Optimizing Individual and Collective Reliability: A Puzzle

Marc-Kevin Daoust, École de technologie supérieure. marc-kevin.daoust@etsmtl.ca

Abstract: Many epistemologists have argued that there is some degree of independence between individual and collective reliability (e.g., Kitcher 1990; Mayo-Wilson, Zollman, and Danks 2011; Dunn 2018). The question, then, is: To *what extent* are the two independent of each other? And in *which contexts* do they come apart? In this paper, I present a new case confirming the independence between individual and collective reliability optimization. I argue that, in voting groups, optimizing individual reliability can conflict with optimizing collective reliability. This can happen even if various conditions are held constant, such as: the evidence jurors have access to, the voting system, the number of jurors, some independence conditions between voters, and so forth. This observation matters in many active debates on, e.g., epistemic dilemmas, the wisdom of crowds, independence theses, epistemic democracy, and the division of epistemic labour.

Keywords: reliabilism, epistemic justification, group reliability, jury theorem, independence thesis

1. Introduction

How can we improve our reliability? Or, how can we optimize our ratio of true to false beliefs?[1]

We can answer these questions from an *individual* as well as from a *collective* point of view. That is, we can improve the reliability of individual agents, like you and I, or we can improve the reliability of group agents, like juries, public institutions, etc. In the individual perspective, research on cognition, reasoning, deference, awareness of biases and blind spots or responsiveness to the evidence helps us to identify ways to optimize reliability.[2] In the collective perspective, research on aggregation of opinions, diversity and complexity helps us identify ways to optimize collective reliability.[3]

---

1   See Goldman and Beddor (2016) for a survey of the debates on reliabilism. See, e.g., Cohen (2002) and Conee and Feldman (1998) on problems for reliabilist theories of justification.

2   See Beebe (2004), Goldman (1986; 1994; 2001), Roberts and West (2015) or Samuelson and Church (2015).

3   See Dietrich and List (2004), Estlund (1994; 1997), Estlund and Landemore (2018), Hedden (2017), Hong and Page (2012), Landemore (2012a; 2012b; 2013), List (2005), List and Goodin (2001), Page (2007; 2010), Pallavicini, Hallsson, and Kappel (2018), Sunstein (2006) or Surowiecki (2005). Jury Theorems are also relevant here. See Condorcet (1976 [1785]), and see Bachrach *et al.* (2012), Dietrich and Spiekermann (2013a, 2013b), Fey (2003), Kaniovski (2010), List and Goodin (2001), Romeijn and Atkinson (2011) or Stone (2015) on various generalizations or applications of Condorcet's Jury Theorem.

Most contemporary philosophers now admit that there is some degree of independence between individual and collective reliability. That is, recent research suggests that there are several contexts in which individual reliability optimization and collective reliability optimization are relatively independent of each other.[4] The question, then, is: To *what extent* are the two independent of each other? And in *which contexts* do they come apart?

In this paper, I present a new case confirming the independence between individual and collective reliability optimization. In some cases, jurors who change their epistemic standards can improve their individual reliability, but decrease group reliability. This can happen even if various conditions are held fixed, such as: the evidence jurors have access to, the voting system, the number of jurors, the degree of (problem-conditional) independence between voters, and so forth. Call this the Puzzle of the Erratic Juror. The Puzzle also has implications in a number of active debates, such as epistemic dilemmas, the Wisdom of Crowds, Independence theses, the relevance of shared epistemic standards in communities, and the relationship between epistemic arguments for democracy and voter competence.

In section 2, I introduce and formalize the Puzzle. In section 3, I discuss a partial solution to the Puzzle, and argue that it does not always succeed. In section 4, I discuss the implications of the Puzzle in a number of active debates.

2. The Puzzle

2.1. An Intuitive Illustration of the Counterexample
The Puzzle of the Erratic Juror is a case in which the optimization of individual reliability and the optimization of group reliability come apart, even if various conditions are held fixed. An intuitive illustration of the puzzle goes as follows. A judge is in charge of forming a jury for several trials. Given the evidence the jurors acquire during the trials, they ought to determine if the defendants are guilty. All jurors are presented with the same evidence and they ought to vote on the basis of the shared evidence only. They are faced with a binary choice (such as 'Guilty' or 'Not Guilty'). They do not discuss with each other before casting their vote.

---

4   See, e.g., Mayo-Wilson, Zollman and Danks (2011) and Dunn (2018) on Independence theses. See also Hong and Page (2012), Landemore (2012a, 2012b), and Kitcher (1990, 1995).

The judge picks two jurors with distinct but equally commonsensical standards or methods of reasoning[5]—call them William and Harry. While they do not have the same methods of reasoning, William and Harry reach the right answer 60% of the time. So, they are fairly reliable. The judge also picks Melania, an 'Erratic Juror' with unorthodox (but often misplaced) standards or methods of reasoning. Melania is less reliable than the other jurors—she reaches the right answer 40% of the time. This is why most people will think that Melania is erratic.

A simple example might help to understand this. Suppose that the following table gives us a 'representative' sample of the votes and reliability levels of jurors:

Table 1. Verdicts

Juror/Case	P_1	P_2	P_3	P_4	P_5
Juror 1 (William):	Not guilty	Guilty	Guilty	Guilty	Not guilty
Juror 2 (Harry):	Not guilty	Guilty	Guilty	Not guilty	Guilty
Erratic Juror (Melania):	Not guilty	Not guilty	Not guilty	Guilty	Guilty
Verdict:	Not guilty	Guilty	Guilty	Guilty	Guilty

Note: The recourse to tables for analyzing votes and reliability levels is less than ideal. They can be misleading. Clearly, the argument should not rest on the analysis of such tables. But at this point, I merely want to give readers an *intuitive* and *accessible* idea of the Puzzle. We'll do better in section 2.3.

Let's assume that the right verdict in each case is 'Guilty.' As we can see, the group reaches the right answer 80% of the time. Also, Table 1 confirms that jurors 1 and 2 reach the right answer 60% of the time, and that the Erratic Juror reaches the right answer 40% of the time.

After a series of verdicts, the Erratic Juror is informed that her odd methods of reasoning are less reliable than the ones entertained by Juror 1 or Juror 2. Even worse, she is informed that her methods of reasoning lead her to the right answer less than 50% of the time. Being reliable matters to her. So, the Erratic Juror is unsatisfied. One day, she discovers new methods of

---

5   See Goldman (2010), Schoenfield (2014), and Titelbaum and Kopec (2019) on standards (or methods) of reasoning.

reasoning that are more reliable. She decides to improve her individual reliability by changing her methods of reasoning.

However, this improvement in individual reliability affects group reliability. Imagine, for instance, that the jurors were to re-vote on P_1 to P_5 (on the same evidence). Then, the votes and the group verdicts would be:

Table 2. Verdicts after the Erratic Juror changed her methods of reasoning

Juror/Case	P_1	P_2	P_3	P_4	P_5
Juror 1:	Not guilty	Guilty	Guilty	Guilty	Not guilty
Juror 2:	Not guilty	Guilty	Guilty	Not guilty	Guilty
Erratic Juror:	Guilty	Not guilty	Guilty	Guilty	Not guilty
Verdict:	Not guilty	Guilty	Guilty	Guilty	Not guilty

As we can see in Table 2, Melania's individual reliability is now up to 60%. However, group reliability is down to 60%. Of course, the Erratic Juror's reliability has improved. However, the group is now less reliable. Apparently, optimizing one's individual standards of reasoning can sometimes reduce group reliability. What can explain this?

