

# **Epistemic Perspectives on Democratic Participation, Freedom and Empowerment**

Shira Ahissar

---

---

A thesis submitted to the department of Philosophy, Logic and Scientific Method in the London School of Economics

## Abstract

The thesis aims to tackle two fundamental questions about epistemic aspects of democratic theory. First, which epistemic goals should democratic decision-making have and can they be pursued in a way that allows for decision-making that is both inclusive and competent? Second, how can these epistemic goals be pursued in a democracy in a manner that respects and promotes voters' epistemic autonomy? In the first two chapters, I discuss the types of truth democracies can track and explain for which of these types more inclusive forms of decision-making will outperform less inclusive ones. In chapter one, I compare direct and representative voting on this basis, and in chapter two, deliberative mini-publics and their alternatives. Chapters three and four aim to tackle the question of which changes to voters' informational environments can improve their epistemic autonomy and empower them. In chapter three, I coin the term "freedom of information choice", understood as the ability to form evaluative judgments autonomously, and explain which sets of information options allow for it. In chapter four, I provide an interpretation of the epistemic empowerment of voters and explain which types of interventions will promote it. Together, these chapters provide a framework for the type of interventions in voters' epistemic environments that could improve both inclusivity and epistemic autonomy while maintaining quality decision-making.

## **Declaration**

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly cited others. The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

The work was conducted under the supervision of Prof. Alex Voorhoeve and Prof. Kai Spiekermann.

The final word count of this thesis is: 50,223 words

## Acknowledgements

First and foremost, I would like to thank my two dissertation supervisors, Alex Voorhoeve and Kai Spiekermann, for their insightful, dedicated, and supportive guidance, which was invaluable to both my dissertation and my wellbeing.

I would like to thank the department faculty, and particularly Richard Bradley, Liam Kofi Bright, Jingyi Wu, Anna Mahtani, Remco Heesen, Lewis Ross, Jonathan Birch, and Jonathan Parry, for their constructive feedback and helpful advice on parts of the dissertation, additional projects and general directions.

I would like to thank my friends and colleagues for providing great companionship and stimulating conversations, and especially Belle Wollesen, Lukas Beck, Shai Agmon, Somayeh Tohidi, Suzanne Bloks, Zhongwei Xu, and Nicolas Côté for providing helpful feedback on my work.

I would like to thank participants in various workshops and conferences, including, but not limited to, BSPS's annual conference, the annual Oxford Graduate Political Theory Conference, the Formal Democratic Theory Workshop, the Annual Conference of the Society for Applied Philosophy, and the LSE Choice Group, for their helpful inputs on my work in various stages.

Finally, I would like to thank my husband, son, and soon to be born baby for taking this journey to a foreign land with me, being my home, and enriching my existence in numerous ways, including my perspective on the world.

---

## Contents

---

|   |           |
|---|-----------|
| <b>Introduction.....</b>  | <b>8</b>  |
| <b>1. Direct or Representative Voting? It Depends on the Type of Truth We Track .....</b> | <b>15</b> |
| 1.1 Introduction .....  | 15        |
| 1.2 Arguments for the superiority of representative democracy .....                       | 16        |
| 1.3 The jury theorem challenge to RD superiority.....                                     | 17        |
| 1.4 Distinguishing two types of truth .....   | 19        |
| 1.5 The universal truth scenario .....  | 21        |
| 1.5.1 The case of a single individual.....  | 21        |
| 1.5.2 Multiple voters .....   | 25        |
| 1.6 Relative truth.....   | 31        |
| 1.6.1 The case of a single individual.....  | 32        |
| 1.6.2 Multiple voters .....   | 37        |
| Conclusions .....   | 40        |
| Appendix: when multiple constituencies track relative truths .....                        | 43        |
| <b>2. The Limitations of Deliberative Mini Publics .....</b>                              | <b>46</b> |
| 2.1 Introduction .....  | 46        |
| 2.2 What are DMPs? .....  | 47        |
| 2.3 Existing criticisms of DMPs .....   | 48        |
| 2.4 The failure of a promising route to explain the advantage of DMPs .....               | 50        |
| 2.4.1 Justifying premise 1 .....  | 51        |
| 2.4.2 Justifying premise 2 .....  | 53        |
| 2.4.3 Justifying premise 3 .....  | 59        |
| 2.4.4 Justifying premise 4 .....  | 62        |
| 2.4.5 Justifying premise 5 .....  | 69        |

|   |           |
|---|-----------|
| Conclusions .....   | 74        |
| <b>3. Freedom of Information Choice .....</b>   | <b>76</b> |
| 3.1 Introduction .....  | 76        |
| 3.2 Defining freedom of information choice.....   | 78        |
| 3.3 What is distinct about freedom of information choice? .....                             | 81        |
| 3.3.1 Not primarily relying on pre-existing evaluative judgments .....                      | 81        |
| 3.3.2 A “cognitive budget set” .....  | 84        |
| 3.3.3 Using these distinctions to explain our intuitions.....                               | 86        |
| 3.4 What criteria of freedom of choice get wrong about freedom of information choice ...    | 87        |
| 3.4.1 The preference criterion.....   | 88        |
| 3.4.2 The reasonableness criterion.....   | 88        |
| 3.4.3 The size and volume criteria .....  | 90        |
| 3.4.4 The accessibility level criterion .....   | 90        |
| 3.5 A new criterion for improving freedom of information choice .....                       | 91        |
| 3.6 A distinct property of the cognitive budget set and improving intra-bundle diversity .. | 95        |
| Conclusion.....   | 97        |
| <b>4. Epistemic Empowerment of Voters.....</b>  | <b>99</b> |
| 4.1 Introduction .....  | 99        |
| 4.2 Voter ignorance .....   | 100       |
| 4.3 Empowerment .....   | 103       |
| 4.4 The epistemically empowered voter .....   | 105       |
| 4.4.1 What limits available capabilities .....  | 106       |
| 4.4.2 The relevant sub-capabilities .....   | 109       |
| 4.4.3 Justifying the conditional subjectivist approach .....                                | 109       |
| 4.4.4 The minimal set of objective capabilities .....                                       | 113       |
| 4.4.5 Information people find relevant .....  | 120       |

|   |            |
|---|------------|
| 4.6 Revisiting the voter ignorance literature ..... | 122        |
| 4.7 Practical interventions .....                   | 125        |
| Conclusion.....                                     | 126        |
| <b>References.....</b>                              | <b>127</b> |

---

## Introduction

---

When I started writing this dissertation, my choice of topic was guided by puzzles I found particularly interesting and timely regarding the role information plays and should play in current democracies.

First, democracy is generally thought to mean rule by the people. Yet, there is a longstanding tradition within democratic theory of refraining from giving the electorate too much power, due to the worry that uninformed voters will engage in bad decisions. Arguments in this vein include the thought that there should not be direct voting due to voter ignorance, which features in Rousseau (1762) and that better educated voters should have more votes, dating back at least to Mill (1865, Chapter VIII). It is also suggested that voters' role should be limited to deciding on societal values, while policy makers or, for instance, prediction markets, should be assigned more complex tasks needed to promote those values (e.g., Christiano, 1996, 2018; Hanson, 2013). Even the founding fathers of the US preferred to give decision-making power to the ruling elite. For instance, Hamilton's comments from the Federalist No. 68 (1788/1961: 458) include the note:

“It was equally desirable that the immediate election should be made by men most capable of analyzing the qualities adapted to the station, and acting under circumstances favourable to deliberation. A small number of persons, selected by their fellows from the general mass, will be most likely to possess the information and discernment requisite to so complicated an investigation.”

Two questions, then, arise: (a) Is the worry about voters' limited ability to engage in good decision-making warranted? (b) Can democracies, which set out to allow rule by the people, improve their ability to foster both inclusivity and good decision-making?

Another puzzle that intrigued me is the change exhibited within democracies in recent years, with the development of the internet and social media. Theoretically, voters now have much more information than they ever had before easily accessible to them online. Such a change might be considered an improvement in democratic standards, as people with better access to multiple sources of information can more easily be self-governing and might be less prone to

## *Introduction*

state manipulation. We might have thought that such a change would bring about a rise in trust between voters and governments in democratic systems. Voters, perhaps, would place more trust in governments that allow such freedom of information, and governments would place trust in voters that have better means to become well-informed. However, voters' trust in democratic institutions has been decreasing in recent years (The Electoral Commission, 2025; University of Southampton, 2025). Additionally, voters are often thought to be biased and misinformed, rather than well informed (e.g., Brennan, 2022, Guerrero, 2024, Ecker et al., 2024). These changes are part of a larger trend, where, worldwide, democracies seem to be in trouble, and clear trends of democratic backsliding have been identified (Nord et al., 2024).

Even though access to information has improved, two factors seem to have prevented improved relations between state and voters. First, the information voters have access to is not always quality information, and it includes misinformation and biased information. Furthermore, it is often controlled by private companies, who take advantage of human weaknesses in a way that creates non-ideal information consumption habits. Additionally, having multiple informational sources has allowed voters to view democratic institutions with a more critical eye and they are therefore better able to identify corruption, biases, lack of neutrality, and the unequal distribution of power within these institutions (Guriev et al., 2021). Consequently, they place less trust in statements, and even factual information, provided by state institutions, especially when a rival political party controls those institutions.

This gives rise to a new puzzle. On the one hand, it seems the state should intervene in voters' epistemic environments, to improve the quality of information consumption. On the other hand, voters are less likely to accept such interventions, particularly if they are initiated by an opposing party and appear non-neutral with respect to values in dispute among the population. Such criticisms were made in the past regarding various proposals for regulation or content moderation of social media (see, for example, Nover, 2025; Vile, 2025), for instance. The question then arises: can there be a sweet spot of state intervention that will aid the relationship between voters and the state, while preserving self-governance?

These puzzles have led me to want to address several important questions about the role of information curation in modern democracies. (a) What type of information do voters need to possess for governments to place a justified trust in them and increase inclusivity in decision-making? (b) Is it possible for states to intervene in voters' epistemic environments in a way that

## *Introduction*

improves, rather than reduces, epistemic autonomy and self-governance, while providing voters with the relevant type of information for inclusive decision making to be feasible?

Thinking particularly about changes that technologies such as recommender systems and large language models have brought to voters' epistemic environments, I also wondered what changes would be beneficial, accepted by voters, and autonomy preserving in a reality where people do not consume the same information as one another, and do not always agree even on basic facts. Additionally, I wondered whether it is possible to capitalize on recent changes in voters' access to information and technological developments in a way that strengthens, rather than weakens, democracies.

I believe that to answer these questions, we must clarify several important underlying concepts. For instance, the information we should expect voters to have to be included in decision-making hinges on the question of what truth we believe democracies should aim to track. The question whether specific mechanisms could increase inclusivity while preserving quality decision-making also hinges on the epistemic benefits such mechanisms could have in tracking specific truths and whether those stand in tension with the particular conception of inclusivity they foster. Answering questions about interventions in voters' epistemic environments that would cause a positive change in democracies also requires understanding human cognition. This is because, while theoretically accessible, information portrayed to people is largely being curated by private companies, often in ways that appeal to certain cognitive biases. An updated theory of epistemic autonomy of voters in modern democracies should take those into account. Finally, the literature has been preoccupied with questions of voter ignorance and competence, whereas a more democratic approach that could guide legitimate interventions would entail studying voters' epistemic empowerment.

These underlying questions are, in this order, what each of my chapters aims to address. Overall, my dissertation explores the role of information in democracies, from the type and level of information voters require for inclusive decision making that leads to good decisions, to the types of interventions states can conduct in voters' epistemic environments. The dissertation comprises four chapters.

The first chapter shows that under plausible conditions, the tension between more inclusive and better decision-making is limited to a specific understanding of the type of truth democratic voting is meant to track. In it, I compare the truth-tracking abilities of direct and indirect voting. I examine an overlooked aspect of this question, showing that the nature of truth we track,

whether it is universal or relative, significantly influences the answer to this question. I use the term “universal truths” to describe cases where the correct vote is the same for all voters, such as when voting on the guilt or innocence of a person. I use the term “relative truth” to describe cases where the correct vote differs for different voters, for instance, since each of them is voting for what is correct given their interests, values or goals.

Using plausible conditions that do not rely on jury theorem assumptions, I compare direct and representative voting, revealing different outcomes for each type of truth. When we track a universal truth, for a relatively incompetent electorate, representative democracy reaches higher competence. Yet, as voter competence improves in both electing representatives and voting on issues, direct democracy will eventually surpass it. By contrast, when tracking relative truths, direct democracy tends to lead to higher competence even when individual voters’ competence is only slightly above random. This chapter offers a new challenge to the idea that voters’ power should be limited due to voter ignorance. It shows that if, as many have supposed (e.g., Schumpeter, 1942, Downs 1957), the aim of democracy is to track relative truths, then voters do not need to be very competent for their inclusion to improve decision making.

The style of this first chapter is the most formal out of my four chapters, however, the ideas in it are also explained intuitively and informally.

In the second chapter, I examine the potential of deliberative mini publics (DMPs), such as citizen assemblies, to lead to inclusive and epistemically advantageous decision-making within democracies. I chose to examine DMPs in particular because of their rising popularity in recent years and their extensive usage within real democracies. These bodies have been recommended by leading world organizations (e.g., WHO, 2021) and are often thought to raise citizens’ trust in existing democracies (Norheim, et al., 2021). I explain why a potentially promising way of explaining their advantage likely fails. In particular, I show that arguments aiming to demonstrate that DMPs can outperform alternatives in tracking a truth relative to the interests or values of the majority of the population fail. I lay out two potential ways of tracking such a truth – (i) where each DMP participant aims to promote their own interests and values, and (ii) where each participant aims to promote the interests and values of the general population. I show that neither way allows us to attribute clear epistemic advantages to DMPs over alternatives. I conclude that other arguments are needed for justifying the use of DMPs in existing democracies.

Within this chapter, I also explain how various means of providing individuals with information could be biased, potentially harming their ability to pursue their own interests, values and goals. I conclude that when aiming to inform voters, it is crucial to take into account means of preserving their epistemic autonomy. This requires the development of an independent theory of what type of information portrayal preserves voters' epistemic autonomy, a question I delve into in the following chapters.

In the third and fourth chapters, I examine the type of interventions governments can have in voters' epistemic environments that would preserve and cultivate epistemic autonomy rather than curtailing it. My examination emphasizes the role cognitive limitations play in current harms to voters' epistemic environments.

In the third chapter, I explore an overlooked aspect of freedom of choice, which I term "freedom of information choice". This is the ability to form evaluative judgments autonomously, which hinges on the information options ones has. I elucidate two distinct characteristics of freedom of information choice. First, it cannot be measured primarily based on pre-existing evaluative judgments. Additionally, instead of the monetary budget set commonly used for economic freedom, in its evaluation, we should primarily consider what I call a "cognitive budget set". Considering freedom of information choice in the light of this cognitive budget set, I argue, enables us to account for harms to such freedom caused by the combination of certain types of information curation with our limited cognitive resources. I clarify why these distinctions render prominent suggestions for enhancing freedom of choice problematic for measuring freedom of information choice. I then propose one way of enhancing freedom of information choice, by increasing what I term the "intra-bundle diversity of options". This involves enhancing the diversity of information within *jointly possible* combinations of information (i.e., within the same bundle) instead of enhancing diversity between mutually exclusive options. What makes information options jointly accessible in my account is that they can be jointly consumed within the limits of a person's cognitive budget. This measure can help us conceptualize, and measure as a threat to freedom, typical problems with online information consumption. For instance, when one is exposed to one-sided information, the cognitive cost of consuming opposing views becomes higher. Hence, often, both cannot be jointly consumed within the cognitive budget. I illustrate how intra-bundle diversity can be achieved by highlighting a distinct quality of the cognitive budget set, namely, that prior information consumption affects the cognitive costs of additional items. Among other

things, this chapter offers a normative explanation of the problem with the current structure of communication in digital spaces.

In the fourth chapter, I offer a particular interpretation of epistemic empowerment of voters. The term “voter empowerment” is often used in public discourse. However, insufficient philosophical thought has been devoted to what such empowerment entails. In the chapter, I develop an understanding of empowerment with respect to the use of information in the decision-making process preceding a vote. Drawing on literature on empowerment in development studies, I argue that epistemic voter empowerment should be measured by particular capabilities. While Sen’s capability approach (e.g., Sen, 1980) is often used to examine what options should reside within a monetary budget set, I argue that for the epistemic empowerment of voters, we need to develop a list of capabilities to determine what information should reside within the cognitive budget set introduced in the previous chapter.

I argue that these capabilities should include a voter’s control over their epistemic process and their ability to use their vote to promote their interests. I develop what I call a “conditional subjectivist” approach to the relevant capabilities. I argue that all voters should possess certain capabilities essential for epistemic autonomy, and once they do so, their own views should determine which information is considered relevant to promoting their interests and well-being. I then explain the advantages this approach has over prominent alternative proposals that are directed at improving voters’ knowledge, including composing tests of citizen knowledge and encouraging citizens to pass them in various ways (Brennan, 2016, 2022). Particularly, I argue that the proposed form of empowerment better protects voters’ epistemic autonomy.

At the end of this fourth chapter, I briefly discuss two important topics that I hope to further develop in the future. The first includes reasons we should expect epistemic empowerment of voters to raise voters’ competence in tracking relative truths. I explain why the rise in competence will be on metrics similar to the ones used to assess the existing proposals for raising the electorate’s competence, such as knowledge tests (Brennan, 2016, 2022), as well as the metric upon which deliberative mini-publics might have epistemic advantages. As empowering the entire population using this method is, I argue, more inclusive and autonomy-preserving than these other methods, it may have an advantage over such methods.

In the final section of this chapter, I briefly discuss how the proposed type of empowerment is timely, in the sense that it can capitalize on existing capabilities of recommender systems and other recent technologies to empower voters. This is because the suggested type of

## *Introduction*

empowerment requires interventions that rely on information that recommender systems currently collect and use to personalize people's epistemic environments, but according to standards that differ from the ones currently employed in such systems. My usage of the capabilities approach allows me to explain the advantage of existing capabilities of recommender systems. It is well known within the capabilities literature that cultivating the same capabilities across different people requires different means, including, in this case, different types of information. In this light, the personalization capacities of such technologies could be used to curate information that fosters the development of these capabilities across diverse citizens.

Even though I do not explicitly address the topic within the thesis, I have written my dissertation while thinking about interventions that could potentially reduce political polarization among voters in current democracies. I believe, in particular, that the proposals in the final two chapters, concerning freedom of information and the epistemic empowerment of voters, might help reduce polarization without resorting to problematic interventions. They could do so, I conjecture, by promoting tolerance for opposing views (allowing those to fall within the cognitive budget set) and giving voters greater control over their informational environment than they currently enjoy online. Of course, developing this idea requires further work.

Although each chapter discusses a different topic, the dissertation completes, in a sense, a full circle, laying out some foundations for what I view as the best way forward for democracies in the present context. Establishing a complete way forward for democracies requires, of course, covering many more related issues. However, I hope these four chapters allow the reader to see the beginning of a larger view of the role of information curation by the state in current democracies, which can help us identify interventions that are inclusive, autonomy-preserving, and likely to lead to better decision-making. The aim of the voter empowerment proposal outlined in the final chapter is to offer a way for democracies to look ahead, embracing rather than resisting current technologies, and engaging with rather than suppressing the growing diversity of opinions and informational habits among voters. Instead of trying to put the genie back in the bottle, it aims to use the positive aspects of these developments to strengthen the relations between voters and governments in contemporary democracies.

---

# 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

---

## Abstract

The literature exploring the epistemic aspects of democracy aims to answer a fundamental question: which form of government is most effective in making decisions that track the truth? In this chapter, I examine an overlooked aspect of this question, showing that the nature of truth we track, whether it's universal or relative, significantly influences the answer to this question. Using plausible conditions that do not rely on jury theorems, I compare direct and representative voting, revealing different outcomes for each type of truth. When we track a universal truth, for a relatively incompetent electorate, representative democracy reaches higher competence. Yet, as voter competence improves in both types of voting, direct democracy will eventually surpass it. In contrast, when tracking relative truths, where the correct vote differs for different voters, direct democracy tends to lead to higher competence even when average voter competence is only slightly above random. Yet, in specific cases, representative voting will be superior for all voter competence levels. Hence, both the epistemically better option and the option with the higher epistemic potential if voter competence improves differ under universal and relative truths.

## 1.1 Introduction

Whether democracy represents the optimal form of government and, if so, which type of democracy is superior, is a longstanding subject of debate (Harrington, 1747; Rousseau, 1762; Madison, 1787; Manin, 1997; Urbinati; 2000; Tormey, 2015; Leemann & Wasserfallen, 2016). In recent years, there has been extensive literature on the epistemic evaluation of democracy, aiming to justify (e.g., List & Goodin, 2001; Estlund, 2008, Goodin & Spiekermann, 2018) or reject (e.g., Brennan, 2016, 2022) democracy based on its capacity for good decision-making.

So far, this literature has neglected to pay much attention to the influence the type of truth we track has over the epistemic comparison of direct and representative democracies. In this chapter, I distinguish between two types of truth, which I term “universal truth” and “relative

### *1. Direct or Representative Voting? It Depends on the Type of Truth We Track*

truth”. This distinction significantly influences the epistemic superiority of each form of democracy.

I identify a truth-tracking threshold (a “crossover point”) after which any higher truth-tracking ability requires more competence of voters in representative democracy. I show that for universal truths, when the voters possess only minimal competence, representative voting often outperforms direct voting. Nevertheless, there always exists a level of voter competence after which direct voting becomes the more effective choice, a crossover point. In contrast, for relative truths, the expected mean of this crossover point is chance level. This indicates that voters require only minimal competence for direct voting to yield superior results. However, there are situations where no crossover point lower than 1 exists, meaning that, regardless of the level of voter competence, representative voting consistently leads to better outcomes.

These results undermine common assumptions of the literature on epistemic aspects of democracy. As I will show, prominent theories take political voting to track relative truths. In such contexts, the results provide epistemic reasons for democratic decision-making to often be done via direct voting. Additionally, the results highlight the dynamics that affect epistemic potential, providing guidelines for when to use each type of voting and undermining the assumption used in many jury theorems, that the context, which determines the type of truth a group aims to track, can be ignored.

I proceed as follows: In section 1.2, I survey literature arguing for the epistemic superiority of representative democracy. In section 1.3, I survey the literature on jury theorems. None of these accounts analyse how the comparison between direct and indirect voting differs given the type of truth we track. Additionally, jury theorem arguments for the superiority of direct voting make unrealistic assumptions. In section 1.4, I distinguish between the two types of truth. In section 1.5, I discuss the location of the crossover point for universal truth, and in section 1.6, I do so for relative truths. In section 1.7, I discuss conclusions regarding which form of government is epistemically more desirable.

## **1.2 Arguments for the superiority of representative democracy**

Two primary forms of democracy are direct democracy (DD) and representative democracy (RD). In DD, individuals vote directly on issues. In RD, they conduct indirect voting, where voters elect representatives to vote on their behalf.

DD is often thought to be epistemically inferior to RD. It is commonly thought to lead to uninformed decision-making. This perspective dates back to Rousseau (1762, III, XV), and continues in contemporary discussions (e.g., Christiano, 1996, 2004, 2018; Guerrero, 2014). Elected representatives are thought to possess better decision-making abilities (Hume, 1754; Condorcet, 1785, 1789; Madison, 1788; Hamilton, 1788; Christiano, 1996, 2018; Landa & Pevnick, 2020). In this chapter, I offer a new challenge to this line of thought. To see that, let us first explore an existing, yet idealized, challenge.

### **1.3 The jury theorem challenge to RD superiority**

Much of the recent literature justifying democracy on epistemic grounds (Cohen, 1986; Coleman, 1986; Estlund et al., 1989; List & Goodin, 2009; Goodin & Spiekermann, 2018), builds on Condorcet's jury theorem (henceforth CJT; Condorcet, 1785), which is used to formalize the notion of the "wisdom of crowds". Roughly, CJT assumes voters are faced with a binary decision where one answer is objectively correct. They aim to track it through a majority vote, where (1) all voters have a fixed probability above 0.5 of voting correctly, and (2) votes are independent.<sup>1</sup> Consequently, (a.) the probability the majority vote aligns with the correct answer monotonically grows with group size<sup>2</sup> and (b.) it converges to 1 as the group size converges to infinity.

The conditions of CJT (1 and 2) allow for a clear trade-off between voter competence and group size. That is, the larger the group, the higher its competence. Hence, for larger groups to reach the same competence, their members can be individually less competent. It follows that electorates, which usually include millions of voters, can reach higher group competence with lower individual competence than a smaller group of representatives. Robert Goodin and Kai Spiekermann (2012), for example, use the assumptions of CJT to compare the competence of large electorates to that of groups the size of existing parliaments, indicating the latter need higher individual competence to surpass the former. Although, notice that according to the unrealistic assumptions of CJT, both groups would easily reach near infallibility, making the difference between them negligible.<sup>3</sup>

---

<sup>1</sup> Conditional on the true state of the world.

<sup>2</sup> CJT is either restricted to odd group numbers, or we assume a probabilistic tie-breaking rule where, in the case of a tie, both options have an equal (and independent of votes and state of the world) probability of winning.

<sup>3</sup>E.g., 650 representatives with 0.58 individual competence yield 0.99997 group competence.

### 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

However, the assumptions of CJT are often thought to be an unrealistic portrayal of electorates, where voters are dependent, heterogeneous, incompetent, and overall fallible (e.g., Mulligan and Hunter, 2002; Somin, 2004, p. 1; Dietrich & Spiekermann, 2013; Achen & Bartels, 2017, p. 57; Brennan, 2016, p. 179-180; Barnett, 2020).<sup>4</sup>

CJT and most of its extensions assume what I term a “universal truth”. Yet, *given the same idealized assumptions*, the conclusions of CJT have been shown to hold for what I call “relative truths”, where different votes are considered correct for different voters (Miller, 1986; List and Spiekermann, 2016, I further explain these terms in the next section). Since under these assumptions, group competence monotonically grows with group size, we can assume that larger groups, such as electorates, will again require lower individual competence than representatives for similar results. However, these assumptions are unrealistic.

In recent years, alternative jury theorems have been developed with the intent of relaxing the idealized assumptions of CJT (e.g., Grofman et al., 1983; Owen et al., 1989; Boland et al., 1989; Dietrich & List, 2004; Dietrich & Spiekermann, 2013; Goodin & Spiekermann, 2018). For instance, in various jury theorems, the independence assumption has been substituted for independence conditional on common causes.<sup>5</sup> Boland et al (1989) discuss a case where all voters are influenced by the same leader, without causally influencing each other. Here, given a sufficiently low influence of the leader, group competence will converge to 1 with group size. Otherwise, it will converge to the competence of the leader (Spiekermann and Goodin, 2012). Dietrich and Spiekermann (2013) show that with independence conditional on features of the questions at hand, group competence converges to the probability of facing a problem where voters are better than chance plus half the probability of facing a problem where voters’ competence is at chance level. Additionally, Dietrich and List (2004) show that if votes are independent conditional on shared evidence, group competence converges to that of an ideal agent who possesses the same evidence.

However, these models are still idealized. For instance, they do not account for the causal influences of voters on one another. Additionally, the conclusions of those jury theorems do not necessarily indicate that electorates outperform representatives. Representatives might be less influenced by a certain leader, influenced by better leaders, perform better than chance on more questions, or have better evidence. In all these cases, they will surpass the electorate’s

---

<sup>4</sup>For example, Binomial models like CJT yield unrealistic conclusions about probability of an election upset (Barnett, 2020) and close elections frequency (Mulligan & Hunter, 2002; Edlin et al., 2007).

<sup>5</sup> See Goodin and Spiekermann (2018, pp. 67-82) for an overview.

competence. Additionally, in medium-sized groups, jury theorems do not necessarily point to an advantage of a direct vote. In fact, they have been used to show that often smaller groups of more competent individuals track the truth better than medium-sized, less competent groups (List, 2008, Bradley & Thompson, 2012). Hence, the questions of when electorates can outperform representatives, and when medium-sized groups can outperform them, remains open.

In this literature, the way different truth types affect the comparison of RD and DD has been overlooked. Moreover, analyses employing jury theorems that offer insights into this comparison (e.g., Boland et al., 1989, List, 2008, Bradley & Thompson, 2012, Goodin & Spiekermann, 2012) do not consider the fact that competent representatives must be selected, and this selection process also demands competence from the public. My chapter aims to fill these gaps. Focusing on these aspects reveals an overlooked separate challenge to the superiority of RD that does not rely on jury theorems.

#### **1.4 Distinguishing two types of truth**

The epistemic comparison of direct and indirect voting requires evaluating their respective capacity to lead to decisions that accurately track a specific truth. We should therefore ask what type of truth those votes aim to track.

One possibility is that all individuals are striving to track the same state of the world. Consequently, the answers given by different individuals will be deemed correct according to an identical set of criteria. I call the state the individuals aim to track in this context a “universal truth”. When tracking a universal truth, every voter’s choice is deemed correct if it aligns with the actual state of the world, i.e., voting 1 when the true state is 1, or voting 0 when the true state is 0. An example of this type of truth-tracking is tracking a person’s innocence or guilt in a specific crime.

Another possibility is that each individual’s vote will be considered correct relative to different set of criteria, i.e., on the same issue, some voters’ vote will be considered correct if they vote 1 and others’ if they vote 0. I term this a “relative truth”. This term simply signifies that not all voters try to track the same thing when casting their vote. They can aim to track different things and still be right or wrong about them. This concept does not in itself determine what each voter should aim to track, and we can interpret in the light of whatever we find acceptable. For

### *1. Direct or Representative Voting? It Depends on the Type of Truth We Track*

example, they could be voting according to different interests, values, goals or priorities, or a subset of these we deem permissible to track. One voter might prioritize affordable healthcare while another prioritizes employment opportunities. Of course, they may be wrong about the actions of which candidate will lead to these results. When aiming to make a decision relevant to a specific group of people, a voting procedure might want to aggregate these different considerations to find out what is best for the group.

Some take political voting to be of the second type (e.g., Miller, 1986, Goldman, 1999, List & Spiekermann, 2016, Goodin & Spiekermann, 2018, Spiekermann, 2020). Alvin Goldman, for instance, argues that in democratic decision-making, each person needs to answer the question of which result would be best from their own point of view (Goldman, 1999). According to these approaches, voters can be wrong when voting on relative truths, just as they can be wrong when voting on universal truths. Others adopt a “relative truth” understanding of a majority vote by taking democratic voting to aggregate the interests of voters. Those voters vote according to those interests (e.g., Schumpeter, 1942, Downs 1957, Mansbridge, 1983, 1990).<sup>6</sup> According to these approaches, majority rule allows taking the interests of most people into account (Mansbridge, 1990, Christiano, 2018, sections 2.2.3, 3.2.1, 3.3.2).

Importantly, prominent literature that illustrates voter incompetence in real democracies often discusses incompetence at tracking relative truths (e.g., Bartels, 1993, Althaus, 1998, Brennan, 2016, 2022). This literature illustrates voters’ votes do not accurately track their interests (e.g., Mansbridge, 1983, p. 24-6, Althaus, 1998, Somin, 2004, Brennan, 2022). For example, Jason Brennan presents the following example: “In June 2016, a slight majority of British voters voted to leave the EU. Economists widely believed – and still believe – this will harm the very citizens who voted to leave” (Brennan, 2022, p. 391). Somin also writes: “Such ignorance also raises doubts about democracy as a means of serving the interests of a majority. Voters … may also demand policies that contravene their own interests” (Somin, 2004, p. 1).

I do not advocate for adopting one interpretation of truth over the other. Plausibly, both types are relevant in different contexts, and both merit discussion in the political context. Therefore, my aim is to compare direct and indirect voting for each type of truth to reveal interesting differences between these cases.

---

<sup>6</sup> In Mansbridge, 1983, 1990, this refers to what she calls “adversary democracy”.

## 1.5 The universal truth scenario

I will show that if a group of voters is trying to track a universal truth, the following holds. When voters' competence is low, an indirect vote will often track the truth better. Nevertheless, enhancing voter competence in selecting a representative for indirect voting does not improve truth tracking as rapidly as enhancing voter competence to directly vote on the issue. If we can equally enhance both types of competence, a crossover point always exists. Beyond this point, the same level of truth-tracking ability requires more competence of voters in indirect voting. The crossover point exists even in the case of a single voter. If crowds possess some collective wisdom, then under plausible conditions, for multiple voters, the average crossover point across voters will be lowered, preserving the case of a single voter as the upper limit of the average crossover point.

### 1.5.1 The case of a single individual

The upper limit of the average crossover point can be found by looking at a simple model. Let there be a true state of the world  $X$  that takes the values 0 or 1, and only one individual,  $I$ , who needs to decide to either track that truth herself, or choose between two possible proxies to track it for her.

#### 1.5.1.1 Option 1: Direct decision-making

Let  $V_I$  represents the option that  $I$  deems correct.  $I$ 's probability of tracking this truth herself, her reliability  $r_I \in [0, 1]$ , is:

$$r_I = \Pr(V_I = x \mid X = x), \text{ for any } x \text{ in } \{0, 1\}.$$

For simplicity, I assume that  $\Pr(V_I = x \mid X = x)$  is the same for any  $x$  in  $\{0, 1\}$ . Since the state of the world  $X$  is a binary variable,  $I$ 's probability of deciding in favour of the wrong state is  $1 - r_I$ .

#### 1.5.1.2 Option 2: Indirect decision-making

Suppose, instead,  $I$  delegates the task. Let us assume there are two candidates,  $C_1$  and  $C_2$  who could serve as  $I$ 's proxy in an indirect decision-making. Let us assume one candidate is more reliable than the other. If  $I$  opts for a proxy,  $I$ 's task is to select the better candidate, the one with a higher probability of tracking the truth.

### *1. Direct or Representative Voting? It Depends on the Type of Truth We Track*

The reliability of the better candidate (B) is  $r_B = \Pr(V_B = x | X = x)$ , and that of the worse candidate (W) is  $r_W = \Pr(V_W = x | X = x)$  for  $x$  in  $\{0,1\}$ , where  $V_B$  and  $V_W$  represent what the better and worse candidates deem correct, respectively.

Consequently, there exists a true state of the world  $Y$ , which takes the values 1 or 2, indicating either that candidate  $C_1$  is the better one, and  $C_2$  is the worse one, or vice versa. This state of the world represents the universal truth that  $I$  aims to track when opting for a proxy. Let  $r_I^s \in [0, 1]$  denote  $I$ 's reliability in selecting the better candidate, i.e.,  $\Pr(V_I^s = y | Y = y)$  for  $y$  in  $\{1,2\}$  where  $V_I^s$  represents the option  $I$  deems correct.

The truth-tracking ability of the indirect decision-making depends on the competence of the selected proxy, which, in turn, relies on  $I$ 's competence in choosing the better proxy. This relationship is expressed as follows:

The expected competence of a proxy ( $P$ ) is then –

$$r_P = (r_I^s \times r_B) + ((1 - r_I^s) \times r_W) \quad (1)$$

That is, the chance that  $I$  selects the better proxy times the better proxy's chance of being right (i.e., their reliability level), plus the chance she selects the worse one times their chance of being right.

#### **1.5.1.3 Proxies as a potential “safety net”**

In some instances, indirect voting can serve as a “safety net”, where even an incompetent  $I$  reaches relatively competent decision-making. This is the case when both candidates are more competent than  $I$ , hence, even if  $I$ 's competence at selecting a candidate is poor, the overall competence of the decision-making will improve. Importantly, candidates can serve as a safety net even if only one candidate is relatively competent. For example, if  $I$  is no better than chance at either type of decision, the overall competence of direct decision-making is 0.5. Conversely, opting for a proxy will lead to a  $(0.5 \times r_B) + (0.5 \times r_W)$  chance of tracking the truth, which will be higher, even if the worse candidate is at chance level, but the better one has higher competence. If candidates are experts or devote substantial amounts of time to studying relevant topics, it may be plausible that at least one is relatively competent.

#### **1.5.1.4 Competence loss in indirect decision-making**

As  $I$  becomes more competent at both types of decision-making, opting for a proxy becomes a less appealing option. This is because *part of any added competence gets lost in the proxy*

## 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

*option.* To see this, consider how improvement in  $I$ 's competence affects both decision-making processes. Suppose by educating herself,  $I$  adds  $d$  to her original competence level  $r_I$  in deciding directly ( $0 < d \leq 1 - r_I$ ). Her new competence level on a direct decision is  $r_I + d$ . Suppose she manages to add  $d'$  to her competence in choosing proxies ( $0 < d' \leq 1 - r_I^s$ ). Then, by equation (1), the competence of an indirect decision will be:  $r_P = ((r_I^s + d') \times r_B) + ((1 - (r_I^s + d')) \times r_W)$ .

