Deliberation and epistemic democracy*

Huihui Ding and Marcus Pivato†

May 11, 2021

Abstract

We study the effects of deliberation on epistemic social choice, in two settings. In the first setting, the group faces a binary epistemic decision analogous to the Condorcet Jury Theorem. In the second setting, group members have probabilistic beliefs arising from their private information, and the group wants to aggregate these beliefs in a way that makes optimal use of this information. During deliberation, each agent discloses private information to persuade the other agents of her current views. But her views may also evolve over time, as she learns from other agents. This process will improve the performance of the group, but only under certain conditions; these involve the nature of the social decision rule, the group size, and also the presence of "neutral agents" whom the other agents try to persuade.

JEL classification: D71, D83.
Keywords: Deliberation; Epistemic social choice; Condorcet Jury Theorem; Probabilistic belief aggregation; Multiplicative pooling

1 Introduction

In many collective decisions, there is an objectively correct answer, and the group wants to find it. For example, a criminal trial jury must decide whether a defendant is innocent or guilty. The Supreme Court must determine the constitutional validity of laws or lower-court decisions. The directorate of the Central Bank must evaluate the risks of inflation

*Supported by CHOp (ANR17-CE26-0003) and Labex MME-DII (ANR11-LBX-0023-01). Ding gratefully acknowledges the support of I-SITE, PIA2, Université de Cergy-Pontoise. We thank Pierpaolo Battigalli, Richard Bradley, David Budescu, Alain Chateauneuf, Franz Dietrich, Simona Fabrizi, Maria Gallego, Mario Roberto Gilli, Ariane Lambert-Mogiliansky, Antoine Mandel, John McCoy, Klaus Nehring, Drazen Prelec, Peter Sørensen, Emily Tanimura, and Bauke Visser for helpful comments.

†THEMA, Université de Cergy-Pontoise, 33, boulevard du Port, 95011 Cergy-Pontoise Cedex, France. Email: huihui.ding@u-cergy.fr and marcuspivato@gmail.com

and recession over the coming year. The senior management of a firm must determine which business strategy will maximize long-term profits. Finally, blue ribbon commissions and scientific committees must advise policy-makers on questions of scientific fact. *Epistemic social choice theory* studies the conditions under which voting rules or opinion aggregation methods can deliver accurate answers to such questions (see Pivato, 2013, 2017 or Dietrich and Spiekermann, 2017 for summaries). This has inspired a parallel literature in political philosophy on *epistemic democracy* (Cohen, 1986; Landemore and Elster, 2012; Schwartzberg, 2015; Goodin and Spiekermann, 2018). Starting with the Condorcet Jury Theorem, most models in epistemic social choice theory have assumed that the opinions of different voters are stochastically independent, conditional on the true state of nature.[1] But in reality, the opinions of voters are correlated, because they deliberate with one another. Indeed, there is now an extensive literature on *deliberative democracy* which argues that deliberation should *improve* the epistemic competency of groups (Estlund and Landemore, 2018). But it is not clear that deliberation is always beneficial in this regard. Deliberation enables agents to pool information, but they might double-count this information (Berg, 1997). Deliberation can also lead to "groupthink" (Mayo-Wilson et al., 2013), informational cascades (Banerjee, 1992), and other pathologies. On the balance, does deliberation really lead to more epistemically reliable group decisions?

Suppose that any initial heterogeneity of beliefs amongst the agents arises from heterogeneity of private information. And suppose that during deliberation, every agent truthfully reveals *all* of her private information to the group, and every other agent correctly understands this information and updates her beliefs accordingly. Then deliberation will generally improve the reliability of the group, because the collective decision will fully incorporate the pooled information of all group members. So if deliberation were this simple, then the answer to the question in the previous paragraph would be trivially affirmative. But in reality, things are not so simple, for two reasons. First, agents may either lie or withhold information to manipulate the group decision to their own advantage. Second, even amongst honest agents, communication can be costly: it may take time and effort for

[1] For exceptions, see Pivato (2017) and the references cited there.

an agent to clearly and credibly convey her private information to other group members so that they fully understand it and update their beliefs accordingly. Indeed, if "private information" is interpreted in the broadest sense, to encompass all of an agent's prior education and professional experience, then it is simply infeasible for her to communicate *all* her information to the group; she must select some small subset to disclose. Over the past two decades, a considerable literature on *strategic deliberation* has developed in response to the first problem. But there has been very little examination of the second problem. This paper aims to fill this gap.