2.2. An Initial Formalization of the Puzzle

How can Melania's individual improvement reduce collective reliability? In order to answer this question, we will begin by making a distinction between two categories of problems. First, there are the ordinary problems for which the Erratic Juror is incompetent. We will denote an ordinary problem π by 'O_π'. In Table 1, P_1, P_2 and P_3 are ordinary problems. Second, there are the specialty problems for which the Erratic Juror is very competent. We will denote a specialty problem π by 'S_π'. In Table 1, P_4 and P_5 are specialty problems.

Accordingly, here is an initial formalization of the reliability levels of our jurors, conditional on the type of problem they face:

Erratic Juror (initial). Consider an erratic juror e. Let $Pr_e(P|\pi)=X$ denote the probability that the erratic juror e will reach the right answer on whether P, conditional on the type of problem π

(i.e. conditional on whether P is an ordinary or a specialty problem π). Then, $Pr_e(P|O)=0$ and $Pr_e(P|S)=1$.

Regular Juror (initial). Consider a regular (or 'non-erratic') juror r. Let $Pr_r(P|\pi)=X$ denote the probability that such a regular juror r will reach the right answer on whether P, conditional on the type of problem π. Then $Pr_r(P|O)\approx0.667$ and $Pr_r(P|S)=0.5$.

It should be noted that agents might not be aware that their reliability varies depending on the type of proposition they consider. It might be a fact that Erratic Jurors are less reliable with respect to ordinary propositions, but they might be unaware of such a fact, or they might be unaware of which propositions are ordinary. Relatedly, perhaps jurors *sometimes* know whether P is an ordinary problem without *always* knowing whether P is an ordinary problem.

The jurors are independent in a specific sense: Conditional on the truth/falsity of the verdict and the type of problem (since the reliability levels vary depending on the type of problem, we need to conditionalize on this factor), the jurors vote independently from each other. That is, the jurors are independent *relative to the type of problem*.[6] Again, consider the above example: Suppose that P is true and P is a specialty problem. Then, the probability that a Regular Juror r will vote for P is 0.5, and the probability that an Erratic Juror e will vote for P is 1. These probabilities remain unchanged even if we learn that another juror has voted for P, or has voted for ~P.

With the above specifications in mind, we see more clearly what's going on in the Puzzle. When Melania changed her standards of reasoning, she became more reliable with respect to ordinary propositions (like P_1, P_2, and P_3), but she became less reliable with respect to specialty propositions (like P_4 and P_5). In other words, Melania abandoned her *erratic* standards, and took instead the standards of reasoning of a *Regular Juror*. If she becomes a Regular Juror, she becomes less reliable with respect to specialty propositions. Of course, she gains reliability with respect to ordinary propositions. However, improving individual reliability for ordinary propositions does not necessarily contribute to collective reliability.

---

6   See Dietrich and Spiekermann (2013a, 99; 2013b, 666) on Problem-Conditional Independence, or New Independence.

2.3. A Generalization of the Puzzle

Having made the distinction between ordinary and specialty problems, we can generalize the intuitive case introduced in section 2.1. First, there is no need to assume that the Erratic Juror takes new standards. We could simply assume that the Erratic Juror is replaced with a Regular Juror who entertains different standards (or replaces his or her standards with the ones of a random Regular Juror). Second, there is no need to assume some specific reliability thresholds. In the initial description of the Puzzle, it is assumed that, when a Regular Juror is faced with ordinary propositions, he or she reaches the right answer 66.7% of the time. The Puzzle works even if we do not assume some specific reliability thresholds for jurors. Third, there is no need to assume that there are only three jurors. There could be more jurors than that. Finally, we will not analyze tables. As I said previously, analyzing votes and reliability levels in tables is less than ideal, since they can be misleading. From now on, we focus on the probability that jurors (and juries) will reach the right answer, conditional on the type of problem they face.

In order to provide a generalized version of the Puzzle of the Erratic Juror, we begin by putting a constraint on ordinary and specialty problems, as in the following:

Types of Problems Condition. The probability that jurors will face an ordinary problem (O) is Q, for $0.5<Q<1$. The probability that jurors will face a specialty problem (S) is 1-Q, for $0<(1-Q)<0.5$.

The constraint roughly states that it is more probable that jurors will be faced with ordinary problems than with specialty problems.

Then, we need to generalize the notions of jurors stated above, as in the following:

Erratic Juror. Let $Pr_e(P|\pi)=X$ denote the probability that the erratic juror e will reach the right answer on whether P, conditional on the type of problem π. Then, $Pr_e(P|O)<0.5$, $Pr_e(P|S)=Z$ (for $Z>0.5$), and $Pr_e(P)<0.5$.

Regular Juror. Let $Pr_r(P|\pi)=X$ denote the probability that a regular juror r will reach the right answer on whether P, conditional on the type of problem π. Then $Pr_r(P|O)>0.5$, $Z>Pr_r(P|S)\geq0.5$ and $Pr_r(P)>0.5$.

We can then calculate the probability that a group of jurors under simple majority will reach the right answer. We are familiar with jury theorems for groups of voters that have different competence levels, and jury theorems for groups of voters that face different types of problems.

We just need to combine these two ideas.[7] Assume that the group's reliability for ordinary propositions when there are n Erratic Jurors is denoted by β_n, and that the group's reliability for specialty propositions when there are n Erratic Jurors is denoted by σ_n. If there are m Regular Jurors, β_n and σ_n are equal to:

$$\beta_n = \sum_{j=1}^{n+1} \left[\sum_{i=\frac{m-1-n}{2}+j}^{m} \binom{m}{i} Pr_r(P|O)^i \cdot (1 - Pr_r(P|O))^{(m-i)} \cdot \binom{n}{n+1-j} Pr_e(P|O)^{(n+1-j)} \cdot (1 - Pr_e(P|O))^{(j-1)} \right]$$

$$\sigma_n = \sum_{j=1}^{n+1} \left[\sum_{i=\frac{m-1-n}{2}+j}^{m} \binom{m}{i} Pr_r(P|S)^i \cdot (1 - Pr_r(P|S))^{(m-i)} \cdot \binom{n}{n+1-j} Pr_e(P|S)^{(n+1-j)} \cdot (1 - Pr_e(P|S))^{(j-1)} \right]$$

Naturally, the above sums presuppose that n>0 (otherwise, the binomial coefficient on the right will be undefined). If n=0, our functions will look like:

$$\beta_0 = \sum_{i=\frac{m+1}{2}}^{m} \binom{m}{i} Pr_r(P|O)^i \cdot (1 - Pr_r(P|O))^{(m-i)}$$

$$\sigma_0 = \sum_{i=\frac{m+1}{2}}^{m} \binom{m}{i} Pr_r(P|S)^i \cdot (1 - Pr_r(P|S))^{(m-i)}$$

Given the Types of Problems Condition, the group's global reliability when there are n Erratic Jurors equals $(Q \cdot \beta_n) + ((1-Q) \cdot \sigma_n)$. In view of the foregoing, here is a generalized reformulation of the Puzzle:

Generalized Puzzle of the Erratic Juror. Assume that a group of jurors containing n Erratic Jurors under simple majority rule are faced with a series of binary choices. Assume also that jurors vote independently of each other and that the Types of Problems Condition is satisfied. Turning an Erratic Juror into a Regular Juror (e.g., going from n to $n-1$ Erratic Jurors) improves his or her individual reliability. But if $(Q \cdot \beta_{n-1}) + ((1-Q) \cdot \sigma_{n-1}) < (Q \cdot \beta_n) + ((1-Q) \cdot \sigma_n)$, turning an Erratic Juror into a Regular Juror decreases collective reliability.