The *added* competence will be:

$$(r_B - r_W) \times d' \quad (2)$$

As long as the better candidate is not infallible and the worse candidate is not always wrong, if  $I$  improves to the same extent in both types of competence, the overall improvement in indirect decision-making will always be smaller than in direct decision-making. Hence, the proxy strategy suffers from what we might call *competence loss*, since part of the improvement is “lost” in the process.

To have the same level of improvement in both types of decision-making,  $d'$  needs to equal  $\frac{d}{r_B - r_W}$ , which often requires a significantly higher  $d'$  than  $d$ . For example, if the better and worse candidates have 0.75 and 0.55 chances of tracking the truth respectively, then  $d'$  needs to be five times bigger than  $d$ . In some instances, the required competence improvement in  $d'$  will not even be possible.

### 1.5.1.5 The existence of a crossover point

Hence, even if indirect decision-making initially leads to better results, any added competence is partially wasted in it. So, if  $I$  begins with equal competence levels in both types of decision-making and improves equally in both, there may be some point where direct decision-making becomes more effective than using a proxy, resulting in better outcomes.

I will now show that there is always such a point, and  $I$ 's competence level at that point is always less than 1, which means there always exists a practical *crossover point*. Due to the competence loss, as  $I$  continues to enhance her competence beyond this crossover point, she will need significantly higher levels of competence in choosing a proxy than in deciding directly for both types of decision-making processes to lead to equally accurate results. If we assume that both types of competence are roughly equally demanding, then beyond the

### 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

crossover point, it becomes more advantageous for  $I$  to opt for direct decision-making. For now, I will work with this assumption.

Let me now define the crossover point more precisely. The crossover point is the individual ( $I$ 's) reliability level ( $r_I^{cross}$ ) where (a) to reach the same decision-making accuracy overall, i.e., reliability in deciding directly equals that of indirect decision-making ( $r_I = r_P$ ), the same reliability is required of the individual ( $I$ ) in both ( $r_I = r_I^s = r_I^{cross}$ ). Also, (b) for any decision-making accuracy higher than the one achieved at the crossover point, the individual ( $I$ ) will need higher reliability in choosing proxies than in deciding herself (that is, when  $r_I^s, r_I > r_I^{cross}$ , for  $r_I = r_P$ ,  $r_I^s > r_I$  is required).

Given equation (1), we can see that in the case of a single individual  $I$  (henceforth  $I$ 's case), the crossover point, i.e., the point where  $r_I = r_I^s = r_P$ , is:

$$r_I^{cross} = \frac{r_W}{1-r_B+r_W} \quad (3)$$

In  $I$ 's case, there is always a crossover point where  $r_I^{cross}$  is smaller than 1. To demonstrate this, we need to (a) identify an upper limit smaller than 1 for  $r_I^{cross}$  where  $r_I = r_I^s = r_P = r_I^{cross}$ , and (b) show that for any decision-making accuracy level beyond this point (where  $r_I = r_P$ ), necessarily  $r_I^s > r_I$ . (b) holds, since, as shown above, larger improvement is required in  $I$ 's reliability in choosing a proxy ( $r_I^s$ ) than in deciding directly ( $r_I$ ) for the same level of improvement overall in decision-making accuracy.<sup>7</sup>

(a) holds as well, since for any  $r_B \neq 1$  and  $r_B \neq r_W$ :

$$r_I^{cross} < r_B.$$

Let's assume the opposite and see that we encounter a contradiction. Due to (3), we assume that  $r_I^{cross} = \frac{r_W}{1-r_B+r_W} \geq r_B$ . Solving this inequality,  $r_W(1 - r_B) \geq r_B(1 - r_B)$ . By our assumptions:  $r_B \neq 1$  and  $r_B \neq r_W$ , it must follow that  $r_W > r_B$ . In other words, the worse candidate is better than the better candidate, which is a contradiction.

Hence, if the better candidate is not infallible and there is indeed a better and a worse candidate, the competence required of  $I$  for deciding herself to be more advantageous is smaller than the

---

<sup>7</sup> This holds true since when  $r_I = r_I^s$ , the derivative of the difference  $r_I - r_P$  with respect to  $r_I$  is always positive, i.e.,  $1 + r_W - r_B$ , signifying that the difference increases with  $r_I$  (and  $r_I^s$ ). This assures us that at any point beyond  $r_I = r_P$ ,  $r_I$  will exceed  $r_P$ .

competence of the better candidate. Therefore, there is always a crossover point smaller than 1.

From (3), we can also deduce the lower limit of the crossover point, i.e.,  $r_I^{cross} > r_W$ .<sup>8</sup> Hence, in  $I$ 's case, there is always a crossover point that falls consistently between the competences of the better and worse candidates. A practical implication of this finding is illustrating where educational efforts are best spent. When  $I$ 's current competence falls below that of the worse candidate,  $I$  is clearly beneath the crossover point, as both candidates would do better. However, if education can raise  $I$ 's competence (in either type of decision-making) above the crossover point, then it is worthwhile to invest in  $I$ 's ability to vote directly. Conversely, when  $I$ 's competence exceeds that of both candidates,  $I$  is clearly beyond the crossover point and should choose direct decision-making. In such cases, the crossover point can also guide whether to prioritize improving  $I$ 's competence or that of the candidates.

### 1.5.2 Multiple voters

$I$ 's case is a toy example of one individual choosing between two candidates. When thinking about larger groups or actual democracies, we need to factor in any effects crowds might have on the comparison between direct and indirect decision-making. Suppose crowds possess some collective wisdom. Then, with multiple voters, in direct decision-making, the accuracy of the direct vote on the issue should improve. In indirect decision-making, having multiple voters will increase the probability of electing more competent candidates and, for multiple representatives, the accuracy of their majority vote. Hence, we can ask which form of voting will benefit more from crowd wisdom.

I will now illustrate how the competence loss exhibited in RD causes any wisdom we attribute to crowds to favour DD. Specifically, I will look at two cases – (i) one unified election, and (ii) a case of multiple constituencies. For each, I will show that when plausible conditions hold, the required average voter competence for DD to surpass RD is lower than for the single individual (i.e.,  $I$ 's case). That is, the average crossover point lowers when considering multiple voters. The required conditions serve to preserve the *competence loss* pattern that already exists

---

<sup>8</sup>  $\frac{r_W}{1-r_B+r_W} > r_W$  since  $1 - r_B < 1$ . We can also see that when  $r_B + r_W = 1$ ,  $r_I^{cross} = 0.5$ . That is, if by chance the two candidates would never agree on any topic, then the individual need not be better than chance for it to be better for her to vote directly. This would mean direct voting is always the safer bet. However, this assumption is implausible.

in the single voter case, rather than cause monotonic competence growth with group size. Consequently, they differ from what is usually assumed in jury theorems, as I will explain.

### 1.5.2.1 Unified election

Let us compare direct voting to a case where the entire population elects a leader out of two candidates. This could be a president, a prime minister, or another type of shared proxy. Here, wisdom of crowds lowers the average crossover point.

#### *Direct voting*

Let  $G$  be a group of  $1, \dots, n$  voters (small groups to millions). Let  $r_G \in [0, 1]$  be the reliability of the majority vote of this group of voters in voting directly, i.e.,  $r_G = \Pr(V_G = x \mid X = x)$ , for  $x \in \{0, 1\}$ , and  $r_G^S \in [0, 1]$  their reliability in selecting the better candidate,  $\Pr(V_G^S = y \mid Y = y)$ , for  $y \in \{1, 2\}$ . Let  $I_G$  be the averagely competent voter in that group, and  $B_G, W_G$  the better and worse candidates to represent that group, with reliabilities defined similarly to  $I$ 's,  $W$ 's, and  $B$ 's in  $I$ 's case.

Let  $g$  be a certain threshold number of voters beyond which crowd wisdom begins to take place.<sup>9</sup> I claim that when the following condition holds, the average competence required of voters for direct to outperform indirect voting (henceforth average crossover point) is lowered when considering groups of size  $n > g$ .

*General Wisdom of the Crowds (GWOC):* A group  $G$  with  $n > g$  voters will have higher probability of reaching the correct decision than the average voter in that group,  $I_G$ , i.e.,  $r_G > r_{I_G}$  and  $r_G^S > r_{I_G}^S$ .<sup>10</sup>

Different jury theorems assign particular values to  $r_G$  given particular assumptions about voter competence and dependence. I do not rely on such assumptions. I treat the source of any wisdom attributed to crowds as a black box. Empirical literature illustrates multiple cases where “the mean of a crowd’s forecasts will typically prove superior in quality to the forecast of the crowd’s average member” (Mannes, et al., 2014, p. 277; Surowiecki, 2012). Stock markets, prediction markets, and Google’s ranking of web pages relying on views have also been listed as indications of collective crowd wisdom (Scoles, 2007; Surowiecki, 2012). Hence, we might suppose there is *some* added wisdom to crowds beyond their average member, even

---

<sup>9</sup> We can select  $g$  according to our assumptions about crowd competence. For  $G < g$ , we might assume majority competence remains the same as the average voter’s competence.

<sup>10</sup> Note that the second equation is not required for bounding the crossover point; omitting it would lower it.

if we don't know precisely what drives it and how extensive it is. To the extent the empirical examples are realistic, the condition also often holds.

We can denote the difference between the group's and  $I_G$ 's probability of reaching the correct decision as:<sup>11</sup>

$$\Delta r_G = r_G - r_{I_G} \quad (4)$$

Given  $GWOC$ ,  $0 < \Delta r_G \leq 1 - r_{I_G}$ . Hence, the group's added competence is  $\Delta r_G$ .

### Indirect voting

Similar to any added wisdom  $I_G$  gains herself, whatever level of wisdom we attribute to crowds, some of it will be lost in an indirect vote. As a result, the average crossover point lowers with crowds wisdom. Let us see why.

Suppose the group elects a representative to vote on their behalf. Again, given  $GWOC$ , the group's competence in selecting the better candidate is  $r_G^s = r_{I_G}^s + \Delta r_G^s$ , where  $0 < \Delta r_G^s \leq 1 - r_{I_G}^s$  is the difference between the group's and  $I_G$ 's competence in selecting the better candidate.

Then, by equation (1), the proxy's expected competence becomes:  $r_{P_G} = ((r_{I_G}^s + \Delta r_G^s) \times r_{B_G}) + ((1 - (r_{I_G}^s + \Delta r_G^s)) \times r_{W_G})$ . Thus, *the expected competence the group adds beyond the average voter competence is:*

$$(r_{B_G} - r_{W_G}) \times \Delta r_G^s \quad (5)$$

Unless the better candidate is infallible and the worse always wrong, this value is smaller than  $\Delta r_G^s$ . That is, *only a fraction of the expected competence added to  $r_G$  is retained in  $r_{P_G}$* .<sup>12</sup>

Given (1),(4) and (5), the new crossover point, i.e., the average competence required of voters for direct to surpass indirect voting, will be:<sup>13</sup>

$$r_{I_G}^{cross} = \frac{r_{W_G} + (r_{B_G} - r_{W_G}) \times \Delta r_G^s - \Delta r_G}{1 - r_{B_G} + r_{W_G}} \quad (6)$$

---

<sup>11</sup> This definition does not assume wisdom of the crowds has additive properties and is thus invariant to  $r_I$ . On the contrary,  $\Delta r$  could be different for each  $G$ .

<sup>12</sup> As long as it is not both the case that the better candidate is infallible and the worse always wrong.

<sup>13</sup> The crossover point is when if  $r_{I_G}^{cross} = r_G = r_G^s$  then  $r_G = r_{P_G}$ . Given equation (1) and  $GWOC$ , it is when  $r_{I_G}^{cross} + \Delta r_G = (r_{I_G}^{cross} + \Delta r_G^s) \times r_{B_G} + ((1 - (r_{I_G}^{cross} + \Delta r_G^s)) \times r_{W_G})$ . Solving for the equation, we get  $r_{W_G} + \Delta r_G^s(r_{B_G} - r_{W_G}) - \Delta r_G = r_{I_G}^{cross}(1 - r_{B_G} + r_{W_G})$ .

This is smaller than the crossover point in  $I$ 's case as long as the following holds:

$$(r_{B_G} - r_{W_G}) \times \Delta r_G^s < \Delta r_G \quad (7)$$

If we assume both candidates are better than chance (hence  $r_{B_G} - r_{W_G} < 0.5$ ), then  $\Delta r_G^s$  can be at least twice the value of  $\Delta r_G$  while maintaining this upper limit.

Therefore, as long as  $GWOC$  and (7) hold, the average crossover point will be lowered, i.e., when there are multiple voters, the average competence required of individuals will be lower than in  $I$ 's case.

### 1.5.2.2 Multiple constituencies

Suppose, alternatively, voters are geographically divided into multiple constituencies. The group of voters needs to decide between opting for a direct vote of all voters or a two-stage process where each constituency elects a representative, and representatives conduct a majority vote.

Here, I introduce another condition:

*Competence preservation (CP):* Let a population of  $n$  voters be divided to  $k$  constituencies  $G_1, \dots, G_k$ , where  $|G_j| \geq 1$ .<sup>14</sup> If, for each constituency  $G_j$ , the probability of a correct majority vote of the voters in the constituency<sup>15</sup> is higher than the probability of a correct vote of their elected representative,<sup>16</sup> then, overall, the probability of a correct majority vote of all  $n$  voters will also be higher than that of their  $1, \dots, k$  representatives.<sup>17</sup>

Roughly, this states that if all constituency groups have higher expected majority competence than the competence of their respective representative, then the majority competence of all voters is also higher than that of all representatives.

I conjecture that, given  $CP$ ,  $GWOC$ , and if (7) holds for each constituency, then the average competence required of voters for the overall direct vote to be epistemically superior to the overall representative's vote is lower in the case where there are multiple voters in each constituency, compared to the case with a single voter in each constituency (one-to-one ratio).

---

<sup>14</sup> Constituencies here are mutually exclusive and collectively exhaustive subsets of the group of voters.

<sup>15</sup> For each constituency  $G_i$ , this probability,  $r_{G_i}$ , is defined similarly to  $r_G$  above.

<sup>16</sup>  $r_{p_j}$  for constituency  $G_j$ , where  $r_{p_j}$  is defined similarly to  $r_p$  above. Hence, the condition entails – for each constituency,  $G_i$ ,  $r_{G_i} > r_{p_j}$ .

<sup>17</sup> These probabilities are defined similarly to  $r_G$  above.

## 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

Hence, the upper limit of the average crossover point can be derived from averaging over crossover points for one individual in each constituency according to *I*'s case, as I explain.

### *The one-to-one ratio scenario*

Suppose an electorate  $G_1$  is composed of  $k$  constituencies. Each constituency includes only one voter,  $I_j$ , for  $j$  in  $1, \dots, k$  voters, voting themselves to select their representative. *If each voter surpasses their crossover point as specified in I's case* (hinging on the competence of their respective candidates, equation (3)), then according to *CP*, the majority competence of the group of voters in a direct vote ( $r_G$ ) *will not be less reliable than the vote of the representatives* ( $r_P$ ). This is because, beyond the crossover point, each voter has higher expected competence than the candidate they elected. Hence, according to *CP*, the probability of a correct vote of all voters is higher than that of their representatives. Therefore, the average required competence for direct to surpass indirect voting in such a case is the average among these crossover points.

### *Higher voter-representative proportions*

Suppose, now, that more voters join voter  $I_j$  in each of the  $1, \dots, k$  constituencies. Then, if *GWOC* and (7) hold in all constituencies, the average required competence for direct to outperform indirect vote within each constituency will be lowered to the crossover point shown in equation (6).<sup>18</sup> This is because, as in the unified elections case, for any wisdom attributed to crowds within each constituency, some of it will be lost in the indirect vote. If the average is lowered in each constituency, the general average across constituencies is also lowered.

### *The plausibility of CP*

The *CP* condition is not equivalent to assumptions made in existing jury theorems. Hence, the argument portrayed in this chapter offers a new challenge to the superiority of RD that does not rely on jury theorems. *CP* might not hold in all cases, but it plausibly holds in many, thereby expanding the cases where the superiority of RD is challenged beyond the particular cases jury theorems can address, whose limitations were explained in section 1.3. The competence comparison in *CP* relies on the competence of group parts, rather than on group size. Regarding group size, *CP* implies only that if the more competent group is at least as large as the less competent one, this will not undermine the preservation of competence. This differs from jury theorems, where group size drives crowd's overall wisdom. Consequently, *CP* does not imply

---

<sup>18</sup> Where group  $G$  is taken to be constituency  $j$  for  $j$  in  $1, \dots, k$ .

### *1. Direct or Representative Voting? It Depends on the Type of Truth We Track*

competence increases monotonically with group size, nor that it converges to 1 as group size approaches infinity.

As with *GWOC*, the sources of the preservation of competence are treated as a black box. It is not necessary for this preservation to have vote independence or low vote dependence. Notably, in contrast to jury theorems, *CP* can hold even under maximal dependence, that is, when all voters' votes are perfectly correlated and all representatives' votes are as well. Suppose that all voters and all representatives are fully dependent within their groups. Then, the overall competence of all the voters matches that of any individual constituency, and the competence of the group of representatives matches that of any individual representative. Hence, *CP* is trivially true. That is, if each constituency is more competent than its corresponding representative, then the entire electorate is more competent than the group of representatives.<sup>19</sup>

Conversely, dependence structures where the votes of different constituencies exhibit higher dependence than those of representatives may violate *CP*. Yet, given that representatives regularly interact with one another, while voters of different constituencies do so less frequently, plausibly this will not happen often.

A more plausible challenge to *CP* arises when taking the effect of deliberation into account. It could be argued that deliberation improves competence more in the group of representatives. Hence, less competent representatives could together produce a more competent group. Yet, deliberation among representatives might come at the expense of deliberation among citizens. When voters have more influence, they tend to be better informed (e.g., Benz, 2004). Hence, in indirect voting, citizens might deliberate less on political issues. Importantly, the deliberation among citizens is likely to occur primarily within constituencies, without significantly increasing dependence levels among them. Hence, when factoring in deliberation, the lower levels of dependence between constituencies than between representatives may still weigh in the direct vote's favour.

To conclude, under the aforementioned conditions, it is always possible to establish a maximum level smaller than 1 of average competence required for voters to benefit more from direct voting. This can be explained intuitively. When tracking universal truths, candidates competing to represent can serve both as a “safety net” and as a “limiting ceiling”. When voters have no competence, they serve as a “safety net”, yet, when voters gain competence, indirect voting

---

<sup>19</sup> Under maximal dependence, the crossover point is as in *I*'s case.

exhibits competence loss and becomes less desirable. Any wisdom attributed to crowds over individuals will also suffer such competence loss in indirect voting.

## 1.6 Relative truth

When tracking a relative truth, a different vote may be considered correct for different voters, depending, for instance, on their values, goals or interests. As explained in section 1.4, within this type of voting, voters may still often vote incorrectly. They might support a policy that doesn't align with their values, interests or goals, as they misunderstand the policy's consequences. Similarly, they can mistakenly vote for a candidate who doesn't promote those values, interests or goals. Consequently, there can be varying levels of voter competence in both direct and indirect voting on relative truths, raising the need for a comparison between their accuracy levels. In discussions of RD, it has been claimed that voters often vote against their interests (e.g., Althaus, 1998, Somin, 2004, Brennan, 2016, 2022), leading to proposals for less democratic systems such as epistocracy (e.g., Brennan, 2016, 2022). This direction overlooks the possibility that direct voting will often be superior to indirect voting for such truths, which I will show in this section.

To conduct the comparison, we must first determine what truth the majority vote aims to track in the context of relative truths. A democratic option would be to choose to adopt a policy that is right for the majority, given their goals (henceforth standing for any relative standards voters may aim to track, such as values, goals, and interests). This intuitive proposal has been argued for and adopted in the literature (Miller, 1986, Goldman, 1999, List & Spiekermann, 2016, Goodin & Spiekermann, 2018), and I will adopt it without arguing for it further.<sup>20</sup>

In the scenario of relative truths, then, the purpose of elections is for the public to indicate which goals they support by choosing a candidate who promotes those goals. This aligns with the idea that in representative democracy, voters contribute by determining general goals, while elected representatives utilize their expertise to advance these objectives (Christiano, 1996,

---

<sup>20</sup> In the literature, some have assumed that the proportion of the majority in the population stays constant as the population grows (e.g., List & Spiekermann, 2016). This assumption is used as part of the jury theorem stating that chance of voting for what is right for the majority monotonically grows with group size and converges to 1 as group size converges to infinity, which is not a condition of my model. Such an assumption is compatible with my model. However, it is also compatible with the model to simply assume that whenever there are multiple voters, the relevant majority is the majority within that finite group of voters. This assumption tracks the idea that what is correct in this context is what is correct for the particular voting group, as, for instance, they are making a decision regarding their own lives.

2018). If this holds true, different candidates should advocate for different goals, allowing people to choose their representatives based on their own goals.

In this context, the expectation that the candidate should be relatively competent is replaced by the expectation that the candidate be relatively competent *given specific goals*. i.e., having a relatively high probability of voting in favour of the outcome that aligns with their goals.

Given these definitions, in the context of relative truths, I take the crossover point to be the point when direct voting becomes more accurate than indirect voting at tracking what is right for the majority, given people's goals. It's important to distinguish between a voter voting competently and a voter contributing to the accuracy of the majority vote. A competent vote is one where the voter chooses the policy that aligns with their own goals. If their goals don't align with the majority, their voting competently can lead the collective vote away from what we consider the most accurate outcome.

Let us now explore the upper limit for the crossover point in situations involving relative truths. For that purpose, we will look at the crossover point between when a single voter can better contribute to a majority vote by voting directly or by selecting a representative to act as their personal proxy, i.e., vote on their behalf.

### 1.6.1 The case of a single individual

Suppose a specific group of  $n$  voters are again facing a binary decision between rejecting or accepting a policy, labelled as 0 and 1, respectively. For any voter  $i$  of the  $1, \dots, n$  voters, their reliability  $r_i$  is measured as the probability of choosing the answer that aligns with their goals. Random variable  $X_i$  refers to the correct answer for voter  $i$  termed  $i$ 's voter-specific truth.  $X_i$  can take two values, 0 or 1, depending on which of the options is the correct one given the voter's goals. This implies the existence of a profile of voter-specific truths  $\langle x_1, x_2, \dots, x_n \rangle \in \{0,1\}^n$ .<sup>21</sup>  $V_1, \dots, V_n$  denote the options the  $n$  voters deem correct, taking the values 0 or 1. For each voter  $i$ , the voter's competence  $r_i \in [0,1]$  in this case is the probability their vote will reflect their truth,  $r_i = \Pr(V_i = x_i \mid X_1 = x_1 \& X_2 = x_2 \& \dots \& X_n = x_n)$ . As explained above, the truth that the group of all voters needs to track is the result that is correct given the goals of the majority of the whole group of voters. This is denoted as  $X_{overall}$ . It will be 0 if more than half the goals in  $\langle x_1, x_2, \dots, x_n \rangle$  are 0, and 1 if more than half are 1.<sup>22</sup> In other words, the truth

---

<sup>21</sup> Here I follow the notation introduced in List and Spiekermann, 2016 to discuss relative truths.

<sup>22</sup> I assume there is always a majority for either 0 or 1.

of the majority equals the truth of more than half the voters. The proportion of the individuals who align with the majority's truth out of the entire population is termed  $q$ , where  $0.5 < q \leq 1$ .<sup>23</sup>

Let us now compare the contribution of a single voter to overall group accuracy when voting directly or when selecting a representative to vote on their behalf.

### 1.6.1.1 Option 1: Direct decision-making

As mentioned above, the overall correct vote is one that aligns with the goals of the majority. Therefore, each specific voter can directly contribute to the majority vote's accuracy either by having a voter-specific truth aligned with the majority and voting correctly or by having a voter-specific truth that doesn't align with the majority but voting incorrectly. Given the probability of any randomly selected individual to hold majority values is  $q$ , the probability of these events occurring is:

$$(q \times r_i) + ((1 - q) \times (1 - r_i)) \quad (8)$$

### 1.6.1.2 Option 2: Indirect decision-making

Now, let's assume there are two candidates,  $C_1$  and  $C_2$ , to be a voter's representative, a candidate promoting the goals of the majority (henceforth the majority candidate) and one promoting those of the minority (henceforth the minority candidate), with reliabilities of  $r_{maj}$  and  $r_{min}$ , respectively. The majority candidate is the one more likely to have a vote that aligns with the voter-specific truth of  $q$  of the population. We can define a profile of candidate-specific truths with two variables  $\langle x_{maj}, x_{min} \rangle$  each taking values 0 or 1, where  $x_{maj}$  aligns with  $X_{overall}$  and  $x_{min}$  aligns with the alternative. Let  $X_{maj}$  and  $X_{min}$  be the correct answers for the majority and minority candidates, respectively. Let  $V_{maj}$  and  $V_{min}$  denote the options the majority and minority candidates deem correct, respectively, taking values 0 or 1.

Then, the competence of each candidate, i.e., their chance of voting correctly according to their goals, is:

$$r_{maj} = \Pr(V_{maj} = x_{maj} \mid X_1 = x_1 \& X_2 = x_2 \& \dots \& X_n = x_n), \text{ where } r_{maj} \in [0,1]$$

$$r_{min} = \Pr(V_{min} = x_{min} \mid X_1 = x_1 \& X_2 = x_2 \& \dots \& X_n = x_n), \text{ where } r_{min} \in [0,1].$$

---

<sup>23</sup>  $q = \frac{\sum_{i=1}^n x_i}{|\{x_1, x_2, \dots, x_n\}|}$  if it is greater than 0.5. Otherwise  $q = 1 - \frac{\sum_{i=1}^n x_i}{|\{x_1, x_2, \dots, x_n\}|}$ .

Consequently, for each voter, there is a voter-specific truth  $Y_i$  that either candidate  $C_1$  or candidate  $C_2$  is the right choice given their goals. This is the relative truth a voter aims to track in an indirect vote. That is,  $Y_i$  can take two values, 1 or 2, denoting the index of the candidate that promotes the voter's goals ( $C_1$  or  $C_2$ ). This also implies a profile of voter-specific truths  $\langle y_1, y_2, \dots, y_n \rangle$  in  $\{1, 2\}^n$ . For any voter  $i$ ,  $V_i$  represents the option the voter deems correct out of the two candidates, taking values 1 or 2. Let  $r_i^s \in [0, 1]$  denote a voter's reliability in selecting the candidate that promotes their goals, i.e.,  $\Pr(V_i = y_i | X_1 = x_1 \& X_2 = x_2 \& \dots \& X_n = x_n)$  for both values of  $y$ .

The probability that a selected candidate will advance the majority vote in the direction that is correct given the goals of the majority is composed of two components: (a) the probability that the voter selects the majority candidate and that this candidate votes correctly based on their goals, and (b), the probability that the voter chooses the minority candidate and that this candidate votes incorrectly given their goals. This can be denoted as

$$\alpha \times r_{maj} + (1 - \alpha) \times (1 - r_{min}) \quad (9)$$

where  $\alpha$  and  $1 - \alpha$  symbolize the probability of selecting the majority and minority candidate respectively. The chance of a voter selecting the majority candidate ( $\alpha$ ) is composed of the chance their goals align with the majority (making the majority candidate the right choice for them) and voting correctly, and the chances they do not, and to vote incorrectly.

$$\alpha = q \times r_i^s + (1 - q) \times (1 - r_i^s) \quad (10)$$

### 1.6.1.3 Competence loss in indirect voting

As with universal truths, there is a competence loss in indirect voting. For both types of voting to contribute the same amount to a majority vote, any improvement in the voter's competence in choosing a candidate needs to equal  $\frac{d}{r_{maj} + r_{min} - 1}$ ,<sup>24</sup> where  $d$  is the improvement in the voter's competence voting directly on the issue. Hence, as long as both candidates are not infallible, higher improvement will be needed in indirect voting for the same result. For example, if one of the candidates has a competence of 0.75 and the other of 0.55, then  $d'$  needs to be 3.33 times

---

<sup>24</sup>When the voter's reliability in voting directly increases by  $d$ , their contribution to the majority vote is  $q \times (r_i^s + d) + (1 - q) \times (1 - (r_i^s + d))$ , hence, the added competence is  $2 \times q \times d - d$ . When voting indirectly, according to equations (9), (10), the expected contribution to an indirect vote can be expressed as  $(q \times (r_i^s + d') + (1 - q) \times (1 - (r_i^s + d')))(r_{maj} + ((1 - q) \times (r_i^s + d') + q \times (1 - (r_i^s + d'))))(1 - r_{min})$ . Therefore, the added competence is  $(2 \times q \times d' - d')r_{maj} - (2 \times q \times d' - d')(1 - r_{min})$ .  $2 \times q \times d' - d'$  equals  $(2 \times q \times d' - d')r_{maj} - (2 \times q \times d' - d')(1 - r_{min})$  when  $d' = d/(r_{maj} + r_{min} - 1)$ .

larger than  $d$  for the same improvement in the contribution of the voter to the majority vote.<sup>25</sup> Conversely, if the voter improves the same amount in both, a larger overall improvement will be reached in direct decision-making.

#### 1.6.1.4 The crossover point

The crossover point, as before, is the point where the same truth-tracking level is achieved by the same level of voter competence in both types of voting, and above which direct voting is more advantageous. That is, when  $r_i^s = r_i$  and the contribution of a voter to the direct vote (8) is the same as the probability of a representative contributing to the representative vote (9). Above this point, direct voting is guaranteed to be more advantageous for the same voter competence, as a part of any growth in the voter's competence is lost in the indirect vote.

Using equations (8), (9), (10) the probability of a representative and that of a voter contributing to alignment with the majority are equal when  $(q \times r_{cross}) + ((1 - q) \times (1 - r_{cross})) = (q \times r_{cross} + (1 - q) \times (1 - r_{cross})) \times r_{maj} + ((1 - q) \times r_{cross} + q \times (1 - r_{cross})) \times (1 - r_{mn})$ . Solving this equation, we get

$$r_{cross} = \frac{2q - 1 - qr_{maj} - qr_{min} + r_{maj}}{2(2q - 1 - qr_{maj} - qr_{min}) + (r_{maj} + r_{min})} \quad (11)$$

We can see that unless  $r_{maj} = r_{min} = 1$ , i.e., both candidates are infallible, which is unrealistic, or  $q = 0.5$ , which is false by our assumptions, the following holds:

if  $r_{maj} = r_{min}$  then  $r_{cross} = 0.5$ .

if  $r_{maj} > r_{min}$  then  $r_{cross} > 0.5$  and if  $r_{maj} < r_{min}$  then  $r_{cross} < 0.5$ .

That is, if both candidates are equally competent at promoting their respective goals, the crossover point is at chance level. Hence, for any voter competence above chance level, direct voting would be more advantageous than representative voting. If the majority candidate is more competent, the crossover point is above chance level. Conversely, if the minority candidate is more competent, the crossover point falls below chance level.

---

<sup>25</sup>We can also see that when  $r_i^s = r_i$ , the derivative of the difference between the expected contribution of a voter to that of a representative is always positive, signifying that the difference increases with  $r_i$  (and  $r_i^s$ ). This assures us that at any point beyond when the expected contributions are equal, direct voting will exceed indirect voting. The derivative of  $diff = q \times r_{cross.} + (1 - q) \times (1 - r_{cross.}) - [(q \times r_{cross.} + (1 - q) \times (1 - r_{cross.})) \times r_{maj} + ((1 - q) \times r_{cross.} + q \times (1 - r_{cross.})) \times (1 - r_{min})]$  with respect to  $r_{cross.}$  is  $4q + r_{maj} + r_{min} - 2 - 2qr_{maj} - 2qr_{min}$ . For it to be positive, we need  $2q(2 - r_{maj} - r_{min}) > (2 - r_{maj} - r_{min})$  which is always true because  $q$  is larger than 1.

If we assume that candidates are not consistently better at promoting either majority or minority goals, then this analysis suggests that for relative truths, direct voting will likely lead to better results. This is because if we consider  $r_{maj}$  and  $r_{min}$  as random variables that are picked out of the same distribution of decision successes, then  $\langle r_{maj} \rangle = \langle r_{min} \rangle$ , where  $\langle x \rangle$  is the mean value of  $x$ . Consequently, the mean of the distribution of  $r_{cross}$  is expected to be 0.5.<sup>26</sup> Hence, if voter competence tends to be better than chance, direct voting is expected to lead to better results.

One might argue that the majority candidate is likely more competent, raising the crossover point. This might be based on the idea that due to the law of large numbers, there is a high chance of finding competent individuals in larger groups. However, since both majority and minority groups typically include millions, this statistical effect applies equally to both. Additionally, Figure 1 below shows that as the majority size increases, the rise in the crossover point becomes slower. Consequently, even if larger majorities enhance the competence of the majority candidate, they simultaneously lower the crossover point. Therefore, our best bet still appears to be direct voting in cases involving relative truths.

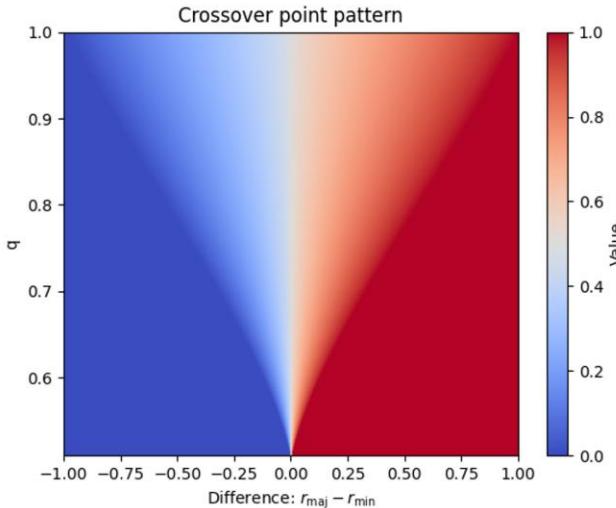
One might also suggest that RD still exhibits the best epistemic potential as educating candidates competing to represent will be more effective than educating voters. However, even if candidates are very competent, their average ability to vote accurately with respect to the goals of the majority will be low, as they promote different goals.

Notice that, in contrast to the case of universal truths, the crossover point does not necessarily fall between the competences of the majority and minority candidates. *I* could be worse than both candidates and still above the crossover point, or more competent than both candidates and still below it. For example, as shown above, if both candidates have the same competence level, i.e.,  $r_{min} = r_{maj}$ , then  $r_{cross} = 0.5$ , whether both candidates have a higher competence than 0.5 or a lower one.

Figure 1 depicts the epistemic comparison between direct and indirect voting for relative truths.

---

<sup>26</sup> These conclusions are supported by various simulations under the assumption that both candidates draw from identical distributions of decision success rates. The simulations included truncated normal distributions with varying parameters:  $q=0.51, 0.55, 0.6, 0.7, 0.8, 0.9, 1$ , candidate mean competence:  $0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1$ , and standard deviations:  $0.01, 0.02, 0.03, 0.05, 0.06, 0.1, 0.2, 0.3, 0.4, 0.5$ , and Beta distributions with:  $q=0.51, 0.55, 0.6, 0.7, 0.8, 0.9, 1$ ,  $\alpha, \beta = 0.1, 0.3, 0.5, 0.8, 1.0, 1.5, 2.0, 3.0, 5.0, 10.0$ . Across all scenarios where  $q > 0.51$ , the mean and median of the crossover point ( $r_{cross}$ ) consistently approximates 0.5. At  $q=0.51$  there were a few exceptions for the Beta distribution. However, the mean value of these exceptions again approximates 0.5.



**Figure 1:** The crossover point as a function of both majority proportion ( $q$ ) and the difference between the reliability of the majority and minority candidates ( $r_{maj} - r_{min}$ ). The value of the crossover point is represented by the colour (colour bar). Values below 0 are clamped to 0 and values above 1 are clamped to 1. The graph illustrates that as the value of the difference increases the crossover point rises, and when the difference is 0 the point is at chance level.

Figure 1 highlights another important distinction between universal and relative-truths. Specifically, in the case of relative truths, there are situations where there is simply no crossover point, i.e., no voter competence level between 0 and 1 where direct voting replaces indirect voting as the more epistemically attractive option. In those cases, for all competence levels, either direct voting is consistently superior (dark blue areas in Figure 1) or representative voting is consistently superior (dark red areas). Hence, in specific cases where the majority proportion is small and the majority candidate is more competent, it may be better to opt for RD.

Also, Figure 1 shows that as the proportion of the majority ( $q$ ) increases, there are more cases where a crossover point can be reached.

### 1.6.2 Multiple voters

Just as in the case of universal truths, under similar conditions, the competence loss exhibited in RD causes wisdom attributed to crowds to favour DD, usually lowering the average crossover point. As for universal truths, this result does not rely on jury theorems.