We will construct a stylized model of deliberation with costly communication amongst Bayesian agents, and investigate whether deliberation improves the reliability of the group's decision. In our model, the true state of the world is unknown; the group wants to determine this state. The agents have each received a set of private signals ("evidence") that are informative about the true state. In each time period, each piece of evidence is either *private* or *public*. *Private* evidence is known by only one (or a few) agents. *Public* evidence is known by everyone. Agents are Bayesian; at any moment in time, each agent's beliefs are obtained by combining her prior beliefs with her currently available evidence —that is, her own private evidence and all the publicly available evidence. Deliberation is a way for agents to "disclose" some of their private evidence, turning it into *public* evidence, so as to modify the beliefs of the other agents.[2] But evidence disclosure takes time. An agent might have many pieces of private evidence, but she can only disclose one piece at a time. Thus, unlike most deliberation models in the literature, our model is *diachronic*; we track the evolution of the agents' beliefs (and evidence base) over time.

Furthermore, evidence disclosure is costly: communication involves an expense of time and effort, both for the sender and the receivers. So an agent will not disclose her private evidence without an incentive. Her incentive is to convince the other agents of the correctness of her current views. Unlike the literature on strategic deliberation, we assume that

[2]This is similar to the *hidden-profiles paradigm* of Stasser and Titus (1985), which has played a prominent role in the social psychology literature on deliberation; see Lu et al. (2012) and Maciejovsky and Budescu (2019) for recent reviews.

the agents deliberate in *good faith* —that is, each agent simply wants the group decision to converge to the correct answer (or at least, what she currently *believes* to be the correct answer). Strategic deliberation is typically driven by heterogeneity of preferences. In our model, there is no heterogeneity of preferences —only heterogeneity of *beliefs*.

Agents take turns disclosing pieces of evidence. Each agent only discloses evidence that will move the beliefs of other group members closer to her own current beliefs. But at the same time, her own beliefs also evolve, as a consequence of the evidence revealed by other agents. This process continues until the group reaches a *deliberative equilibrium* in which every agent is either unwilling or unable to disclose further evidence. This raises two questions. First, do such equilibria exist? Second, how accurate are the resulting group decisions?

The rest of the paper is organized as follows. Section 2 presents the model. Section 3 considers deliberation in a binary decision, and contains Theorems 1. At last, there is a conclusion.

2 Framework

Let $\mathcal{X}$ be a finite set of possible states of the world. Let $\mathcal{Y}$ be a space of possible signal values. Let $\Delta^*(\mathcal{X})$ be the set of all probability measures on $\mathcal{X}$ with *full support* —i.e. such that every element of $\mathcal{X}$ receives nonzero probability. Define $\Delta^*(\mathcal{Y})$ in the same way. Let $\rho : \mathcal{X} \longrightarrow \Delta^*(\mathcal{Y})$ be a function; if the true state of the world is $x \in \mathcal{X}$, then the signals will be conditionally independent, identically distributed (i.i.d) random variables drawn from the probability distribution $\rho(x)$, which we will write as ρ_x.