In other words, the Puzzle is triggered if and only if $(Q \cdot \beta_{n-1}) + ((1-Q) \cdot \sigma_{n-1}) < (Q \cdot \beta_n) + ((1-Q) \cdot \sigma_n)$.

---

7   See Grofman *et al.* (1983) and Stone (2015) on reliability for groups of voters that have different competence levels. See Dietrich and Spiekermann (2013a, 99; 2013b, 666) on reliability for groups of voters that face different types of problems.

Here is a specific case illustrating this. Suppose that a jury is composed of five jurors. They are under simple majority rule and are faced with a series of binary choices. We are given the following opportunities: Form a jury of five Regular Jurors, or form a mixed jury of three Regular Jurors and two Erratic Jurors. Mathematically, we can represent the reliability levels of such possible juries as follows:

Jury #1: Five Regular Jurors and No Erratic Juror.

$$\beta_0 = \sum_{i=3}^{5} \binom{5}{i} Pr_r(P|O)^i \cdot (1 - Pr_r(P|O))^{(5-i)}$$

$$\sigma_0 = \sum_{i=3}^{5} \binom{5}{i} Pr_r(P|S)^i \cdot (1 - Pr_r(P|S))^{(5-i)}$$

$$Total\, reliability = Q \cdot \beta_0 + (1-Q) \cdot \sigma_0$$

Jury #2: Three Regular Jurors and Two Erratic Jurors.

$$\beta_2 = \sum_{j=1}^{3} \left[\sum_{i=j}^{3} \binom{3}{i} Pr_r(P|O)^i \cdot (1 - Pr_r(P|O))^{(3-i)} \cdot \binom{2}{3-j} Pr_e(P|O)^{(3-j)} \cdot (1 - Pr_e(P|O))^{(j-1)} \right]$$

$$\sigma_2 = \sum_{j=1}^{3} \left[\sum_{i=j}^{3} \binom{3}{i} Pr_r(P|S)^i \cdot (1 - Pr_r(P|S))^{(3-i)} \cdot \binom{2}{3-j} Pr_e(P|S)^{(3-j)} \cdot (1 - Pr_e(P|S))^{(j-1)} \right]$$

$$Total\, reliability = Q \cdot \beta_2 + (1-Q) \cdot \sigma_2$$

Given the above equations, many values of Q, $Pr_r(P|O)$, $Pr_e(P|O)$, $Pr_r(P|S)$ and $Pr_e(P|S)$ will give rise to a conflict between individual and collective reliability. For instance, suppose that, when P is an ordinary problem, Regular Jurors reach the right answer 99% of the time and Erratic Jurors reach the right answer 10% of the time. When P is a specialty problem, Regular Jurors reach the right answer 50% of the time and Erratic Jurors reach the right answer 90% of the time. The probability that jurors will face an ordinary problem is 65% (and so the probability that they will face a specialty problem is 35%).

A jury composed of five Regular Jurors has ≈99.9% chance to reach the right answer on all the ordinary propositions. However, they have 50% chance to reach the right answer on the specialty propositions.[8] Given the proportion of ordinary and specialty propositions, the

---

8   Formally: $\beta_0 \approx 0.999$ and $\sigma_0 = 0.5$. We get these reliability levels by calculating:

$$\beta_0 = \binom{5}{3} 0.99^3 \cdot 0.01^2 + \binom{5}{4} 0.99^4 \cdot 0.01^1 + \binom{5}{5} 0.99^5 \cdot 0.01^0$$

reliability of the group will be ≈82.5%.[9] By way of contrast, consider a jury composed of three Regular Jurors and two Erratic Jurors. With respect to ordinary propositions, there is ≈97% chance that such a jury will reach the right answer. So, even if Erratic Jurors are very bad at reaching the right answer on ordinary propositions, it is highly unlikely that their votes will lead the jury to the wrong answer on such propositions. However, with respect to specialty propositions, a jury composed of three Regular Jurors and two Erratic Jurors should reach the right answer 80% of the time.[10] So, given the proportion of ordinary and specialty propositions, the mixed jury should reach the right answer ≈91% of the time.[11] Since 91%>82.5%, it is optimal to include Erratic Jurors in the jury.[12] This also means that turning an Erratic Juror into a Regular Juror will decrease the group's reliability.

Again, this confirms the conflict we have described so far. When an Erratic Juror like Melania becomes a Regular Juror, she gains reliability with respect to ordinary propositions, but she also becomes less reliable with respect to specialty propositions. However, this might not contribute to group reliability. A change in Melania's standards can improve her individual reliability, but it can also reduce collective reliability.

2.4. A Concrete Application for the Puzzle

One might worry that the Puzzle I describe will not find concrete applications. Is this just a mathematical Puzzle with limited real life applications?

First, a methodological remark: The literature on jury theorems often makes simplifying assumptions. For example, it is assumed that voters do not interact with each other, that they have

$$\sigma_0 = \binom{5}{3} 0.5^3 \cdot 0.5^2 + \binom{5}{4} 0.5^4 \cdot 0.5^1 + \binom{5}{5} 0.5^5 \cdot 0.5^0$$

9   Formally: Given that $\beta_0 \approx 0.999$ and $\sigma_0 = 0.5$, $(0.65 \cdot \beta_0) + (0.35 \cdot \sigma_0) \approx 0.825$.

10   Formally: $\beta_2 \approx 0.97$ and $\sigma_2 = 0.8$. We get these numbers by calculating:

$$\beta_2 = \sum_{j=1}^{3} \left[\sum_{i=j}^{3} \binom{3}{i} 0.99^i \cdot 0.01^{(3-i)} \cdot \binom{2}{3-j} 0.1^{(3-j)} \cdot 0.9^{(j-1)} \right]$$

$$\sigma_2 = \sum_{j=1}^{3} \left[\sum_{i=j}^{3} \binom{3}{i} 0.5^i \cdot 0.5^{(3-i)} \cdot \binom{2}{3-j} 0.9^{(3-j)} \cdot 0.1^{(j-1)} \right]$$

11   Formally: Given that $\beta_2 \approx 0.97$ and $\sigma_2 = 0.8$, $(0.65 \cdot \beta_2) + (0.35 \cdot \sigma_2) \approx 0.91$.