#### 1.6.2.1 The unified elections scenario

Any added competence we attribute to crowds over individuals will be spent less effectively in indirect voting, due to the competence loss exhibited in indirect voting, under the following condition (for a group size threshold for crowd wisdom  $g$ ):

### 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

*General wisdom of the crowds for relative truths (GWOCRT):* if group G's majority is aligned with the tracked majority, and  $n > g$ , group G has a higher probability of producing the correct outcome for the general majority than the average voter in it ( $I_G$ ).<sup>27</sup>

Since the majority in general population aligns with the general majority by definition,<sup>28</sup> for unified elections, the crowds have higher expected competence than  $I_G$ , according to *GWOCRT*.

As with universal truths, the improvement caused by crowds can be depicted as  $\Delta r_G^s$ .<sup>29</sup> In the case of relative truths, this improvement is in the majority vote's ability to track what is right given the goals of the majority of voters. Now, whatever wisdom we attribute to crowds, some will be lost in the indirect vote. Given equation (9), a representative's expected ability to contribute to overall accuracy now becomes  $(\alpha + \Delta r_G^s) \times r_{maj_G} + (1 - (\alpha + \Delta r_G^s)) \times (1 - r_{min_G})$ , where  $0 < \Delta r_G^s \leq 1 - \alpha$  is the improvement caused by the crowds in selecting the majority candidate, and  $r_{maj_G}, r_{min_G}$  the competences of the majority and minority candidates competing to represent group G, defined similarly to  $r_{maj}, r_{min}$ .

Thus, *the overall improvement in the indirect vote* is:

$$(r_{maj_G} + r_{min_G} - 1) \times \Delta r_G^s \quad (12)$$

which is smaller than  $\Delta r_G^s$ , as long as it is not the case both candidates are infallible given their goals, indicating a *competence loss*.

---

<sup>27</sup> Formally: Let  $X_{overall}^G$ , taking values 0 or 1, be what is correct for the majority in group  $G$ , defined as what is right for more than half the voters in  $G$ , i.e., 1 if more than half the values in  $\langle x_1, \dots, x_l \rangle$  are 1 and 0 otherwise, where  $1, \dots, l, (1 < l \leq n)$  are the voters in group  $G$ . Let  $X_{overall}$  represent the option that is right for the majority in the general population of  $n$  voters, 1 if more than half  $\langle x_1, \dots, x_n \rangle$  are 1 and 0 otherwise. Let there be two candidates competing to represent group  $G$ ,  $C_1^G$  and  $C_2^G$ , one promoting the values of the overall majority, those aligning with  $X_{overall}$ , and one promoting the values of the overall minority, with reliabilities  $r_{maj}^G$  and  $r_{min}^G$  respectively. Let  $Y_G$  taking values in 1 and 2 denote the index of  $r_{maj}^G$  out of  $C_1^G$  and  $C_2^G$ . Let  $V^G$  be the outcome of a collective majority vote of group  $G$ . Let  $I_G$  be the averagely competent voter in group  $G$ . Then, *General wisdom of the crowds for relative truths (GWOCRT): if  $X_{overall}^G = X_{overall}$ ,  $Pr(V^G = X_{overall} | X_1 = x_1 \& X_2 = x_2 \& \dots \& X_m = x_m) > Pr(V^{\{I_G\}} = X_{overall} | X_1 = x_1 \& X_2 = x_2 \& \dots \& X_m = x_m)$ , and  $Pr(V^G = y_G | X_1 = x_1 \& X_2 = x_2 \& \dots \& X_m = x_m) > Pr(V^{\{I_G\}} = y_G | X_1 = x_1 \& X_2 = x_2 \& \dots \& X_m = x_n)$ .*

<sup>28</sup> Again, recall, I take the relevant majority to be a majority out of the entire finite group of the voting population. See footnote 20 for expansion.

<sup>29</sup> As before, the value of  $\Delta r_G$  may be different for any competence level of a single voter.

Hence, when *GWOCRT* holds, the average competence required of voters for direct to outperform indirect voting, i.e. the average crossover point, is lowered, as long as there is not a significant difference between  $\Delta r_G^s$  and  $\Delta r_G$ , or more specifically:<sup>30</sup>

$$(r_{maj_G} + r_{min_G} - 1) \times \Delta r_G^s < \Delta r_G \quad (13)$$

### 1.6.2.2 Multiple constituencies

When considering multiple constituencies, in the vast majority of cases, the average competence required of voters for direct to outperform indirect voting, i.e., the average crossover point, lowers when considering wisdom of crowds. This holds under similar conditions to the case of universal truths, i.e., (1) *GWOCRT*, (2) *CP* (the same *CP* introduced for universal truths), and (3) there is no vast difference between the competence crowds add when electing a candidate and when voting directly (i.e., equation (13) above).

Let there be  $n$  voters divided into  $k$  constituencies. I assume that in a majority of constituencies ( $m > k/2$ ), the majority of voters in each such constituency ( $i$  in  $1, \dots, m$ ) share the same goals as the overall majority of the  $n$  voters. If this assumption doesn't hold, direct voting has a built-in advantage, as explained in the appendix. In the appendix, I also discuss an alternative definition of the relevant majority.

Given this assumption, I follow the same reasoning used in the universal truth case. I begin with a one-to-one ratio of voters to representatives. In this setup, when each voter crosses the crossover point, they have a higher competence than the expected competence of their representative. Then, according to *CP*, the group of voters is more competent than the group of representatives.

Next, I consider what happens when each constituency consists of more voters. In constituencies aligned with the overall majority (which, by assumption, comprise the majority of constituencies), adding voters will lower the average crossover point required for direct voting to outperform indirect voting, as long as equation (13) holds. This effect arises, again, because the larger group of voters benefits from crowd wisdom (per *GWOCRT*), while some of this competence is lost in the process of electing a representative.

---

<sup>30</sup> Given equations (8),(9),(10),(12) the new crossover point of average voter competence is:  $\frac{2q_G - 1 - q_G r_{maj_G} - q_G r_{min_G} + r_{maj_G} + (r_{maj_G} + r_{min_G} - 1) \Delta r_G^s - \Delta r_G}{2(2q_G - 1 - q_G r_{maj_G} - q_G r_{min_G}) + (r_{maj_G} + r_{min_G})}$  ( $0.5 < q_G \leq 1$  is the proportion of the local majority within group G).

Now, suppose each constituency contains the same number of voters.<sup>31</sup> Then, since the majority of voters are located in the majority of constituencies, where the average crossover point decreases, it is reasonable to expect that the overall average crossover point also decreases in most cases. I support this claim with extensive simulations in the appendix, which show that the overall crossover point decreases in the vast majority of cases (e.g., for most spreads of candidate competence across constituencies).

Notably, the exceptions where the average crossover point increases after adding multiple voters, primarily occur when the majority proportion is small. This is consistent with the results in the single voter tracking relative truths' case, where small majority proportions are cases where indirect voting might be preferable.

Therefore, when the mentioned conditions hold, and we have no further knowledge, such as the precise size of the majority, then the best bet for achieving good epistemic results when tracking relative truths is a direct voting system, such as the one used in referenda and DD.

I will conclude this section with an intuitive explanation of the problem indirect voting faces when tracking relative truths. If we believe different people are entitled to have different values or goals, or that they are permitted to vote according to their interests, which naturally differ, then for people to be able to express these differences between them in an indirect voting system, different candidates need to promote different goals. Therefore, the average ability of the candidates to promote the goals of the majority is expected to often be low, as the one promoting the goals of the minority is often worse than random at promoting the goals of the majority. Hence, the aforementioned “safety net” of the average high competence of the representatives, which makes representative voting superior when voters are relatively ignorant, does not exist in such a case, and direct voting is the best option.

## Conclusions

I have analysed the dynamics of how the type of truth being tracked affects the epistemic comparison between direct and indirect voting. The relations uncovered here can serve us in the evaluation of the epistemically best form of government, direct, semi-direct or representative democracy, given our beliefs about the type of truth being tracked, voters'

---

<sup>31</sup> Simulations in the appendix show the claim holds also when constituency sizes differ, as long as they don't consistently favour minority aligned constituencies.

### 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

competence levels, the proportion of the majority, and candidate competence. This can also indicate whether votes on certain topics should be referendums or votes conducted by representatives. The results also apply to smaller groups deciding between a direct or an indirect vote.<sup>32</sup>

The results might be particularly relevant for liquid democracy, in which voters choose to either vote directly or delegate their vote to a particular individual (e.g., Blum & Zuber, 2016). The results serve not just to suggest whether opting for liquid democracy is advisable on epistemic grounds, but also as a guideline for voters within such a system. The crossover points presented in this chapter illustrate the point after which it may no longer make sense for the voter to delegate their vote, as higher competence will be required of the voter in deciding to whom to delegate than in voting directly. Notably, this chapter focuses on the cases where there are two options for whom to delegate the vote. However, similar patterns may occur for a larger number of candidates, as there is still expected to be a competence loss, caused by any added voter competence only factoring towards a chance of selecting a fallible candidate.<sup>33</sup>

Let us assume, for the moment, that tracking the truth when voting directly on policies is not vastly more difficult than electing representatives to make policy decisions, and that *GWOC(RT)* and *CP* hold. Then, my results suggest the following.

When the goal is to track a *universal truth*, RD can provide a form of protection, a “safety net”, against generally incompetent voters, since even in that case, the expected competence of representatives can still be relatively high. However, the epistemic potential of RD when voters become more competent is ultimately lower than that of DD, because there exists a threshold of voter competence beyond which DD outperforms RD. Once that threshold is crossed, DD

---

<sup>32</sup> One might argue that the cognitive load of voting repeatedly reduces competence in DD. However, empirical studies have shown that political systems allowing for more opportunities for political participation lead to increased political knowledge (e.g., Tolbert et al. 2003; Benz & Stutzer, 2004). Generally, for both types of decision-making, voters need to either increase the chance that for every subsequent vote of representatives, or for every direct vote of their own, the vote will be correct. They can either do so by voting according to the issues and votes directly, which is as complicated in both cases, or by using heuristics, such as the credentials of candidates to represent. However, heuristics can also be used in direct-voting by asking the advice of people around with impressive credentials. Direct voting has the advantage that there are more trustworthy people to use as heuristics, and that there isn't an added layer of concealment of information about how the candidate will vote. Either way, even if higher competence is less likely in direct democracy, the presented results would still apply to direct voting in the form of referendums.

<sup>33</sup> For example, the crossover point will remain the same when there are four candidates under certain conditions. For instance, if (1) the original probability of selecting the better candidate is now split equally between the two better candidates, and the same is true for the worse candidates, and (2) the average competence of the better candidates equals that of the better candidate in the original election, and the same for the worse candidates. Then

$$r_p = \left(\frac{r_l^s}{2} \times r_B\right) + \left(\frac{r_l^s}{2} \times r_{B'}\right) + \left(\left(1 - \frac{r_l^s}{2}\right) \times r_w\right) + \left(\left(1 - \frac{r_l^s}{2}\right) \times r_{w'}\right)$$

will remain the same.

### *1. Direct or Representative Voting? It Depends on the Type of Truth We Track*

achieves better results with less individual voter competence. Hence, RD might be considered a risk-averse option, but with limited epistemic *potential* relative to DD.

In contrast, when the goal is to track a *relative truth*, DD is generally better, as the expected mean of the upper limit competence threshold for DD to outperform RD is near chance level, given plausible assumptions about candidate competence. If voters are even slightly better than random at promoting their own goals, DD is epistemically preferable. This suggests that, if in political contexts the truths being tracked are relative, as some have argued, DD may outperform RD, contrary to common thought. Yet, for relative truth, there are specific cases where the epistemic *potential* of DD will be inferior to RD, as there is no crossover point. Specifically, that is the case when majority candidates are significantly better than minority candidates and the majority is close to half.

These conclusions suggest that the intuition that voting on policies requires more competence can be explained with the idea of the “safety net”. That is, we mistake the fact that we don’t *need* to be highly competent to choose representatives for the idea that we *are* competent in doing so. We may feel that current democracies function adequately even when voters are ignorant. However, this is primarily due to the safety net representatives provide even when voters are incompetent at selecting good candidates, rather than voters having high level of competence in selecting such candidates.

Now, let’s suppose that tracking the truth in voting on policies differs in difficulty from electing representatives to vote on policies. My central claim about a significant difference between universal and relative truths still holds. When tracking relative truths, direct voting is most likely still the best choice, as usually the expected mean of the upper limit of the crossover point is chance level, equivalent to no competence at all, which is likely equally undemanding for both types of competences. For universal truths, RD will still be the better choice for low levels of voter competence. For higher levels, it will depend on the difference between two types of competences. Evidently, then, distinguishing between these two types of truth is still important.

Hence, I pose a new challenge to the claim that voter ignorance is a problem for direct voting and the better option is representative democracy. It can be explained as an intuition about universal truths. However, this intuition is misguided if we aim to track relative truths in the political context. Also, even for universal truths, direct voting often has more potential, as it improves more with any added competence.

## Appendix: when multiple constituencies track relative truths

### *Number of constituencies in the majority*

When elections are organized across multiple constituencies, a question arises: Which majority should the vote aim to track? There are two main possibilities: (1) the majority of the general population, or (2) the majority across constituencies. That is, the majority of local-level majorities. Focusing on option (1), two further scenarios are possible: (a) in the majority of constituencies, the local majority aligns with the overall majority. (b) in the majority of constituencies, the local majority does not align with the overall majority.

In case (b), indirect voting has a structural disadvantage. Dividing the electorate into constituencies, which is only needed for indirect voting, can distort the results. If each constituency vote accurately tracks its local majority, then the majority of representatives will be misaligned with the overall majority, leading to an incompetent overall vote. Direct voting does not suffer from this distortion. Hence, in such cases, it is likely the better option. In my argument, I focus on case (a). This allows me to tackle the more challenging scenario to the superiority of direct voting, for which I argue. This setup also allows me to tackle option (2), tracking the majority across constituencies.

### *Simulation Results*

I use simulations to show that when the required average competence for direct to surpass indirect voting lowers within the majority of constituencies, in most cases, the general required average competence across all constituencies also lowers, i.e., the overall average crossover point is lowered.

Simulations confirm this pattern, assuming there is no consistent bias in favour of minority aligned-constituencies. That is, I assume constituencies to be statistically symmetric in the following ways:

1. Majority and minority candidate competence levels across constituencies are drawn from the same distribution,
2. The proportion of the local majority in each constituency is drawn from the same distribution.

### 1. Direct or Representative Voting? It Depends on the Type of Truth We Track

3. In each constituency, the competence added by majority voting over the averagely competent individual within that constituency is drawn from the same distribution (the added competence can be due to more competent individuals, a larger constituency size, or any other reason. Sources of crowd wisdom remain a black box).

Given this lack of bias, the average competence required of all voters for direct to surpass indirect voting, i.e., the overall average crossover point, decreases in the vast majority of cases. Specifically, out of 264,688 simulated cases, in 93.7% of cases (tested cases listed below), the average crossover point for multiple voters in each constituency was lower than the average crossover point in the case of a single individual per constituency.

Formally:<sup>34</sup>

$$(\sum_{G_i=1}^n r_{G_i}^{cross})/n < (\sum_{G_i=1}^n r_{I_{G_i}}^{cross})/n$$

For  $G_1, G_2, \dots, G_n$  constituencies, where  $r_{G_i}^{cross}$  is the average crossover point with multiple voters and  $r_{I_{G_i}}^{cross}$  with an individual, in consistency  $G_i$ .

Notably, the 6.3% of cases where the average crossover point increased with the addition of multiple voters mostly occurred when the overall majority proportion was small, which is expressed in a small majority of constituencies aligning with the overall majority.

#### Tested cases

The tested cases included both truncated normal<sup>35</sup> and beta distributions for key model parameters. Parameters for truncated normal distribution included:

1. *Candidate competence*: for the competence levels of majority and minority candidates within each constituency, truncated normal distributions were used with means sampled

---

<sup>34</sup>  $r_{G_i}^{cross} = \frac{2q_{G_i} - 1 - q_{G_i}r_{maj_{G_i}} - q_{G_i}r_{min_{G_i}} + r_{maj_{G_i}} + (r_{maj_{G_i}} + r_{min_{G_i}} - 1)\Delta r_{G_i}^s - \Delta r_{G_i}}{2(2q_{G_i} - 1 - q_{G_i}r_{maj_{G_i}} - q_{G_i}r_{min_{G_i}}) + (r_{maj_{G_i}} + r_{min_{G_i}})}$ , where  $\Delta r_{G_i}^s$ ,  $\Delta r_{G_i}$  are the added

competence due to crowd wisdom in constituency  $G_i$  over the averagely competent individual in constituency  $G_i$  (see equation 12 in the body of the text).  $r_{I_{G_i}}^{cross} = \frac{2q_{G_i} - 1 - q_{G_i}r_{maj_{G_i}} - q_{G_i}r_{min_{G_i}} + r_{maj_{G_i}}}{2(2q_{G_i} - 1 - q_{G_i}r_{maj_{G_i}} - q_{G_i}r_{min_{G_i}}) + (r_{maj_{G_i}} + r_{min_{G_i}})}$ , where  $q_{G_i}$ ,

$r_{maj_{G_i}}, r_{min_{G_i}}$  are the proportion of the local majority, the competence of the majority and minority candidates respectively within constituency  $G_i$ , defined similarly to  $q$ ,  $r_{maj}$ ,  $r_{min}$  (see equation 10 in the body of the text).

<sup>35</sup> The distribution was truncated to the interval  $[0, 1]$  for competence levels, and  $[0.51, 1]$  for majority size.

### *1. Direct or Representative Voting? It Depends on the Type of Truth We Track*

at five evenly spaced values between 0 and 1, and standard deviations sampled at five evenly spaced values between 0 and 0.5.

2. *Local majority proportion*: for the proportion of voters constituting the local majority within a constituency, truncated normal distributions were tested with means ranging from 0.51 to 1 (five evenly spaced values), and standard deviations ranging from 0 to 0.5 (also five evenly spaced values).
3. *Crowd wisdom gain*: for the additional competence the group decision has over the average individual, truncated normal distributions were used with means ranging from 0 to 0.5 and standard deviations ranging from 0 to 0.5, both in five evenly spaced steps.
4. *Alignment across constituencies*: the proportion of constituencies in which the local majority aligns with the overall majority was tested using 10 evenly spaced values between 0.51 and 1.

For each simulation, the parameter values for each constituency were drawn randomly from the corresponding distributions.

In addition to truncated normal distributions, beta distributions were also tested for parameters (1)–(3). The values tested for both parameters ( $\alpha$  and  $\beta$ ) included 0.1, 0.5, 5, and 10. For the alignment rate across constituencies (4), beta distributions were tested using alignment probabilities of 0.51, 0.55, 0.6, 0.7, 0.8, 0.9, and 1.

---

## 2. The Limitations of Deliberative Mini-Publics

---

### Abstract

In this chapter, I critically assess a seemingly plausible argument supporting the use of Deliberative Mini-Publics (DMPs) in existing democracies and show that it fails. I show that DMPs do not have a superior ability to track a truth relative to the interests or values of the majority of the population. I lay out two potential ways of tracking such a truth: (i) where each DMP participant aims to promote their own interests and values, and (ii) where each participant aims to promote the interests and values of the majority of the population. I show that neither way allows us to attribute clear epistemic advantages to DMPs over alternatives. I conclude that other arguments are needed to justify the use of DMPs in existing democracies.

### 2.1 Introduction

Deliberative Mini-Publics (DMPs) are a popular proposal for increasing democratic inclusiveness, offering advice to policymakers, and aiding decision-making. Many instances of such mini-publics have been put into practice (e.g., Norheim et al., 2021; Humanists UK, 2022; Casassus, 2022; UK Parliament, 2022).

In this chapter, I question the enthusiasm surrounding DMPs. I assess a potential epistemic benefit of such DMPs and show that this benefit is unlikely to materialize, leaving proponents of DMPs with a need to find alternative justifications. The chapter is structured as follows. In section 2.2, I provide a brief introduction to DMPs. In section 2.3, I survey existing criticisms of DMPs, which are consistent with seeing them as having an epistemic advantage over methods of decision-making that do not involve DMPs. In section 2.4, I question this purported advantage. I lay out an argument for why a potentially appealing explanation of their epistemic advantage fails.

## 2.2 What are DMPs?

DMPs typically follow a standard design: a randomly selected group of citizens, intended to represent the public, convenes to receive information on a given topic, engages in deliberation, and arrives at conclusions designed to influence decisions. These conclusions usually take the form of policy recommendations. A concluding majority vote among participants is sometimes conducted. DPMs are currently used to supplement existing democracies, by providing recommendations to policymakers or by providing their conclusions to the public, to aid with decision-making. It has also been suggested that DMPs can replace parliaments as legislators (e.g., Guerrero, 2024). Additionally, sometimes the conclusions of DMPs are voted on in referendums (e.g., Involve, 2018). In this chapter, I will focus on DMPs that provide policy advice to legislators or to the public within existing democracies, as existing DMPs do. Notably, DMPs are often combined with other institutions, such as subsequent referenda. My goal, however, is not to assess the epistemic performance of such combinations of institutions. Rather, I focus exclusively on the epistemic contribution of the DMP, as this will provide analytical clarity about the DMPs function.

Existing DMPs take different forms. Citizen juries usually include 12-16 participants, citizen panels 30-40 participants (Dampster, 2020; Willis, 2020; WMRCG, 2020; Casassus, 2021), citizen assemblies can include 100-150 participants (e.g., House of Commons et al., 2021), and deliberative polls can engage 500 participants. Most of these forums last for one to a few days, but some last for a series of weekends over several months (Setälä & Smith, 2018).

Since DMPs are not elected and hence cannot be controlled by the public through the threat of withdrawing electoral support, they are usually not thought of as a form of responsive representation. Philip Pettit, for example, defines responsive representation as a case where the representer's action are a causal response to the mindset of the represented, different from a form of representation where the fact that the representer acts in a certain way serves as evidence that the represented has a corresponding mindset (Pettit, 2010). The justification for DMPs, and more specifically, the claim that policymakers or the general public should assign substantial weight to the results of DMP deliberations in deciding which policy to implement or support, has most often been attributed to the second kind of representation, i.e., their ability to be a descriptive representation of the population. That is, to mirror the population's experiences and points of view, and hence the epistemic process they would have gone through had they had the time (Khoban, 2021; Mansbridge, 1999; Guerrero, 2014; Fishkin, 2009, 2013,

2018; Curato et al., 2021). As Fishkin puts it, they support “a hypothetical inference—these are the conclusions the population would come to if they could somehow consider the issue in depth under the same good conditions” (Fishkin, 2018, p. 322).<sup>36,37</sup>

DMPs have been enthusiastically advocated for in the philosophical literature (e.g., Guerrero, 2014, 2024; Fishkin, 2009, 2013, 2018; Landemore, 2013, 2020; Norheim et al., 2021), by governments, such as the governments of the UK and France (Humanists UK, 2022; Casassus, 2022; UK Parliament, 2022), in the media (Rice-Oxley, 2022; Casassus, 2022; Harvey, 2022), in analyses of existing DMPs (e.g., Fournier et al., 2011), and by different organizations such as the World Health Organization (e.g., WHO, 2021) and the European Climate Foundation (ECF, 2022).

Given the enthusiasm surrounding DMPs, and specifically their extended use as supplements in contemporary representative democracies, we may wish to ask whether they in fact have a clear and strong advantage when used in such a role. This is the question I tackle in this chapter.

### 2.3 Existing criticisms of DMPs

Some problematic aspects of deliberation have been noted, in a manner that does not rule out the potential usefulness of including DMPs in the decision-making process. Critics have pointed out, for example, that deliberation can sometimes lead to group polarization (Sunstein, 2009) and that despite the agreed-upon importance of equal participation in deliberation (Thompson, 2008; Mutz, 2008; Steiner et al., 2004), equal participation and influence often do not occur in practice (Jackman & Sniderman, 2006; Mendelberg & Oleske, 2000; Sunstein, 2002; Thompson, 2008; Setälä & Smith, 2018; Elstub et al., 2021). Yet, these are empirical contingencies that could, in theory, be solved using an expert moderator. Others have pointed

<sup>36</sup> A few other suggestions have been made for the type of representation that is to be found in DMPs (Mansbridge, 2003, Dryzek, 2010, Reuchamps et al., 2023). These include *discursive* representation, namely, representing types of discourse. Dryzek argues that “collective outcomes are responsive to the balance of competing discourses in the public sphere” (p. 24). It also includes *surrogate representation* – where representatives are not only committed to their own constituency, and their traits cause them to represent people with similar traits from a different constituency. Another suggested form of representation is *gyroscopic*. Mansbridge suggests this happens when representatives can act in ways we expect voters would approve of due to their shared traits.

<sup>37</sup> Notice that descriptive representation is necessary even if DMPs are used to provide advice to the general public. For instance, Warren and Gastil propose that mini-publics can function as “trusted information proxies” (Warren and Gastil, 2015, p. 562), supplementing the democratic process. Such an example is when mini-publics support democracy by providing recommendations directly to the public, as was done, for example, in the Oregon State Voters’ Pamphlet (Setälä & Smith, 2018). Warren and Gastil explain why the public should trust DMPs in such instances due to their ability to mirror the population’s interests, values, etc. This, again depends on their success in being descriptively representative.

out shortcomings of their descriptive representation. Jane Mansbridge, for example, argues that a random sample is likely to select less talented individuals than elections, and that using descriptive representation leads to the assimilation of minority interests to those of the dominant subgroup within each group (Mansbridge, 1999). Yet, she argues, such problems can be mitigated, and that these drawbacks are often outweighed by the benefits of having such a representation, such as enhanced communication with members of the public with shared properties and a better representation of interests.

Cristina Lafont adds that accepting the conclusions of a DMP without further deliberation among the wider public requires blind deference, which is incompatible with the democratic ideal of self-governance (Lafont, 2015, 2019). She concludes that DMPs should not replace deliberation in the public sphere, though she finds them useful as a tool within it (2019). Simone Chambers (2009) also argues that, although they have advantages, for instance due to “adding a citizen perspective to many questions and bridging the gap between lay citizens and policy experts” (p. 330), DMPs should not be thought to replace the need for deliberation within the mass public. She gives the example of DMPs that reached conclusions that later did not pass in referendums, suggesting deliberation in the latter stage deserves more attention (2009).

Additionally, analyses of specific citizen assemblies have noted vulnerabilities of the deliberative process. Fournier et al. (2011), for example, assert that it is important to ensure that the structure of the deliberation in DMPs does not provide undue influence to just a few people. Nonetheless, they come to favourable conclusions regarding the three study cases they investigate and the general usefulness of DMPs, claiming that participants “worked intensely and produced sensible decisions and recommendations” (p. 152). Elstub et al. suggest that democratising the agenda-setting stage of DMPs could help tackle problems that stem from having assemblies address complicated topics such as climate change (Elstub et al., 2021).

Despite the discussion of these potential weaknesses, this literature is generally supportive of the DMP project. Their popularity both in practice and in philosophical literature seems to suggest a widespread view that such weaknesses can be overcome or outweighed by DMPs’ strengths. In this chapter, I wish to more closely examine whether DMPs can truly be said to offer clear advantages, when used, as they currently are, within existing democratic systems. I examine this by looking at a seemingly plausible way of arguing for their advantage and showing that it fails.

## 2.4 The failure of a promising route to explain the advantage of DMPs

I examine one account of the unique advantage of DMPs. I show that this line of thought, which corresponds to common claims in the literature regarding the virtues of DMPs, does not withstand scrutiny. Hence, justification of an authoritative<sup>38</sup> role for DMPs in informing the decision-making of the public or elected representatives and other public decision-makers (e.g., civil servants in agencies entrusted with key policy decisions) must take a different form.

In the remainder of this chapter, I will defend the following argument:

*Premise 1:* If DMPs possess a clear and distinct advantage over alternative supplements to existing decision procedures within democracies, it is because their process generates some type of epistemic advantage over these alternatives.

*Premise 2:* One natural candidate for the type of epistemic advantage that DMPs have is in more accurately tracking what policy would best promote the values, interests, and goals of the general population (call this a population-relative truth). Moreover, a plausible interpretation of what would best promote these values, interests, and goals is the majoritarian one—what would command a majority vote if each were to vote according to an accurate picture of their own interests, values, and aims.

*Premise 3:* The two most natural ways of attempting to track this population-relative truth within DMPs are: (1) each person acts as an agent promoting their own values, interests, etc. (e.g., votes according to them, advocates for them in discussions and in authoring concluding documents), or (2) each person aims to promote the interests, values, and goals of the majority.

*Premise 4:* (1) is not a promising way to finding a clear advantage to the DMP process over alternatives.

*Premise 5:* (2) is also not a promising way to finding a clear advantage to the DMP process over alternatives.

*Conclusion:* This seemingly plausible line of thought to justifying the epistemic authority of DMPs fails, and proponents of DMPs need to find alternative means of justifying them.

---

<sup>38</sup> The term “authority” here is meant in a sense similar to Allan Gibbard’s description depicted here: “When we say that a person accepts something *on authority*, we mean that he takes someone else’s acceptance of it as his own reason for accepting it... Suppose you tell me it made no sense for Cleopatra to be angry at [a] messenger [who brought her bad news]. I am ignorant of history, perhaps, and confident that you know your history, and that you and I share the same basic norms for anger. In that case, I can take your normative reasoning as proxy for my own. I think that you are reasoning just as I would if I knew the facts.” (Gibbard, p. 174).

### 2.4.1 Justifying premise 1

Premise 1 states that if DMPs possess a clear and distinct advantage over alternatives, it is because *the process they employ* generates some *epistemic* advantage over these alternatives.

The idea of this premise is that potential advantages that DMPs may possess are epistemic. This is because the product DMPs produce is an epistemic product – they produce arguments, lists of recommendations, votes, etc. For instance, the West Midlands citizens' panel, which aimed to address the effects of COVID-19 (WMRCG, 2020), resulted in a list of priorities and policy guidelines. In the Oregon State Voters' Pamphlet, the DMP produced a pamphlet with recommendations to the public (Setälä & Smith, 2018). In such cases, which are characteristic of the results of many DMPs, any influence of DMPs over people's lives is through providing information that feeds into decision-making. Hence, for such DMPs to have an advantage, these epistemic outputs need to be, in some sense, better than the ones that already exist without the input of DMPs, or the ones that could exist if we gather alternative groups.

One might suppose certain advantages DMPs possess go beyond an epistemic advantage. One natural thought is that such a method provides greater inclusivity in deliberations than other methods because random selection picks out members of minority groups that are underrepresented in parliaments or in expert groups (e.g., World Bank, 2023). However, if such inclusivity is to support a claim to DMPs' advantage over alternatives, explanations of this type should also be translatable to some form of epistemic advantage this process produces. Think of different explanations for the benefits of inclusivity. A primary benefit is providing underrepresented groups with more influence over decision-making. Due to the small size of the participating group, providing only the participants of DMPs with larger influence due to their participation in the process would only increase exclusivity, and not achieve this goal.<sup>39</sup>

---

<sup>39</sup> This response applies to other potential non-instrumental considerations like non-domination (elections create a political class that effectively rules over others; elections skew in terms of who comes to hold power), and individual rights to participate in decisions that will affect one's life. When only a small group is elected to participate, we cannot claim these considerations hold. The small select group still dominates the larger group, unless there is an epistemic claim for producing accuracy in representation of the larger group. It also applies to instrumental considerations, such as encountering others in person in these ways plausibly affects how and the extent to which people care about, empathize with, and are motivated to work with others. When these only affect the small participating group, they cannot be claimed to benefit the larger population. Finally, one non-instrumental value of such random selection is respect for political equality, achieved through everyone having equal chance of being selected in the lottery. However, this type of a-priori equality, combined with the a posteriori lack of equality of only some participating in practice, is inferior to procedures that enable both types of equality, such as elections and referendums. They might also be inferior in value to procedures that do not guarantee equal participation, but are more epistemically accurate in the sense that they improve chances of equal consideration of interests in outcomes. Hence, the epistemic accuracy in reaching such equal consideration is again essential for justifying DMPs.

## 2. The Limitations of Deliberative Mini-Publics

The way the entire underrepresented group could have influence is by the epistemic outputs of the DMPs more accurately tracking what policies fit with the interests, values, viewpoints or goals of the underrepresented group. This is an epistemic advantage; an added accuracy in tracking something about underrepresented groups.

Another potential advantage could be solving polarizing societal issues. This would be an epistemic advantage as well, offering policy advice or other outputs that more accurately track a societal consensus that would exist under more ideal conditions. For this reason, it might give the general public and policymakers a reason to assign substantive weight to its input. Other advantages of inclusivity in DMPs would similarly require the epistemic outcomes of DMPs to be more accurate in tracking something about underrepresented groups or society than alternative decision-making structures.

Plausibly, the epistemic advantage attributed to DMPs needs to result from the process employed in DMPs. It could be that in specific cases, the results produced, for example, policy recommendations, sound promising for reasons independent of the process that produced them. However, to attribute epistemic authority to DMPs more generally, beyond whether they managed to convince a policymaker at a particular instance of a particular policy, something about their *process* should produce valuable epistemic outputs. Epistemic authority, in this context, means taking the mere fact that DMPs reached certain conclusions to in itself be a reason for accepting them. This resembles the notion of authority introduced by Allan Gibbard, who says: “When we say that a person accepts something *on authority*, we mean that he takes someone else’s acceptance of it as his own reason for accepting it” (Gibbard, p. 174). Justifications of DMPs are often of this type. For instance, as explained above, some argue that the DMP process allows its decisions to reflect what the general population would think were they well informed (e.g., Fishkin, 2018). Proponents of such claims suppose that both the public and policy makers should attribute importance to the outputs of DMPs regardless of whether they are independently convinced by them, due to this unique process.

DMPs include a randomly selected group coming together for a limited time to learn certain information presented to them, deliberate, and produce some such epistemic outcome. What epistemic advantage can be attributed to having such a process, in addition to the usual representative system? In particular, how would DMPs do better than having referendums, having a group of experts deliberate and produce similar products, or having randomly selected people each be paid to conduct independent learning, without deliberating with others, and then

vote? The answer needs to hinge on the epistemic advantage produced by the unique combination of random selection, short information briefing period, and deliberation.

Hence, if DMPs possess a clear advantage over alternatives, it is because *the process they employ* generates some *epistemic* advantage over these alternatives. Then, to assess whether the enthusiasm surrounding them is warranted, we need to examine whether the process they employ produces an epistemic advantage over such alternatives.

#### 2.4.2 Justifying premise 2

Premise 2 states that a natural candidate for the type of epistemic advantage that DMPs have is in more accurately tracking what policy would best promote the values, interests, and goals of the general population. I will call this an advantage in tracking a “population-relative truth”.

Before I defend premise 2, let us explore the concept of a population-relative truth a bit further. To see what characterises a population-relative truth, it might be most helpful to examine the types of truths that are not included in it. The population-relative truth that I claim DMPs can potentially have an advantage in tracking refers to a truth relative to the interests, values, etc. of the group composing that society at the moment in time when the random sample of participants was selected. This contrasts, for instance, with other types of truths that could be tracked, such as what would promote the interests, values, etc. of the members that society in ten years; what would promote the interests and values of society when taking into account future generations; or what would be morally right given universal moral principles the demands of which do not hinge on the particular interests, values and aims of a society’s population.

The population-relative epistemic advantage is limited. For some questions, tracking a population-relative truth might seem fitting, while for others less so. For some questions posed to DMPs, for instance, whether assisted dying should be legalized (e.g., Nuffield Council 2025), one might suppose more universal standards are required. We might suppose that what matters most in this case is universal individual rights, such as the right to self-determination in the face of grave suffering and medical professionals’ right to assist competent patients in their exercise of this form of self-determination. For other questions posed to DMPs, such as combatting climate change (e.g., Elstub et al., 2021), we might also suppose other non-population-relative factors related to future generations or global interests are of high importance. Hence, one might suppose that answers to such questions should not be assessed solely by population-relative truth tracking. Consequently, my argument entails that

alternatives to DMPs should most likely be used to tackling such questions, as DMPs will not have an epistemic advantage with respect to them. Such alternatives include, for instance, referendums or panels of experts, as discussed below.

There are various potential interpretations of what it means to track a population-relative truth, i.e., to promote the interests, values, etc. of the population at large. A plausible interpretation is the majoritarian one: tracking what would promote the interests, values, etc. of the majority of the population. This view, which was introduced in chapter 1, is prevalent in the literature on the epistemic value of democracy, for instance, in the jury theorem literature (e.g., Miller, 1986; Goldman, 1999; List and Spiekermann, 2016).

