” *Group Processes & Intergroup Relations* 8 (4): 331–54.

Bright, Liam Kofi, Haixin Dang, and Remco Heesen. 2018. “A Role for Judgment Aggregation in Coauthoring Scientific Papers.” *Erkenntnis* 83 (2): 231–52.

Chambers, Simone. 2004. “Behind Closed Doors: Publicity, Secrecy, and the Quality of Deliberation.” *Journal of Political Philosophy* 12 (4): 389–410.

Cialdini, Robert B., and Vladas Griskevicius. 2010. “Social Influence.” In *Advanced Social Psychology: The State of the Science*, 385–417. New York, NY, US: Oxford University Press.

Elster, Jon, ed. 2015. *Secrecy and Publicity in Votes and Debates*. New York: Cambridge University Press.

Fazelpour, Sina, and Daniel Steel. 2022. “Diversity, Trust, and Conformity: A Simulation Study.” *Philosophy of Science* 89 (2): 209–31.

Forsyth, Donelson R. 2019. *Group Dynamics*. Seventh edition. Australia ; Boston, MA: Cengage.

Gibbard, Allan. 1973. “Manipulation of Voting Schemes: A General Result.” *Econometrica* 41 (4): 587.

Gosseries, Axel, and Tom Parr. 2022. “Publicity.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2022. Metaphysics Research Lab, Stanford

Granberg, Donald, and Brandon Bartels. 2005. “On Being a Lone Dissenter.” *Journal of Applied Social Psychology* 35 (9): 1849–58.

Hartmann, Stephan, and Soroush Rafiee Rad. 2020. “Anchoring in Deliberations.” *Erkenntnis* 85 (5): 1041–69.

Imbert, Cyrille, Thomas Boyer-Kassem, Vincent Chevrier, and Christine Bourjot. 2020. “Improving Deliberations by Reducing Misrepresentations Effects.” *Episteme* 17 (4): 403–19.

Janis, Irving L. 1983. *Groupthink: Psychological Studies of Policy Decisions and Fiascoes*. 2nd ed., Rev. Boston

Kahneman, Daniel, Paul Slovic, and Amos Tversky, eds. 1982. *Judgment under uncertainty: heuristics and biases*. Cambridge: United States.

Kerr, Norbert L., and R. Scott Tindale. 2004. “Group Performance and Decision Making.” *Annual*

Review of Psychology 55 (1): 623–55.

Kuran, Timur. 1997. *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. United States of America.

Mallinson, Daniel J., and Peter K. Hatemi. 2018. "The Effects of Information and Social Conformity on Opinion Change." *PLOS ONE* 13 (5)

Martini, Carlo, and Manuela Fernández Pinto. 2017. "Modeling the Social Organization of Science." *European Journal for Philosophy of Science* 7 (2): 221–38.

Meade, Ellen E., and David Stasavage. 2008. "Publicity of Debate and the Incentive to Dissent: Evidence from the US Federal Reserve*." *The Economic Journal* 118 (528): 695–717.

Mercier, Hugo, and Olivier Morin. 2019. "Majority Rules: How Good Are We at Aggregating Convergent Opinions?" *Evolutionary Human Sciences* 1 (January).

Mohseni, Aydin, and Cole Randall Williams. 2021. "Truth and Conformity on Networks." *Erkenntnis* 86 (6): 1509–30.

Ottaviani, Marco, and Peter Sørensen. 2001. "Information Aggregation in Debate: Who Should Speak First?" *Journal of Public Economics* 81 (3): 393–421.

Šešelja, Dunja. 2023. "Agent-Based Modeling in the Philosophy of Science." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta and Uri Nodelman, Winter 2023. Metaphysics Research Lab, Stanford University.

Stasser, Garold, and William Titus. 1985. "Pooling of Unshared Information in Group Decision Making: Biased Information Sampling during Discussion." *Journal of Personality and Social Psychology* 48:1467–78.

Sunstein, Cass R. 2003. *Why Societies Need Dissent*.

Tindale, R. Scott, and Jeremy R. Winget. 2019. "Group Decision-Making." In *Oxford Research Encyclopedia of Psychology*, by R. Scott Tindale and Jeremy R. Winget. Oxford University Press.

Urfalino, Philippe, and Pascaline Costa. 2015. "Secret-Public Voting in FDA Advisory Committees." In *Secrecy and Publicity in Votes and Debates*, edited by Jon Elster, 165–95.

Weatherall, James Owen, and Cailin O'Connor. 2021. "Conformity in Scientific Networks." *Synthese* 198 (8): 7257–78.