12   Formally: $(0.65 \cdot \beta_2) + (0.35 \cdot \sigma_2) > (0.65 \cdot \beta_0) + (0.35 \cdot \sigma_0)$.

access to the same evidence, and that they must make a choice between only two options. These assumptions simplify the formal results discussed in this literature. However, the reality is often more complex. It is difficult to find situations that *perfectly* fit the conditions stated in these theorems. Nevertheless, we can find situations that are highly similar to the ones stipulated in jury theorems.

Here is one of them. Private foundations and public agencies fund scientific research. Scientists try to get funding from these organizations, by submitting proposals. These proposals then go through various mechanisms to determine whether or not funding is granted. Lamont (2009) has documented how these mechanisms work in the United States, in the fields of humanities and social sciences.

Lamont observes that the main mechanism for evaluating applications is interdisciplinary deliberation among experts (ibid., 116-120). Typically, different experts from different fields discuss and rank the proposals. But there is often a prior mechanism for screening proposals (ibid., 28). This mechanism aims to determine whether a proposal meets certain 'minimum quality thresholds,' a notion defined in broad strokes by the founding agencies (ibid., 38). In other words, screening is meant to 'weed out unpromising proposals' (ibid., 37).

The experts commissioned at the screening stage have the same evidence (the proposals written by scientists). They also rarely meet or talk to each other. Instead, the screeners might be asked to vote independently of each other.[13] For instance, for a given scientific proposal, they can either vote for 'the proposal is promising' or vote for 'the proposal is not promising.'

Screeners do not evaluate proposals in the exact same way. They have different standards of reasoning. They use a variety of criteria, like originality, clarity, diversity, or significance (ibid., chaps. 5-6). Experts do not always mention the same criteria for their decisions. When they do, they do not give the same weight to these different criteria (ibid.). This will lead them to be (un)reliable in different ways. For instance, some screeners give a lot of importance to clarity, and will reject promising proposals that are slightly confusing. Other screeners give little weight to feasibility, and will support some proposals even though those proposals are unpromising in virtue of being wildly unrealistic.

---

13 Screeners can also be asked to give a numerical score, or a rank, to the proposals (Lamont 2009, 28, 37).

In line with the above, one could imagine the following situation. A program officer (i.e., the person in charge of these different mechanisms for funding) could recruit some screeners to evaluate proposals. A proposal passes screening if a majority of votes supports it. The program officer might realize that one of the screeners has unreliable standards of reasoning.[14] The screener's standards lead him or her to make the wrong call on a high proportion of proposals. The officer will then have to decide whether the screener should be replaced from the (current or future) process with another one.

Presumably, program officers will want to exclude unreliable screeners from actual or future processes. However, if we take the lessons of the Puzzle seriously, there can be a good reason to include some unreliable screeners in these processes. A person who is less reliable than others can, on some occasions, be epistemically beneficial to the group. They can improve group reliability. This depends on the type of proposals for which this person is (un)reliable.

In sum, screening mechanisms have features that are similar to the ones described in the Puzzle: (i) Screeners vote independently of each other on various proposals; (ii) they base their decisions on the same evidence (the proposals submitted by scientists); (iii) the screeners have different standards of reasoning; (iv) some of these screeners may be less reliable than others. When these conditions are met, program officers have a decision to make—namely, whether to replace unreliable screeners. Here, the Puzzle's lesson can be helpful. If unreliable screeners are Erratic Jurors (i.e., if they are more reliable than other voters for evaluating certain types of 'specialty' proposals), it may be reliable to include them in the evaluation process.

2.5. Old Wine in New Bottles?

One might object that the Puzzle doesn't bring anything new to the table. First, we already know that groups of inquirers or problem solvers do better when their members think differently, which might include agents that are less reliable than others (Hong and Page 2012; Landemore 2012a, 2012b; Kitcher 1990, 1995). Suppose agents take part in a pub quiz in mathematics. If Melania is only good in algebra while Harry is only good in geometry, it is likely that they will reach the

---

14 In the puzzle, I focus on an extreme case where a juror reaches the correct answer less than 50% of the time. Lamont's research tells us that some screeners were perceived as unreliable, not that they reached the right answer less than 50% of the time (see Lamont 2009, 38-9). However, the extreme case described in the Puzzle is instructive for understanding the more nuanced situations described by Lamont. Even if a screener is profoundly unreliable, we might still have a good reason to include him or her in the group.

right answers in algebra and geometry. But if Melania and Harry are both only good in algebra, it is unlikely that they will reach the right answers in geometry. So, diversity of expertise can help groups improving their collective reliability. Second, we know that a jury composed of jurors who think differently of each other can outperform a jury composed of jurors with the same methods of reasoning (Estlund 1994; Dietrich and List 2004). The satisfaction of some independence conditions among reliable jurors is often an essential condition for improving collective reliability.[15] For instance, a group of independent but moderately reliable jurors can outperform a homogeneous group of highly reliable jurors. In view of the foregoing, one could argue that the Puzzle of the Erratic Juror is old wine in new bottles—we already knew that individually unreliable agents can contribute to group reliability (e.g., by improving diversity within the group).

Yet the Puzzle is different from the above results. There is an important difference between groups of problem solvers and groups of voters under simple majority, namely, the way agents can defer to each other and divide the epistemic labour among themselves. Suppose Melania is good in algebra but not good in geometry. In a group of inquirers or problem solvers (like a pub quiz), she might defer to Harry for questions in geometry. This is a good mechanism for reaching more right answers. But in a group of independent voters under simple majority, Melania might have to vote on propositions in geometry without being able to consult Harry. This difference matters. Showing that unreliable agents can contribute to the reliability of voting groups under simple majority is different than showing they can contribute to the reliability of teams of inquirers or problem solvers.

Also, we know that the satisfaction of independence conditions among jurors is often an essential ingredient of collective reliability in voting groups. So, reducing independence might affect collective reliability. But the Puzzle of the Erratic Juror assumes that there is no opinion leader in the group and that, conditional on the type of problem they face, jurors are independent of each other. So, the Puzzle concerns a different problem.

Thus, the Puzzle is different from familiar conflicts between optimizing individual and collective reliability.

---

15 This observation goes back to Condorcet (1976 [1785]).

3. A Partial Solution to the Puzzle

The possibility of an Erratic Juror poses challenges for meeting certain requirements. For illustration purposes, take the putative requirement of epistemic Immodesty. According to many philosophers, agents fall under an obligation to be epistemically immodest.[16] That is, they should take their own standards and attitudes to be among the most truth-conducive ones available to them. Now, from which point of view should agents be immodest? We can make a distinction between an 'individual' and a 'collective' interpretation of Immodesty. Under the individual interpretation, you should think that your standards and attitudes are among the most truth-conducive ones available to you *personally*. Under the collective interpretation, you should think that your standards and attitudes are among the most truth-conducive ones available to you *qua member of a group*. However, the problem seems to be that Erratic Jurors can't have it both ways. They might know that they have to make a choice between what serves them personally and what serves the group. So, either one of the interpretations of Immodesty is not a genuine requirement, or agents sometimes can't fully comply with the demands of Immodesty.