Let $\widetilde{x}$ be the (unknown) true state of the world. Suppose a Bayesian agent has prior beliefs about $\widetilde{x}$ given by $\pi \in \Delta^*(\mathcal{X})$, and she receives a sequence $\mathbf{y} = (y_1, y_2, \ldots, y_M)$ of i.i.d. random signals drawn from $\rho_{\widetilde{x}}$. Let $B_{\mathbf{y}}$ be her posterior beliefs about $\widetilde{x}$, given $\mathbf{y}$. Then $B_{\mathbf{y}}$ is the following probability distribution over $\mathcal{X}$:

$$B_{\mathbf{y}}(x) \quad = \quad \frac{1}{(\text{SNC})}\, \pi(x) \cdot \prod_{m=1}^{M} \rho_x(y_m), \qquad \text{for all } x \in \mathcal{X}. \tag{2A}$$

Here and throughout the paper, "(SNC)" refers to "Some Normalization Constant", needed to ensure that the expression in question defines a probability distribution. These normalization constants are not important to the analysis, so we will not specify them explicitly.[3]

Now let $\mathcal{I}$ be a set of agents. For all $i \in \mathcal{I}$, let $\pi_i \in \Delta^*(\mathcal{X})$ be a probability distribution describing i's prior beliefs, before acquiring any information. Let $\mathcal{M}$ be a finite indexing set, and let $\{y_m\}_{m\in\mathcal{M}}$ be a set of i.i.d. random signals drawn from ρ_x, where x is the (unknown) true state of the world. We refer to these signals as ***evidence***. For all $i \in \mathcal{I}$, let $\mathcal{M}_i \subset \mathcal{M}$ be the set of evidence received by agent i —this is i's initial private information. (We do not necessarily assume that these sets are disjoint.) Meanwhile, let $\mathcal{C}^0 \subset \mathcal{M}$ be the set of evidence that is common information at time 0. (We assume it is disjoint from the sets $\mathcal{M}_i$.) Thus, before deliberation, formula (2A) says that i's initial beliefs are given by

$$B_i^0(x) \quad = \quad \frac{\pi_i(x)}{\text{(SNC)}} \prod_{c\in\mathcal{C}^0} \rho_x(y_c) \cdot \prod_{m\in\mathcal{M}_i} \rho_x(y_m), \qquad \text{for all } x \in \mathcal{X}. \tag{2B}$$

We can assume without loss of generality that

$$\mathcal{M} \quad = \quad \mathcal{C}^0 \sqcup \bigcup_{i\in\mathcal{I}} \mathcal{M}_i, \tag{2C}$$

because any evidence which is not in this union will never be learned by the group after any amount of deliberation, and is therefore irrelevant to our analysis.

Deliberation takes place in a series of "rounds", indexed by the set of natural numbers $\mathbb{N} = \{0, 1, 2, \ldots\}$. For any $t \in \mathbb{N}$, let $\mathcal{C}_i^t \subseteq \mathcal{M}_i$ be the part of agent i's evidence which she has disclosed (i.e. made into common information) at time t. Thus, $\mathcal{P}_i^t := \mathcal{M}_i \setminus \mathcal{C}_i^t$ is the part of agent i's evidence which remains "private" at time t. Let

$$\mathcal{C}^t \quad := \quad \mathcal{C}^0 \sqcup \bigcup_{i\in\mathcal{I}} \mathcal{C}_i^t. \tag{2D}$$

This is the set of evidence that is common information at time t. Let B_i^t be the probabilistic beliefs of agent i at time t. This is a combination of her prior, the publicly available

[3]To be precise, in equation (2A), $\text{(SNC)} = \sum_{x\in\mathcal{X}} \pi(x) \left(\prod_{m=1}^{M} \rho_x(y_m) \right)$.

evidence, and her own (undisclosed) private evidence. By applying (2A), we obtain:

$$B_i^t(x) \quad = \quad \frac{\pi_i(x)}{(\mathrm{SNC})} \cdot \prod_{p \in \mathcal{P}_i^t} \rho_x(y_p) \cdot \prod_{c \in \mathcal{C}^t} \rho_x(y_c), \qquad \text{for all } x \in \mathcal{X}. \tag{2E}$$

Thus, when agent i discloses evidence (i.e. transfers it from $\mathcal{P}_i^t$ to $\mathcal{C}^t$), her own beliefs will not change, but this disclosure may affect the beliefs of other people. Conversely, agent i's own beliefs can evolve as $\mathcal{C}^t$ accumulates more and more evidence from other people. At any moment in time, her deliberation behaviour (i.e. the sort of evidence she discloses) depends on her current beliefs. So as her beliefs evolve, her behaviour may change. She may stop advocating one position (i.e. stop disclosing evidence supporting this position), and even start advocating a position she previously opposed. In Sections 3, we will explore two special cases of this model.