Notably, the majority interpretation does not necessarily require only taking the *interests* of the majority into account. It could be interpreted as what is correct given the values or viewpoints of the majority of the population. If a DMP conducts a majority vote at the end of the deliberation and this DMP is representative of the population, such votes could capture something like what is morally right, according to the moral conceptions or viewpoints of the majority of the population. Notice also that the majority-relative interpretation does not rule out, for instance, the vote reflecting what is best for the least well off, if according to the values of the majority, the best course of action is what is best for the least well-off.

In this chapter, I adopt this majoritarian interpretation. Of course, there could be other interpretations of what population-relative truths should track. Those might include always prioritising the interests of a particular societal group, such as the least well-off, regardless of the values held by the majority of individuals, or attaching weighted importance to different groups. For concreteness of argumentation, I will not focus on such alternative interpretations. Although, in section 2.4.4 I briefly explain why such interpretations are less relevant to the first method of truth tracking I survey, and in section 2.4.5 I will briefly explain why the second method of truth tracking I survey is also unlikely to give DMPs an epistemic advantage in tracking such interpretations.

Having clarified the concept of a population-relative truth, to justify premise 2, let us see why tracking a population-relative truth is plausibly a more promising justification for the potential epistemic advantage of DMPs than several leading alternative explanations.

First, the population-relative interpretation of their advantage is compatible with prominent justifications of DMPs in the literature. As explained, DMPs are often taken to be a representative sample of the population, representing the experiences and points of view typical

## *2. The Limitations of Deliberative Mini-Publics*

of members of the general population (Khoban, 2021; Mansbridge, 1999; Guerrero, 2014; Fishkin, 2018; Curato et al., 2021). This suggests that an advantage DMPs are meant to have is the ability to collectively possess evidence about the general population – typical experiences, typical interests, typical values and points of view that exist in this population. This is thought to be the source of their ability to be descriptively representative. Hence, DMPs would seem to be better placed than certain other entities in finding answers to questions that hinge on such properties. Such questions might include what policy is most compatible with the interests of the population. A common question DMPs are thought to be able to answer is what the population would want if they were well informed (Fishkin, 2018). The answer to this question hinges on properties of the general population, i.e., it is a population-relative truth.

Since DMPs are randomly selected with the intent of representing the typical properties in the population, the assortment of individuals is meant, by definition, to be typical of the general population in their ability to handle different tasks. That is, some will be better than average, some worse, but the representativeness of the group implies that their mean abilities are typical. We might ask, then, for what topic does such a group produce an epistemic advantage over other groups whose mean is beyond what is typical in some way. One plausible answer is: when the subject matter studied is the group itself, and hence what is typical serves in itself as helpful evidence. This, again, suggests that DMPs' epistemic advantage arises when studying questions whose answer hinges on properties of the members of the general population, such as their interests, values, or goals.

Some alternative arguments have or could be made for the advantage of DMPs. I will now go over several key arguments and explain why these are either translatable to an advantage in tracking a population-relative truth, or we can find alternative supplements to representative democracy that will perform better in achieving this advantage, such as groups with particular expertise or referendums. This discussion will help justify premise 2, as it shows epistemic advantages of DMPs are likely within the realm of tracking population-relative truths.

An epistemic advantage prominently attributed to DMPs is that deliberation among a randomly selected group enhances competence due to the cognitive diversity of such groups (Landemore, 2013). This suggestion also relates to literature in standpoint epistemology, where it is claimed that people who occupy different standpoints in society encounter different types of evidence, providing them with unique epistemic advantages (Hartsock, 1983; Collins, 1990). This, again,

## 2. The Limitations of Deliberative Mini-Publics

leads to helpful diversity within the population, and has also been linked to the epistemic advantage of DMPs (Guerrero, 2024).

However, the idea that a group mirroring the general public is helpfully diverse cannot provide a strong justification for DMPs as an epistemic authority. Rather, assuming such diversity might provide the general public with an epistemic advantage over DMPs. Different jury theorems and other studies in Machine Learning illustrate that diversity in the population might be better utilized for reaching good decisions via a direct vote of the entire population (e.g., de Condorcet, 1785; Owen et al., 1989; Ladha, 1992; Kunapuli, 2023). They show that if people are on average better than chance at solving problems, then direct democracy with vote independence or low vote dependence would lead to a relatively competent majority vote.<sup>40</sup> The small number of participants and the fact that this type of interpersonal diversity is diminished during deliberation (Dietrich & Spiekermann, 2022) do not allow for such results within DMPs.

The diversity argument for DMPs has been most prominently suggested by Hélène Landemore (e.g., Landemore, 2013). Landemore offers a formal explanation for why the diversity in the population is helpful for decision-making, and can trump the ability of experts. However, when examining Landemore's argument for such cognitive diversity, more problems arise. Landemore relies on a formal model presented by Hong and Page to illustrate the advantages of cognitive diversity within a group (Hong & Page, 2001; Page, 2007). This model suggests that diversity can outperform ability under highly specific and idealized conditions. Landemore's use of this model has therefore been criticized due to the model's limited applicability (e.g., Quirk, 2014; Grim et al., 2019). Moreover, to reach the conclusion that the diversity in the population trumps the ability of experts in dealing with problems, this model specifically assumes that people are, on average, better than chance at problem-solving, by postulating participants' intelligence (Page, 2007). As explained, together with the diversity claim, these assumptions suggest direct voting, such as in referendums, could yield competent decision-making (e.g., de Condorcet, 1785; Owen et al., 1989; Ladha, 1992).

Consequently, justifying DMPs based on diversity is not a promising route. Justifying randomly selected assemblies requires demonstrating clear and strong advantages over involving the entire population in the decision-making process, as the latter approach has the

---

<sup>40</sup> For complete vote independence, jury theorems have shown infallibility of the electorate (de Condorcet, 1785). However, other jury theorems, as well as machine learning literature have also shown high reliability of the electorate for high voters competence combined with low vote dependence (e.g., Ladha, 1992; Kunapuli, 2023).

significant advantage of being more inclusive. Proponents of DMPs would argue that the organized deliberation, which cannot happen in the same way within the general public, also increases competence. However, without a convincing argument that the epistemic advantages of DMPs cause it to outperform direct voting, the inclusivity of direct voting gives it an edge. Hence, this explanation of a non-population-relative epistemic advantage of DMPs fails, as the referendum alternative is likely better under the assumptions required for attributing such an epistemic advantage to DMPs.

Another possible challenge to the claim that DMPs are best suited to track population-relative truths (premise 2) is claiming that DMPs are particularly suited to tackle moral questions beyond population-relative ones, since the general population possesses a perspective that biased experts lack.<sup>41</sup> To the extent that the moral question at hand is not a population-relative truth, the benefit of DMPs in such moral questions will not come from having insight into the interests, values, experiences of this particular population. Rather, it will need to come from an ability to track independent moral truths. Admittedly, there might be bias among experts with regards to such questions. However, if we suppose the diversity among the population sufficiently cancels out biases, leading to a low vote dependence, then the claim regarding the advantages of direct voting resurfaces.

Moreover, there are reasons to think both DMPs and expert groups will be biased. In that case, I claim that the most plausible reason to prefer the bias that exists in the general population is if we aim to track a population-relative type of moral truth. Hence, epistemic advantage of DMPs is once again linked to tracking population-relative truths. Let me explain. Certain types of biases with regards to moral questions are known to exist in the general population. For example, in literature discussing juries, another group in which random (initial) selection plays a part, a prominent claim is that many jurors believe rape myths, and that those plausibly influence jury decision making (e.g., Laverick, 2025). It has also been argued that the public, and consequently members of juries, experience a bias called “system justification”, where they are motivated to defend the status quo of social order (e.g., Jost, 2020). This bias is, research suggests, less prominent in people with more extensive knowledge regarding societal inequalities (Saguy et al., 2008; Samuelson & Zeckhauser, 1988). Hence, particular groups of experts, which are non-typical in some respect, will plausibly do better than a DMP with respect

---

<sup>41</sup> Alex Guerrero, for instance, writes “those who are experts about non-moral issues are not thereby experts about moral ones and might share various biases or perspectives with respect to moral issues.” (2024, p. 222). Although, Guerrero argues particularly for a lottocratic model where DMPs function as legislators.

to at least some relevant biases. For example, for both these biases, a group particularly selected for their information level on relevant societal issues will likely do better.<sup>42</sup> Additionally, groups that possess an above-average ability or training in composing and critically examining arguments about morality may plausibly be better able to observe biases, also within their own arguments, and plausibly do better overall. To further reduce biases, such a group can be particularly selected to represent a variety of views on the relevant moral topic. For instance, for examining a bill on assisted dying, it could be selected to include both advocates and people against these ideas. It could be that the selection of the particular type of opinion diversity in experts' views might itself incorporate certain biases, as the decision of what counts as diverse might not be value-neutral (see Garnett, 2016). However, that type of bias is not clearly worse than the biases that exist in the general population. It seems the most plausible reason to favour the latter type of biases is if we are aiming to track a population-relative moral truth. That is, find out what is compatible with the values, viewpoints, interests of the population, rather than tracking an independent moral truth. Hence, an epistemic advantage attributed to DMPs over expert groups in tracking moral questions would indicate these questions are translatable to population-relative truths, aligning with premise 2.

We might wish to challenge premise 2 by considering additional types of epistemic outcomes that could be beneficial for the population's decision-making process, where DMPs might be beneficial. However, for some such outcomes, alternatives to DMPs will again be better, and for others, any advantage to DMPs will be due to tracking a population-relative truth. For instance, producing high quality, logically sound arguments, providing information supported by historical or scientific facts, etc. However, for such benefits, plausibly, there would be groups better suited to produce these epistemic goods than the randomly selected group. For instance, for logically sound arguments, we might want philosophers or logicians, for scientifically rooted arguments, scientists, and so on. Some other potential advantages of DMPs are translatable to tracking population-relative truths. For instance, one may claim the deliberation among the randomly selected group produces arguments that are responsive to reasoning. This would only clearly make the *randomly selected group* deliberation better than that of alternative groups if the reasoning we would like the arguments to be responsive to are ones this group is best suited to produce, due to being representative of experiences, values,

---

<sup>42</sup> Recall I am specifically addressing DMPs as advisory bodies here, who generally tend not to last for very long durations, hence, not enabling citizens to become experts.

interests of the general population. That, again, leads to having an advantage specifically in tracking population-relative truths.

One might object, however, that DMPs do not possess a single epistemic advantage, but are the best compromise among different potential advantages. However, to claim that the epistemic outcomes of DMPs have epistemic authority due to the process that produces them, and hence, that people and policy makers should listen to them, we should be able to identify the precise set of epistemic benefits this process is able to produce. As long as we can clearly define that, it is not clear why we would take a randomly selected group to better produce it than a group particularly selected for specific combinations of skills. Additionally, if different benefits are needed for tackling different questions, we can supplement existing democracies with different types of advisory groups for each question at hand.

So far, I have explained why an advantage in tracking a population-relative truth is compatible with common claims in the literature about DMPs and how it plausibly relates to the concept of random selection of participants. I also explained why for alternative explanations of DMPs' epistemic advantage, either they are actually translatable to tracking population relative truths, or, if not, then referendums or group of experts will plausibly do better. Hence, it seems that the most promising route for finding a clear advantage for DMPs is as being superior in tracking population-relative truths, namely, inquiries that hinge on the interests, values, opinions, perspectives, or societal experiences or members of that particular society. Notably, this refers to those characteristics of that society at the particular time when the randomly selected group was chosen, as that is as far as its claim for a descriptive representation of the population can go. The randomly selected group's advantage in deliberating on such matters emerges from their capacity to offer relevant evidence – their personal experiences, values, and viewpoints.<sup>43</sup>

Hence, it seems premise 2 is justified. That is, that a primary type of epistemic advantage where DMPs seem to potentially have promise is in tracking a population-relative truth. As explained, in this chapter, I focus on the majoritarian interpretation of such a truth.

#### 2.4.3 Justifying premise 3

According to premise 3, two key ways of tracking a population-relative truth include (1) each person acting as an agent promoting their own values, interests, etc. (e.g., voting according to

---

<sup>43</sup> This notion is supported by claims in the literature such as the following: "By bringing people from widely divergent groups, backgrounds, and experiences together, and having them talk and engage with each other, reflective lotterocracy holds out the promise of creating genuine understanding, empathy, and political community that embraces the full spectrum of experiences and interests of that community" (Guerrero, 2024, p. 229)

them, advocating for them in discussions and in writing inputs to, or conclusions of, the DMP), or (2) each person acting as an agent promoting the interests, values, etc. of the entire population. Given the majoritarian interpretation I adopt, that entails promoting the interests, values, etc. of the majority of that population.

Approach (1) to tracking a population-relative truth (what best aligns with the populations' values, interests, and aims) is for each person to be tasked with tracking what, in the previous chapter, I termed "relative truth". For, on this approach, each person's decisions are considered correct according to different criteria—criteria determined by each person's own interests, values, and goals. The idea of approach (1) is that if DMPs are descriptively representative of the population, by tasking each individual in a DMP with gathering information, arguing and voting according to their own interests, values, and aims, the collective decisions of the DMP as a whole will track the population-relative truth of what is best given the values, interests, and aims of the majority of the population.<sup>44</sup>

In contrast, approach (2) tasks each individual in the DMP with gathering information on, deliberating about, and then voting for the option that best promotes the values, interests, etc. of the majority of the population. In a sense, this means each member of the DMP is tasked with tracking a specific type of what, in chapter 1, I called a "universal truth", since the decisions of all participants are considered correct according to the same criterion. In this context, universal truth-tracking simply implies that the same criteria are applied to assessing the correctness of the decisions of different participants within the same DMP, mirroring the same society. That is, the same decision on a policy, for instance, will be either correct for all participants of that DMP, or incorrect for all, depending on whether it is what is correct given the values, interests, etc. of the majority of the population. Note that, despite being "universal" in this sense, the truth is also population-relative, in the sense that the same decision on a policy

---

<sup>44</sup> The interpretation of population-relative truths as what promotes the interests, values, etc. of the majority seems a particularly plausible interpretation of what individual-relative truth tracking will be best suited to track. For instance, suppose a DMP makes their decisions by conducting a majority vote in the end of the deliberative process. If each person votes according to their conception of their own interests, values, and aims and their votes indeed accurately track these interests, values and aims, then a majority vote will reveal the interests of the majority of the DMP. If, furthermore, the DMP is representative of the population, then the majority vote in the DMP will track what would command majority support in the population at large, if they were voting in line with what would, genuinely, further their interests, values, and aims. It might be that for other decision-making means, such as concluding the deliberation only when there is consensus, individual-relative truth tracking could overall track something other than the majority. However, if we take the relative-truth tracking idea that each person continuously advocates for their own interests, values, etc. seriously, then consensus would be extremely hard to reach, and could only be reached on policies that are good according to what everyone considers to be their interests.

could still be considered correct for participants of a DMP in society A but not for participants of a DMP in society B. In what follows, I will simply refer to this as each participant “tracking a population-relative truth”. Since each participant themselves aims to track the population-relative truth, it is clear why this form of truth tracking is also a plausible way for the group as a whole to track the same population-relative truth.

To summarize, what defines population-relative truths is that correctness hinges in some way on the properties of the population (in particular, their interests, values, and aims). For both options for what individuals should track, the individual-relative and the population-relative ones, these properties of the population are being taken into account as part of the decision-making process. Consequently, both will track a form of correctness that hinges on those properties. Hence, both seem like potential ways to track population-relative truths, and premise 3 is established.

I will now justify the focus on these two options for truth tracking rather than on the possibility of so-called mixed-motive voting (Wolff 1994). Mixed-motive voting occurs when some participants engage in individual-relative truth tracking, and others in population-relative truth tracking. First, some cases that are initially perceived as mixed-motive voting will actually be cases of individual-relative voting. For instance, if people vote according to their own values, but to do so, they take the interests of others into account, that is still an individual-relative voting according to my definitions, as it is relative to values of individuals. Hence, the arguments below regarding the lack of advantage of DMPs for individual-relative voting will apply to such cases.

Second, mixed-motive voting is not a great way to track what would promote the values, interests, etc. of the majority (Wolff 1994). Wolff (1994) shows that if some voters in a democracy vote according to their own interests, and others according to what is good for all, then it is possible that the overall vote neither reflects what is in the interests of the majority, nor what the majority believes to be what is good for all. This result undermines the value of mixed-motive voting as a way of reliably tracking population-relative truths. One might argue that although it is not the best way of truth tracking, mixed-motive voting is likely to occur in practice. However, in that case, expert groups that could be better taught not to conduct mixed voting will likely outperform DMPs that engage in mixed voting in tracking population-relative truths.

Finally, when mixed-motive voting is common, DMPs have important disadvantages also compared to direct votes. Particularly, mixed-motive voting in DMPs is likely less legitimate than in the general population. This is because in mixed-motive voting, the interests, values, etc. of the general population are unequally considered. Some are overrepresented in their own individual-relative voting and in the votes of others that vote according to the group, while others are underrepresented. While that may be legitimate when the people themselves choose whether to prioritize their own interests and values or not, it seems far less legitimate if unelected representatives conduct such an unequal treatment.<sup>45</sup>

I therefore focus on two options – either the goal we attribute to all participants is to track an individual-relative truth, or the goal we attribute to all<sup>46</sup> is to track the population-relative truth that the entire group is meant to track. As premise 3 states, these are two key ways of tracking population-relative truths.

#### 2.4.4 Justifying premise 4

According to premise 4, the first type of decision-making, which tasks each individual with tracking their individual-relative truth, is not a promising way to finding a clear advantage to DMPs. To justify premise 4, I will make the following argument.

*Premise A:* Either participants are already good at tracking what will best serve their own interests, values, and aims prior to the start of the DMP process or they are not.

*Premise B:* If they are already good at it, it is likely better to have direct votes of the population than DMPs.

---

<sup>45</sup> This criticism also applies to a case where, rather than mixed voting, where each voter chooses to *either* track their own interests, etc. *or* to track those of the general population, each person tracks both, deciding how to balance between both considerations. Some might take DMPs to be best placed with allowing people to balance and modulate between the questions of what would be good for them and what would be good for others. One might learn, for example, that although policy A would be better for them, it would actually be incredibly harmful for others, and so even if there were more people like them, they should still decide not to go with policy A. However, such an intra-personal mixed voting leads to the same problem that interests and values of different people are considered unequally in the process. One participant may choose to prioritise their own interests, or weigh them more heavily against others, while others choose not to do so. If we want an objectively preferable weighing of interests, experts have an advantage. Alternatively, if we want each to be allowed to decide how much to weigh their own interests, that would be more legitimate in an election of the entire population than when decided by unelected representatives.

<sup>46</sup> Attributing such a goal to participants does not necessarily entail they all actually use this rational when voting. However, it can be helpful in assessing how competent they can be relative to this goal.

*Premise C:* If they are not already good at it, the particular structure of DMPs is not the most promising way to improve this ability and hence is not the best route to competence in individual-relative truth tracking.

*Conclusion:* Neither assumption about people's ability to track what best promotes their interest, values, and aims leads to this form of truth tracking being a promising justification of DMPs over alternatives.

Premise A is straightforward. Let me now discuss premise B. Suppose participants of DMPs are already good at tracking what would best align with their interests, values, and aims before the DMP process starts. As explained, DMPs are meant to be descriptively representative of the population. Hence, participants already being competent in tracking their interests, values, etc. is an indication that the entire population is good at doing the same.

In this case, direct voting will plausibly be more advantageous also in tracking population-relative truths. Let us see why. First, the general population is competent at tracking their interests, values, etc. Second, to the extent there is diversity among their views within the DMP, which is a major justification for its necessity, such diversity reflects diversity in the general population. Of course, it could be that all people share the same values and only make decisions according to those. But then DMPs would lead to similar results as only having one or a few people make decisions, and hence the randomly selected group will not have an epistemic advantage. Therefore, this condition likely entails high competence in tracking interests, values, etc. and low vote dependence. However, under these two conditions, majority votes have been shown to function quite well, as explained above. This has also been shown particularly when tracking individual-relative truths, and where each person votes according to their own interests, values, etc. (Miller, 1986; List & Spiekermann, 2016).<sup>47</sup>

In chapter 1, I also established that when voters are better than chance at tracking such individual-relative truths, a direct vote usually outperforms a representative vote. For this claim, I also do not rely on independence or low vote dependence, as explained in the previous chapter. This is also an indication of a certain level of competence of individuals voting in referendums, and the ability of referendums to be epistemically advantageous as supplements of representative democracy. At the very least, this line of argument challenges proponents of DMPs to show the

---

<sup>47</sup> Proponents of DMPs will usually turn to the deliberation conducted in them as providing an important advantage. However, within this type of truth tracking, where each participant aims to track their own interests, the deliberation within DMPs does not seem particularly helpful in raising competence, as will be explained in the justification of premise C.

superiority of a DMP over a referendum, rather than simply over representative democracy. As direct votes of the entire population are more inclusive than DMPs, if they reach high competence under the same conditions needed for DMPs' success, we may wonder why we should opt for DMPs.

Let me now justify premise C. Suppose participants are not good at tracking their own interests, values, etc. Then, it is unlikely that the DMP will significantly improve their ability to do so, and other means of improving their ability will likely be a more effective way of engaging in individual-relative truth tracking. Let me explain.

DMPs foster deliberation among a randomly selected group, which is meant to represent a diverse assortment of interests and viewpoints. The way the deliberation of such group in DMPs usually goes, and what it is usually praised for, is that people, coming from different communities and worldviews share their views with one another, each discussing their own thoughts, interests, relevant life experiences, etc.

Now, in some circumstances, it may be that understanding others' experiences, viewpoints and interests aids in establishing which policies truly align with one's own interests, values, and goals. For example, if a person needs to establish which position on assisted dying best accords with her values, and according to those values, it is important how others will view the opportunity to choose assisted dying (e.g., whether they see it as essential to their autonomy), then she will need to learn others' perspectives on the matter. However, even for a values based decision-making on assisted dying, if a person deems the testimonies of elderly or terminally ill people as more relevant, or if they aim to decide based on universal moral values rather than the opinions that characterise this particular society at a given moment in time, other processes, e.g., discussions with selected groups rather than a random selection of individuals, might well be more useful for figuring out what best fits their own values. Also, others will likely discuss the topics based on their own differing moral convictions that the person will often not share. For instance, hearing a religious person discuss the topic from a religious perspective will likely not help an atheist make up their minds about it given their own values. More generally, the particular input received in the deliberation of a randomly selected group will not be the most productive way for each to learn about what would best serve their own interests, values, and aims.

Let us look at another example of a DMP convening to provide recommendations for COVID-19 regulations (e.g., WMRCG, 2020). Suppose one is a small business owner whose most

## 2. The Limitations of Deliberative Mini-Publics

relevant interests include how regulations will affect their business over time, and one aims to decipher which policies about employee protections and particular types of governmental support would aid them. To promote one's own interests, it is plausibly more useful to discuss possibilities with other small business owners and to examine together particular facts relevant to their situation than to hear the testimonies of other interest groups, such as elderly people or large business owners. While these might be helpful in reaching a compromise among different interests, that is irrelevant for this type of voting.

Hence, to make decisions based on a clear view of what would serve one's own interests, values, and goals, a process of learning about the topic from the point of view of a diverse group of others who do not share the same interest, values, etc., may well be confusing and counter-productive. At the very least, we could find alternative structures better suited for this goal. For instance, individual learning without deliberation or deliberation among smaller groups of people who share similar interests and values, may often promote such goals better. Hence, on an approach that tasks each person with tracking their individual-relative truth and then taking the majority opinion to represent the view of the population, there will often not be an advantage to the particular combination of characteristics of DMPs.

One might argue that the information provided to participants at the start, and the ability to process it through deliberation, are the beneficial parts when tracking individual-relative truths.<sup>48</sup> However, those are achievable in alternative structures that do not include a deliberation among a randomly selected group, such as deliberations among similar interest groups and individual learning followed by a vote. Hence, this claim does not in itself justify the DMP structure.

Additionally, there are inherent tendencies towards bias in the portrayal of information in DMPs. This bias causes the result of the learning phase to often shape participants' opinions in a way that reflects the values and opinions of the DMP organizers, rather than aiding them better pursue their own interests and values. I will now explain this in more detail.

### *Biased information portrayal in DMPs*

DMPs that serve as supplements of a representative democracy, advising policy makers, for instance, do not usually last long. A long duration would be hard to justify given their limited

---

<sup>48</sup>When justifying lottocracy, Alex Guerrero, for instance, attributes importance to deliberation in aiding participants process the learned materials (Guerrero, 2024).

role and the cost to both participants' personal lives and public funds. The need to provide concise information on complex topics due to their short duration and the non-expert nature of participants causes information to often be portrayed in a biased way, as I will show. Consequently, the learning process participants go through is substantially shaped by the opinions of the organizers and experts.

The nature of many DMPs necessitates presenting information in an easily digestible manner. This introduces the risk of biasing participants.<sup>49</sup> There could be many causes of potential bias in information portrayal. It may arise, for example, from omitting intricate yet critical information, neglecting to communicate uncertainty regarding specific data, failing to present all possible alternatives to a given action, framing discussion questions with bias, and phrasing information in a way that appeals to known psychological cognitive biases. To illustrate this vulnerability, I will now discuss the existence of such biases in several DMPs that were considered successful (Fournier et al., 2011; Warren and Gastil, 2015; Setälä & Smith, 2018; Participedia, 2021; Norheim et al., 2021; Casassus, 2021).

One example of not listing all solution options can be found in an information booklet distributed at the outset of an online deliberation panel in British Columbia: "Mandating quarantines can be challenging to enforce. There might be technologies enabling quarantine enforcement, but public acceptance remains uncertain" (BC-Data-Deliberation-research-team, 2020, p. 8). This phrasing primarily highlights the role of invasive technologies in enhancing quarantine adherence. However, studies show that better public education about quarantine effects and fostering trust can also increase public cooperation (e.g., Webster et al., 2020).

A further example concerns a discussion question in a French citizen panel aimed at promoting vaccination (Casassus, 2021). Rather than allowing participants to debate the best approach to the handling the pandemic, the panel's objective presupposed that vaccination is the correct option. However, the panel was intended to be representative of the population. Given that the government initiated the panel due to their awareness of the population's mistrust of the vaccine, assuming a representative panel's agreement with vaccination is problematic. To ensure the panel's recommendations genuinely mirrored their opinions and, consequently, those of the public, alternative options should have also been debated.

---

<sup>49</sup> This is not needed in other methods such as direct votes, and will probably be less problematic for other deliberative groups, such as groups of experts.

Another case of omitting solution options is evident in the British Columbia assembly on electoral reform (CAER, 2004a, 2004b). For instance, the “Assessing electoral systems” fact sheet notes: “Voter turnout rates are declining in most western democracies. Aspects of an electoral system that could impact participation include clarity, simplicity, and choice” (CAER, 2004b). Notably absent are other significant factors contributing to low voter turnout in Canada, such as public dissatisfaction with politicians and political institutions (Elections Canada, 2000).

A different instance involves information phrased in a way that triggers cognitive biases. Short-duration DMPs are especially susceptible to these biases as participants encounter a limited range of phrasings for relevant information and lack the knowledge to critically assess them. For example, the West Midlands citizens’ panel, which aimed to address the effects of COVID-19, participants received six information cards with statistics on relevant topics. Three cards were entirely positively framed, portraying only good effects of COVID-19 (e.g., “1 in 5 expressed intent to cycle more,” “Almost half plan to increase walking,” WMRCG, 2020, p. 83), while three were entirely negatively framed (e.g., “Less than half of parents without higher education feel confident guiding home learning,” “Only 23% of the most deprived schools have online learning platforms,” WMRCG, 2020, p. 82). When participants were asked to choose three topics as priorities, they chose the three negatively framed topics (WMRCG, 2020, p. 26). This may indicate the influence of the framing effect on their responses (Druckman, 2001).

In addition to information packets, participants are often exposed to expert testimony. A study of DMPs reveals that information about the witnesses and experts involved in DMPs is often lacking (Lindell, 2011, p. 20), making it difficult to evaluate biases in their testimonies. Generally, the same problematic tendencies identified in information packets can be expected to occur in expert testimonies. For instance, it is challenging for experts not to inadvertently convey their personal opinions about certain facts through specific framing. They may also fail to communicate the uncertainty surrounding their claims or fail to mention which statistical tools can be used to assess the strength of their claims. Furthermore, differences in conveying uncertainty among different experts could lead some claims to mistakenly sound more certain than others. Additionally, experts often behave as goal-oriented motivated reasoners. That is, they tend to favour facts aligning with their own agendas and opinions (Chong 2013, Druckman 2012). These biases may influence their testimonies. This is especially problematic if the process is too short for participants to effectively challenge experts’ testimonies.

## 2. The Limitations of Deliberative Mini-Publics

The examples above illustrate that DMP decisions are often influenced by the perspectives of organizers, who prepare information packets, and experts. While this might not be universally true for all DMPs, the brief gatherings of randomly selected non-experts make the process vulnerable to such influence. This is because conveying complex issues in nuanced ways within a short timeframe is inherently challenging. Consequently, the information stage of DMPs also does not seem to provide a strong reason to suppose they are well suited to advance people's ability to track their *own* interests, values, and aims rather than those of organizers and experts. Also, as mentioned, even if an information stage could enhance such abilities, in individual-relative truth tracking, there is no clear reason to attach it particularly to a DMP structure that also involves deliberation among a randomly selected group.

Finally, I wish to briefly explain why this approach to tracking population-relative truths is unlikely fruitful in tracking non-majoritarian interpretations of such truths, for instance, what is best for the least well off. It seems that if we wish the decision-making to reach recommendations that are best for the least well-off, then people should vote according to the least well off, rather than according to their own interests, values, etc. Otherwise, poor accuracy is guaranteed. While some may take the advantage of DMPs to come from enabling the group's judgments to be responsive to the minority's complaints, leading to a better tracking of some non-majoritarian truth, those should only be taken into account by each participant if those participants are not each tasked with tracking their own interests and values. Unless such reason responsiveness can be captured by each person still voting according to their own values, in which case the arguments above still hold. Hence, such potential advantage fits better with the type of task discussed in the next section.<sup>50</sup>

Hence, it seems the conclusion of our sub-argument, i.e., premise 4, follows. No assumption about people's ability to track what best promotes their interests, values, and aims leads to an approach that tasks each individual with doing that being a promising route towards establishing DMPs' epistemic authority on population-relative truths.

---

<sup>50</sup>Also, at the end of the DMP process, generally, either a unanimity is reached, or there should be some way to end the process, such as a majority vote. Generally, if each person is advocating for their own interests, values, etc., then the majority vote would track what a majoritarian view aims to track. Also, if we take the idea that each advocates for their own interests, etc. seriously, reaching a consensus would entail that the result is right according to the interests, values, etc. of everyone, including the majority. Hence, it again aligns with the majoritarian interpretation. It seems that wishing a DMP group to track non-majoritarian population-relative truths would entail asking for each to track what is best for all, rather than for themselves. This comes with its own set of problems discussed in the next section.

### 2.4.5 Justifying premise 5

According to premise 5, tasking each DMP participant with tracking a population-relative truth, i.e., where each participants tries to promote what they believe is right for a majority, given the whole population's values, interests, etc., is also not a promising way to establishing the advantage of DMPs. To justify this premise, I present three arguments.

#### 2.4.5.1. *Argument 1: No DMP size allows participants to perceive an accurate representation of the population*

Ensuring that deliberation enhances people's ability to track what aligns with the majority's interests, values, and aims requires a deliberation process where members encounter testimonies that are sufficiently representative of the general population. I argue that participants are unlikely to receive the right type of exposure, making their decisions significantly inaccurate according to the discussed standard. This is due to an inherent tension between small and large-sized DMPs, causing neither to sufficiently allow the individual participant to be exposed to an accurate descriptive representation of the population. Particularly, the smaller the DMP, the less overall descriptively representative it is. The larger the DMP, the less likely it is that the deliberation will allow participants to accurately perceive the relevant properties of all other participants.

Let me start with small DMPs. Random selection of a small-scale panel tends not to precisely mirror the proportional ratios of various groups within the population. Consequently, to increase the chance of sampling relevant properties, such DMPs often employ stratified random sampling. This form of sampling divides the population into sub-groups according to relevant properties and then randomly samples from within each group (e.g., Lacelle-Webster, 2020; Scully, 2020). However, the challenge with stratified random sampling is that it requires organizers to pre-determine the traits that should be sampled. But it is not clear which traits are jointly sufficient to mirror the population (Morone & Marmore, 1981).

Furthermore, one reason why especially large groups are needed for descriptive representation is that it requires representing not only individual traits but *combinations* of traits present in society. Sampling combinations of traits that are found in society requires many participants. For example, even in the limiting case where only five characteristics matter for descriptive representation (e.g., religion, age, gender, ethnicity, and financial status), each restrictively limited to four possibilities, we would already need  $4^5 = 1024$  participants to represent all combinations.

If we adopt the reasoning behind descriptive representation, only sampling individual traits without accounting for trait combinations is insufficient. As mentioned above, descriptive representation is often justified by the assumption that if a randomly selected individual comes to hold a certain opinion it is “evidence that members of the political community who share contextually salient characteristics with that individual would also come to have those views, had they gone through the same experience” (Guerrero, 2014b, p. 159). This stems from the assumption that individuals with similar traits share similar worldviews. However, it is more plausible that such shared worldviews hold for individuals with the same trait combinations. People with specific trait combinations share distinct experiences that those lacking the relevant combinations do not have. This argument aligns with intersectionality literature (Crenshaw, 1990, 2005). Sally Haslanger, for instance, notes that “the gender norms for Black women, Latinas, and White women differ tremendously, and even among women of the same race, they differ depending on class, nationality, sexuality, [and] religion” (Haslanger, 2012, p. 9). Thus, combinations of multiple traits likely lead to specific worldviews. Consequently, to accurately represent the worldviews prevalent in society, combinations of traits need to be represented, necessitating a larger number of representatives.

Importantly, this may lead to drastic changes in the potential outcomes of DMPs. Suppose, arbitrarily, as in the example above, that 1024 participants were required to represent all relevant experiences in society. Then DMPs that include less than 500 participants, which is most of existing and proposed DMPs, would not even represent half of the viewpoints in society. Hence, even if participants can perceive the interests, etc. of all other participants, it will not allow them to have an accurate perception of those found in society. This will likely not allow conceiving of which policy aligns with the interest, values, etc. of the majority. In cases where the majority is small, this problem will be even more severe.

The limitations of small DMPs in representing the broader public might lead us to consider larger DMPs. As explained, this would require significantly larger groups than are common. Yet, while in very large DMPs, it might be possible for the entire group to accurately represent the electorate, it is still challenging for each particular individual within it to gain an accurate perception of the electorate from their personal interactions, which, as explained, is needed for them to know what will promote the interest, values, etc. of the majority of participants. This is because the large group size makes deliberation where all voices are heard equally by all others highly challenging (see Sunstein, 2009; Thompson, 200; Jackman & Sniderman, 2006; Mendelberg & Oleske, 2000; Sunstein, 2002; Thompson, 2008; Setälä & Smith, 2018; Elstub

et al. 2021). While expert facilitators may help create a more balanced combination of testimonies (e.g., Elstub & Escobar, 2019), large groups (for instance, 1000 participants) cannot all effectively deliberate within a single forum. For this reason, it has been suggested (Guerrero, 2024), and implemented in various DMPs (e.g., Campbell et al. 2011; Giraudeau, 2021), that smaller groups form, informing the larger group. In such structures, participants may be unevenly exposed to different viewpoints, likely learning more about the people in their immediate group. Consequently, worries about not knowing precisely what would best promote the interests, values and aims of the majority resurface in this context, as participants will not be exposed to unfiltered, first-hand testimonies by and debates with all others. Even if people are briefly exposed to such views, they might naturally be less influenced by them than by their immediate group members. Hence, it is not clear whether the deliberation achieves the goal of improving participants' ability to vote in accordance with what is right for the majority.

To conclude, it appears that no participant count in DMPs is likely to cause participants to have an accurate view of what is best for the majority of society. Consequently, the epistemic advantage of DMPs is also put in doubt in aiming to track a population-relative truth by each individual aiming to track this truth themselves, i.e., when each participant promotes what they believe holds according to the interests, values, etc. of the majority of the population.