We could say similar things about other putative requirements. The point is that cases like the Puzzle force us to make uneasy choices between what is beneficial to us personally and what is beneficial to groups. Is there a way for Erratic Jurors to improve individual and collective reliability simultaneously? Perhaps we can figure out a way to solve the Puzzle, so that there is no dilemma between improving individual and collective reliability.

Suppose that Melania, an Erratic Juror, comes to know some regular methods of reasoning as well as her own unorthodox methods of reasoning. And assume furthermore that she is part of a group which would be more reliable if it included an Erratic Juror. Does Melania necessarily face a dilemma between improving individual reliability and improving collective reliability (provided, of course, that she cares about individual and collective reliability)?

According to a potential solution to the Puzzle, Erratic Jurors should combine the regular standards of reasoning they acquire with their own unorthodox standards of reasoning. Specifically, the Erratic Juror should use his unorthodox standards when facing a specialty

---

16 See Lewis (1971). See Daoust (2021) for discussion in jury contexts.

problem, but use the regular standards when facing an ordinary problem. Call this a mixed standard, as in the following:

Mixed Standard Solution. Suppose that A knows the standards employed by a Regular Juror as well as the standards employed by an Erratic Juror. Then, A should employ the standards of a Regular Juror when facing an ordinary problem, and employ the standards of an Erratic Juror when facing a specialty problem.

This way, the agent is guaranteed to be reliable at the individual level without also reducing collective reliability. The standards of a Regular Juror are more reliable for ordinary problems, and the standards of an Erratic Juror are more reliable for specialty problems. From an epistemic point of view, this will necessarily benefit the group as well as the agent.

While this solution is interesting and can solve the problem in *many* contexts, it does not *always* solve the problem. This solution presupposes that agents know the kind of problem they are faced with. If agents do not know when they are facing an ordinary problem (or a specialty problem), the Mixed Standard Solution is unavailable to them. Even if an Erratic Juror's standards are less reliable with respect to ordinary propositions, he or she might be unaware of which propositions are ordinary. At least, one might not be able to neatly classify every problem one faces. So, *on some occasions*, the Mixed Standard Solution will not apply.

Consider the following case. The probability that Melania reaches the right answer on ordinary problems is 20% and the probability that she reaches the right answer on specialty problems is 60%. The probability that she will face an ordinary problem is 70%. She comes to know the standards of a Regular Juror who reaches the right answer 55% of the time on ordinary problems and 50% of the time on specialty problems. Half the time, Melania can tell the type of problem she is facing (and half the time she can't tell). Accordingly, on some occasions, she can satisfy the Mixed Standard. But since she can't always tell what type of problem she is facing, she still has to make a choice between:

MixErr. Take the Mixed Standard half the time and be an Erratic Juror half the time.
MixReg. Take the Mixed Standard half the time and be a Regular Juror half the time.

Her reliability levels for each possibility are:

Table 3. Melania's Choice Between MixReg and MixErr

Standard/Reliability	Ordinary Problems	Specialty Problems	Global Reliability
MixReg	0.55 (0.55·1)	0.55 (0.5·0.5+0.6·0.5)	0.55 (0.55·(0.7+0.3))
MixErr	0.375 (0.2·0.5+0.55·0.5)	0.6 (0.6·1)	0.4425 (0.375·0.7+0.6·0.3)

As we can see in Table 3, we are back to square one. If she picked MixErr, Melania's individual reliability would be suboptimal (and less than 50%). However, when it comes to specialty propositions, MixReg is less reliable than MixErr. This means that, in some cases, taking MixErr will optimize collective reliability. So, Melania can still face a conflict between optimizing individual reliability and optimizing collective reliability.

Hence, the Mixed Standard Solution is fully effective only if an Erratic Juror (i) comes to know some regular standards of reasoning and (ii) can always tell which propositions are ordinary. If such conditions are satisfied, there will be no dilemma between individual and collective reliability. But in some cases, such conditions won't be satisfied. So, the Mixed Standard Solution doesn't provide a full solution to the Puzzle.

So far, I have focused on the type of solutions jurors can implement for solving the conflict between individual and collective reliability. That is, I have focused on what *individual members of a group* can do to solve the problem. But perhaps a better way to solve the puzzle will come from an institutional perspective. For instance, institutions could tell us which standards to use when we are faced with a problem. If we face an ordinary problem, the institutions could invite us to use the standards of a Regular Juror, and if we face a specialty problem, the institutions could invite us to use the standards of an Erratic Juror. However, this putative solution is subject to the same caveats we just discussed. Institutions can do this if they have fine-grained knowledge of the types of problems jurors face. But they might not have this kind of knowledge.

This is, in part, what makes the Puzzle of the Erratic Juror interesting. Institutions and jurors can know that they are in a conflict between optimizing individual and collective reliability. Unlike some epistemic conflicts discussed in print, this one can be explicit from the

agent's point of view.[17] Except if jurors or institutions have fine-grained knowledge of the types of problem they face, there is no easy solution out of the Puzzle. They have to make a choice between their own individual 'epistemic interests' and what serves the group.

4. Discussion: Other Implications of the Puzzle

4.1. The Wisdom of Crowds

One implication of the Puzzle of the Erratic Juror is that adding individually unreliable agents to a group can improve collective reliability. This observation allows us to extend the scope of the 'Wisdom of Crowds' literature. Let me explain.

The Wisdom of Crowds literature was initiated with a study made by Galton (1907). During the West of England Fat Stock and Poultry Exhibition, participants were invited to guess the weight of an ox. Galton analyzed the individual guesses and found out that the median guess was almost perfectly accurate. Yet, the individuals who made these guesses were fairly inaccurate. The correct answer was 1,198 pounds, but the guesses were ranging from 1,074 pounds to 1,293 pounds. So, how can individually unreliable guesses be accurate once they are aggregated? Galton noted that, under some conditions, gross overestimations and underestimations tend to *cancel each other out*. This observation has been generalized in later studies by, e.g., Page (2007, chaps. 7-8) and Davis-Stober *et al.* (2014).

The lesson of these arguments is that having inaccurate individuals in a group is not an obstacle to collective reliability. In fact, they can contribute to collective reliability. I reach the same conclusion with the Puzzle. However, the formal model underlying the Puzzle of the Erratic Juror is different than the kind of model we typically find in the literature on the Wisdom of Crowds. The formal model underlying the Puzzle has different applicability conditions, and this matters.

Take the ox's weight example. The reason why groups tend to be accurate in such cases is that we are looking for a cardinal variable (e.g., the weight). This allows individual overestimations and underestimations to counterbalance each other out. For instance, suppose

---

17 See Hughes (2019) on epistemic conflicts from a third-personal point of view.

that Anna overestimates the ox's weight by 120 pounds, Bob underestimates the ox's weight by 60 pounds, and Carol underestimates the ox's weight by 30 pounds. Then, their average estimation will be off by 10 pounds only.