3 Deliberating binary decisions

In this section, we will focus on the simplest nontrivial example of the model, where $\mathcal{X} = \mathcal{Y} = \{\pm 1\}$. Fix $p \in (\frac{1}{2}, 1)$, and suppose that $\rho(x|x) = p$ and $\rho(-x|x) = 1 - p$, for both $x \in \mathcal{X}$. This is essentially the set-up of the Condorcet Jury Theorem. We will describe $+1$-valued signals as "positive" and -1-valued signals as "negative". We assume all agents have a common prior π, which assigns probability $\frac{1}{2}$ to each state.

For all $i \in \mathcal{I}$ and $t \in \mathbb{N}$, the *opinion* of agent i at time t is the variable $s_i^t \in \{-1, 0, 1\}$ describing whether i believes the positive state, or the negative state to be more probable, according to the probabilistic beliefs B_i^t defined by equation (2E). Formally,

$$s_i^t \quad := \quad \begin{cases} 1 & \text{if} \quad B_i^t(1) > B_i^t(-1); \\ 0 & \text{if} \quad B_i^t(1) = B_i^t(-1); \\ -1 & \text{if} \quad B_i^t(1) < B_i^t(-1). \end{cases} \tag{3A}$$

Using equation (2E) and the common prior π, it is easily verified that s_i^t is entirely determined by the amount of positive and negative evidence available to i at time t. Formally,

$$s_i^t \quad = \quad \mathrm{sign}\left(\sum_{p \in \mathcal{P}_i^t} y_p + \sum_{c \in \mathcal{C}^t} y_c \right). \tag{3B}$$

Dissent and disclosure. We assume that it is easy for agents to learn their peers' opinions during each round of deliberation, simply by asking them yes/no questions or holding a straw vote. But in this section, we assume that it is *not* possible for agents to learn the underlying beliefs of their peers. Furthermore, it is difficult to learn what evidence —or even *how much* evidence —their peers have to justify these opinions. We have abstractly represented each "piece of evidence" as a single binary signal. But in reality, it may by a complex corpus of facts, analysis, interpretation and arguments that may take considerable time and effort for an agent to clearly explain to her peers. Thus, an agent will not "disclose" this evidence (i.e. explain these facts and arguments) unless she has an incentive to do so. We assume that each agent wants the collective decision to be correct. At any time during deliberation, she believes that her *current* opinion is correct; thus, she seeks to persuade other group members to agree with her current opinion.[4] She will disclose evidence only if it advances this goal. To be precise, she will only disclose evidence if her current opinion *disagrees* with the current collective decision —in this case, we say that she ***dissents*** from the group. Agents who already agree with the current collective decision will not disclose evidence, because such disclosure is costly, and they have no reason to incur this cost.

Although it will not play any formal role in the model that follows, it might be helpful to rationalize these behavioural assumptions with a stylized utility function. Let $\epsilon > 0$. Suppose agent i starts with M_i pieces of private evidence, and during the course of deliberation, she discloses C_i of them. Suppose that the group decision is $x \in \{-1, 0, 1\}$, but agent i's opinion is y. Then i's final utility will be $-|x-y| - \epsilon\, C_i/M_i$. So disclosing each piece of evidence is costly, but this cost is outweighed by i's desire for the group decision to agree with her opinion. Thus, at any time t, she is always willing to disclose more evidence, *if* she thinks it will increase by more than ϵ the probability of the group agreeing with her. If ϵ is very small, then even a tiny gain in the probability of the group agreeing with her opinion is sufficient incentive for i to disclose more evidence. But the group already agrees

[4] Note that $s_i^t = 0$ does not mean that i has "no opinion" —it means that i thinks ± 1 are *equally probable*, and hence the correct answer for the group is to remain ambivalent.

with her, then she will not incur the cost of disclosing further evidence.