It is worth noting that the ability of DMPs to track a population-relative truth by each participant aiming to track that truth themselves can also be called into question for interpretations of population-relative truths that are not majoritarian. This is because many such interpretations also require an accurate representation of the population. If, for instance, we suppose that what is right for the electorate is what is best for the least well off in that society, and hence, that each participant needs to vote according to this standard, participants could get such a result wrong if they don't encounter the least well off, and they mistakenly take the least well off out of those they encountered to be the least well off in society at large. In an unrepresentative sample, there is a non-negligible chance that this will occur. For instance, if they encounter less than 50 percent of societal viewpoints, as in the example above, they will have less than a fifty percent chance of encountering the relevant testimony to answer this question. Hence, a representative sample is needed. The same is true for questions such as:

What is right for the population most affected by this policy? Or: what is best for those most passionate about it?<sup>51</sup>

Of course, it is hard to orchestrate situations where participants are exposed to the wide range of viewpoints plausibly needed for achieving an accurate perception of the population at large. For this reason, it might be better to employ certain types of experts, such as people who specialize in learning about societal issues, rather than attempting to gather a deliberative group where each individual is meant to serve as evidence of the views, interests, values of the group they descriptively represent. Alternatively, due to the complicated nature of each participant tracking a population-relative truth, it might be better to task each individual with tracking an individual-relative truth. I will delve more into this point in the next argument.

*2.4.5.2. Argument 2: Tasking people with directly voting on their understanding of what would promote the majority's interests and values is likely outperformed by tasking them with voting on what would best promote their own interests and values.*

Another reason why participants engaging in tracking what is best from the standpoint of a majority is unlikely to provide a strong epistemic advantage to DMPs is that such a method is likely often outperformed by individual-relative truth tracking. I claim this for two reasons. First, individuals will likely be less competent in the former. This is because tracking what is correct for one person (the individual themselves) is likely less complicated than tracking what is correct for a majority of many people. The latter requires engaging in the same procedure as the former, i.e., examining particular interests and figuring out which policy best promotes them, but many times over. Also, the latter requires figuring out the prevalence of varying interests in society, which is an additional complex task not needed in the former. This more complicated procedure for truth-tracking introduces many more possibilities for fallacies and misconceptions. Additionally, given that people have a reason to make decisions according to their own interests and values, they might also be less motivated to put an effort into accurately

---

<sup>51</sup> One might suggest that participants' encounter with experts will enable them to gain the relevant information for tracking a population-relative truth, for instance, information about the majority or the least well-off. However, in such a case, the truth of what is best for the least well off in society is best tracked using a group of experts rather than the DMP. One might suggest that there is an advantage in both tracking this truth and making decision-making more inclusive by using DMPs rather than experts. However, DMPs are a small, select group participating in decision-making. As explained, they are only inclusive to the extent that their decisions reflect the general public. If their decisions aim to reflect the public by tracking the interests of the least well off in it, as is the case in the discussed scenario, we might as well opt for the method that best tracks this, hence best reflecting the public according to the same measure.

deciding according to the majority's interests and values, which may also affect their competence.

Second, the assumption that people can be competent in figuring out what would best promote the interests, etc. of the majority relies on the assumption that they are each at least somewhat competent in doing this with regards to their own interests, values, etc. This is because, as part of understanding where the majority's interests, values, etc. lies, they would need to understand also where the interests, values, etc. of the group they themselves are meant to represent lie. Also, for participants to get an accurate portrayal of what such promotion for society at large looks like, each of the other participants they encounter needs to be able to portray that information regarding the group they are meant to represent. They would need to be able to bring up the relevant interests, values, and goals, and share the most relevant experiences on a given topic, indicating competence in making the needed links. Hence, it seems that in scenarios where people would be competent at tracking the majority's interests and values, they would also be competent in tracking their individual interests and values, while the opposite is not true.

It follows that, within a DMP, tasking individuals with uncovering what is best for the majority is likely inferior to tasking them with figuring out what is best from their perspective. But as explained in section 2.4.4, the latter is, in turn, worse than alternative methods of truth tracking that do not involve DMPs.

#### 2.4.5.3. *Argument 3: Biased information portrayal*

The third argument is that, just as in the case of individual-relative tracking, the inherent vulnerability to biases that exists in information portrayed to participants is likely to harm the truth tracking process. In section 2.4.4, I explained why information provided to participants of DMPs is likely to be biased and tend to represent the opinions of organizers and experts. I also illustrated how that influences the decision-making and thought process of participants. This is also an obstacle for DMP participants' ability to get an accurate perception of what is right given the interests, etc. of the majority of the population, as the "informed" opinions participants encounter will not necessarily represent the population but the organizers.

I have now justified premises 1-5, leading to the conclusion that a seemingly plausible line of thought to justifying DMPs is not actually promising. Hence, alternative means of justifying DMPs are required.

## Conclusions

In this chapter, I have questioned the enthusiasm surrounding DMPs as supplements of current democracies. I have presented an argument showing how a potentially promising way of justifying this enthusiasm likely fails. Hence, proponents of DMPs will need to come up with alternative means of justifying their unique value.

Specifically, I claimed that advantages of DMPs over alternatives are likely the result of the process they employ having some epistemic advantage. A natural candidate for such an advantage is tracking a population-relative truth, which I interpreted as what is right given the interests and values of the majority of the population. I then examined two means of tracking such a truth: (1) where each participant tracks their own interests and values, and (2) where they each track those of the majority. I showed that no matter which approach is adopted, DMPs are unlikely to offer clear epistemic advantages over alternative supplements to existing representative governments. I concluded that this seemingly plausible way of justifying their advantages is not actually promising.

If my argument is successful, it highlights, among other things, the need for alternative means of making democratic systems more inclusive. Both this and the previous chapter potentially support the idea of enhancing current democracies using referendums, for instance. However, whichever form of inclusivity is chosen, it will likely benefit from changes in voters' informational environments. Within this chapter, I have also touched upon the topic of how various means of providing individuals with information could be biased, potentially harming their ability to pursue their own interests, values, and goals. I argued that the DMP structure is particularly prone to providing information that leads to biased decision-making. I showed that DMPs cannot be considered descriptively representative of the general population when their conclusions are largely influenced by information biases reflecting the views of organizers. However, to some extent, such worries exist with any intervention in voters' epistemic environments. Generally, in informing voters, it is crucial to take into account means of preserving their epistemic autonomy. This requires the development of a theory of what type of information portrayal preserves voters' epistemic autonomy, a question I pursue in the following chapters.

*2. The Limitations of Deliberative Mini-Publics*

---

### 3. Freedom of Information Choice

---

#### Abstract

This chapter explores an overlooked aspect of freedom of choice, which I term “freedom of information choice”. This is the ability to autonomously form evaluative judgments, such as opinions and beliefs, which hinges on the options one has regarding which information to consume. I claim this is the freedom harmed by the design of many contemporary epistemic environments, including ones prevalent online. I identify two distinct characteristics of freedom of information choice. First, it cannot be measured primarily based on pre-existing evaluative judgments. Additionally, as cognitive limitations are the most crucial constraint in processing information, I suggest that instead of the monetary budget set that is common in the evaluation of economic freedom, we should primarily consider a “cognitive budget set”. I therefore propose a new measure of enhancing freedom of information choice, namely, increasing the intra-bundle diversity of options. This involves a diversity of information within *jointly possible* combinations of information (i.e., within the same bundle) being accessible *within the limits of one's cognitive budget*, instead of diversifying between mutually exclusive options. I illustrate how this can be achieved by highlighting a distinct quality of the cognitive budget set.

#### 3.1 Introduction

A defining feature of the digital age is the unprecedented availability of information options provided to individuals. We might expect this abundance to provide individuals with a particular form of freedom: freedom of choice. This prevalence, however, is not always helpful to individuals' epistemic process. In this chapter, I argue that the problem lies in the absence of a particular form of freedom of choice.

Choices we face between information options influence our beliefs, and in turn, broader life choices. Determining which information option-sets enhance and restrict our freedom, however, is not an easy task. There are properties of information options that make evaluating the freedom they provide more elusive than in other more straightforward contexts of choice,

such as choosing which goods to buy in a supermarket. While in more straightforward contexts, it is often assumed people hold some values or opinions that allow them to have established attitudes towards options, information option sets are meant to shape a person's evaluative judgments by providing evidence and arguments for and against them. Hence, relying primarily on preexisting evaluative judgments to assess such option sets defeats the purpose of offering them. Also, because information choices directly engage with a person's cognitive capacities, considering cognitive limitations becomes a primary concern in the context of information options.

In this chapter, I introduce the concept of "freedom of information choice", the ability to autonomously form evaluative judgments, such as opinions and beliefs, which hinges on the options one has regarding which information to consume. I propose that, due to these distinct properties, the measurement of freedom of information choice cannot primarily rely on preexisting evaluative judgments. Also, to measure freedom of information choice, we should adopt the idea of a "cognitive budget set", i.e., the set of all alternatives of jointly accessible information given our limited cognitive capacities. I argue that for freedom of information choice, this budget set should be our primary concern rather than a monetary one.

As existing literature on freedom of choice has failed to take these properties into account, it cannot capture the nature of information option sets. As a result, I argue in this chapter, it does not account for our intuitions regarding which option sets offer more freedom of information choice. I therefore propose a new idea: that such freedom increases with an increase in the diversity of perspectives exhibited *within* each of the information bundles *included in a person's cognitive budget set*. I term such diversity "intra-bundle diversity". This contrasts with enhancing diversity among mutually exclusive alternatives or different bundles, i.e., inter-bundle diversity. Inter-bundle diversity alone, I claim, cannot account for our intuitions about prominent cases of lack of freedom of information choice.

The chapter is structured as follows. In section 3.2, I define freedom of information choice and offer some examples. In section 3.3, I explain the two properties that make it distinct from the paradigmatic cases we usually have in mind when thinking about freedom of choice. In section 3.4, I explain why prominent theories of freedom of choice cannot accurately rank information option sets. In section 3.5, I explain why diversity within information bundles improves freedom of information choice. In section 3.6, I point to a distinct quality of the cognitive

### *3. Freedom of Information Choice*

budget set and use it to suggest a few adjustments to option set structuring that can enhance such diversity.

#### **3.2 Defining freedom of information choice**

Let us examine several scenarios.

The “Rabbit Hole Scenario”:

Martha opens her news website before the U.S. elections and sees an enticing partisan headline like “Democrats Will Rig the Elections!” or “Republicans Are Undermining Democracy!”. Without much thought, she clicks on the article. She reads a short piece with colourful images and strong statements that resonate with her views, making her angry about the situation. This leads her down a rabbit hole where she keeps receiving enticing recommendations to articles, each one slightly more extreme than the last. She keeps clicking these “click baits” until she is finally pulled out. By then, the radical ideas she encountered seem familiar and convincing.

The “Intolerance Scenario”:

Continuous exposure to one-sided opinions leads Bob to develop intolerance toward the ‘other side’, experiencing anger at Fox News if Democratic or at the New York Times if Republican. Without reading these sources, he believes everything they say is misleading or biased. Opposing views cause him to experience cognitive dissonance and avoid further exposure. The anger he experiences discourages him from considering the arguments he encounters, resulting in a vicious cycle where past lack of exposure prevents future exposure.

The “Bite-Size Information” scenario:

Hannah becomes accustomed to one-sided information from her favourite partisan sources, making opposing opinions mentally challenging to process, as they require more concentration, active searches and consolidating conflicting views. Consequently, when she wants to both relax and stay informed about a current topic, which happens several times a day, she consumes the one-sided articles she is accustomed to. Despite her abstract desire for exposure to opposing viewpoints, she finds it too demanding during daily moments when she typically consumes information, such as short work breaks, and her opinions are shaped accordingly.

### 3. Freedom of Information Choice

Such scenarios are common occurrences in people's everyday lives (e.g., Iyengar & Leelkes, 2012; Iyengar & Westwood, 2015; Milano et al., 2020; Boucher, 2022).<sup>52</sup> In each scenario, there appears to be an issue with how the available information options shape people's judgments. Judgments are caught in vicious cycles, where lack of exposure to certain opinions perpetuates itself. Beliefs are formed on the basis of contingent circumstances, like the initial data one was exposed to that initiated the described vicious cycle. Additionally, people's epistemic process is riddled with psychological biases and well-known threats to autonomy, like addiction, indoctrination, manipulation, and ignorance (e.g., Dworkin, 1988; Christman, 1991; Buss, 2018). In the Rabbit Hole Scenario, for example, the content one clicks on is designed to be addictive, to manipulate one into spending more time online. Such an option-set limits people's control over their epistemic process and pushes them to form judgments and opinions in a non-ideal way. I claim such information option-sets importantly limit people's *freedom of information choice*.<sup>53</sup>

To define the notion of freedom of information choice more carefully, I will first introduce the general concept of freedom of choice. The formal freedom of choice literature, advanced by scholars like Amartya Sen, Prasanta Pattanaik, and Yongsheng Xu, evolved around the understanding in economics literature that option sets have a value beyond the benefit one receives from the chosen option. This is the intrinsic value of the freedom to choose (Sen, 1988, 1995; Pattanaik & Xu, 2015), which is derived from the value of being an autonomous decision-maker with control over one's actions (Sugden & Jones, 1982; Sen, 1988; Raz, 1988; Pattanaik & Xu, 2015) and of free development of one's individuality through choices (Mill, 1859). Freedom of choice has been taken to be different from freedom simpliciter in that the former involves the possibility of making a *selection* out of more than one option, while the latter only requires the possibility of *acting* (Carter, 2004).

The freedom of choice literature explores means of comparing different sets of mutually exclusive options available to a person based on their ability to provide this freedom. They offer criteria for such a comparison that provides a ranking of option sets (Sen, 1993A; Pattanaik and Xu, 2015). Theorists have investigated various criteria that potentially enhance freedom of choice, i.e., criteria that can be used to rank option sets (e.g., Pattanaik & Xu, 1990,

---

<sup>52</sup> While some researchers find echo chambers less concerning (Bruns, 2019) or beneficial (Erickson, 2024), I appeal to those who specifically find examples of the type described here to often be problematic.

<sup>53</sup> While such scenarios may not be *equally* problematic regardless of people's specific worldview, individuals lack autonomy in important aspects of their belief formation process regardless of initial viewpoint, due to insufficient self-control in the epistemic process, and threats such as addiction, manipulation and ignorance.

### 3. Freedom of Information Choice

2000, 2008, 2015; Puppe, 1996; Sen, 1993B; Xu, 2004; Côté, 2020; Costella, 2023). Such criteria include, for example, whether the options in the set are preferred, and the number of options that could be reasonably valued. Such suggestions attempt to determine what makes for a more meaningful selection menu. While the aim of these suggestions has been to come up with general criteria that fit different types of options, the focus has been mainly on various options for what to do, be, or buy. Examples presented in the literature include different possible careers (Pattanaik & Xu, 2015; Costella, 2023), deciding between taking a bus or a car to work (e.g., Pattanaik & Xu, 2000), and mutually exclusive bundles of commodities (e.g., Pattanaik & Xu, 1990, 2000; Xu, 2004). Importantly, this literature has neglected option sets composed of *information* options, whose purpose is developing one's ability to form judgments. As a result, prominent suggestions for such measurements cannot account for the lack of freedom of choice in the scenarios above, as I will show in section 3.4.

I now define the concept of freedom of information choice. As mentioned, the freedom of choice literature compares the freedom offered by different option sets (Pattanaik and Xu, 1990; Sen, 1993B; Pattanaik and Xu, 2000, 2008; Côté, 2020). Accordingly, a measurement of freedom of information choice should compare the relevant type of freedom of different information option sets. By “information options”, I mean items from which one can extract informational content, including evidence and justifications for evaluative judgments. Relevant sources include articles, books, papers, speech, and social media posts.

Now, what is the relevant type of freedom that information options provide? In this chapter, I am concerned with information options that serve as a basis for forming beliefs and evaluative judgments. As mentioned, measurements of freedom of choice often aim to determine what is a meaningful set of options to choose from, and its value is thought to be derived from the value of autonomous decision-making, exercising control, and self-determination. Likewise, measurements of freedom of information choice should determine what is a meaningful set of information options. Information option sets that provide more freedom of information choice are ones that promote one's ability to form evaluative judgments, such as opinions and beliefs, autonomously. They do so by allowing greater self-control and self-determination in the process of forming such judgments.<sup>54</sup>

---

<sup>54</sup> Providing freedom of information choice means allowing autonomy within one's epistemic process. Therefore, freedom of information choice relates to the concept of “epistemic autonomy”. In recent years there have been several different interpretations of this concept (e.g, Prichard, 2013, Matheson, 2024). To the extent we understand “epistemic autonomy” to mean having autonomy, governance and self-control over one's judgment formation

### 3. Freedom of Information Choice

Importantly, in keeping with the economic tradition in this area (e.g., Sen, 1988, 1993B), I take the freedom to  $x$  to imply the ability to  $x$ , focusing on people's *real set of options* rather than what we might call their nominal set or their theoretically available set. I am not taking freedom to only imply not having one's choices be actively interfered with by others. Some prominent theories of freedom, and particularly with regards to judgment formation, adopt a similar understanding of the concept (e.g., Chrisman, 2024). I believe this focus is particularly helpful when considering typical contemporary structures of informational environments, such as the ones in the aforementioned scenarios. These scenarios show that typical cases of lack of autonomy in people's epistemic process, especially when people have uncensored access to the internet, do not necessarily involve intentional active interference by others.

Notably, there could be other uses to epistemic sources, including books and newspapers, that can facilitate other types of freedom of choice. For example, one could wish to have better options for a book to pass the time, for enjoyment, as an intellectual challenge, etc. For any of these uses, it might not matter, for example, if a piece of information contains false evidence or if one is only exposed to one-sided views. However, these are not the types of freedoms I am concerned with. Instead, I will ask what information option sets allow for a meaningful selection when forming evaluative judgments. To answer that question, we first need to understand what is distinct about freedom of information choice.

#### 3.3 What is distinct about freedom of information choice?

##### 3.3.1 Not primarily relying on pre-existing evaluative judgments.

One thing that sets freedom of information choice apart is promoting a fundamental type of freedom: the freedom of forming judgements. This freedom is foundation for other types of freedom. Typical examples of cases where a person lacks freedom of choice often involve people's choices not aligning with their evaluative judgments. Such an example is a person missing an appointment they value due to their cigarette addiction, which they do not (Carter, 2022, section 1). Another example is not having the option of following a certain career a person values, for instance, being a doctor (e.g., Costella, 2023, p. 4). In such scenarios, we

---

process, information option sets that better promote epistemic autonomy will be ranked as providing more freedom of information choice.

### *3. Freedom of Information Choice*

evaluate people's lack of freedom based on their evaluative judgments and resulting attitudes toward different options.

However, to assess people's ability to autonomously form evaluative judgments, we cannot primarily rely on judgments and attitudes towards options, as these are meant to be formed in the process. In the Rabbit Hole and Intolerance scenarios above, for example, Martha and Bob's judgments are shaped by the homogenous and radicalized information to which they have been exposed. Their beliefs were formed as part of the vicious cycle process and therefore cannot serve as the primary criterion for assessing it. To the extent prior judgments were not formed freely, using them as assessment criteria creates cycles where past unfreedoms cause future ones. Establishing that these prior judgments were formed freely cannot itself rely on those judgments. Hence, it brings us back to the initial question.

For instance, deciding whether an abortion-related option set should include two pro-choice items or one pro-choice and one pro-life items cannot rely primarily on pre-existing judgments of the individual towards these options. As the purpose of the option set is to allow a person to learn about a new topic and form evaluative judgments about it, we should not assume the prior existence of firmly held evaluative judgments of this type. People might be new to a topic and not have firmly held beliefs, or they may change their beliefs in response to exposure to new information and arguments. For instance, they might firmly hold that pro-life views are better. However, to shape information options accordingly without allowing them the opportunity to explore alternative views would plausibly significantly undermine the purpose of gaining new information on the topic.

Individuals may also have evaluative judgments towards sources of information, preferring certain sources over others. However, these views are also influenced by the information to which people have been exposed. Hence, we need to determine whether such views about sources were freely formed. We cannot assess the freedom of a certain judgment, such as the judgment that a particular partisan news source is best, given an option set that is assessed using that judgment itself. If we were to define freedom in terms of an option set that already aligns with that judgment (i.e., news source  $X$  is best), then we would take a judgment about  $X$  that is solely based on information from  $X$  to be freely formed. However, this, again, is a vicious

### 3. Freedom of Information Choice

cycle.<sup>55</sup> Hence, evaluating the option set in forming the judgment about news source  $X$  would require determining another criterion for an option set that allows freely forming that judgment.

One might suggest using attitudes towards the option set itself, such as preferring a diverse or uniform set. Let's call such views "meta views" for convenience. However, these meta views are also influenced by the information to which people have been exposed. For instance, if prior lack of exposure caused intolerance to certain views, as in the Intolerance Scenario, it will affect these meta views, as the person will not want to have information pieces supporting such views in their option set. Hence, these meta views may not be freely formed. Even if these meta views were independent of the content of specific options, they would still need to have been formed somehow. Then, again, either they have been freely formed or not, meaning we need a measurement for which option sets enable the free formation of such meta-views. One might suggest using what we may call "meta meta views", i.e., views about the information option set best suited to freely form an opinion about the right design for an option set. However, we would then need to determine if those meta meta views were freely formed and that would lead to an infinite regress.

Hence, when it comes to freedom of information choice, we need to answer the question – how can we measure the degree to which an option set supports the free formation of judgments without primarily relying on pre-existing judgments? This is essential at least for assessing the free formation of key judgments. Of course, this does not imply that people cannot use pre-existing judgments as part of the evaluation of evidence they encounter. Such assessment, which does not prevent *mere possibility of exposure* to evidence, does not cause the same type of vicious cycles.

Another reason why we may not wish to rely primarily on attitudes towards options is that often people are not initially aware of all options. Lack of awareness prevents them from being able to form intentions to pursue those options. For instance, searching relevant keywords on Google is necessary to reach a particular article. Furthermore, regular consumption of news sources that present alternative viewpoints as insubstantial and fail to mention the existence of contrary evidence causes people to be unaware of the evidence's existence, and be deprived of reasons to seek it. For instance, if a news source presents a bombing as being caused by A, neglecting to mention that whether A caused it is in dispute, its consumers are unaware of the

---

<sup>55</sup> Note that due to initially only being exposed to source  $X$ , people cannot have a correct prior distribution of what they are likely to find in each source.

dispute and lose grounds for seeking further evidence. Therefore, relying on pre-existing judgments creates vicious cycles also due to lack of awareness. Hence, to break such cycles, additional criteria are required.

### 3.3.2 A “cognitive budget set”

Freedom of information choice also differs from traditional ways of thinking about freedom of choice since the limited budget we should be primarily concerned with in freedom of information choice is cognitive. At the root of the idea of freedom of choice held by many economists, at least when economic activities are concerned, resides the notion of a budget set. This is the set of alternatives available to a person given their income and the prices of the commodities in the market. Any combination of existing commodities whose aggregate cost is smaller than or equal to the person’s budget is considered a separate bundle, i.e., one alternative (Pattanaik & Xu, 1990; Sen, 1995; Xu 2004; Böhm & Haller, 2008; Miyagishima, 2010). As the bundles are mutually exclusive, a person’s freedom of choice is usually measured by looking at a one-off choice from an option set containing these bundles (e.g., Pattanaik & Xu, 1990, 2000). This original notion has been vastly expanded, for example, by Amartya Sen’s capabilities approach (Sen, 1974, 1988, 1995). Still, using a budget set to determine a person’s freedom of choice remains prevalent in this literature.

For freedom of *information* choice, the primary budget constraint is cognitive rather than monetary. The abundance of information available through the internet and various devices comes at little to no monetary cost but is highly mentally taxing. Most individuals do not have a monetary limit preventing them from reaching vast amounts of content, but their limited cognitive capacities restrict what they are able to process. Their cognitive resources are often used as a currency with which they “pay” for information. For example, the cost of receiving information is often an implicit agreement to have constant attempts to distract their attention with sponsored content, which exhausts limited cognitive resources. A cognitive budget is generally the most relevant constraint on information options as cognitive limitations vastly influence the way opinions are formed. Enabling free opinion formation requires acknowledging our limitations rather than idealizing us as epistemic agents.

The analogy between a monetary and a cognitive budget set works well. Like the relationship between commodities and money, the different information options we have each come at a different cognitive cost. In cognitive psychology “cognitive cost” refers to the mental effort required to perform a particular task (Christie & Schrater, 2015). It is widely accepted that tasks

### *3. Freedom of Information Choice*

have varying effort levels, ranging from automatic processes that are effortless, to tasks demanding high cognitive control, which often involve active concentration and suppressing default processes (Székely & Michael, 2021). For instance, the process of giving in to distracting commercials is often automatic. Conversely, reading while suppressing distractions requires more cognitive control (Goldstein, 2014). Familiar tasks are less cognitively costly, as they draw on long-term memory, rather than working and short-term memory (Chase & Simon, 1973). High cognitive costs are also associated with active attention 85rticle85n, selective attention, task switching and suppressing strong emotions (Zenon et al., 2018; Christie & Schrater, 2015; Székely & Michael, 2021).

Also, analogously to money and commodities, the cognitive cost of processing information comes from a limited budget. It is well established in cognitive psychology that every usage of cognitive resources creates opportunity costs at other tasks' expense (e.g., Kurzban et al., 2013). The level of cognitive effort people invest in a task directly influences the accuracy and completion time of subsequent tasks (e.g., Marcus et al., 1996; Sweller et al., 2011). According to the "law of least mental effort", people seek less cognitively demanding strategies, wishing to preserve limited cognitive resources (Kahneman et al., 1991; Christie & Schrater, 2015). Cognitively demanding tasks cause a feeling of fatigue (e.g., Borragán et al., 2017).

I therefore suggest treating our cognitive capacities at a given timeframe as a limited budget determining a budget set. Behavioural economics and bounded rationality show that considerations of cognitive cost significantly influence how individuals make choices within a monetary budget set (Kahneman et al., 1991; Kahneman, 2011; Thaler & Sunstein, 2009; Thaler, 2016). For instance, whether people need to opt in or out of a pension plan changes the chance they will participate in it, due to the cognitive cost associated with the process (Thaler, 2016). However, I suggest that when dealing with information options, the cognitive cost becomes the primary constraint, and we should treat it as determining a cognitive budget set.

As with a monetary budget set, the cognitive budget set is the set of all informational alternatives that are feasible within the limitations of the cognitive budget. Each alternative is itself a bundle of information pieces, and any combination of available information pieces whose aggregate cognitive cost is smaller than or equal to the person's cognitive budget will be considered a separate bundle, i.e., one alternative. For example, a commodity bundle will include both bread and eggs only if a person has money for both. Similarly, an information bundle will include both a pro-life and pro-choice article only if the person has sufficient

cognitive resources to open, read, and comprehend both articles. Consequently, to improve freedom of information choice we should focus on improving people’s “cognitive budget set”, the mutually exclusive combinations (i.e., bundles) of information pieces attainable within their cognitive budget for a given timeframe.

Notice that the cognitive budget set does not need to be modelled in precisely the same way as the classic modelling of the monetary budget set (e.g., Mas-Colell et al., 1995). As explained, cognitive costs exhibit important similarities to monetary costs, allowing them to be depicted as falling within the limits of a budget, with the consumption of some information pieces falling outside the budget. Yet the way the cost of each item is determined also differs in some crucial respects from that of a monetary price, suggesting the need for a modified model for this budget set. I discuss this more in section 3.6.

### 3.3.3 Using these distinctions to explain our intuitions.

Using a cognitive budget set to describe freedom of information choice helps explain the scenarios above, where a person theoretically has many options but ends up in an informational “echo chamber”, i.e., an environment where she only encounters views similar to her own. In such cases, opposing views may reside outside the person’s cognitive budget set. This could be because opposing views require diverting attention from distracting recommendations (the Rabbit Hole Scenario), suppressing strong emotions and overcoming cognitive dissonance (the Intolerance Scenario), or tackling task switching and unfamiliar tasks (the Bite-Size Information scenario). In each case, exposure to opposing views has a higher cognitive cost, possibly leaving options requiring it outside people’s cognitive budget within a specific timeframe.

This may often be the case if we consider the limitations put on a budget set by the combined cognitive cost of daily life *and* learning about a new topic. For instance, if a person only finds time to consume new information after a demanding workday, their cognitive resources may already be exhausted, preventing them from consuming cognitively demanding information. It is important to account for the demands of daily life when considering the cognitive budget a person can allocate to learning new information. These demands can be considered a part of what needs to fit within the same cognitive budget or as a factor influencing budget allocation, which limits the resources a person can dedicate to specific topics at certain times.

Importantly, the cognitive limitations are not the only constraint on freedom of information choice. Even though there is a hard limit to a person’s mental capacity, within this limit, the

### 3. Freedom of Information Choice

allocation of a specific budget to a certain topic at a particular time is partly within the person's control. However, the vicious cycles that arise from relying on pre-existing judgments deprive people of the opportunity for deciding on a budget allocation that best matches their most crucial knowledge gaps. If, for example, people are only exposed to one-sided views that are dismissive of opposing opinions, or unaware of the existence of certain evidence, they do not have the opportunity to fully realise that the topic at hand includes substantial unresolved questions. Consequently, they are deprived of reasons to suffer the opportunity cost to devote more of their limited cognitive resources to studying it.

Hence, together, the two distinct qualities of information options explain why people are not free to explore alternative points of view on specific topics. Namely, they have a limited cognitive budget and are deprived of reasons to suffer opportunity cost when appropriate. As mentioned, the type of freedom I am concerned with here is people's *real set of options* rather than the theoretically available one. I am not taking freedom only to imply not having one's choices be actively interfered with by others.

Cognitive budget sets also help explain other intuitive characteristics of freedom of information choice. For example, people have differing cognitive capacities, which causes differences in the amount and complexity of information they can process (Eysenck et al., 1994). Similar to monetary differences in individual budgets, this causes disparities in the level of freedom of information choice they have. Some bundles or specific information pieces are outside the reach of some people but not others. For example, no matter how hard they try, not everyone can understand quantum physics. Discussing people's freedom of information choice in terms of their cognitive budget set helps explain the difference in the level of freedom caused by such differences between individuals.

Importantly, the two distinct criteria of freedom of information choice affect which modifications to people's option set will increase freedom of choice, as I will explain in the next section.

#### 3.4 What criteria of freedom of choice get wrong about freedom of information choice

In the literature on freedom of choice, various criteria have been proposed to rank option sets by the freedom they provide, but many are unsuitable for measuring freedom of information

choice due to the distinct qualities discussed earlier. I will examine prominent proposals and explain their limitations for freedom of information choice.

### 3.4.1 The preference criterion

A common idea is that better preferred options increase freedom of choice (Sen, 1991, 1993). However, as mentioned, when ranking information options, it is not fitting to primarily rely on pre-existing attitudes towards options, such as preferences. This is because these options are meant to shape those attitudes. The scenarios above illustrate this problem. For example, in the Intolerance Scenario, Bob's original one-sided option-set shapes his preferences, causing a vicious cycle. Bob cannot tolerate opposing views and strongly prefers the one-sided information he is used to consuming. Moreover, known characteristics of cognitive cost may render a preference-based ranking meaningless. There is a typical correlation between the options people will tend to prefer and the less cognitively costly options, which easily fit within a budget. That is, continuous exposure to a certain type of information both causes people to prefer it and to more easily process it. One-sided exposure provides only evidence-based reasons for adopting those opinions, causing a preference for them. Additionally, as these opinions are familiar, they require less task switching, generate less cognitive dissonance, and need less anger suppression, reducing their cognitive costs and raising the likelihood of them fitting within the cognitive budget. This creates a self-reinforcing compatibility between preferences and the option-set, making preferences an inadequate measure of freedom of information choice<sup>56</sup>

### 3.4.2 The reasonableness criterion

Another proposal in the freedom of choice literature is that having additional options that could be reasonably valued increases freedom of choice (Dowding & Van Hees, 2007; Jones, 1982; Pattanaik and Xu, 1990; Sen, 1993). According to Pattanaik and Xu, reasonable is determined by “the values prevailing in the society to which the individual under consideration belongs”, and “the preferences of a reasonable person in  $I$  (the person faced with the option set)’s community” (Pattanaik & Xu, 2015: p. 373).

---

<sup>56</sup>The idea that the range of choices available to a person can influence their preferences is also discussed in studies on adaptive preferences (e.g., Elster, 1983; Nussbaum, 2001, Costella, 2023). This research highlights how oppression can cause individuals to develop preferences that are not genuinely their own. This phenomenon is not treated as inherent to all option sets, in contrast to what I claim to be the case for information options due to their unique qualities. Also, unlike the scenario described in adaptive preferences research, where pre-existing genuine preferences are assumed, our goal when considering information choices is to identify the conditions that enable the free development of preferences.

### *3. Freedom of Information Choice*

The reasonableness measure is also problematic, as it again relies on pre-existing evaluative judgments, though not necessarily those of the individual. This undermines the goal of allowing people to form opinions about what they find valuable, reasonable, etc. Even if the evaluative judgment is made by the person's community, an option set guided by it can still limit their freedom. Consider a person from a closed community functioning as an echo chamber. That person's community is strongly opinionated against certain views and does not value them. If we were to take values prevailing in this community as a benchmark for reasonableness, we would wrongfully consider information options presenting similar views as providing more freedom of choice than those offering opposing views. Hence, we would again allow a lack of sufficient exposure to perpetuate itself. Broader societal values can also be problematic, because society at large could also sometimes function like an echo chamber, suppressing certain views and promoting others. The Catholic Church's ban on books supporting the theory that the earth revolves around the sun exemplifies this. We wouldn't want to claim that option sets that limit people's information according to such standards enhance freedom of information choice. Hence, there is something problematic about freedom of information choice criteria being primarily determined by current societal values.

One might interpret the reasonableness criterion as a rationality criterion, i.e., valuable options are ones that a person can rationally prefer. However, it is commonly thought that rational beliefs are updated based on available evidence. Insofar as the options a person may rationally value depend on the evidence they are exposed to, using such a measurement creates vicious cycles. Given common types of initial evidence bases, people would be deprived of important reasons to rationally value information promoting the same line of thought as what they have previously seen, as explained above. For example, an evidence base presenting opposing views as biased or manipulative makes it somewhat irrational to want to consume them. Additionally, if people are exposed to information presenting a certain fact as certain – for instance that side A of a conflict caused a certain bombing, neglecting to mention that in other sources that information is treated as dubious and contrary evidence are presented, people will be deprived of reasons to look out for such opposing views. Such cases, which are prevalent in polarised media and online platforms (e.g., Iyengar & Westwood, 2015), make the reasonableness measure compatible with lack of exposure to opposing views perpetuating itself, and hence with the scenarios described above.

### 3.4.3 The size and volume criteria

Another idea in the literature is that increasing the number of options in a set is one of the factors enhancing freedom of choice (Pattanaik & Xu, 1990, Sen, 1993B).<sup>57</sup> According to Michael Garnett, for instance, a cardinality criterion is essential for preserving a commitment to value neutrality, as other criteria, such as option variety, are value-laden (Garnett, 2016). Monetary budget sets are often listed as a central example of the type of option sets that theorists have in mind. It has also been proposed that increasing the volume of one's monetary budget set enhances freedom of choice (Xu, 2004). However, increasing the number or volume of options that fit within one's monetary budget will often not increase freedom of information choice, and might even reduce it. This is because it will not necessarily increase the number or volume of options that fit within one's cognitive budget, as these two types of budgets have different constraints. Moreover, increasing the options within one's monetary budget set might actually reduce the amount of options within one's cognitive budget set. This is because too many options can cause choice overload, i.e., being overwhelmed and unable to process options effectively for informed decisions (e.g., Park & Jang, 2013; Schwartz, 2012). This can lead to choice paralysis, where not choosing becomes the only manageable option. When a person experiences choice paralysis, the option of not choosing remains the only option available within the cognitive budget set. Opting for specific choices, such as opening and reading a specific article, becomes too costly. This clearly offers less freedom of choice than a more varied set. Hence, increasing the size of one's theoretically available options may decrease the size of one's real option set, depicted by the cognitive budget set. Therefore, this criterion should be revised considering the limitations of cognitive budget sets. Since for freedom of information choice our primary budget concern is a cognitive one, we should only consider an increase in the number or volume of options that fall within people's cognitive budget to enhance freedom of information choice.