However, in jury scenarios, this condition is often not satisfied. That is, jurors are often *not* estimating the value of a cardinal variable. Suppose, for instance, that there are 13 jurors under simple majority. They need to figure out who killed Smith. As it happens, Bob is guilty. There are three suspects: Anna, Bob and Carl. 7 jurors overestimate the evidence for Anna's guilt, and vote for 'Anna is guilty.' 6 jurors overestimate the evidence for Carl's guilt, and vote for 'Carl is guilty.' Unfortunately, these overestimations will not result in the right conviction. Anna will be convicted, since a majority of jurors think she is guilty. Thus, the reason why groups of voters under simply majority are wise is not that, in jury settings, overestimations and underestimations counterbalance each other out. Arguments for the collective intelligence of juries track a separate phenomenon.

But this is good news. This means that we have different models, *with different applicability conditions*, confirming that crowds composed of unreliable agents can be wise. For instance, we can generalize this conclusion beyond contexts in which individuals estimate the value of a cardinal variable.

4.2. Independence Theses
The Puzzle of the Erratic Juror provides a novel confirmation of Independence theses, which roughly state that reliable groups can include some unreliable agents (and vice versa).

An influential version of the Independence Thesis has been discussed by Philip Kitcher (1990) in a paper devoted to diversity in science.[18] Kitcher begins by assuming that epistemic communities should aim at getting significant truths and avoiding error. Then, he argues that scientific research using distinct incompatible methods serves this goal, even if some of these methods are less plausible (or truth-conducive) than others. This leads him to conclude that the norms that optimize the advancement of science as a whole can conflict with the individual

---

18 See footnote 4 and Palmira (2018, sec. 6) on the Independence Thesis and doxastic attitudes.

epistemic norms that apply to scientists. Individual scientists who employ unreliable methods can contribute to the global reliability of science.

The Puzzle provides a new perspective on Independence theses. We already know that there is some independence between individual and collective reliability. The interesting question is: How much independence is there between the two, and in which contexts? That is, what we don't know exactly is *how much* independence there is between individual and collective reliability, and *which situations* lend support to Independence theses. As I have explained in section 2.5, an independence result for group of inquirers or problem solvers who deliberate with each other might not have import for groups of independent voters under simple majority. Kitcher has shown that there can be independence between individual and collective reliability in groups of inquirers who want to contribute to the advancement of science. The Puzzle makes a similar point, but in a completely different context—namely, groups of independent voters under simple majority.

Independence theses say that group reliability and individual reliability are independent from each other. However, this doesn't mean that there is absolutely no connection between individual and group reliablility. Some philosophers have argued that, all things being equal (e.g., the number of agents involved, the decision process, the evidence available to agents, etc.), an increase in individual reliability contributes to collective reliability. And this claim is compatible with Independence theses. Goldman (2014), for instance, has endorsed the following principle:

Group Justification. 'If a group belief in P is aggregated based on a profile of member attitudes toward P, then *ceteris paribus* the greater the proportion of members who justifiedly believe P and the smaller the proportion of members who justifiedly reject P, the greater the group's grade of justifiedness in believing P.' (Goldman 2014, 28; Goldman and Beddor 2016, §4.2)

Since Goldman understands justification in terms of reliability, his Group Justification principle can be reformulated as follows:

Group Reliability. If a group belief in P is aggregated based on a profile of member attitudes toward P, then *ceteris paribus* the greater the proportion of members who are reliable in believing P and the smaller the proportion of members who are reliable in rejecting P, the greater the group's reliability in believing P.

At first sight, Goldman's Group Reliability principle makes sense. Most of the time, reliable individuals contribute positively to the reliability of groups. So, it is natural to assume

that, if we want to achieve collective reliability, we should promote individual reliability. However, the Puzzle of the Erratic Juror is in tension with Goldman's thesis. When Melania, an Erratic Juror, improves her individual reliability by taking the standards of a Regular Juror, this leads to lower collective reliability. We can observe this even if we hold several factors fixed, such as the number of jurors, the degree of independence among them, the decision-making mechanism, etc. The only thing Melania needs to change is her standards of reasoning.

Now, perhaps the worry is merely apparent, and some clarifications or adjustments will dissipate it. First, Goldman's Group Reliability Principle has a *ceteris paribus* clause. He says that, *ceteris paribus*, improving individual reliability also improves collective reliability. However, Goldman doesn't tell us how to interpret the *ceteris paribus* clause. A strong interpretation of the *ceteris paribus* clause can save his Group Reliability principle: We could say that, in order to satisfy the *ceteris paribus* clause, we must exclude changes in a juror's standards. This strong interpretation of the *ceteris paribus* clause can save the Group Reliability principle. So, minimally, the Puzzle reveals some of the conditions under which Goldman's thesis holds (or not).

Second, Goldman could specify that individual and group reliability is relative to the specific type of problem jurors face, as in the following:

Group Reliability (Revised). Suppose that P is a problem of type X (e.g., an ordinary or a specialty problem). If a group belief in P is aggregated based on a profile of member attitudes toward P, then *ceteris paribus* the greater the proportion of members who are reliable with respect to problems of type X, the greater the group's reliability in believing P.

This revised formulation also avoids the worry. When Melania, an Erratic Juror, improves her individual reliability by taking the standards of a Regular Juror, this leads to lower collective reliability *for specialty problems*, but not to lower collective reliability *for ordinary problems*. Again, the Puzzle is helpful for clarifying Goldman's Group Reliability principle.

4.3. Shared Reliable Epistemic Standards
The Puzzle also has implications in the debate on the significance of shared reliable epistemic standards. Several philosophers think that, in order to optimize the division of epistemic labour,

epistemic communities should entertain shared epistemic standards. After presenting a recent argument in favour of such a conclusion, I will explain how the Puzzle compromises some of its assumptions.

Dogramaci and Horowitz think that members of an epistemic community like us regularly argue and discuss with each other because it is valuable to evaluate each other's doxastic attitudes (Dogramaci and Horowitz 2016, 132). The epistemic evaluations we make notably include promoting rational beliefs and discouraging (or criticizing) irrational beliefs (ibid., 131). But when we say that one has an irrational belief, what do we criticize, exactly? Judging that a belief is irrational typically means that such a belief was formed in accordance with an unreliable belief-forming process.[19] So, for Dogramaci and Horowitz, when we promote and criticize each other's attitudes, we are in fact evaluating the *epistemic standards* licensing certain beliefs relative to a body of evidence (ibid.).

The question, then, is whether agents should have the *exact same* standards. Dogramaci and Horowitz think there are clear advantages to having the same standards. If members of epistemic communities reason from the same set of reliable standards, they can treat each other as epistemic surrogates, namely as agents with sufficiently similar modes of reasoning. For example, if I reliably conclude that P and you reliably conclude that P implies Q, we could share our respective conclusions and reliably conclude that Q. In view of the goal of getting significant truths, having epistemic surrogates is valuable, since it allows agents to efficiently 'divide the labor of *collecting evidence* and the labor of *reasoning.*' (ibid., 138). That is, epistemic surrogates can provide reliable information to each other through testimony. However, for Dogramaci and Horowitz, when agents entertain distinct standards, we can't take advantage of such a division of epistemic labour. If agents reason from distinct standards, they will not be able to treat each other as epistemic surrogates. Members of epistemic communities would then constantly have to review each other's standards to reach a conclusion. This compromises the goal of getting significant truths through reliable mechanisms. So, for Dogramaci and Horowitz, we should have the same epistemic standards.[20]

---

19 Dogramaci and Horowitz argue that, while there is a strong connection between rational standards and reliable processes, reliability is not a sufficient condition for epistemic rationality (ibid., 135).