Deliberative equilibrium. For most of this section, we will assume that the collective decision is made by *majority vote*. Formally, for any $t \in \mathbb{N}$, the ***majority opinion*** at time t is defined

$$\mathrm{Maj}^t \quad := \quad \mathrm{sign}\left(\sum_{i \in \mathcal{I}} s_i^t\right), \tag{3C}$$

where $\{s_i^t\}_{i \in \mathcal{I}}$ are the individuals' opinions from formula (3B). Let $i \in \mathcal{I}$. If $s_i^t = \mathrm{Maj}^t$, then we assume that i remains silent during round t. If $s_i^t > \mathrm{Maj}^t$, then i may disclose one positive signal from $\mathcal{P}_i^t$, so that it becomes part of $\mathcal{C}^{t+1}$.[5] On the other hand, if $s_i^t < \mathrm{Maj}^t$, then i may disclose one negative signal from $\mathcal{P}_i^t$, so it becomes part of $\mathcal{C}^{t+1}$.

After each disclosure, everyone in the group will update their beliefs based on the newly revealed evidence. This process of deliberation must end in finite time (because $\mathcal{C}^1 \subset \mathcal{C}^2 \subset \mathcal{C}^3 \subset \cdots \subset \mathcal{C}^T \subseteq \mathcal{M}$ and $\mathcal{M}$ is finite). When it ends, all the agents are silent, either because they have no evidence left to disclose, or because they have no incentive to disclose their remaining evidence. We then say that the group is in a ***deliberative equilibrium***.

Reliability and full disclosure. If the agents hold a majority vote at time t, then the decision is given by formula (3C). The ***reliability*** of this decision is the probability that it correctly identifies the true state.[6] We say there is ***full disclosure*** if all agents have complete information —that is $\mathcal{C}^t = \mathcal{M}$. In this case, all agents update to identical posterior beliefs via formula (3B); thus, the group is unanimous.

Proposition. *The (unanimous) majority decision under full disclosure achieves the maximum reliability possible given the information in $\mathcal{M}$.*

[5] We say she *may* disclose a signal because the actual timing of disclosures depends on the deliberation protocol, as explained below. Also, note that $\mathcal{P}_i^t$ must contain positive signals, because otherwise an examination of equation (3B) shows that it would be impossible for $s_i^t = 1 > \mathrm{Maj}^t$.

[6] A collective decision of "0" is assigned reliability 0.5. Also, we neglect the possibility of strategic voting raised by Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996, 1997, 1999).

Since communication is costly, full disclosure may be unachievable. Fortunately, it may also be unnecessary. We will say that a deliberative equilibrium is *full-disclosure equivalent* if the collective decision in the equilibrium is the same as the consensus that *would be* reached under full disclosure. So it is sufficient to seek full-disclosure equivalence.

Deliberation protocols. We have not yet described the timing of evidence disclosure —what we call the *protocol*. We will consider two possible protocols. In the *serial* protocol, only *one* dissenting agent can speak during each round of deliberation. In the *parallel* protocol, *all* dissenting agents can speak during each round of deliberation. The distinction between the two protocols is best understood as a question of the relative speed at which agents can transmit information, versus the speed at which they can incorporate new information into their beliefs. In the serial protocol, agents can incorporate new information very quickly: every time an agent reveals information, the other agents immediately update their beliefs. In the parallel protocol, agents update their beliefs more slowly. Thus, *all* dissenting agents have a chance to disclose new information, and then (perhaps during a brief interlude), all agents update their beliefs before the next round of deliberation begins.

This small difference in timing causes a big difference in the outcome. In the *serial* protocol, deliberation leads to full-disclosure equivalence under broad conditions (Theorem 1). But in the *parallel* protocol, deliberation is guaranteed to yield full-disclosure equivalence only under very special conditions (Theorem 2). These results depend on two assumptions.

(A) Agents have disjoint information sets. That is, $\mathcal{M}_i \cap \mathcal{M}_j = \emptyset$ for all distinct $i, j \in \mathcal{I}$.