### 3.4.4 The accessibility level criterion

Some theories also take the ease or difficulty of achieving certain options to affect the level of freedom they provide (Sunstein, 2019; Côté, 2020). Nicolas Côté, for example, suggests that

---

<sup>57</sup> Pattanaik and Xu suggest increasing the number of options that could be reasonably valued increases freedom of choice. This proposal still takes an increase in the *number* of (reasonable) options to increase freedom of choice, therefore taking cardinality to be an important factor in increasing freedom of choice. In Sen's proposed measurements, one of the determining factors for A to provide at least as much freedom of choice as B is if there is some "one-to-one correspondence" between set B and a subset of A, putting an emphasis on size. The criticism here applies to such accounts.

option set A provides more freedom of choice than set B if they contain the same options, but some options are more accessible in A, without any options being less accessible. He interprets accessibility as the probability of choosing an option conditional on intending to do so. For example, the option to study is highly accessible to Sarah if there is a high probability that she will study if she intends to do so. Although not explicitly categorized as literature on freedom of choice, literature on nudging aims to describe a choice architecture where the difference in the accessibility levels of different options benefits individuals (Thaler & Sunstein, 2009). Sunstein argues that a good measure of a freedom-enhancing choice architecture is one that influences choices in a way that will make people better off “as judged by themselves” (Sunstein, 2019: p. 72). For example, a smoker who desperately wants to quit judges a choice architecture that makes quitting easier to be better. Hence, such choice architecture provides them with more freedom (p. 63).

Both these suggestions are often not sufficient for freedom of information choice, as they rely on pre-existing evaluative judgments and attitudes to options. Such interpretations only tell us what freedom consists of *given certain intentions or judgments*. Hence, it cannot account for cases where we have not yet formed such intentions or judgments, or when harms to our freedom are caused by the alteration of our intentions and judgments.<sup>58</sup> For example, in the Rabbit Hole Scenario, overexposure to a single-minded line of thought influenced Martha’s judgments. Accordingly, her intentions changed. She now intends to read more articles that agree with this radical line of thought. Such problematic changes cannot be accounted for by such theories.

### 3.5 A new criterion for improving freedom of information choice

Having examined how unique characteristics of information options make certain criteria unsuitable for assessing freedom of information choice, I now propose a new method. I argue that one aspect contributing to freedom of information choice is the extent of diversity of information within the same bundle, which I term “intra-bundle diversity”. As mentioned earlier, an information bundle is a combination of information pieces that can be jointly consumed given the limits of one’s cognitive budget. Hence, intra-bundle diversity is the diversity within the information that can be jointly consumed by an individual, given their cognitive limitations. Notably, just like general freedom of choice criteria, there can be multiple

---

<sup>58</sup> In this regard, there is some similarity between information options and the idea of a transformative experience (Paul, 2014), as in both cases, a person’s attitudes towards options change dramatically after the intervention, leaving judgments about its value less useful.

### 3. Freedom of Information Choice

criteria relevant to increasing freedom of information choice. It could be, for instance, that certain evaluations of the quality of information provided within the cognitive budget set are also relevant to a freedom of information choice ranking. In this section, I claim that intra-bundle diversity is one overlooked criterion that is essential for ranking information option sets according to their ability to enable autonomous judgment formation, and hence the freedom of information choice they provide.<sup>59</sup>

Free judgment formation requires diverse information. Individuals must be exposed to various opinions for the scenarios outlined in the first section, which result from underexposure to opposing views, not to occur.<sup>60</sup> Exposure to diverse views is thought to allow people to reflect on their views, plausibly switch views if they find others more convincing, and to justify their views if they don't. According to Mill, without the challenge of opposing views, the “grounds of the opinion are forgotten” (Mill, 1859: p. 55). Thomas Scanlon claims that autonomous agents need to consider different reasons for judgments (Scanlon, 1972). Recent epistemic autonomy literature claims autonomous agents rely on their own judgments when forming beliefs (Goldman, 1991; Fricker, 2006, p. 226). This plausibly requires awareness to opposing views. Others posit that an autonomous person not only acts in accordance with their beliefs but also reflects and reasons about them (Christman, 1991; Friedman, 2003; Meyers, 2005). Many other theories posit that an autonomous person not only acts in accordance with their beliefs and preferences but also reflects on those beliefs and preferences (Christman, 1991; Friedman, 2003; Meyers, 2005) and can reason about them (Buss, 2018). This plausibly requires some familiarity with opposing views. For example, Christman maintains that: “[a] person who endorses his decisions ... having been denied minimal education and exposure to alternatives, does not adequately reflect in this way” (2009: p. 147). Also, Meyers notes, “When individuals have little opportunity to explore alternative value systems ... function as indoctrination, which precludes critical reflection on the values and desires that shape one's choices” (2005: p. 30).

Considering diversity's essential role in autonomous agency, option sets that facilitate it should count as enhancing freedom of choice. Hence, the suggestion from the literature on freedom of

---

<sup>59</sup> Note also that, in this paper, I focus on obstacles to our *freedom* caused by problematic option sets, rather than on obstacles to epistemic improvement.

<sup>60</sup> In these examples, an essential part of the problem is the vicious cycle caused by lack of exposure perpetuating itself. Hence, in those examples, lack of exposure is an essential and major part of the problem.

### 3. Freedom of Information Choice

choice that comes closest to rightly ranking information option sets<sup>61</sup> is that diversity of alternatives enhances freedom (e.g., Klemisch-Ahlert, 1993; Pattanaik & Xu, 2000; Rosenbaum, 2000; Bervoets & Gravel, 2007). It has been proposed that option sets providing more types of options or a wider range of alternatives should be ranked as offering more freedom of choice (Pattanaik & Xu, 2000; Klemish-Ahlert M, 1993; Rosenbaum, 2000). However, this suggestion ranks the freedom offered by option sets according to the diversity between different *mutually exclusive* alternatives, such as different commodity bundles (Pattanaik & Xu, 2000., 2008; van Hees, 2007; Rosenbaum, 2000; Klemish-Ahlert M, 1993), or having the option to travel to a desired location by *either* bus *or* train (Pattanaik & Xu, 2000).

This criterion still falls short in capturing freedom of information choice. Freedom of information choice requires not just the option to be exposed to either side A or, instead, side B but mainly the option to be exposed to both.<sup>62</sup> Instances in our daily lives lacking freedom of information choice are precisely those where we have the option to choose between side A or B but not both. For example, we could have the option to immerse ourselves in a left-wing echo chamber or a right-wing echo chamber, providing a diversity of mutually exclusive alternatives, but intuitively not allowing the diversity needed to freely form judgments. Consider scenarios like the Rabbit Hole, Intolerance, and Bite-Size Information introduced earlier. In each case, the person had the initial option to choose from a diverse set of online sources, i.e., inter-bundle diversity. However, once they made a choice, without much thought, switching to the opposing view seemed outside the person's cognitive budget. The initial diversity did not offer true freedom of information choice, as it did not enable diversity within the same information bundle—diversity that is not mutually exclusive.

Therefore, I propose a new criterion for ranking information option sets, focusing on intra-bundle diversity rather than inter-bundle diversity. That is, rather than diversity in mutually exclusive information options, we rank them according to diversity in non-mutually exclusive options that their combination resides within a person's cognitive budget. This depicts the level of diversity within the pieces of information a person can jointly consume given their cognitive

---

<sup>61</sup> Besides the simple cardinality rule, which is not considered sufficient for ranking option setoption sets (Pattanaik & Xu, 1990).

<sup>62</sup> Simple cardinality ranking is insufficient. It would correctly rank the set {homogeneous opinions of type A, homogeneous opinions of type B, diverse opinions bundle 1} above {homogeneous opinions of type A, homogeneous opinions of type B}, Yet, we would also need the set {diverse opinions bundle 1, diverse opinions bundle 2} to rank higher than {homogeneous opinions of type A, homogeneous opinions of type B}, even though these sets are the same size. Also, of course, a model of a monetary budget set *allows* the existence of diversity, in the sense that it is possible to buy several different commodities within the budget. However, this has not been proposed as a measurement of freedom of choice.

### *3. Freedom of Information Choice*

limitations. According to an inter-bundle diversity ranking, a budget set with two bundles, where the first contains only right-wing views and the other only left-wing views, would rank higher than a budget set with two bundles, each containing both views. However, it would rank lower according to the proposed intra-bundle diversity ranking.

Despite the difference between inter and intra-bundle diversity, we can still find inspiration for a more precise definition of the freedom-enhancing diversity we refer to in intra-bundle diversity in the freedom of choice literature's discussions on inter-bundle diversity. One simple approach, suggested by Pattanaik and Xu (2000), involves categorizing items into groups where two items can be either similar or dissimilar, i.e., either in the same group or not. Diversity among items refers to such dissimilarity between them. This suggests classifying a set of information pieces as diverse or non-diverse based on whether they are in favour of a similar or dissimilar opinion. For example, two pro-life articles would count as similar, and a pro-choice and pro-life article would count as dissimilar. As explained above, this ability to be exposed to opposing views represents a basic form of diversity, prominently thought to be necessary for autonomous decision-making. Hence, enhancing this minimal intra-bundle diversity seems crucial for improving freedom of information choice. I therefore assume such a definition of diversity as a starting point.

Some literature also suggests that the definition of freedom-enhancing diversity should be sensitive to the degree of similarity among options (e.g., Klemisch-Ahlert, 1993; Bervoets & Gravel, 2007; Rosenbaum, 2000). To illustrate, choosing between a red and a white car is considered less diverse than choosing between a red car and a white train. Suppose we accept that a higher degree of dissimilarity further enhances freedom of information choice. Then, we might rank an option set with bundles of information pieces from The New Yorker and CNN as relevantly less intra-diverse than a Fox News-CNN bundle, for instance. Besides degree of similarity, we might also wish to rank option setoption sets according to the proportion of intra-diverse bundles, or the variety of information pieces that can be found in intra-diverse bundles. Notably, intra-bundle diversity is not necessarily the only relevant measure for freedom of information choice. Other measures of information quality might also be relevant, and could affect the ranking of option sets that include such bundles (e.g., Fox News-CNN vs. The New Yorker-CNN) according to the freedom of information choice they provide. I leave all these matters as topics for further investigation.

### *3. Freedom of Information Choice*

Now, let us examine in more detail why intra- rather than inter-bundle diversity is potentially helpful for the scenarios above. Consider the Rabbit Hole Scenario, where, intuitively, the option to either fall into a left or right-wing rabbit hole is problematic. The challenge here is that once an initial choice is made, subsequent recommendations provided by a recommender system portray only similar views. The option to continuously have access to diverse information, whatever initial choice one makes, is what may be helpful. That is, non-mutually exclusive diversity. Similarly, in the Intolerance Scenario, the possibility to choose between inter-diverse bundles translates to the possibility to either develop intolerance towards left wing views or towards right wing views, but not to the possibility of not developing such intolerance at all. Tolerance implies intra-bundle diversity within one's budget set. In the Bite-Size Information Scenario, inter-bundle diversity does not ensure the integration of diverse sources into one's daily routine. In all these cases, the problem is forming opinions without exposure to opposing viewpoints. Thus, what is lacking is intra-bundle diversity—having diverse pieces of information with an aggregated cognitive cost low enough to fit within one's cognitive budget.

### **3.6 A distinct property of the cognitive budget set and improving intra-bundle diversity**

How can we improve the intra-bundle diversity within a cognitive budget set? With monetary budget sets, one way to enable new bundles within the set is by lowering the costs of commodities. Similarly, we can change cognitive costs to fit diverse information within a cognitive budget. This requires acknowledging a distinct quality of cognitive costs, which affects how available bundles are formed.

Within the classical literature on monetary budget sets and some related discussions of freedom of choice, each commodity has a fixed price that is not influenced by one's choices (e.g., Pattanaik & Xu, 1990, 2000).<sup>63</sup> This might be a reasonable simplification for monetary budget sets, where all items could theoretically be purchased together. However, spending cognitive resources necessarily occurs over time. Information that was consumed in the past affect the cognitive costs of information consumed later on. For instance, familiar information costs less,

---

<sup>63</sup> Exceptions include literature on adaptive preferences (e.g., Bowles, 1998). However, the relationship between past consumption and cost differs in freedom of information choice because it follows cognitive patterns.

while information challenging existing views or causing cognitive dissonance costs more, due to tasks like switching focus or learning new information.

Therefore, the order of information consumption within the bundle affects whether that bundle resides within the cognitive budget set. For instance, let A denote a pro-life article and B denote a pro-choice article. The information consumption sequences  $(A, A, A, B, B, B)$  and  $(A, B, A, B, A, B)$  may have different cognitive costs. Consequently it is possible that only one of them will reside within the cognitive budget set.<sup>64</sup> This is due to the different forms of task switching, levels of familiarity, and possibly discomfort and cognitive dissonance in exposure to opposing views exhibited in each sequence, due to the differing order of consumption.

While the budget set analogy is useful, cognitive budget sets differ from the classical models of monetary ones. Both budgets sets include combinations that are jointly affordable given the budget. However, the modelling of the cognitive budget set needs to account for the fact that the order of consumption of information affects its cognitive costs. While in both cases it is possible to simplify series of item choices in a choice between bundles, in the case of cognitive bundles, the prices of those bundles should be influenced by order of information consumption. Hence, the price function should take into account previous items in the consumption sequence.<sup>65</sup>

It follows from this distinct quality of the cognitive budget set that achieving intra-bundle diversity it requires different techniques than achieving it within a monetary budget set. While presenting a precise model of the cognitive budget set is beyond the scope of this chapter, understanding this distinction already enables us to identify possible interventions that can increase intra-bundle diversity within a *cognitive* budget set. As explained, the cognitive cost of consuming a certain information piece  $i$  is not determined independently for each piece, but rather is a function of the sequence of pieces consumed prior to this piece  $(1, \dots, i - 1)$  as well as the current piece. Hence, we can influence the cost of diverse options at the  $i$ th position in the sequence either by (1) modifying that information piece, or by (2) altering something about

---

<sup>64</sup> Consequently, the price function, which may assign a cognitive cost to each sequence, must reflect these order-dependent effects. For instance, while the cost of the first instance of  $B$  in the former sequence depends on the preceding subsequence  $(A, B)$ , the cost of the first  $B$  in the latter sequence depends on the full preceding subsequence  $(A, A, A, B)$ , thereby capturing the differing cognitive burden as a function of the consumption history.

<sup>65</sup> One way to incorporate the significance of consumption order is to model bundles as finite sequences, and define cognitive costs via a function that assigns a “price” to each item based on the entire consumption sequence up until that item. In this chapter, I do not lay out a precise model of this type, and leave the different possibilities open. My goal is to explain the criterion a plausible model must uphold, namely, that order of consumption matters to costs.

the prior pieces of information consumed, or (3) by altering the function determining precisely how prior information influences the present cost. Let us examine an example of each.

Modifying the information piece at point  $I$  (1) can include, for instance, modifying the complexity level of that piece. For instance, it is possible to use an introductory pieces supporting a certain opinion rather than a more complex one. Reducing the complexity can balance the added costs of having intra-diverse information pieces. This also lowers the cost of processing subsequent complex pieces with similar views as it increases familiarity, preventing scenarios like the bite-size information problem.<sup>66</sup>

Modifying prior selected information pieces (2) can involve restricting the order of selections. For example, requiring opposing opinion selections after  $m$  same-opinion selections. This can reduce intolerance from lack of exposure and lower the cost of consuming diverse options. Thus, a larger proportion of initial choice combinations within one's budget set could be found in intra-diverse bundles, solving the Intolerance Scenario issue by ensuring initial diversity and leaving a larger proportion of the series of choices compatible with an intra-diverse bundle, keeping those options open.

It is also possible to adjust the precise way prior information influences the present cost, i.e., to achieve 3. Although prior information consumption inherently affects current costs, its effect can either be weakened or amplified. Recommender systems, for example, reduce the cognitive cost of consuming similar items to what was previously consumed by directing the attention of users' to such items. Similarly, measures that would direct attention to dissimilar options would reduce the cost of consuming intra-diverse bundles. This also prevents a lack of awareness from limiting options. Directing attention to alternative views is different from "attention capture" (Milano et al., 2020), where platforms are designed to be addictive precisely because attention is directed towards dissimilar options, hence it rarely becomes highly costly to resist a recommendation.

## Conclusion

In this chapter, I have discussed what I term "Freedom of Information Choice", the ability to autonomously form evaluative judgments. I explained why contemporary epistemic environments often harm this freedom. I then outlined two distinct characteristics of freedom

---

<sup>66</sup> This translates, then, to replacing a certain item  $X$  with a less costly yet similar item  $X'$ .

### *3. Freedom of Information Choice*

of information choice. First, ranking option-sets cannot primarily depend on evaluative judgments. Second, we should use what I termed a “cognitive budget set” to assess such freedom. I clarified how these distinctions render well-known measurements of freedom of choice less relevant to freedom of information choice. I then proposed that increasing intra-bundle diversity enhances freedom of information choice, and suggested a few ways in which that could be achieved, given that cognitive prices depend on the history of consumption. These measures, I claim, can help avoid different harms to freedom people exhibit in their daily lives.

---

## 4. The Epistemic Empowerment of Voters

---

### Abstract

The term voter empowerment is often used in public discourse. However, insufficient philosophical thought has been devoted to what such empowerment entails. In this chapter, I will develop an understanding of empowerment with respect to the use of information in the decision-making process preceding a vote, which I term “epistemic voter empowerment”. Drawing on literature on empowerment in development studies, I argue that epistemic voter empowerment should be measured by the development of particular capabilities, including a voter’s control over their epistemic process and their ability to use their vote to promote their interests. I develop a conditional subjectivist approach to these capabilities. I argue that all voters should possess the capabilities essential for epistemic autonomy, and once they do so, their own views determine which information is relevant to promoting their interests. I then explain the advantages this approach has over prominent alternative proposals that are directed at improving voters’ knowledge and briefly discuss practical implications for possible interventions with voters’ epistemic environments.

### 4.1 Introduction

Campaigns, government acts, and projects conducted by NGOs use the term “voter empowerment” (e.g., Environmental Health Coalition, 2025; Independent Sector, 2025). That term is commonly thought to refer to giving voters the power to use their votes and participate in politics. Voter empowerment is generally understood to be important and is presented as the goal of different interventions. However, within philosophical discussions, the precise meaning of such empowerment is underdiscussed. In this chapter, I discuss a particular type of voter empowerment, which I term “epistemic voter empowerment”. Adopting a particular understanding of such empowerment may aid understanding what type of interventions would empower voters in practice.

Voters’ epistemic deficits have mostly been discussed within evaluations of the ability of democratic decision-making to produce good outcomes. It is widely thought that voters are

generally uninformed (e.g., Bartels, 1996; Althaus, 1998; Brennan, 2016, 2022). However, this literature does not in itself tell us what is required for voters to be epistemically empowered, that is, whether the lack of knowledge of certain facts results in a lack of empowerment. This distinction is important for determining which of the interventions proposed for increasing the competence of majority votes would also empower voters. For that purpose, we need a better understanding of voter epistemic empowerment.

The chapter proceeds as follows. In section 4.2, I summarize common claims about voter ignorance and proposals for informing voters, explaining why those are not equivalent to questions of empowerment. In section 4.3, I extract some pertinent ideas from the existing literature on empowerment in the field of development studies. In section 4.4, I propose a specific understanding of epistemic voter empowerment, which I term a “conditional subjectivist” approach. In section 4.5, I discuss the normative stance of epistemically empowering interventions compared to methods suggested for raising the competence of voters. In section 4.6, I briefly discuss some practical implications.

## 4.2 Voter ignorance

Ample empirical literature discusses voter incompetence (e.g., Palfrey and Poole 1987; Lupia 1994; Delli Carpini and Keeter 1996; Bartels, 1996; Althaus, 1998; Caplan 2007; Somin 2013; Brennan, 2016, 2022). Some suggest that to vote competently, voters need to be well informed (e.g., Brennan, 2016). Others say it is enough for voters to mimic other well-informed voters with similar opinions (Lupia, 1994). This literature usually measures a competent vote using interviews, tests of knowledge of political facts, or by comparing voters’ vote to the vote of other voters with similar demographic characteristics that rank highly in such tests. For example, Delli Carpini and Keeter (1996) and Althaus (1998) tested voters’ ability to vote competently using data collected by the American National Elections Studies (NES), with questions such as identifying an office held by a public figure and the party that held a majority status in both houses of Congress (Althaus, 1998). Such studies usually conclude that voters are not well-informed (e.g., Somin, 2013, p. 192; Achens & Bartels, 2017)

However, the fact that voters do not know certain facts does not necessarily mean that they are not epistemically empowered. For example, voters might not wish to know certain facts they do not find relevant. Insofar as their judgments of what is relevant matter, part of their empowerment is allowing them to exercise their agency in this respect, since, as I will explain,

empowerment is linked to agency and self-governance. Hence, the question of whether voters are epistemically empowered remains open.

Following such analyses of voter ignorance, several theorists have suggested means of causing the electorate to produce more competent voting results. For example, Jason Brennan suggests creating a knowledge test, which includes about 50 basic questions, out of 200 overall questions posted online. Such questions can include – “who are their representatives, which party controls which house, what are recent bills that were passed, how much money is spent on this or that, questions about recent events, questions about social indicators (such as the unemployment rate)” (Brennan, 2022, p. 96).<sup>67</sup>

Brennan also suggests different ways of motivating voters to educate themselves sufficiently to pass such a test. Citizens could either be compensated monetarily for knowing the answers to the questions (“voter achievement day”, p. 96) or they could vote, but only their demographic characteristics will be considered, and their preferences will be simulated based on people with similar demographic characteristics and higher scores (“enlightened preference voting”, pp. 97-102), or, they could not be allowed to vote without knowing these facts (“epistocracy”, p. 102). The latter proposal bears some similarity to Mill’s suggestion that the better educated should receive more votes (Mill, 1865, Chapter VIII).

While such proposals may increase the chance that voters, or the people whose vote matters, will learn the specified facts, it still does not mean such interventions empower voters. For some such methods, the opposite may be true. For example, not allowing voters to vote without knowing such facts will disempower certain voters by denying them voting power. Additionally, proposals such as simulating the well-informed vote prevent self-governance by the people. Those opinions might not actually match what the people themselves would want were they well-informed. There might also be disputes regarding who is well-informed in the relevant sense. Furthermore, as argued in Chapter 1, it is challenging to identify demographic resemblance sufficient for a claim for true representation. Additionally, even if such preferences reflect what people would choose were they well-informed, it does not reflect their current choice. Hence, such methods provide less self-governance than actual voting. To the extent that autonomy and self-governance are needed for empowerment, such proposals will

---

<sup>67</sup> Of course, others have claimed that “being informed is not simply a matter of having been immersed in high-quality information, but also of having attained specific cognitive styles and strategies, and acquired a rich set of political information that has implications for how new information is subsequently processed” (Ahlstrom-Vij, 2022). However, such literature still relies on a basic set of facts for knowledge measuring.

therefore disempower voters. Such interventions also particularly harm voters' epistemic autonomy,<sup>68</sup> as voters are strongly encouraged to go down a particular epistemic path. This, again, may clash with epistemic voter empowerment.

Another approach to overcome voter ignorance is deliberative opinion polls (Fishkin, 2018). In those, a random group meant to represent the demographic characteristics in society receives information and deliberates for a short amount of time, after which they express their updated opinions. Some take those to simulate the opinion of the population were they well-informed (e.g., Fishkin, 1995). As James Fishkin suggests, "the deliberative poll can be thought of as an actual sample from a hypothetical society-the deliberative and engaged society we do not have" (Fishkin, 1995, p. 171). However, as explained in Chapter 2, this judgment is often inaccurate, because facilitators may lack objectivity and fail to present unbiased summaries of information, and because the claim for descriptive representation of the epistemic outcomes of the randomly selected group is often false.

Hence, we cannot rely on the literature on voter ignorance and consequent proposals for improving collective decision-making as a guide for interventions that would empower voters. We require a better understanding of what voter epistemic empowerment entails. Developing such an understanding deserves more attention than it has been given so far, as it can enable a path to more legitimate interventions in voters' epistemic environments.

Moreover, proposals that encourage citizens to learn specific political facts potentially attributes too much importance to memorizing information, especially when thinking of voter empowerment. Voters today have information easily accessible to them. When they begin their epistemic process to decide how to vote, they often turn to online resources if they realize that they are missing relevant knowledge that they find relevant. Therefore, to enable voters to reach a well-informed voting decision, it may be sufficient to provide them with easy access to relevant information and raise their awareness of its existence. The fact that voters cannot answer factual questions at a given point in time might not be a sign of lack of empowerment. It could be, for instance, that they would not have considered those facts relevant had they known them. Alternatively, they might realize their relevance later on, look them up and vote accordingly, even if those facts do not enter their long-term memory. In such cases, knowing

---

<sup>68</sup> By "epistemic autonomy", I simply mean autonomy in epistemic processes, that is, self-control and self-governance in using information to form opinions. I do not attach my account to a specific conception of epistemic autonomy, out of the ones discussed in recent literature (e.g., Goldman, 1991; Fricker, 2006; Pritchard, 2013; Matheson, 2024).

that fact in advance likely would not have changed how the voters voted. This consideration will factor into the approach I develop in the following sections.

### 4.3 Empowerment

As mentioned, the concept of voter or citizen epistemic empowerment is underdiscussed in the literature on epistemic aspect of democracy, a lacuna which I aim to correct. To do so, I will first survey the discussion of the term “empowerment” within the context where it has been most thoroughly discussed, namely, development.

Development work often sets as its goal to empower oppressed and deprived people (e.g., Khader, 2011). When considering the type of interventions that would lead to such empowerment, there is a shared understanding that empowerment is not about making decisions on behalf of people or forcing change in their lives. Rather, it is about enabling them to take control over their own lives.

In unpacking this idea,<sup>69</sup> a resemblance has been identified between empowerment and freedom of choice. Various theorists have taken empowerment to be primarily linked to improving people’s option sets (Kabeer, 1999; Hadi, 2001; Khader, 2011; Alsop and Heinsohn, 2005). Naila Kabeer (1999), for example, takes disempowerment to be a form of lack of choice. She describes empowerment as “the process by which those who have been denied the opportunity to make strategic choices acquire such an ability” (Kabeer 1999, p. 435). Following a similar line of thought, Ruth Alsop and Nina Heinsohn (2005) suggest empowerment relates to the question “if you wished to make decisions with respect to X, could you?” Serene Khader also claims that choice-based definitions of empowerment illuminate the important fact that people who are disempowered face very limited option sets (Khader, 2011, p. 180).

In the freedom of choice literature, Amartya Sen’s capability approach (e.g., Sen, 1987) has been highly influential. According to the capabilities approach, what truly matters in an option set is the actual capabilities the options provide. That is, being able to achieve valuable goals, such as being well-nourished, good health and self-respect, is what matters to freedom of choice, rather than the commodities one can buy, which compose one’s immediate options.

---

<sup>69</sup> I will not survey economic empowerment in the sense of poverty reduction, even though it is sometimes discussed in these contexts (see Pettit, 2012), as it is irrelevant to my debate.

The capability approach has been adopted as a central conception of empowerment by many in the development literature (Clark, 2005; Alkire and Deneulin, 2009; Byskov, 2016, 2018; Sen, 1999; Nussbaum, 1999; Nussbaum, 2011; Comin, 2018; Biggeri et al., 2019). This understanding of freedom is particularly important in this context, as it enables discussing the true capabilities people need to be empowered, rather than the mere options they have. For instance, people should possess the ability not to starve, rather than have the option to buy bread, which may or may not be sufficient for them not to starve. Additionally, the transformation to discussing capabilities rather than options allows acknowledging the differences between people, and the fact that a fulfilment of those capabilities would require different types of interventions for different people. For instance, for women, different dietary options are required for good health than for men (Sen, 2017, pp. 198-9).

However, to apply the empowerment approach in practice via interventions, it is necessary to decide what the relevant capabilities are that people require to be empowered. A common conception in development work is that empowerment involves enabling people to make choices aligned with their interests. Khader, for example, states that people “are disempowered to the extent that choices that would allow them to unambiguously advance their interests are unavailable” (p. 187). Similarly, empowerment can be thought of as the freedom to promote one’s well-being, a focus of the capabilities approach (e.g., Sen, 2017).

Finally, among the capabilities relevant to justice is the freedom to collaborate with others in pursuit of social change. Sometimes, taking control over one’s life cannot be done individually. Some social changes require collaboration among social groups. This has been noted in the empowerment literature. Jo Rowlands claims there is a personal, relational, and collective dimension to empowerment (Rowlands, 1997, p. 14). The collective dimension involves people working together to reach a more extensive impact than they could have had individually (p. 15). One aspect of empowerment, she writes, is “a sense of the whole being greater than a sum of the individuals, especially when a group tackles problems together” (p. 13). Iris Marion Young similarly discusses two types of empowerment. One involves developing individual autonomy, self-control, and confidence. The other involves a group of people collectively having influence over the social conditions that affect their lives (Young, 1997, p. 89). According to Julian Rappaport (1984) empowerment is the mechanism that causes people, organizations, and communities to gain control over their lives (p. 3).

These interpretations of the concept of “empowerment” do not yet tell us what type of access to and curation of information is needed to create a voter who is epistemically empowered. In the next section, I will use them as a basis for proposing such requirements.

#### **4.4 The epistemically empowered voter**

The capabilities approach has several important advantages. First, it acknowledges individual differences, which is also essential for epistemic empowerment, as people’s interests, values, goals, and consequently the information they require to promote their well-being, naturally differ. Additionally, the approach’s focus on enhancing agency is essential for any type of empowerment. Hence, I adopt the capabilities approach as a framework for epistemic voter empowerment, asking what minimal capabilities voters need to be epistemically empowered. Answering this question will determine what information should be offered to different people to enable such empowerment.

As we have seen, conceptions of empowerment and of capability sets focus on strengthening people’s agency, autonomy, and self-control, while equipping people with the ability to promote their well-being and their other interests, alone and with others. Following this logic, for epistemically empowering voters, the purpose of the capability set would be to allow voters control over their epistemic process and improve their epistemic autonomy, particularly within the process of gathering information to make up their mind about their vote. Additionally, epistemically empowering voters entails providing them with information that will enable them to vote in accordance with their interests and well-being.

For practical interventions, it is essential to determine which sub-capabilities are included under those overarching capabilities. These sub-capabilities help us determine which informational options, when made available to people, will enable their empowerment. Within the capability literature, there is a longstanding debate regarding the best way to determine essential capabilities. Two prominent proposals are Sen’s proposal to focus on options that people “value and have reason to value” (Sen, 2002, p. 5), and Nussbaum’s proposal for a set list of essential capabilities, which are necessary for living a life “worthy of the dignity of a human being” (2000, p. 72). Other theorists have also suggested combining the derivation of a list of capabilities from moral principles with procedural processes, such as democratic deliberation (Byskov, 2018).

I address this question specifically in relation to the capabilities required for an epistemically empowered voter, in light of the overarching capabilities discussed above. I propose a *conditional subjectivist* approach, arguing that from the standpoint of the role of the state all voters should be provided with a minimal set of capabilities that guarantee epistemic autonomy across time, and that once this condition is met, individuals should determine for themselves what information they find relevant and what their interests are. I elaborate on this account in the sections that follow.

#### 4.4.1 What limits available capabilities

In Sen's capability approach, the capabilities arise from the alternatives available to a person.<sup>70</sup> That is, helpful interventions open options that give rise to the desired capabilities. However, in what sense should options be available? In classic examples provided by Sen, such as the option of being well-nourished, the option needs to be available in the sense that one can physically access the relevant products. Also, the person should be able to purchase the relevant products given the limitations of their income.<sup>71</sup> A common approach in economics is depicting monetary restrictions using a budget set –the set of alternative combinations of commodities, i.e., different bundles, that are jointly affordable given the budget (e.g., Pattanaik & Xu, 1990; Xu 2004; Böhm & Haller, 2008).

For such modelling, a capability set (or just “capability”, Pattanaik & Xu, 2020) can be envisioned as the set of all the bundles of functionings that are available to the individual given their budget, where functionings are things the individual chooses to do or be, such as taking a walk or being well-nourished (e.g., Kaushik et al., 2011, Pattanaik & Xu, 2024). A “utilization function” builds on the person's unique “conversion factors”, i.e., characteristics of the person or their surroundings that enable the conversion of commodities found in the budget set into functionings. For example, to turn a car into the functioning of traveling to a certain place, one needs both good roads and legal permission to drive. Hence, those serve as conversion factors. Thus, the utilization function can transform a commodity bundle into a functioning bundle. In this manner, the bundles of available functionings still depend on the original budget set. All

---

<sup>70</sup> Some capabilities, such as breathing clean air, are available across all alternative. Hence, they are not a choice that one can opt out of.

<sup>71</sup> In other classical examples, such as the capability for political influence, laws and regulation also limit the capability set (Clark, 2005).

functioning bundles feasible for the person in this way are included in the capability set (Sen, 1999; Leßmann, 2004; Pattanaik & Xu, 2020).<sup>72</sup>

When it comes to information options, a different approach is needed for what options count as available. This is because money is often not the primary constraint on voters' epistemic empowerment. To provide people with real epistemic power, it is necessary to consider human cognitive limitations, as explained in Chapter 3.

Today, in wealthier nations, the main reason voters are not epistemically empowered is not a lack of access to information or monetary constraints. I have defined epistemic empowerment of voters as the ability to control their epistemic process and vote in accordance with their interests and well-being. As explained, the literature on voter ignorance does not in itself address the question of voter empowerment, since voters can choose to purposefully be ignorant of specific facts they do not deem important. However, voters do seem to lack empowerment as they are not only ignorant but seem to lack control of their own epistemic process. Many voters report they perceive themselves as lacking information they themselves deem essential (Graham, 2020; Asano, 2023). This, as opposed to failing to meet someone else's standards, shows lack of control and clear lack of ability to pursue their interests.

Access to the internet allows voters theoretically have access to vast amounts of information that could enable them to pursue their interests and well-being. However, this abundance makes it cognitively demanding to identify the most relevant options to casting their vote. In such cases, the abundance of information becomes "a hindrance rather than a help" (Bawden & Robinson, 2020, p. 2) to voters' control over their epistemic process and ability to pursue interests. Researching a topic thoroughly usually involves searching for the most relevant information on multiple platforms and aggregating information from multiple sources. The obstacles voters experience are not due to lack of existing information, but rather due to a limited cognitive capacity to process it, and the opportunity cost of doing so.

Moreover, voters are also disempowered due to finding themselves in filter bubbles and rabbit holes (Milano et al., 2020). Their attention is being guided by algorithms online, sometimes distracting them from information they would have valued. They also encounter biases and misinformation, which they cannot assess critically due to lack of sufficient background knowledge (Milano, 2020). Cognitive biases such as familiarity and confirmation bias also

---

<sup>72</sup> This is somewhat simplified, as Pattanaik and Xu point out, the choices of many consumers can affect the price vector and hence the budget set. However, this is the base comparison.

weaken their control over their epistemic process (e.g., McKee & Stuckler, 2010; Boston, 2017; Jacobs, 2016). These sources of lack of empowerment are also connected to cognitive, rather than monetary, limitations. Due to a limited cognitive capacity to process information, voters often use heuristics such as aligning with party lines, and hence are more vulnerable to populist rhetoric and political manipulation (e.g., Sabl, 2010; Noggle, 2021; Rozīna & Karapetjana, 2009). Unsurprisingly, it is commonly claimed that voters often unknowingly vote against their own interests (e.g., Mansbridge, 1983, pp. 24–26; Althaus, 1998; Somin, 2004; Brennan, 2022).

Hence, I suggest the best way to epistemically empower voters is to review the capability set that stems from their “cognitive budget set”. As mentioned, an individual’s monetary budget set represents the set of alternatives available to them given their income and the prices of the goods in the market, where each alternative is a bundle of commodities whose aggregate price is equal to or lower than the individual’s income (Pattanaik & Xu, 1990; Sen, 1993; Xu 2004; Miyagishima, 2010). Analogously, their cognitive budget set is the set of combinations (i.e., bundles) of information options that can be jointly consumed within the limits of a person’s cognitive budget. As explained in Chapter 3, in a manner analogous to the relationship between commodities and money, the different information options one has come at a different cognitive cost. Also, cognitive resources come out of a limited budget in any given timeframe. Hence, it makes sense to treat our cognitive capacities as determining a limited budget for investment in particular topics, such as a particular vote.

Therefore, the question we should ask is what interventions would allow sets of bundles of information-related functionings to fall within the cognitive budget set in a way that allows voters to have epistemic autonomy, control over their epistemic process, and the ability to become informed about which policy would best advance their wellbeing and interests, given their unique conversion factors. This epistemic capability set needs to arise from the set of information options available to the voter before casting their vote.

This definition entails specific types of sub-capabilities that need to reside within the cognitive budget set. Although importantly, just like other capability sets, the information needed for providing these capabilities differs between people. I view this as a significant advantage compared to the lists of facts voters are required to know to be considered “well informed” according to the prominent theories discussed in Section 4.2 of this chapter, which do not take into account voters’ differing interests, abilities for crucial thinking and information processing, etc.

#### 4.4.2 The relevant sub-capabilities

In this section, I outline the sub-capabilities that should be promoted to ensure that voters can control the epistemic process leading up to their vote and vote in a way that promotes their interests and well-being.