20 A similar argument can be found in Greco and Hedden (2016). See Daoust (2017) for other objections.

The Puzzle of the Erratic Juror conflicts with the above line of reasoning. First, the strength of Dogramaci and Horowitz's argument depends on the type of reliability we are concerned with. The above argument can be successful for individual reliability. It could be difficult for individuals to treat each other as epistemic surrogates if they do not share the same standards. However, entertaining distinct incompatible standards of reasoning can optimize an aggregation procedure. The Puzzle of the Erratic Juror confirms this, since an Erratic Juror with distinct standards can improve the group's reliability. So, it is false that the absence of homogeneous standards of reasoning in a group is an obstacle to the goal of acquiring significant truths.

Second, there is no straightforward connection between (i) being reliable and (ii) participating in an efficient division of epistemic labour.[21] Erratic Jurors are unreliable individuals. However, they can participate in an efficient division of epistemic labour. Specifically, their vote on various issues can help groups to get significant truths. So, it might be unwise to criticize unreliable members of an epistemic community, or to invite them to revise their standards. In fact, if we want to optimize collective reliability, it might be preferable to promote (or at least tolerate) their erratic standards of reasoning.

4.4. The Epistemic Toleration of Unreliable Agents
The Puzzle of the Erratic Juror suggests that there exists a rationale for the epistemic toleration of individuals with some individually unreliable methods of reasoning. There is not *necessarily* something wrong with including individually unreliable agents in collective belief-forming mechanisms. In fact, including Erratic Jurors in collective belief-forming mechanisms can optimize collective reliability. So, if we want to optimize collective reliability, some Erratic Jurors should be allowed to participate in these mechanisms.

This allows us to reconsider some common claims in the literature on epistemic democracy. For instance, according to Maskivker, citizens have a duty to vote well (Maskivker 2016, 1). This is so, because not voting well would deprive societies of the epistemic benefits that come out of collective decision mechanisms. More specifically, she says that 'by voting without

---

21 This observation is also confirmed by Kitcher's decision-theoretic argument. See §4.2.

regard to the quality of our vote, we contribute to denying democracy the epistemic properties that come with the aggregation of (good) votes' (ibid. 2).

Now, what does Maskivker mean by 'voting well,' exactly? According to her, voters must attain a certain (probabilistic) threshold of competence.[22] She says:

> [M]y notion of a duty to vote minimally well is subject to an important constraint: The epistemic properties of majority rule are unlikely to come about if individual voters fail to attain a certain threshold of competence, which we can label 'better than random' (i.e., better than a coin flip). (Maskivker 2016, p. 4)

The Puzzle of the Erratic Juror allows us to refine Maskivker's claim. Including agents that are not 'better than random' in collective decision mechanisms can sometimes be epistemically beneficial to the group. Unreliable agents that are good at solving specialty problems, like the Erratic Juror, can help us get the epistemic benefits that come out of collective decision mechanisms. In fact, it can be *optimal* to include them in collective decision mechanisms.

Yet, this comes with a warning. Small changes in the number of Erratic Jurors can greatly affect the group's reliability. Here is an example. Suppose that, when P is an ordinary problem, Regular Jurors reach the right answer 70% of the time and Erratic Jurors reach the right answer 10% of the time. When P is a specialty problem, Regular Jurors reach the right answer 51% of the time and Erratic Jurors reach the right answer 90% of the time. The probability that jurors will face an ordinary problem is 65% (and so the probability that they will face a specialty problem is 35%). If the jury is composed of fifteen voters, here are the reliability levels for all the possible compositions of the jury:

Table 4. Group reliability depending on the ratio of Regular Jurors to Erratic Jurors

Reliability/ Ratio of jurors	15:0	14:1	13:2	12:3	11:4	10:5	9:6	...	0:15
	0.803	0.807	0.797	0.768	0.716	0.642	0.556	...	0.350

Here is how to interpret Table 4. The ratio 15:0 means that there are fifteen Regular Jurors and no Erratic Juror in the group. When there are fifteen Regular Jurors and no Erratic Juror, the group

---

22 Maskivker's 'better than random' condition concerns binary choices (i.e., choices between two options, like 'Guilty' and 'Not guilty'). I have also focused on binary choices throughout this article.

reaches the right answer 80.3% of the time. When there are fourteen Regular Jurors and one Erratic Juror, the group reaches the right answer 80.7% of the time, etc.

As we can see, having an Erratic Juror on the jury is epistemically beneficial to this group. In fact, with respect to the parameters of the case, the inclusion of one Erratic Juror in the group *optimizes* collective reliability. However, if there are more than two Erratic Jurors in the group, collective reliability starts going down very quickly, and can go as low as 35%.

So, it can be epistemically beneficial to include Erratic Jurors in collective belief-forming mechanisms *in small doses*. In other words, the Puzzle of the Erratic Juror provides a rationale for the epistemic toleration of some individually unreliable agents, but this comes with a warning: Having too many Erratic Jurors can lead groups to a bad ratio of true to false beliefs.

References

Bachrach, Yoram, Graepel, Thore, Kasneci, Gjergji, Kosinski, Michal and urgen Van Gael. 2012. "Crowd IQ: Aggregating Opinions to Boost Performance." In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, 535–542.

Beebe, J. R. 2004. "Reliabilism, Truetemp And New Perceptual Faculties." *Synthese* 140 (3): 307–29.

Cohen, Stewart. 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65 (2): 309–29.

Condorcet, Nicolas de. 1976. "Essay on the Application of Mathematics to the Theory of Decision-Making." In *Condorcet: Selected Writings*. Indianapolis: Bobbs-Merrill.

Conee, Earl, and Richard Feldman. 1998. "The Generality Problem for Reliabilism." *Philosophical Studies* 89 (1): 1–29.

Daoust, Marc-Kevin. 2017. "Epistemic Uniqueness and the Practical Relevance of Epistemic Practices." *Philosophia* 45 (4): 1721-1733.

Daoust, Marc-Kevin. 2021. "Should agents be immodest?" *Analytic Philosophy* 62 (3): 235-251.

Davis-Stober, Clintin P., Budescu, David V., Dana, Jason, & Broomell, Stephen B. 2014. "When is a Crowd Wise?." *Decision* 1 (2): 79-101.

Dietrich, Franz, and Christian List. 2004. "A Model of Jury Decisions Where All Jurors Have the Same Evidence." *Synthese* 142 (2): 175–202.

Dietrich, Franz, and Kai Spiekermann. 2013a. "Epistemic Democracy With Defensible Premises." *Economics & Philosophy* 29 (1): 87–120.