(B) There is some agent $i \in \mathcal{I}$ who is *internally neutral*, by which we mean $\sum_{m \in \mathcal{M}_i} y_m = 0$.

3.1 The serial protocol

For any $t \in \mathbb{N}$, let $\{s_i^t\}_{i \in \mathcal{I}}$ be the individuals' opinions from formula (3B), and let $\mathcal{I}_+^t$ be the set of agents with *positive* opinions at time t. Let $\mathcal{I}_-^t$ be the set of agents with *negative* opinions at time t. Let $\mathcal{I}_0^t$ be the set of agents with *neutral* opinions at time t. Let

$E^0 := \sum_{c \in \mathcal{C}^0} y_c$ be the balance of public evidence at time zero. Suppose $E^0 \neq 0$; we will say there is *initial disagreement* if at least one agent has an opinion different from $\mathrm{sign}(E^0)$ at time 0. In other words, if $E^0 > 0$, then $\mathcal{I}_-^0 \sqcup \mathcal{I}_0^0 \neq \emptyset$, whereas if $E^0 < 0$, then $\mathcal{I}_+^0 \sqcup \mathcal{I}_0^0 \neq \emptyset$. During each round of deliberation, exactly one dissenting agent will reveal exactly one piece of evidence. To be precise, during round t, if $\mathrm{Maj}^t = 1$, then exactly *one* (randomly chosen) agent in $\mathcal{I}_-^t \cup \mathcal{I}_0^t$ reveals one piece of negative evidence. Likewise, if $\mathrm{Maj}^t = -1$, then exactly one (randomly chosen) agent in $\mathcal{I}_+^t \cup \mathcal{I}_0^t$ reveals one piece of positive evidence. Finally, if $\mathrm{Maj}^t = 0$, then exactly one (randomly chosen) agent in $\mathcal{I}_+^t \cup \mathcal{I}_-^t$ reveals one piece of evidence (either positive or negative). We have not specified the probability distributions by which these dissenting agents are "randomly chosen". If we specified these distributions, then we could describe serial protocol deliberation as a stochastic process. But it turns out that the precise probability distribution doesn't matter.

Theorem 1 *Assume* (A) *and* (B)*, and suppose that either* $E^0 = 0$*, or there is initial disagreement. Then the serial protocol always reaches a deliberative equilibrium that is full-disclosure equivalent.*

3.2 The parallel protocol

In the serial protocol of Section 3.1, only one agent could disclose evidence during each round of deliberation. But in the parallel protocol, *every* agent can disclose evidence in each round. However, an agent will disclose evidence only if she dissents from the majority. For any $t \in \mathbb{N}$, let $\mathcal{N}^t := \{i \in \mathcal{I};\ \sum_{p \in \mathcal{P}_i^t} y_p = 0\}$. (Thus, $\mathcal{N}^0$ is the set of internally neutral agents who appear in Assumption (B).) We will require a third assumption.

(C) For all $t \in \mathbb{N}$, no agent in $\mathcal{N}^t$ discloses information in round t.

If $\mathcal{P}_n^t = \emptyset$ for all $n \in \mathcal{N}^t$, or if $\mathrm{Maj}^t = \mathrm{sign}\left(\sum_{c \in \mathcal{C}^t} y_c\right)$, then Assumption (C) is trivially satisfied at time t. Otherwise, it is a substantive behavioural assumption. In effect it says: even if agents in $\mathcal{N}^t$ disagree with the majority decision at time t, their opinions are not strong enough to motivate them to disclose any further evidence. Let $\mathcal{I}^* := \mathcal{I} \setminus \mathcal{N}^0$ be the set of agents in $\mathcal{I}$ who are *not* internally neutral. Here is our second result.