As explained, I propose that under the condition that certain capabilities essential for agents to have epistemic autonomy are continuously available within people's cognitive budget sets, people should be allowed to choose what information they deem relevant. This will allow both control over their epistemic process and the pursuit of their conception of their own well-being, under circumstances of epistemic autonomy, avoiding many cases of problematically shaped views.

In this proposal, I offer a new middle ground for the role of the state, which I believe is needed for treatment of contemporary issues with voter empowerment. I claim that this type of intervention is what is justified, at least if we believe that it is outside the role of the state to make value judgments on issues in large dispute among the population. The latter view is prevalent in philosophical thought. Liberal thought commonly supposes the state should maintain a level of neutrality with regards to evaluative questions that are in dispute among the population (Garnett, 2016). Rawls and Sen (e.g., Sen, 2017) claim the state should adopt “a point of view that everyone can adopt on an equal footing” rather than a “personal slant” (Rawls, 1971, p. 516–17). On the other hand, I take the state to still have a role in intervening with individuals' epistemic environments. While private companies, such as social media companies, shape the majority of voters' epistemic environments, harming their freedom of information choice, I consider measures that treat individuals as deserving of epistemic autonomy as justifiable. This is especially true for *epistemic* autonomy, which is arguably the root of any possibly legitimate agreement to forgo other types of autonomy. In the next section, I develop this line of thought, justifying this view specifically with respect to epistemic empowerment of voters.

#### 4.4.3 Justifying the conditional subjectivist approach

To see why a conditional subjectivist approach is needed, let us first see why pure objectivist and pure subjectivist approaches fail when it comes to the capabilities required for being an epistemically empowered voter.

Perhaps the most prominent objectivist account of capabilities was proposed by Martha Nussbaum. Nussbaum's prominent account provides a set of what she takes to be a minimum

threshold of central human capabilities (Dixon & Nussbaum, 2011), which she claims are needed for basic human dignity (Nussbaum, 2000). These include life and bodily health, but also capabilities such as imagination and thought, play, and control over one's environment (Nussbaum 2006, p. 76).

If the definition of the epistemically empowered voter were to follow this view, it would entail that voters are epistemically empowered when information options that explain which policy promotes the particular capabilities Nussbaum lists reside within the cognitive budget set. However, I reject this direction for three reasons.

First, providing information only on those particular capabilities will violate epistemic autonomy according to prominent conceptions of such autonomy. Mill prominently opposed state censorship of particular views, even if they are wrong. This was part of his interpretation of epistemic autonomy as enabling people to determine for themselves which views they would like to accept or reject (Mill, 1859). This applies to interventions favouring particular worldviews over others. Hence, following an extensive list of one particular view of well-being, such as Nussbaum's, would violate this conception of epistemic autonomy. As epistemic autonomy and control over one's epistemic process are essential for epistemic empowerment, this seems to be a futile path for epistemic empowerment.

Second, since voting is the mechanism by which governing members of the state are elected, we should be particularly cautious about the state interfering with voters' epistemic environments preceding their votes in a way that entails particular conceptions of what the good life consists of, and which might favour a particular way of voting over another. It seems this should not be the role of the state in such situations. There is a common conception that elections are meant for voters to determine general goals or values, while the role of elected representatives and other mechanisms is to advance these goals (Christiano, 1996, 2018; Hanson, 2013). This conception implies that values, goals and conceptions of the good should be determined bottom up through the voting process. They should not be determined indirectly by the state shaping voters' epistemic environments prior to their votes in a way that influences this vote. Hence, any state interventions with voters' epistemic environment should not be based on overly specific conceptions of the good life or the relevant values.

Finally, epistemic autonomy involves a fundamental type of autonomy, which is essential for forming evaluative judgments that are later used for achieving other types of autonomy. For instance, a person may be considered autonomous if they are able to pursue the options that

they judge to be valuable, as explained in Chapter 3. Hence, if self-governance is not provided at this stage, it will lead to a graver lack of autonomy, due to all the consequent uses of the evaluative judgments formed. Such autonomy is arguably even more important than not being able to agree to be sold into slavery, which Mill takes to be an important limitation needed to protect freedom (Mill, 1859, chapter 5). This is because, even if we were to take a person consenting to being sold into slavery as legitimate, that would plausibly only be the case if that decision was made in a reasonably epistemically autonomous manner by that person. If we were to suspect it was based on brainwash or manipulation, that decision would seem less legitimate. Since the epistemic capability set is used in shaping evaluative judgments, it is prone to manipulation, where the capabilities given could influence people's basic world conceptions. Hence, it is essential this particular type of autonomy is not being infringed on by interferences adopting a specific conception of well-being.

The case of abortions can exemplify the problem with relying on such an objective list of capabilities. Together with Rosalind Dixon, Nussbaum has claimed that laws restricting abortions violate central capabilities, including bodily integrity and health (Dixon & Nussbaum, 2011). Hence, abortion rights are protected under Nussbaum's conception of central capabilities. This idea has been disputed in the literature, as it does not seem to arise from "overlapping consensus" (Robeyns, 2006), which Nussbaum has claimed as a central justification for her list of capabilities (Nussbaum, 2000). According to Nussbaum's list of capabilities, voters should have access specifically to information regarding which policies will increase their chances of having abortion rights. This would skew the information accessible in favour of pro-choice options, which seems to undermine epistemic autonomy. The three arguments provided above suggest that information should be introduced in a balanced way, including both pro-life and pro-choice options. Notably, such a balanced information sample will be included under the requirement for having access to opposing views within the cognitive budget set, which is one of the objective requirements I list below.

Contrary to Nussbaum, Sen has prominently objected to providing a list of capabilities that will count as objectively promoting well-being. He says that specifying a list of "proper" objectives and values... can be a source of a substantial "unfreedom" preventing people from using reason to determine their values and goals (Sen, 2002, p. 6). Hence, he argues that relevant capabilities involve "what a person would want to have and have reason to value having", focusing on people's own attitudes towards options (Sen, 2002, p. 5). What it means to "reasonably value" an option remains vague in Sen's work, and different interpretations have

been discussed (Clark, 2005). He has sometimes claimed that the preferences people have reason to value are the ones that stand up to scrutiny in individual or public deliberation (Sen, 2002; Byskov, 2020).

Sen's proposals are, in themselves, also insufficient for epistemic empowerment. As explained in chapter 3, when it comes to information options, primarily relying on pre-existing attitudes towards options to assess the option set is problematic. It cannot provide freedom when such judgments have been problematically shaped. In that chapter, I also argued that, for this reason, the prominent interpretation of "reasonably valued options" in the freedom of choice literature, which also relies on pre-existing evaluative judgments, is problematic for evaluating information options. To the extent that people's judgments are already problematically shaped, the interpretation of "reasonably valued" options as what stands up to scrutiny is also problematic. Public deliberation is also based on societal principles and pre-existing judgments. Hence, my criticism of such an understanding of "reasonableness" applies to here as well. That is, if applied to information options, it will prevent information that goes against societal conceptions from being available and will consequently prevent societal openness.

Even though *purely* relying on evaluative judgments is insufficient, taking them into account is essential for epistemic autonomy. This is because, if people cannot use their own judgments to navigate information options, they cannot be said to have true control over their epistemic process. Also, it likely harms their well-being. Hence, the ability to exercise one's judgments is a main aspect of empowerment.

For this reason, I suggest that the most plausible approach is a *conditional subjectivist* approach, where what people find relevant is considered, however, under specific conditions. Those conditions serve to break the vicious cycles of problematically shaped judgments explicated in chapter 3. They include minimal conditions needed for people's judgments follow from autonomous decision-making regarding the capabilities they find valuable and can truly reflect control over their epistemic process.

More specifically, I submit that under certain conditions, promoting epistemic autonomy entails providing, within the cognitive budget set, the information people themselves find relevant to their well-being (as they conceive of it), and to their other interests and goals. These necessary conditions, as I explain below, include the continuous access within one's cognitive budget set to: (1) awareness of the various information options, e.g., topics and types of options that could be relevant; (2) information on opposing views; (3) the ability for epistemic collaborations; and

(4) information on whether a policy promotes or harms capabilities essential for the very possibility of being an epistemically empowered voter (e.g., non-starvation, literacy, the right to vote). Together, these conditions ensure that multiple epistemic paths are genuinely open to voters, and consequently, that their perceptions of relevance are not due to lack of true choice. Interventions making such information available within the cognitive budget set will be epistemically empowering. The need for this list is justified by it being essential for epistemic autonomy.

Both Sen and Nussbaum claim that a given list of capabilities need not be a complete list of all capabilities required for well-being. Nussbaum takes her list to include minimal essential capabilities and to be open to future revisions (Nussbaum, 2000). Sen claims the central aim is not to describe a utopian ideal, but to decrease the level of injustice in current societies (Sen 2006, Clark, 2005). I follow this line of thought in suggesting that the list detailed below is a minimal list of essential capabilities that are required for voters' epistemic empowerment, without arguing they are necessarily sufficient for an epistemic ideal. They are needed as they allow voters to have multiple epistemic paths, which is a crucial part of epistemic autonomy. Hence, they provide good initial guidelines for interventions in voters' epistemic environments to improve on the current situation.

#### **4.4.4 The minimal set of objective capabilities**

The minimal set of objective capabilities enables agents to continuously access the type of information that allows them to have multiple epistemic paths, including the possibility of changing direction during their epistemic process. This ensures that their decision not to choose certain paths is not due to a lack of ability. This, I claim, is essential for control over their epistemic process and for epistemic autonomy.

##### **4.4.4.1 Opposing views**

Improving voters' autonomy and control over their epistemic process plausibly involves enabling voters to have access to diverse points of view and opposing arguments within their cognitive budget. Literature on autonomy has long taken exposure to diversity of information and viewpoints to be a key component of it (e.g., Christman, 1991; Friedman, 2003; Meyers, 2005). This is particularly true for epistemic autonomy. As explained in the previous chapter, according to Mill, exposure to diverse opinions enables individuals to discuss and justify their thoughts. It also enables people to challenge their opinions and have reasons to change them. Consequently, Mill supports free speech and opposes state censorship. Of course, the lack of

censorship does not entail people being force-fed opposing views, only that such views exist in the public sphere and are easily accessible to people. Such access is essential for epistemic autonomy because it enables people to have different possible epistemic paths they can follow, i.e., have the possibility of changing their mind about an issue. Without access to opposing views, people may find themselves conditioned to believe the one-sided view that is available. They might not realize any compelling challenges to it exist. When such access does exist, their choice to adopt, develop, and investigate particular views seems more substantial, given that they had the ability to do otherwise.

An important addition on Mill's original notion, which is essential when many of voters' obstacles to empowerment come from an overabundance of information or strategic curation of information, is that such accessibility is required within the limits of the cognitive budget set. As explained in chapter 3, people often have theoretical access to opposing views, but overexposure to one-sided information from polarized media sources has limited their cognitive capacity to seriously consider opposing views. Hence, such access means a person can entertain opposing views and is not led by vicious cycles shaping their beliefs. Importantly, this requires the ability to jointly access opposing views, rather than having the mutually exclusive ability to access opposing views, which people already have today. Hence, as explained in chapter 3, intra-bundle diversity, i.e., joint exposure to opposing views, should reside within the cognitive budget set. Importantly, such views need to be continuously accessible, so that voters cannot enter informational echo chambers that they later cannot exit. An empowering intervention may hence include limiting certain long-duration one-sided epistemic paths, at least in cases where such paths will result in the consumption of opposing views requiring cognitive costs beyond that individual's budget. Alternatively, it may require directing attention to opposing views.

#### **4.4.4.2 Awareness of the existence of various information options**

Importantly, self-control also implies that people's judgment regarding which information is relevant is based on awareness of various options. Otherwise, people's judgment or prioritization would not reflect their complete ranking, but a random subset that was determined by the previous information they accidentally were exposed to. Suppose, for example, a voter was trying to form an opinion about a new abortion policy. Initially, they might think they solely care about foetal development stages. However, this might be due to their lack of awareness of the fact that many women have illegal abortions when legal ones are not available. This aspect of the question might have never crossed their minds. Once it is

brought to their attention, they may realize that they consider it to be highly relevant. Hence, another guideline for information curation is that various topics related to the question are brought to the voter's attention. Information curation that follows what voters themselves find relevant should only come after such awareness is cultivated.

Interventions to include awareness of options within the cognitive budget set could, for example, include directing people's attention to the existence of various types of information, as is done by recommender systems (Burke et al., 2011). Instead of recommendations for similar material, one could receive suggestions for directions of investigation that cover a wide range of possible topics or summarized information about existing research areas. This is because active direction of attention is cognitively costly (Van Merriënboer, et al., 2002).

#### 4.4.4.3 Epistemic collaboration

Sometimes, the information needed to determine which policy will best promote voters' interests and well-being cannot reside within an individual voter's cognitive budget set. This is because some topics are inherently complex. In such cases, supplying people with the capability to pursue their interests implies enabling epistemic collaboration for the sake of pursuing shared interests. This could be thought of as a form of *distributed epistemic empowerment*.

Epistemic collaborations involve people coming together to aid each other in studying a topic or to divide the epistemic labour between them to collectively reach better decision-making. Examples of epistemic collaboration include co-authoring work, engaging in discourse where each learns of the other's ideas, recommending information sources to one another, etc. Some forms of epistemic collaboration typically occur between voters, who engage in informal discussions about politics or read articles written by others whose salaries are partly funded by their tax money. However, epistemically empowering voters involves enabling and facilitating particular epistemic collaborations that are helpful in promoting these voters' goals, interests and well-being. Those involve *beneficial* epistemic collaborations. A voter may read an article that was recommended to them, but does not relate to their interests or what they find relevant. They may also engage in discussions with members of an opposing political group that are poorly grounded in facts and often result in conflict. These collaborations would not be particularly epistemically empowering. Importantly, empowering collaborations should be beneficial when taking into account voters' cognitive limitations. Hence, a recommendation for an article that might be somewhat relevant to the voter, but not worth the opportunity cost of

exploring more relevant options, hence overall derailing that voter's epistemic process, will also not be an empowering epistemic collaboration.

Epistemically empowering options plausibly include the ability to find other people with shared interests, values and goals and to collaborate with them in a manner that enhances one's ability to pursue those interests, values and goals. This idea reflects a broader understanding of empowerment as a social phenomenon. Historical social changes that empowered people were often caused by collaborations among people sharing similar interests or values. In the Women's Suffrage Movement and the Civil Rights Movement, for example, individuals collaborated by marching together and conducting collective boycotts. The collaboration could achieve what the individual alone could not.

We can understand distributed epistemic empowerment in similar terms. That is, facilitating connections that allow the collective to achieve what the individual cannot. Due to the overabundance and complexity of information, individuals can only access a fraction of the information relevant to their interests. Since decisions based on more complete information tend to be better, this limitation reduces people's ability to promote their well-being and interests. For example, someone trying to vote for the economic policy that would best promote equality, job growth, and higher wages for minorities may find it extremely difficult to determine the most effective policy, due to the complexity of assessing the long-term influence of different policies. If this complexity undermines their ability to pursue their own interest, it constitutes a lack of empowerment. In such cases, a more empowering approach may be to enable them to participate in a collaboration that promotes shared goals through a division of epistemic labour. This way, individuals can rely on knowledge they would wish to include but that is too complex to master alone.

Suppose a group of people sharing similar interests have gathered together to study a topic prior to a vote. What type of division of epistemic labour would enable them to better promote their shared interests? They might decide, for instance, to conduct a pre-vote among themselves, preceding the actual vote. Then, each voter from the group would the vote according to the majority vote in the pre-vote (see Goodin & Spiekermann, 2015 for details of such a collaboration)<sup>73</sup>, or at least use the information revealed in the pre-vote in their decision-making process preceding the actual voting. Or they might decide to raise their group's competence within the final vote by making sure that as a group, their majority vote has

---

<sup>73</sup> Goodin & Spiekermann discuss masses coming together to overcome elites using such a method.

increased competence. For this purpose, they could divide among themselves the information they consume preceding a vote in a way that aims to enhance their group competence, even at the expense of their individual competence levels.

Enhancing group competence through a division of labour could be done using the understanding prevalent in jury theorems (e.g., Ladha, 1995; Goodin & Spiekermann, 2018) that low dependence between votes can increase collective competence on majority votes. This might be achieved by ensuring a diversity of information between people who choose to participate in such an epistemic collaboration. That is, if different information pieces are being made accessible to different people, allowing them to develop different expertise, or study different aspects of the same topic, their collective vote could reach a higher competence. Suppose that all voters believe that the same three pieces of information are the most relevant so that, were they not in a collaboration, they would all consume the same three pieces. Nonetheless, being in a collaboration, they might choose to diversify their sources as a group. This could come from the understanding that, although individually their competence may slightly drop, as a group promoting the same interests, they are more likely to promote those interests during the vote.

Hence, one path to empowering voters is making it easier for them to learn of other individuals with similar interests or values who wish to enter an epistemic collaboration. That could be achieved, for example, by enabling a type of “matchmaking” service for potential collaborators. Additionally, interventions that distribute diverse information between members of the group according to agreed-upon principles can aid in such collaboration. Such a division of labour could reduce dependence between voters and raise their group competence.<sup>74</sup> Importantly, for

---

<sup>74</sup> Let's see why that is the case. Suppose each voter has a competence level of 0.6, that is, each person votes correctly on 60 percent of the issues. If voters were maximally dependent, they would all be right on the *same* 60 percent of issues. In that case, there would be complete consensus on 60 percent of the issues (everyone voting correctly), and complete error on the remaining 40 percent (everyone voting incorrectly). The overall group competence would remain at 0.6. While such alignment may be advantageous under a unanimity decision rule, it is less effective under a majority vote. To maximize group performance under majority rule, it is preferable for correct votes to be spread across different issues, so that on as many topics as possible, a *small majority* of voters is correct. Suppose there are  $n$  voters and  $x$  topics to vote on. Then, across all topics, there are  $xn$  individual votes. Achieving a correct majority outcome on all topics requires at least  $(xn/2) + x$  correct votes, just over half the votes for each topic. Therefore, to maximize the number of correct majority outcomes, it is optimal to distribute correct votes across topics rather than concentrate them. Yet, any improvement in the spread of competence across topics, even if not maximal, will result in the majority vote being correct on more issues, i.e., in higher competence. Such improvements are possible to achieve without complete independence. Hence, a helpful type of diversity is one that allows voters to develop different expertise, making them much more likely to answer certain questions correctly than others, even if their expertise areas are not completely independent from one another. This logic guides uses of diversity in the training data in Machine Learning methods (See, for example, Li, 2021).

a division of labour to work, the group must collectively possess the necessary resources. For example, a financial boycott can only succeed if the group has enough combined economic influence to put pressure on existing institutions. The same is true in our case for mental and epistemic resources. That is, the shared interest group that agreed to participate in the collaboration should collectively be able, within the limits of their respective cognitive budget sets, to handle a complex topic. Hence, interventions for summarising information and fitting the complexity level to group capabilities will also aid in distributed epistemic empowerment.

Distributed epistemic empowerment is also essential for epistemic autonomy. Feminist conceptions of autonomy tell us that this concept should not necessarily be taken to mean being self-sufficient (Jaggar, 1985; Stoljar, 2024). The feminist conception of “relational autonomy” values social relations and places importance on the influence of social context. For instance, in bioethics, some theorists have taken relational autonomy to entail decision-making about major events, such as end-of-life decisions, to involve shared decision-making with loved ones (Gómez-Vírsedaet et al., 2019).

We may also treat epistemic autonomy in this general spirit, measuring it not only in relation to people’s ability for to achieve certain goals individually, but also in relation to their ability to form epistemic relations and collaborations. The latter, we may suppose, is just as important, and if people are deprived of opportunities for such collaborations, their epistemic autonomy is lacking. This is especially the case if many of the epistemic goals they wish to promote cannot be sufficiently achieved when working individually, due to the complexity of the information involved and their personal cognitive limitations. Hence, I suggest extending the notion of control over the epistemic process to incorporate the ability to form epistemic collaborations, where people harness their shared resources to promote their shared interests.

Consequently, if we want to take people’s judgments as to what they find relevant to express true control of their epistemic process, these judgments need to be shaped in a context in which epistemic collaborations are possible. This is because people may deprioritise certain information and deem it insufficiently relevant if they suppose they cannot process it given their limited cognitive resources. However, suppose they could collaborate with others and expand overall resources. Then, perhaps they would have deemed the information relevant and thought the group should collectively become informed of it. Hence, evaluations that are not based on the ability to collaborate may express preferences that have been adapted to a sub-ideal set of options, and that do not reflect an empowered choice.

Importantly, a division of epistemic labour does not contradict improvement in other essential capabilities, such as access to information one deems relevant or awareness of options. There is ample information available regarding each specific topic one deems relevant, as well as about the necessities of epistemic voter empowerment. Suppose, for example, one is facing a vote on a COVID-19-related policy and wishes to compare the risk to one's health vs. the risk to one's economic security caused by a specific policy. The information relevant to such an assessment could be highly complex. For example, a comparative analysis with other countries may be necessary to assess the potential impact of a policy on the spread of the virus and its economic effects. Hence, different people could opt into a collaboration, based on a shared understanding of the importance they attribute to these two goals. One possible division of epistemic labour could involve different people studying the effect in different countries. The epistemic diversity achieved will likely raise the collective competence of the group's votes as long as each voter gains a competence of above chance level, and different voters are competent on different questions.<sup>75</sup>

#### **4.4.4.4 Information about policies that would threaten capabilities that are preconditions for being an epistemically empowered voter**

As mentioned above, a broad list of objective capabilities, such as the one suggested by Nussbaum, should be excluded from the definition of an epistemically empowered voter. However, I maintain that a limited set of objective capabilities that are preconditions for epistemic autonomy should be included. These are capabilities whose absence would be in direct contradiction to someone being an epistemically empowered voter. Those include capabilities necessary for existence (e.g., not starving), capabilities required for being a voter (e.g., access to voting rights), and ones essential for epistemic empowerment itself (e.g., literacy). Because the absence of these capabilities directly contradicts the very possibility of being an epistemically empowered voter, the ability to pursue their continuation must be considered part of what defines such empowerment. Hence, the cognitive budget set should include information about whether the policy or candidates put to vote could infringe upon one of these rights of voters. Such information could be made to reside within the cognitive budget set by, for instance, providing summaries of relevant points in accessible language or directing attention towards such information.

---

<sup>75</sup> See previous footnote.

#### 4.4.5 Information people find relevant

According to the conditional subjectivist approach, if the aforementioned capabilities remain continuously available within the limits of people's cognitive budget sets, interventions to include more information voters themselves find relevant within their cognitive budget set (or distributing it among collaborating groups who agree it is relevant) will also count as epistemically empowering.<sup>76</sup>

When people have continued access to the aforementioned types of information, including being aware of information options in existence and continuously being cognitively able to process opposing views, this means they have power in paving the path of their epistemic process. Hence, their judgments about which information is relevant can be taken to reflect control over this process. Such control is needed for epistemic autonomy. Consequently, we can apply Sen's reasoning and allow people to determine what is relevant to their well-being, interests, and goals.

The idea of empowering individuals by ensuring that information they find relevant to decision-making fits within their cognitive budget also aligns with prominent interpretations of how people pursue their interests. An example is the “enlightened preferences” interpretation used in the literature on voter ignorance (e.g., Althaus, 1998; Brennan, 2016), which has often been used to justify autonomy-violating interventions that disempower voters, as discussed in section 4.2. This prominent definition of people's interests was first given by Jane Mansbridge, who takes “enlightened preferences” to be what option they would prefer if they could see how things would turn out for either option (Mansbridge, 1983, pp. 24-6).

At first glance, that would seem to entail simply giving them full information about potential consequences of each policy. However, notice that what is relevant according to this interpretation is a vote being aligned with people's *preferences* regarding the way the future will look. What matters for people to vote in alignment with their own preferences is only the information about future implications they themselves deem relevant, as that information could potentially have changed their preferences and influenced their votes. Information they would have been exposed to that would have resulted in no change of preference is therefore unneeded for them to vote in alignment with their “enlightened preferences”. Importantly, though, there

---

<sup>76</sup> This implies people can decide what their interests and well-being consists off. It could entail, for instance, that on a policy question regarding lock-up to prevent the spread of an epidemic, people can decide whether the impact on risk to physical health, impact on the economy, taxes and employment, or both, are priorities. Following such decisions, information can be recommended.

could be aspects of the policy consequences that the voter was not aware of in advance, or did not sufficiently reflect on, but once they would have seen the new reality, they would have realized its relevance their preferences. Hence, what is important for discovering people's "enlightened preferences" is what they find relevant *given* that they are made aware of different related topics and questions, and can properly reflect on them to see if those views are stable. An essential part of the latter is the ability to examine challenges to their views. Additionally, as Mansbrige notes, finding out consequences of policies is highly complex, and voters are unlikely to be able to know in advance all the policy's implications, which is needed for fully enlightened preferences. Hence, methods that increase likelihood of processing complex information improve on the ability to reach enlightened preferences. This could be achieved by epistemic collaborations that enhance collective resources. Therefore, the proposed conditions (awareness to options, access to opposing views, epistemic collaborations and basic preconditions of voting, literacy, etc.) will improve voters' ability to track their enlightened preferences. Given such conditions, tracking their enlightened preferences involves providing them with information they find relevant, and hence may have a real influence on those preferences. Thus, the conditional subjectivist approach would work increase people's ability to pursue their interests according to this conception.

Notice that for the objective capabilities proposed to be continuously available within the cognitive budget set, certain paths of investigation might not be possible, even if people find them relevant. For instance, if viewing only one-sided information for a long time makes a person cognitively incapable of viewing and processing opposing views, then the option for opposing views would not be continuously available within the cognitive budget set. Hence, an intervention enabling the continuous viewing of one-sided information within the limits of the cognitive budget set would not count as epistemically empowering, even if the person finds it most relevant. The rationale for this judgment, which differs from the judgment of the individual in question, is that in such scenarios, the evaluative judgment might be problematically shaped and not indicate true control. Additionally, allowing information that people find relevant to fit within the cognitive budget set might mean adjusting its complexity level, for instance.

So far, we have discussed a few characteristics of the information that are likely essential for people to have the desired capabilities. However, to convert such information option sets into the capability set, i.e., to truly provide people with the desired capabilities, it is important to note that people's unique conversion factors will affect what information is most helpful for

their empowerment. For instance, if one has higher abilities in comparing and contrasting different complex views, more nuanced and complex pieces would better promote the desired capabilities. Alternatively, if someone lacks integration skills across various information sources, perhaps fewer but more comprehensive ones would be more empowering.

To conclude, I have listed a few important criteria for the cognitive budget set to be empowering, representing a conditional subjectivist approach to the desired capabilities. I have argued that interventions to empower voters by enhancing their access to information should focus on curating the informational environment to promote these capabilities. Importantly, those capabilities should be enabled when taking into account the limits of people's cognitive budget set. Notice that for voters with access to the internet, almost all these capabilities could be thought to exist already if one does not account for cognitive limitations. Due to people's access to the internet, they have theoretical access to almost all information in existence, including the best information possible for promoting each of the listed capabilities. However, the current curation of information, characterized by constant distractions, algorithmic bias, and the complexity of content that is often not presented in accessible formats, prevents voters from being epistemically empowered.

#### **4.6 Revisiting the voter ignorance literature**

The unpacking of voters' epistemic empowerment provided in this chapter supports interventions with voters' epistemic environment that align with the presented conditional subjectivist approach. As explained, such interventions include making accessible the information people themselves find relevant, conditional on a list of essential capabilities remaining continuously accessible. This list of capabilities is essential for voters' epistemic autonomy. This condition bears resemblance to Mill's idea that it should not be possible for people to agree to be sold into slavery (Mill, 1859). Mill states that the "principle of freedom cannot require that [a person] should be free not to be free" (chapter 5). Similarly, the minimal restrictions on information consumption include only those that allow people to have continuous access to alternative epistemic paths, rather than making one decision that prevents future epistemic freedoms (e.g., as they lack awareness to options, do not have the mental capacity to internalise views that challenge their current path, or unknowingly vote for a policy that will prevent them from voting in the future). Although, as mentioned above, this is

plausibly even more essential when it comes to epistemic autonomy than with being sold into slavery as it is a more fundamental type of autonomy.

Interventions guided by this proposal can be contrasted with the interventions prominently offered to tackle voter ignorance, which were introduced in section 4.2. Recall that such interventions primarily included providing benefits to voters who learned a specified list of facts. These proposals harm voters' epistemic autonomy as they pressure voters to follow a predetermined epistemic path. The proposal in this chapter, on the other hand, ensures the continuous availability of multiple epistemic paths. Additionally, as argued above, these proposals violate democratic principles, such as equality and self-governance. Brennan himself sees "epistocracy" (where the right to vote depends on learning specific facts) as an alternative to democracy. The proposal made in this chapter aims to preserve democratic values. It increases self-governance in the epistemic process, by offering interventions that better allow people to choose and pursue what they find relevant within the cognitive limits, up until the point where those choices contradict their self-governance (e.g., information about whether a certain policy will take away one's right to vote is always supplied). Additionally, by relying on the capabilities approach, it prioritizes actual equality, aiming to provide the same capabilities to different people given their unique conversion factors, rather than merely supplying the same information.

Importantly, there are also reasons to believe that interventions aimed at enhancing voters' epistemic empowerment could improve their collective competence according to the same criteria used to justify these existing proposals. For instance, Brennan's proposed interventions to increase voter knowledge are grounded in the "enlightened preferences" literature. As explained, such literature does not suppose people should end up voting according to the same values or conceptions of the good, but rather according to their well-informed preferences. In section 4 of this chapter, I explain how the conditional subjectivist approach enhances people's ability to pursue their enlightened preferences. This is due to providing people with information that they find relevant to their own preferences, conditional on them being in a better epistemic position to evaluate their preferences. They are more likely to change their vote due to information they personally find relevant to their vote. This approach is more likely to help voters vote in alignment with their enlightened preferences than causing them to learn a list of facts they may not find relevant to their vote (e.g., recent bills passed or government spending on various issues; Brennan, 2022). Such lists may consume cognitive resources at the expense of information that would influence voters' choices.

There are also reasons to think interventions to epistemically empower voters can increase competence according to other similar conceptions of such competence that exist in the literature. In chapter 1, I discussed two types of voter competence, competence with respect to universal truths and competence with respect to relative truths. Competence in voting on relative truths involves voting according to one's interests, values or goals, and prominent literature adopts such understanding of democratic voting, particularly focusing on voters voting in accordance with their interests (e.g., Schumpeter, 1942, Downs 1957, Mansbridge, 1983, 1990; Miller, 1986, Goldman, 1999, List & Spiekermann, 2016, Goodin & Spiekermann, 2018, Spiekermann, 2020). The definition of empowerment provided here includes enhancing voters' ability to do that, and consequently plausibly their competence according to this conception.

In chapter 2, I also argue that a natural explanation for the potential epistemic advantage of DMPs is in tracking population-relative truths. I identify two means of tracking population-relative truths. One of them is by each person making decisions based on their own interests, values, etc, which I also claim is likely to lead to higher competence. Again, the definition of empowerment provided here includes enhancing voters' ability to do that. Hence, such empowerment may enhance their ability to vote competently according to interpretations of voter competence that should be used to assess prominent means of improving democratic decision-making.

Finally, I have explained that the capability of epistemic collaboration, which is part of the offered account, can aid in reaching collective voter competence, due to an epistemic division of labour and lower vote dependence. Although not all individuals will choose to participate in such collaborations, enabling this opportunity increases the likelihood that some will. Such division of labour is likely to raise competence also on universal truths, as is noted in jury theorem literature.

This plausibility of competence enhancement coupled with empowerment and perseverance of epistemic autonomy suggests that epistemic voter empowerment should be further investigated as means of either enhancing various proposals for democratic-decision making, such as representative votes or deliberative mini publics, or aiding in replacing them with more inclusive methods.

#### 4.7 Practical interventions

One might argue that, compared with encouraging voters to learn a specific finite list of facts, the proposal made here for epistemic empowerment is much more complex to execute. This is because the suggested form of empowerment entails the particular information pieces that should be made accessible differ between different individuals, according to their personal traits, what they find relevant, their cognitive capacities, etc.

However, this proposal for epistemic voter empowerment is in alignment with recent technological advancements and the possibilities they bring. The ability to personalize information according to metrics of this type is already in existence.

Many of the listed goals align with capabilities employed by recommender systems to shape informational environments, although these systems currently shape those environments according to different standards. For example, adjusting informational environments to what people find relevant resembles the capability of adjusting people's recommendations according to "liked" posts, except it asks for a different type of feedback from users. Promoting opposing views for the sake of epistemic autonomy is also possible, and recommender systems that promote such diversity have been developed (Zhao et al., 2025), although insufficiently used on large existing platforms. Moreover, it is possible to tailor information to account for the different cognitive costs items impose on different individuals, since current systems already track detailed information consumption histories, which influence the present cognitive cost associated with different information. For example, recommender systems currently hold sufficient information to assess whether a certain content would be familiar to a particular user, which reduces its cognitive cost. Additionally, there exist tests for the cognitive cost of a specific task for a specific person, which are easily automatable (e.g., Marcus et al., 1996; Sweller et al., 2011). Those include, for instance, time spent on next article, which is already tracked (Ge et al., 2024, Facebook Help Centre, 2025).

Additionally, interventions that could be helpful for epistemic collaborations can rely on existing technologies. Matchmaking between people with shared interests involves similar capabilities to current matchmaking algorithms, although the input parameters differ. Means of distributing diverse information across voters who engage in a collaboration differ slightly more for current usages of these technologies, however, capabilities of text classification and clustering information can be helpful to achieve this goal (Di Lascio et al., 2017; Ahmad, et al., 2021; Shahina & Kumar, 2022). Online votes preceding the actual vote are also easily

achievable. Finally, large language models have been making major progress in automatically summarizing relevant information in a way that reduces the cognitive cost of internalizing it.

Hence, the proposed type of empowerment is a natural candidate for providing guidelines as to how online information curation should happen to make it empowering rather than disempowering. It highlights the potential of existing technological capabilities to epistemically empower voters. In this sense, the personalized nature of the suggested empowerment makes the proposal timely and could be seen as an advantage of the theory rather than a burden. Of course, each possible intervention to epistemically empower voters using such technological capabilities should be examined more closely to identify possible AI safety concerns. However, I believe this chapter provides initial justification for the pursuit of such investigations, the plausibility of some justifiable interventions of this type, as well as guidelines for assessing them.

## Conclusion

In this chapter, I have unpacked the concept of an epistemically empowered voter. I have argued that to be epistemically empowered, voters must possess certain capabilities that arise within the constraints of their cognitive budget sets. These capabilities include the ability to access information they find relevant, conditional on the continuous availability, within their cognitive budget, of: (1) awareness of the various information options; (2) information on opposing views; (3) the ability for epistemic collaborations; and (4) information on whether a policy is compatible with the preconditions for being an epistemically empowered voter. I outlined the advantages of interventions guided by this framework over prominent alternatives in the literature, particularly in terms of preserving and enhancing epistemic autonomy and improving competence by similar standards. I have also explained why this approach highlights the potential of existing technological capabilities in empowering voters.

---

## References

---

Achen C. H. & Bartels L. M., (2017) *Democracy for Realists: Why Elections Do Not Produce Responsive Government*, NJ: Princeton University Press.

Ahlstrom-Vij, K. (2022). The Case for Modelled Democracy. *Episteme*, 19(1), 89-110.

Ahmad, Z., Shahid, K. A., Wai, S. C., Abdullah, J., & Ahmad, F. (2021). *Transactions on Emerging Telecommunications Technologies*, London: Wiley.

Alkire, S. and Deneulin, S., (2009), The Human Development and Capability Approach, in Shahani, L. and Deneulin, S. (eds.), *An Introduction to the Human Development and Capability Approach: Freedom and Agency*, London: Earthscan, pp. 22–48.

Alsop, R., & Heinsohn, N. (2005). Measuring Empowerment in Practice: Structuring Analysis and Framing Indicators. *World Bank Policy Research Working Paper No. 3510*.

Althaus, S. L. (1998) Information Effects in Collective Preferences. *American Political Science Review* 92(3), 545-558.

Andersen, V. N. and Hansen, K. M. (2007). How Deliberation Makes Better Citizens: The Danish Deliberative Poll on the Euro. *European Journal of Political Research*, 46(4), 531–56.

Barnett, Z. (2020). Why You Should Vote to Change the Outcome. *Philosophy & Public Affairs*, 48(4), 422-446.

Bartels L. M. (1996). Uninformed Votes: Information Effects in Presidential Elections. *American Journal of Political Science* 40(1), 194-230.

Batliwala, S. (2007) Taking the Power Out of Empowerment. *Development in Practice*, 17(5), 557–565.