Dietrich, Franz, and Kai Spiekermann. 2013b. "Independent opinions? On the Causal Foundations of Belief Formation and Jury Theorems." *Mind, 122* (487): 655-685.

Dogramaci, Sinan, and Sophie Horowitz. 2016. "An Argument for Uniqueness About Evidential Support." *Philosophical Issues* 26 (1): 130–47.

Dunn, Jeffrey. 2018. "Epistemic Free Riding." In *Epistemic Consequentialism*, edited by Jeffrey Dunn and Kristoffer Ahlstrom-Vij, 310–32. Oxford: Oxford University Press.

Estlund, David. 1994. "Opinion Leaders, Independence, and Condorcet's Jury Theorem." *Theory and Decision* 36 (2): 131–162.

Estlund, David. 1997. "Beyond Fairness and Deliberation: The Epistemic Dimension of Democratic Authority." In *Deliberative Democracy: Essays on Reason and Politics*, edited by James Bohman and William Rheg, 173–204. Cambridge: MIT Press.

Estlund, David, and Hélène Landemore. 2018. "The Epistemic Value of Democratic Deliberation." In *The Oxford Handbook of Deliberative Democracy*, edited by André Bächtiger, John S. Dryzek, Jane Mansbridge, and Mark E. Warren, 113–31. Oxford: Oxford University Press.

Fey, Mark. 2003. "A Note on the Condorcet Jury Theorem with Supermajority Voting Rules." *Social Choice and Welfare* 20 (1): 27–32.

Galton, Francis. 1907. "Vox Populi (the Wisdom of Crowds)." *Nature* 75 (7): 450-451.

Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.

Goldman, Alvin. 1994. "Naturalistic Epistemology and Reliabilism." *Midwest Studies in Philosophy* 19 (1): 301–320.

Goldman, Alvin. 2001. "Experts: Which Ones Should You Trust?" *Philosophy and Phenomenological Research* 63 (1): 85–110.

Goldman, Alvin. 2010. "Epistemic Relativism and Reasonable Disagreement." In *Disagreement*, edited by Ted Warfield and Richard Feldman, 187–213. Oxford: Oxford University Press.

Goldman, Alvin. 2014. "Social Process Reliabilism: Solving Justification Problems in Collective Epistemology." In *Essays in Collective Epistemology*, edited by Jennifer Lackey, 11–41. Oxford: Oxford University Press.

Goldman, Alvin, and Bob Beddor. 2016. "Reliabilist Epistemology." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics

Research Lab, Stanford University. https://plato.stanford.edu/archives/win2016/entries/reliabilism/.

Greco, Daniel, and Brian Hedden. 2016. "Uniqueness and Metaepistemology." *The Journal of Philosophy* 113 (8): 365–95.

Grofman, Bernard, Guillermo Owen, and Scott L. Feld. 1983. "Thirteen Theorems in Search of the Truth." *Theory and Decision* 15 (3): 261–278.

Hedden, Brian R. 2017. "Should Juries Deliberate?" *Social Epistemology* 31 (4): 368–86.

Hong, Lu, and Scott Page. 2012. "The Micro-Foundations of Collective Wisdom." In *Collective Wisdom: Principles and Mechanisms*, edited by Jon Elster and Hélène Landemore, 56–71. Cambridge: Cambridge University Press.

Hughes, Nick. 2019. "Dilemmic Epistemology." *Synthese*, *196* (10): 4059-4090.

Kaniovski, Serguei. 2010. "Aggregation of Correlated Votes and Condorcet's Jury Theorem." *Theory and Decision* 69 (3): 453–468.

Kitcher, Philip. 1990. "The Division of Cognitive Labor." *The Journal of Philosophy* 87 (1): 5–22.

Kitcher, Philip. 1995. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.

Lamont, Michèle. 2009. *How Professors Think*. Cambridge (MA): Harvard University Press.

Landemore, Hélène. 2012a. "Democratic Reason: The Mechanisms of Collective Intelligence in Politics." In *Collective Wisdom: Principles and Mechanisms*, edited by Hélène Landemore and Jon Elster, 251–89. Cambridge: Cambridge University Press.

Landemore, Hélène. 2012b. "Why the Many Are Smarter than the Few and Why It Matters." *Journal of Public Deliberation* 8 (1): 1–12.

Landemore, Hélène. 2013. "Deliberation, Cognitive Diversity, and Democratic Inclusiveness: An Epistemic Argument for the Random Selection of Representatives." *Synthese* 190 (7): 1209–31.

Lewis, David. 1971. "Immodest Inductive Methods." *Philosophy of Science 38* (1): 54-63.

List, Christian. 2005. "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective." *Episteme* 2 (1): 25–38.

List, Christian, and Robert E. Goodin. 2001. "Epistemic Democracy: Generalizing the Condorcet Jury Theorem." *Journal of Political Philosophy* 9 (3): 277–306.

Maskivker, Julia. 2016. "An Epistemic Justification for the Obligation to Vote." *Critical Review* 28 (2): 224-247.

Mayo-Wilson, Conor, Kevin J. S. Zollman, and David Danks. 2011. "The Independence Thesis: When Individual and Social Epistemology Diverge*." *Philosophy of Science* 78 (4): 653–77.

Page, Scott E. 2007. *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton (NJ): Princeton University Press.

Page, Scott E. 2010. *Diversity and Complexity*. Princeton (NJ): Princeton University Press.

Pallavicini, Josefine, Bjørn Hallsson, and Klemens Kappel. 2018. "Polarization in Groups of Bayesian Agents." *Synthese*, 1–55. https://doi.org/10.1007/s11229-018-01978-w

Palmira, Michele. 2018. "Inquiry and the Doxastic Attitudes." *Synthese*, 1–27. https://doi.org/10.1007/s11229-018-01955-3.

Roberts, Robert C., and Ryan West. 2015. "Natural Epistemic Defects and Corrective Virtues." *Synthese* 192 (8): 2557–76.

Romeijn, Jan-Willem, and David Atkinson. 2011. "Learning Juror Competence: A Generalized Condorcet Jury Theorem." *Politics, Philosophy & Economics* 10 (3): 237–262.

Samuelson, Peter L., and Ian M. Church. 2015. "When Cognition Turns Vicious: Heuristics and Biases in Light of Virtue Epistemology." *Philosophical Psychology* 28 (8): 1095–1113.

Schoenfield, Miriam. 2014. "Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Influences on Belief." *Noûs* 48 (2): 193–218.

Stone, Peter. 2015. "Introducing Difference into the Condorcet Jury Theorem." *Theory and Decision* 78 (3): 399–409.

Sunstein, Cass R. 2006. *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press.

Surowiecki, James. 2005. *The Wisdom of Crowds*. New York: Anchor.

Titelbaum, Michael G., and Matthew Kopec. 2019. "When Rational Reasoners Reason Differently": 205-31. In *Reasoning: Essays on Theoretical and Practical Thinking*, edited by M. Balcerak Jackson and B. Balcerak Jackson. Oxford: Oxford University Press.