Theorem 2 *In the parallel protocol, there is a unique deliberative equilibrium. If assumptions* (A), (B) *and* (C) *hold, $E^0 = 0$, and $|\mathcal{I}^*| \leq 4$, then this equilibrium is full-disclosure equivalent.*

Being similar to Section 3.1, assumptions (A) and (B) and the hypothesis $E^0 = 0$ are needed for the conclusion of Theorem 2. Assumption (C) and the condition $|\mathcal{I}^*| \leq 4$ are necessary because otherwise the balance of public evidence can lurch to an extreme negative or positive value, causing the group to get “stuck” in a suboptimal equilibrium, as shown in the next two examples.

Remarks. Note that the conditions of Theorems 1 and 2 are sufficient but not *necessary* for full-disclosure equivalence. For example, if all agents begin with the same information (i.e. $\mathcal{M}_i = \mathcal{M}$ for all $i \in \mathcal{I}$), then the group immediately reaches a full-disclosure equivalent deliberative equilibrium, whether or not Assumptions (A), (B) or (C) or the other hypotheses of Theorems 1 and 2 are satisfied. We do not yet know of conditions which are both necessary and sufficient for full-disclosure equivalence.

Conclusion

We have presented three stylized models of deliberation in which Bayesian agents with private information try to persuade one another by selectively revealing some of this information. We have shown that such a process of deliberation can enable the group to make optimal use of the information of its members, but only under certain conditions, which involve the nature of the group decision rule, the presence of internally neutral agents , and the number of opinionated agents.

References

Austen-Smith, D., Banks, J. S., 1996. Information aggregation, rationality, and the condorcet jury theorem. American Political Science Review 90 (1), 34–45.

Bächtiger, A., Dryzek, J., Mansbridge, J., Warren, M. (Eds.), 2018. The Oxford Handbook of Deliberative Democracy. Oxford University Press, Oxford.

Banerjee, A. V., 1992. A simple model of herd behavior. The Quarterly Journal of Economics 107 (3), 797–817.

Berg, S., 1997. Indirect voting systems: Banzhaf numbers, majority functions and collective competence. European Journal of Political Economy 13 (3), 557–573.

Cohen, J., 1986. An epistemic conception of democracy. Ethics 97 (1), 26–38.

Dietrich, F., Spiekermann, K., 2017. Jury theorems. (in press).

Estlund, D., Landemore, H., 2018. The epistemic value of democratic deliberation. In: Bächtiger et al. (2018), Ch. 7, pp. 113–131.

Feddersen, T., Pesendorfer, W., 1996. The swing voter's curse. American Economic Review, 408–424.

Feddersen, T., Pesendorfer, W., 1997. Voting behavior and information aggregation in elections with private information. Econometrica, 1029–1058.

Feddersen, T., Pesendorfer, W., 1999. Elections, information aggregation, and strategic voting. Proceedings of the National Academy of Sciences 96 (19), 10572–10574.

Goodin, R. E., Spiekermann, K., 2018. An epistemic theory of democracy. Oxford University Press.

Landemore, H., Elster, J., 2012. Collective wisdom: Principles and mechanisms. Cambridge University Press.

Lu, L., Yuan, Y. C., McLeod, P. L., 2012. Twenty-five years of hidden profiles in group decision making. Personality and Social Psychology Review 16, 54–75.

Maciejovsky, B., Budescu, D. V., 2019. Too much trust in group decisions: Uncovering hidden profiles by groups and markets. (preprint; presented at the 18th Internationall Conference on Social Dilemmas in Sedona Arizona).

Mayo-Wilson, C., Zollman, K., Danks, D., 2013. Wisdom of crowds versus groupthink: learning in groups and in isolation. International Journal of Game Theory 42 (3), 695–723.

Pivato, M., 2013. Voting rules as statistical estimators. Social Choice and Welfare 40 (2), 581–630.

Pivato, M., 2017. Epistemic democracy with correlated voters. Journal of Mathematical Economics 72, 51–69.

Schwartzberg, M., 2015. Epistemic democracy and its challenges. Annual Review of Political Science 18, 187–203.

Stasser, G., Titus, W., 1985. Pooling of unshared information in group decision making: Biased information sampling during discussion. Journal of Personality and Social Psychology 48 (6), 1467.