Bawden, D., & Robinson, L. (2020). *Information Overload: An Overview*, Oxford: Oxford University Press.

BC-Data-Deliberation-research-team. (2020). *Public Input into Pandemic Planning Deliberating Trade-offs in COVID-19 Policy Making*. Background Information Booklet.

Benz, M. & Stutzer, A. (2004). Are Voters Better Informed When They Have a Larger Say in Politics? Evidence for the European Union and Switzerland. *Public Choice*, 119(1), 31-59.

Berlin, I., (1969), Two Concepts of Liberty, in I. Berlin, *Four Essays on Liberty*, London: Oxford University Press.

Bervoets, S., & Gravel, N. (2007). Appraising Diversity with an Ordinal Notion of Similarity: an Axiomatic Approach. *Mathematical Social Sciences*, 53(3), 259-273.

Biggeri, M., & Frediani, A. A. (2019). *The Capability Approach, Empowerment and Participation*. London: Palgrave Macmillan.

Blum, C., & Zuber, C.I. (2016). Liquid Democracy: Potentials, Problems, and Perspectives. *Journal of Political Philosophy*, 24(2), 162-182.

Boland, P. J. (1989). Majority Systems and the Condorcet Jury Theorem. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 38(3), 181-189.

Borragán, G., Slama, H., Bartolomei, M., & Peigneux, P. (2017). Cognitive Fatigue: A Time-Based Resource-Sharing Account. *Cortex*, 89, 71–84.

Boston, J. (2017). *Governing for the Future: Designing Democratic Institutions for a Better Tomorrow*. Northsprings: Emerald Group Publishing Limited.

Boucher, V. (2022). *Down the TikTok rabbit hole: Testing the TikTok Algorithm's Contribution to Right-Wing Extremist Radicalization*, Doctoral dissertation, Queen's University Canada.

Bowles, S. (1998). Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions. *Journal of Economic Literature*, 36(1), 75–111.

Bradley, R. & Thompson, C. (2012). A (Mainly Epistemic) Case for Multiple-Vote Majority Rule. *Episteme*, 9(1), 63-79.

Brennan J (2016) *Against Democracy*. Princeton, NJ: Princeton University Press.

## References

Brennan J (2022) In Defense of Epistocracy: Enlightened Preference Voting, in Melenovsky C.M. (ed.) *The Routledge Handbook of Philosophy, Politics, and Economics*. NY: Routledge, pp. 391-400.

Brennan, J., & Landemore, H. (2022). *Debating Democracy: Do We Need More or Less?* Oxford: Oxford University Press.

Bruns, A. (2019). *Are Filter Bubbles Real?* New Jersey: John Wiley & Sons.

Burke, R., Felfernig, A., & Göker, M. H. (2011). Recommender Systems: An Overview. *AI Magazine*, 32(3), 13-18.

Buss, S. and Westlund A. (2018). Personal Autonomy. *Stanford Encyclopedia of Philosophy*. Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/>>.

Byskov, M. F. (2016) Democracy, Philosophy, and the Selection of Capabilities. *Journal of Human Development and Capabilities*, 18 (1), 1–16.

Byskov, M. F. (2016) Democracy, Philosophy, and the Selection of Capabilities. *Journal of Human Development and Capabilities*, 18 (1), 1–16.

Byskov, M. F. (2018). Selecting Capabilities for Development, in Comim, F., Fennell, S., Anand, P.B. (eds.), *New Frontiers of the Capability Approach*, Cambridge: Cambridge University Press, pp. 198-231.

Byskov, M. F. (2020). Beyond ‘Having Reason to Value’: Why We Should Adopt a Procedure-Independent and Value-Neutral Definition of Capabilities. *Journal of Economic Methodology*, 27(1), 18-35.

Campbell, T., Sostrin, J., & Masters, B. (2011). How to Educate the Public about the ACA: Recommendations from California Speaks. *Health Affairs Forefront*. Available at: <https://www.healthaffairs.org/content/forefront/educate-public-aca-recommendations-california-em-speaks-em> [accessed September 2025].

Casassus, B. (2021) France Turns to Citizens’ Panel to Reduce Vaccine Skepticism. *Science* 371, 763-764.

Carter, I. (1999). *A Measure of Freedom*. Oxford: Oxford University Press.

## References

Carter, I (2004). Choice, Freedom, and Freedom of Choice. *Social Choice and Welfare*, 22, 61-81.

Carter, I (2022). Positive and Negative Liberty, *The Stanford Encyclopedia of Philosophy* Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2022/entries/liberty-positive-negative/>>.

Cheng, F. F., & Wu, C. S. (2010). Debiasing the Framing Effect: The effect of Warning and Involvement. *Decision Support Systems*, 49(3), 328-334.

Chase, W. G., & Simon, H. A. (1973). Perception in Chess. *Cognitive Psychology*, 4(1), 55-81.

Chong, D. (2013). Degrees of Rationality in Politics, in Sears D., Huddy L., and Levy J (eds.) *Handbook of Political Psychology*, 2nd ed., Oxford: Oxford University Press, pp. 96-129.

Christiano, T. (1996) *The Rule of the Many: Fundamental Issues in Democratic Theory*. NY: Routledge.

Christiano, T. (2018). Democracy. The Stanford Encyclopedia of Philosophy , Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2018/entries/democracy/> 5.

Christiano, T. (2022). Algorithms, Manipulation, and Democracy. *Canadian Journal of Philosophy*, 52(1), 109-124.

Christie, S. T., & Schrater, P. (2015). Cognitive Cost as Dynamic Allocation of Energetic Resources. *Frontiers in Neuroscience*, 9, 289-304.

Christman, J. (1991). Liberalism and Individual Positive Freedom. *Ethics*, 101(2), 343-359.

Citizen's Assembly on Electoral Reform (CAER). (2004a). *Assessing Electoral Systems*. [online] Voterequality.org. Available at: <<http://voterequality.org/wp-content/uploads/2015/12/BCCA-FS02-Assessing-Electoral-Systems.pdf>> [Accessed 11 October 2022].

Citizen's Assembly on Electoral Reform (CAER), (2004b). *Politics in BC: What do we Want?*. Available at: <<http://voterequality.org/wp-content/uploads/2015/12/BCCA-FS01-Politics-in-BC.pdf>> [Accessed 11 October 2022].

## References

Citizen's Assembly on Electoral Reform (CAER) (2004c) *Making every vote count - voter equality*. Available at: <http://www.voterequality.org/wp-content/uploads/2015/12/BCCA-BC-Citizens-Assembly-Technical-Report-2004-English.pdf> [Accessed: October 20, 2022].

Clark, A. & Chalmers D. (1998) The Extended Mind. *Analysis* 58(1), 7-19.

Clark, D. A. (2005). The Capability Approach: Its Development, Critiques and Recent Advances. In Clark, D. A. (ed.), *The Elgar Companion to Development Studies*. Cheltenham: Edward Elgar, pp. 1-18.

Code, L., (1991), *What Can She Know?* Ithaca: Cornell University Press.

Costella, A. (2023). Adaptive Preferences, Self-Expression, and Preference-Based Freedom Rankings. *Economics & Philosophy*, 40(3), 1-22.

Cohen, J. (1986). An Epistemic Conception of Democracy, *Ethics*, 97(1), 26– 38.

Coleman, J.L. & Ferejohn, J. (1986). Democracy and Social Choice, *Ethics* 97(1), 6– 25.

Collins, P. H. (1990). *Black Feminist Thought*. Boston: Unwin Hyman.

Comim, F. (2018). Sen's Capability Approach, Social Choice Theory and the Use of Rankings. In Comim, F., Fennell, S., and Anand, P.DB (eds.) *New Frontiers of the Capability Approach*. Cambridge: Cambridge University Press, pp. 179-197.

Côté, N. (2020). Weakness of Will and the Measurement of Freedom. *Ethics*, 130(3), 384–414.

De Condorcet N (1976) (Originally published in 1785). Essay on the Application of Mathematics to the Theory of Decision-making. in Michael K. (Trans. & Ed.) *Condorcet: Selected Writings*. Indianapolis: Bobbs-Merrill, pp. 33–70.

Converse, P. E. (2006). The Nature of Belief Systems in Mass Publics (1964). *Critical Review*, 18(1-3), 1-74.

Crenshaw K.W. and Bonis O. (2005) Mapping the Margins: Intersectionality, Identity Politics, and Violence Against Women of Color. *Cahiers du Genre* 39(2), 51-82.

Curato, N., Farrell, D. M., Geissel, B., Grönlund, K., Mockler, P., Pilet, J. B., ... & Suiter, J. (2021). *Deliberative Mini-Publics: Core Design Features*. Bristol: Bristol University Press.

## References

Delli Carpini, M. X., & Keeter, S. (1992). An Analysis of Information Items on the 1990 and 1991 NES Surveys. *Board of Overseers for the National Elections Studies*, (17).

Dietrich, F. and Spiekermann, K. Jury Theorems, *The Stanford Encyclopedia of Philosophy* (Spring 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/spr2023/entries/jury-theorems/>>.

Di Lascio, F. M. L., Durante, F., & Pappada, R. (2017). *Copulas and Dependence Models with Applications: Contributions in Honor of Roger B. Nelsen*. NY: Springer International Publishing.

Dietrich, F., & Spiekermann, K. (2024). Deliberation and the Wisdom of Crowds. *Economic Theory*, 79, 603-655.

Dietrich, F. & List, C. (2004). A Model of Jury Decisions Where All Jurors Have the Same Evidence. *Synthese*, 142(2), 175-202

Dietrich, F. & Spiekermann, K. (2013). Epistemic Democracy with Defensible Premises, *Economics and Philosophy*, 29(1), 87-12.

Dixon, R., & Nussbaum, M. C. (2011). Abortion, Dignity and a Capabilities Approach. *The Public Law and Legal Theory Working Paper Series Working Paper no. 345*.

Dowding, K., & Van Hees, M. (2007). Freedom of Choice. In Anand, P., Pattanaik P., & Puppe, C. (eds.) *The Handbook of Rational and Social Choice*, Oxford: Oxford University Press, pp. 374-392.

Downs, A. (1957). *An Economic Theory of Democracy*, Addison Wesley: Boston.

Druckman, J. N. (2001). Evaluating Framing Effects. *Journal of Economic Psychology*, 22(1), 91-101.

Druckman, J. N., (2012). The Politics of Motivation. *Critical Review: A Journal of Politics and Society* 24(2), 199–216.

Dryzek, J. S. (2010). *Foundations and Frontiers of Deliberative Governance*. Oxford: Oxford University Press.

Dworkin, G. (1988). *The Theory and Practice of Autonomy*. Cambridge: Cambridge University Press.

## References

Ecker, U., Roozenbeek, J., Van Der Linden, S., Tay, L. Q., Cook, J., Oreskes, N., & Lewandowsky, S. (2024). Misinformation Poses a Bigger Threat to Democracy Than You Might Think. *Nature*, 630(8015), 29-32.

Edlin, A., Gelman, A., & Kaplan, N. (2007). Voting as a Rational Choice: Why and How People Vote to Improve the Well-Being of Others. *Rationality and Society*, 19(3), 293-314.

Elections Canada (2000) *Explaining the Turnout Decline in Canadian Federal Elections: A New Survey of Non-voters*, – Elections Canada. Available at: <https://www.elections.ca/content.aspx?section=res&dir=rec%2Fpart%2Ftud&document=reasons&lang=e> [Accessed: October 12, 2022].

Elster, J., (1983). *Sour Grapes: Studies in the Subversion of Rationality*, Cambridge: Cambridge University Press: Cambridge.

Elstub, S., & Escobar, O, (eds.), (2019). *Handbook of Democratic Innovation and Governance*. Cheltenham: Edward Elgar Publishing.

Elstub, S., Carrick, J., Farrell, D. M., & Mockler, P. (2021). The Scope of Climate Assemblies: Lessons from the Climate Assembly UK. *Sustainability* 13(20), 1-17.

Environmental Health Coalition, (2020), Voter Empowerment, Accessible at- <https://www.environmentalhealth.org/our-work/voter-empowerment/> [Accessed 15 May 2025].

Erickson, J. (2024). Rethinking the Filter Bubble? Developing a Research Agenda for the Protective Filter Bubble. *Big Data & Society*, 11(1), 1-4.

Estlund, D. Waldron, J. Grofman, B. & Feld, S.L. (1989). Controversy: Democratic Theory and the Public Interest: Condorcet and Rousseau Revisited. *American Political Science Review*, 83 (4), 1317– 40.

Estlund, D. (2008). Introduction: Epistemic Approaches to Democracy. *Episteme*, 5(1), 1-4.

European Climate Foundation (ECF). (2021 October 12). *The Growing Traction of Climate Citizens Assemblies* - European Climate Foundation. Available at: <<https://europeanclimate.org/stories/the-growing-traction-of-climate-citizens-assemblies/>> [Accessed 9 October 2022].

## References

Exploring Public Views on Assisted Dying. (2025). Nuffield Council on Bioethics. Available at: <<https://www.nuffieldbioethics.org/project/exploring-public-views-on-assisted-dying/>> [accessed September 2025].

Eysenck, M. W., Ellis, A. W., Hunt, E. B., & Johnson-Laird, P. N. E. (1994). *The Blackwell Dictionary of Cognitive Psychology*. NJ: Basil Blackwell.

Facebook Help Centre (accessed 2025) *What are Recommendations on Facebook?* Available at: <<https://www.facebook.com/help/1257205004624246>> [accessed September 2025].

Fishkin, J. S. (2009). *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford: Oxford University Press.

Fishkin J. S. (2013). When Citizens Decide: Lessons from Citizens' Assemblies on Electoral Reform; Designing Deliberative Democracy: The British Columbia Citizens Assembly. *Perspectives on Politics*, 11(2), 670.

Fishkin J. S. (2018). Deliberative Polling. In Bächtiger A., Dryzek J.S., Mansbridge J., Warren M.D.(eds.) *The Oxford Handbook of Deliberative Democracy*. Oxford: Oxford University Press, pp. 314-328.

Fournier, P., Van Der Kolk, H., Carty, R. K., Blais, A., & Rose, J. (2011). *When Citizens Decide: Lessons from Citizen Assemblies on Electoral Reform*. Oxford: Oxford University Press.

Fricker, E. (2006). Testimony and Epistemic Autonomy, in Lackey, J. & Sosa, E. (eds.) *The Epistemology of Testimony*. New York: Oxford University Press, pp. 225-50.

Friedman, M. (2003). *Autonomy, Gender, Politics*. Oxford: Oxford University Press.

Fung A. (2004) *Empowered Participation: Reinventing Urban Democracy*, NJ: Princeton University Press.

Garnett, M. (2016). Value Neutrality and the Ranking of Opportunity Sets. *Economics & Philosophy*, 32(1), 99-119.

Ge, Y., Liu, S., Fu, Z., Tan, J., Li, Z., Xu, S., ... & Zhang, Y. (2024). A Survey on Trustworthy Recommender Systems. *ACM Transactions on Recommender Systems*, 3(2), 1-68.

## References

Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Oxford: Oxford University Press.

Giraudeau, L. G., Apouey, B., Arab, H., Baeckelandt, S., Begout, P., Berghmans, N., ... & Tournus, S. (2021). *Deliberating on Climate Action: Insights from the French Citizens' Convention for Climate*. Available at: <https://hal-enpc.archives-ouvertes.fr/hal-03119539> [accessed September 2022].

Goldman, A. I. (1991). Epistemic Paternalism: Communication Control in Law and Society. *The Journal of Philosophy*, 88(3), 113–131.

Goldman, A. (1999). *Knowledge in a Social World*. NY: Oxford University Press.

Goldstein, D. G., Suri, S., McAfee, R. P., Ekstrand-Abueg, M., & Diaz, F. (2014). The Economic and Cognitive Costs of Annoying Display Advertisements. *Journal of Marketing Research*, 51(6), 742-752.

Gómez-Vírseda, C., De Maeseneer, Y., & Gastmans, C. (2019). Relational Autonomy: What Does it Mean and How is it Used in End-of-Life Care? A Systematic Review of Argument-Based Ethics Literature. *BMC medical ethics*, 20(1), 76

Goodin, R. E., & Spiekermann, K. (2012). Epistemic Aspects of Representative Government. *European Political Science Review*, 4(3), 303-325.

Goodin, R. E., & Spiekermann, K. (2015). Epistemic Solidarity as a Political Strategy. *Episteme*, 12(4), 439-457.

Goodin, R. E., & Spiekermann, K. (2018). *An Epistemic Theory of Democracy*. Oxford University Press.

Graham, M. H. (2020). Self-awareness of Political Knowledge. *Political Behavior*, 42(1), 305-326.

Grim P., Singer, D. J. and Bramson, A. (2019). Diversity, Ability, and Expertise in Epistemic Communities. *Philosophy of Science*, 86(1), 98-123.

Grofman, B. O. & Feld S. L. (1983). Thirteen Theorems in Search of the Truth. *Theory and Decision*, 15(3), 261-278.

Grofman, B. O. & Feld, S. L. (1988). Rousseau's General Will: a Condorcetian Perspective. *American Political Science Review*, 82(2), 567-576.

Guerrero, A. (2014) Against Elections: The Lottocratic Alternative, *Philosophy & Public Affairs*, 42 (2), 135-78.

Guerrero, A. (2024). *Lottocracy: Democracy Without Elections*. Oxford: Oxford University Press.

Guriev, S., Melnikov, N., & Zhuravskaya, E. (2021). 3g Internet and Confidence in Government. *The Quarterly Journal of Economics*, 136(4), 2533-2613.

Gustafsson, J. E. (2020). A Paradox for the Intrinsic Value of Freedom of Choice. *Noûs*, 54(4), 891-913.

Hadi, A. (2001). Effects of the Public Role of Bangladeshi Women on Their Reproductive Decisions. *Asia-Pacific Population Journal* 16(4), 21–30.

Hamilton, A. (1788). Federalist No. 68, in J. Cooke (ed.), *The Federalist*, Middletown, CT: Wesleyan University Press, pp. 457–462.

Hamilton, A. & Madison, J. (2003). *The Federalist, with Letters of 'Brutus'*. (Originally published in 1788). Ball T. (ed.). Cambridge: Cambridge University Press.

Hanson, R. (2013). Shall We Vote on Values, But Bet on Beliefs? *Journal of Political Philosophy*, 21(2), 151-178.

Harrington, J. (1747). *The Oceana and Other Works*. London: A. Millar.

Hartsock, N. (1983). The Feminist Standpoint: Developing the Ground for a Specifically Feminist Historical Materialism. In Harding S. and Hintikka M. (eds.) *Discovering Reality: Feminist Perspectives on Epistemology, Metaphysics, Methodology, and Philosophy of Science*. Dordrecht: Kluwer, pp. 283–310.

Harvey, F., (2022, October 5). *Global Citizens' Assembly to be Chosen for UN Climate Talks*. TheGuardian.<<https://www.theguardian.com/environment/2021/oct/05/global-citizens-assembly-to-be-chosen-for-un-climate-talks>> [Accessed 9 October 2022].

Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.

Van Hees, M. (2004). Freedom of Choice and Diversity of Options: Some Difficulties. *Social Choice and Welfare*, 22, 253–266

## References

Hirst, P. (1988). Representative Democracy and its Limits. *The Political Quarterly*, 59(2), 190-205.

Hong, L. and Page, S. (2001). Problem Solving by Heterogeneous Agents. *Journal of Economic Theory*, 97(1), 123–163.

House of Commons, Involve and Sortition Foundation et al. (2021). *The Path to Net Zero Climate Assembly UK Full report*, available at: <https://www.climateassembly.uk/report/read/final-report.pdf> [accessed September 10 2021].

Humanists UK. (2022 September 9th). *France Announces Launch of Assisted Dying Citizens Assembly*. [online] Available at: <<https://humanists.uk/2022/09/09/france-announces-launch-of-assisted-dying-citizens-assembly>> [Accessed 9 October 2022].

Hume, D. (1754). Idea of a Perfect Commonwealth. Essay 16 in *Essays, Literary, Moral and Political*. London: A. Millar.

Independent SAGE. (2022, October 21st). *Media Highlights* Available at: <https://www.independentsage.org/> [Accessed October 18, 2022].

Independent Sector. (2025). *It's On Us: The Voter Empowerment Project*. Available at: <https://independentsector.org/nonprofit-voter-empowerment-project/> [Accessed 15 May 2025].

Involve. (30 May 2018). *The Citizens' Assembly Behind The Irish Abortion Referendum*, Available at: <https://www.involve.org.uk/news-opinion/opinion/citizens-assembly-behind-irish-abortionreferendum#:~:text=These%20resources%20became%20a%20valuable,up%20to%20certain%20gestation%20limits> [accessed October 2022].

Iyengar, S., Sood, G., and Lelkes, Y. (2012). Affect, not Ideology: A Social Identity Perspective on Polarization. *Public Opinion Quarterly*, 76(3), 405–431.

Iyengar, S. and Westwood, S. J. (2015). Fear and Loathing Across Party Lines: New Evidence on Group Polarization. *American Journal of Political Science*, 59(3), 690–707.

Jackman S, Sniderman P. M., (2006). The Limits of Deliberative Discussion: A Model of Everyday Political Arguments. *The Journal of Politics*, 68(2), 272–83.

## References

Jacquet, V. (2017). Explaining Non-participation in Deliberative Mini-Publics. *European Journal of Political Research*, 56(3), 640-659.

Jacobs, A. M. (2016). Policy Making for the Long Term in Advanced Democracies. *Annual Review of Political Science*, 19, 433–454.

Jaggar, A. (1985). *Feminist Politics and Human Nature*, NJ: Totowa.

Jones, P. & S. R. (1982). Evaluating Choice. *International Review of Law and Economics*, 2(1), 47–65.

Jost, J. T. (2020). *A Theory of System Justification*. Cambridge: Harvard University Press.

Kabeer, N. (2001). Conflicts Over Credit: Re-evaluating the Empowerment Potential of Loans to Women in Rural Bangladesh. *World Development* 29(1), 63–84.

Kahneman D. and Tversky A. (1979). Prospect Theory: An Analysis of Decision Under Risk. *Econometrics* 47(2), 263-291.

Kahneman, D. (2011). *Thinking, Fast and Slow*, London: Pinguin Books.

Kaushik, B., & Lòpez-Calva, L. F. (2011). Functionings and Capabilities. In Arrow K. J., Sen, A., Suzumura, K. (eds.) *Handbook of Social Choice and Welfare* Vol. 2, pp.153-187.

Khader, S. J. (2011). *Adaptive Preferences and Women's Empowerment*. Oxford: Oxford University Press.

Kitcher, P. (1990). The Division of Cognitive Labor. *The Journal of Philosophy*, 87(1), 5-22.

Kitcher, P. (2022). *The Main Enterprise of the World: Rethinking Education*. Oxford: Oxford University Press.

Klemisch-Ahlert, M. (1993). Social Choice Welfare Freedom of Choice, A Comparison of Different Rankings of Opportunity Sets. *Social Choice & Welfare*, 10(3), 189-207.

Kramer, M. H. (2008). *The Quality of Freedom*. Oxford: Oxford University Press.

Kunapuli, G. (2023). *Ensemble Methods for Machine Learning*. NY: Simon and Schuster

Kurzban, R., Duckworth, A., Kable, J. W., Myers, J. (2013). An Opportunity Cost Model of Subjective Effort and Task Performance. *Behavioral and Brain Sciences*, 36(6), 661–679.

## References

Lacelle-Webster, A. (2020). *Citizens' Panel – Planning the West Midlands' Recovery*. Available at: <https://participedia.net/case/7085> [accessed September 20 2020].

Ladha, K. K. (1992). The Condorcet Jury Theorem, Free Speech, and Correlated Votes. *American Journal of Political Science*, 36(3), pp. 617-634.

Ladha, K. K. (1995). Information Pooling through Majority-Rule Voting: Condorcet's Jury Theorem with Correlated Votes. *Journal of Economic Behavior & Organization*, 26(3), 353-372.

Lafont, C. (2015). Deliberation, Participation, and Democratic Legitimacy: Should Deliberative Mini-Publics Shape Public Policy?. *Journal of Political Philosophy*, 23(1), 40-63.

Lafont C. (2019). *Democracy Without Shortcuts: A Participatory Conception of Deliberative Democracy*. Oxford: Oxford University Press.

Landa, D. & Pevnick, R. (2020). Representative Democracy as Defensible Epistocracy. *American Political Science Review*, 114(1), 1-13.

Landemore, H. (2013). Deliberation, Cognitive Diversity, and Democratic Inclusiveness: an Epistemic Argument for the Random Selection of Representatives. *Synthese*, 190(7), 1209-1231.

Leemann, L. & Wasserfallen, F. (2016). The Democratic Effect of Direct Democracy. *American Political Science Review*, 110(4), 750-762.

Leverick, F., Chalmers, J., Kane, E. and Kinghan J. (2025). *Rape Myths and Jury Decision Making: an Evidence Review*. Glasgow: University of Glasgow

Leßmann, O. (2004). Some Reflections on Operationalising the Capability-Set and a Definition of Poverty as Capability Failure. In *4th International Conference on the Capabilities Approach "Enhancing Human Security"*.

Levendusky, M. (2013). Partisan Media Exposure and Attitudes Toward the Opposition. *Political Communication*, 30(4), 565-581

Li, W. (2021). *Optimal Ensembles for Deep Learning: Theory and Practice*, Doctoral dissertation, Colorado State University.

## References

Lin, H., Ristic, J., Inzlicht, M., & Otto, A. R. (2022). The Average Reward Rate Modulates Behavioural and Neural Indices of Effortful Control Allocation. *Journal of Cognitive Neuroscience*, 34(11), 2113-2126

Lindell, M. (2011, August). Same But Different—Similarities and Differences in the Implementation of Deliberative Mini Publics. In *ECPR General Conference, Section Democratic Innovations in Europe—A Comparative Perspective, August*, pp. 24-27.

List, C. & Goodin, R.E. (2001). Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy*, 9(3), 277-306.

List, C. (2008). Distributed Cognition: A Perspective from Social Choice Theory, in Albert M., Schmidtchen D. & Voigt S. (eds.), *Scientific Competition: Theory and Policy, Conferences on New Political Economy*, Tübingen: Mohr Siebeck, pp. 285-308.

List, C. & Spiekermann, K. (2016). The Condorcet Jury Theorem and Voter-Specific Truth, in McLaughlin, B. P. and Kornblith, H. (eds.), *Goldman and His Critics. Philosophers and their Critics*, Hoboken: Wiley-Blackwell, pp. 219-234.

List, C. (Winter 2022 Edition). Social Choice Theory, *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta & Uri Nodelman (eds.) URL = <<https://plato.stanford.edu/archives/win2022/entries/social-choice/>>.

Lupia, A. (1994). Short Cuts Versus Encyclopedias: Information and Voting Behavior in California Insurance Reform Elections. *American Political Science Review*, 88(1), 63-76.

Madison, J. (1961). Federalist #10. (Originally published in 1787), in C. Rossiter (ed.), *The Federalist Papers*. NY: New American Library.

Manin, B. (1997). *Principles of Representative Government*. Cambridge: Cambridge University Press.

Mannes, A. E., Soll, J. B., & Larrick, R. P. (2014). The Wisdom of Select Crowds. *Journal of Personality and Social Psychology*, 107(2), 277.

Mansbridge, J. (1983). *Beyond Adversary Democracy*. Chicago: University of Chicago Press.

Mansbridge, J. (1990). Democracy and Common Interests. *Social Alternatives*, 8(4), 20-25.

## References

Mansbridge J. (1999). Should Blacks Represent Blacks and Women Represent Women? A Contingent “Yes”. *The Journal of Politics*, 61(3), 628-657.

Mansbridge, J. (2003). Rethinking Representation. *The American Political Science Review*, 97(4), 515–528.

Marcus, N., Cooper, M., & Sweller, J. (1996). Understanding Instructions. *Journal of Educational Psychology*, 88(1), 49–63.

Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic Theory* (Vol. 1). New York: Oxford university press.

Matheson, J. (2024). The Philosophy of Epistemic Autonomy: Introduction to Special Issue. *Social Epistemology*, 38(3), 1-7.

McKelvey, Richard D., and Peter C. Ordeshook. (1986). Information, Electoral Equilibria, and the Democratic Ideal. *Journal of Politics*, 48, 909-37.

McKee, M., & Stuckler, D. (2010). How Cognitive Biases Affect our Interpretation of Political Messages. *British Medical Journal*, 340, 1-3.

Mendelberg T, Oleske J. (2000). Race and Public Deliberation. *Political Communication* 17(2), 169–91

Van Merriënboer, J. J., Schuurman, J. G., de Croock, M. B., & Paas, F. G. (2002). Redirecting Learners’ Attention During Training: Effects on Cognitive Load, Transfer Test Performance and Training Efficiency. *Learning and Instruction*, 12(1), 11-37.

Meyers, D. T. (2005). Decentralizing Autonomy: Five Faces of Selfhood. In Christman, J & Anderson, J. (eds.), *Autonomy and the Challenges to Liberalism*, Cambridge: Cambridge University Press, pp. 27-55.

Milano, S., Taddeo, M., & Floridi, L. (2020). Recommender Systems and Their Ethical Challenges. *AI and Society*, 35(4), 957–967.

Mill, J. S. (1859). *On Liberty*. London: John W. Parker and Son, West Strand.

Mill, J. S. (1991). *Considerations on Representative Government*. (Originally published in 1865). London: Prometheus.

## References

Miller, N. R. (1986). Information, Electorates and Democracy: Some Extensions and Interpretations of the Condorcet Jury Theorem, in Grofman, B. and Owen, G. (eds.) *Information Pooling and Group Decision-making*. Stamford: JAI Press, pp. 175–94.

Miyagishima, K. (2010). Ranking Linear Budget Sets. *Social Choice and Welfare*, 35, 163-173.

Morone, J. A. and Marmor T. R. (1981). Representing Consumer Interests: The Case of American Health Planning. *Ethics*, 91 (3), 431-50.

Mulligan, C. B., & Hunter, C. G. (2003). The Empirical Frequency of a Pivotal Vote. *Public Choice*, 116(1), 31-54.

Nguyen, C. T. (2020). Echo Chambers and Epistemic Bubbles. *Episteme*, 17(2), 141-161.

Noggle, R. (2021). Manipulation in Politics. *Research Encyclopedia of Politics*. Oxford: Oxford University Press.

Nord, M., Lundstedt, M., Altman, D., Angiolillo, F., Borella, C., Fernandes, T., ... & Lindberg, S. I. (2024). Democracy Winning and Losing at the Ballot: Democracy Report 2024. *V-Dem Institute Report*.

Norheim, O. F., Abi-Rached, J. M., Bright, L. K., Bærøe, K., Ferraz, O. L., Gloppen, S., & Voorhoeve, A. (2021). Difficult Trade-Offs in Response to COVID-19: The Case for Open and Inclusive Decision Making. *Nature Medicine*, 27(1), 10-13.

Noveck, B. S. (2009). *Wiki Government: How Technology Can Make Government Better, Democracy Stronger, and Citizens More Powerful*. Washington: Brookings Institution Press.

Nover S, (Mar 5, 2025) No, Trump Did Not ‘Bring Back Free Speech’, *Time*, Available at: <https://time.com/7264811/trump-free-speech-joint-address-essay/> [accessed September 2025].

Nuffield Council on Bioethics. (2021). *Genome Editing and Farmed Animal Breeding: Social and Ethical Issues*, Available at: <https://www.nuffieldbioethics.org/publications/genome-editing-and-farmed-animals/read-the-short-report/overview-and-key-findings> [accessed December 1st 2021].

Nussbaum, M. (1999). Defense of Universal Values. *Idaho Law Review*, 36(1), 379.

## References

Nussbaum, M. (2000). *Women and human development: The capabilities approach*. Cambridge: Cambridge University Press.

Nussbaum, M., (2001). Adaptive Preferences and Women's Options, *Economics and Philosophy*, 17, 67–88.

Nussbaum, M. (2006). *Frontiers of Justice: Disability, Nationality, Species Membership*, Cambridge: Harvard University Press.

Nussbaum, M. (2011). *Creating Capabilities: The Human Development Approach*. Cambridge: Harvard University Press.

The Original. (7 June 2008). *Senator On-Line*. Available at: <https://web.archive.org/web/20080607140855/http://senatoronline.com.au/> [accessed September 2025].

OECD. (2021). *Trust in Government*. Available at: [www.oecd.org/gov/trust-in-government.htm](http://www.oecd.org/gov/trust-in-government.htm) [accessed December 24, 2021].

OECD. (2020). *Innovative Citizen Participation and New Democratic Institutions: Catching the Deliberative Wave*. Available at: <https://doi.org/10.1787/339306da-en> [accessed October 2022].

Otto, A. R., & Daw, N. D. (2019). The Opportunity Cost of Time Modulates Cognitive Effort. *Neuropsychologia*, 123, 92-105.

Owen, G., Grofman, B., & Feld, S. L. (1989). Proving a Distribution-Free Generalization of the Condorcet Jury Theorem. *Mathematical Social Sciences*, 17(1), 1-16.

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

Palfrey, T. R., and Keith T. P. (1987). The Relationship Between Information, Ideology, and Voting Behavior. *American Journal of Political Science*, 31(3), 511-30.

Park, J. Y., & Jang, S. S. (2013). Confused By Too Many Choices? Choice Overload in Tourism. *Tourism Management*, 35(C), 1-12.

Participedia. (2021). *A global network and crowdsourcing platform for researchers, educators, practitioners, policymakers, activists, and anyone interested in public participation and democratic innovations*. Available at: <https://participedia.net> [accessed October 2022].

Pattanaik, P. K., & Xu, Y. (1998). On Preference and Freedom. *Theory and Decision*, 44(2), 173-198.

Pattanaik, P. K., & Xu, Y. (1990). On Ranking Opportunity Sets in Terms of Freedom of Choice. *Recherches Économiques de Louvain/Louvain Economic Review*, 56(3-4), 383-390.

Pattanaik, P. K., & Xu, Y. (2000). On Diversity and Freedom of Choice. *Mathematical Social Sciences*, 40(2), 123–130.

Pattanaik, P. K., & Xu, Y. (2008). *Ordinal Distance, Dominance, and the Measurement of Diversity*. Berlin: Springer Berlin Heidelberg.

Pattanaik, P. K., & Xu, Y. (2015). Freedom and its Value, in Hirose I. and Olson J. (eds.) *The Oxford handbook of value theory*, Oxford: Oxford University Press, pp. 356-380.

Pattanaik, P. K., & Xu, Y. (2020). On Capability and its Measurement. In Martinetti, C. E. (ed.) *The Cambridge Handbook of the Capability Approach*. Cambridge: Cambridge University Press, pp. 271-291.

Pattanaik, P. K., & Xu, Y. (2024). *Multidimensional Well-Being, Deprivation and Inequality*. NY: Springer International Publishing.

Paul, L. A. (2014). *Transformative Experience*. Oxford: OUP.

Pettit, P. (2010), Representation, Responsive and Indicative, *Constellations* 17(3), 427–28.

Pettit, J. (2012). Empowerment and Participation: Bridging the Gap Between Understanding and Practice. *United Nations Headquaters*, 10(6), 39.

Pollak, R. A. (1978). Endogenous Tastes in Demand and Welfare Analysis. *The American Economic Review*, 68(2), 374–379.

Pritchard, D. (2013). Epistemic Paternalism and Epistemic Value. *Philosophical Inquiries*, 1(2), 9–37.

Puppe, C. (1996). An Axiomatic Approach to “Preference for Freedom of Choice”. *Journal of Economic Theory*, 68(1), 174–199.

Quirk, P. J. (2014). Making It Up on Volume: Are Larger Groups Really Smarter?, *Critical Review*, 26(1-2), 129–50.

## References

Rappaport, J. (1984). Studies in Empowerment: Introduction to the Issue. *Prevention in Human Services*, 3(2-3), 1–7.

Rawls, J. (1971). *A Theory of Justice*, Cambridge: Harvard University Press.

Raz, J. (1988). *The Morality of Freedom*, Oxford: Oxford University Press.

Reuchamps, M., Vrydagh, J., & Welp, Y. (eds.). (2023). *De Gruyter Handbook of Citizens' Assemblies*. Boston: De Gruyter.

Rice-Oxley, M., (2022). Should Citizens Assemblies be Mandatory? *The Guardian*. Available at: <https://www.theguardian.com/world/2020/sep/11/should-citizens-assemblies-be-mandatory> [accessed October 2022].

Robeyns, I. (2006) The Capability Approach in Practice. *Journal of Political Philosophy* 14(3), 351–76.

Robeyns, I. and Byskov M. F. (Summer 2023 Edition). The Capability Approach, *The Stanford Encyclopedia of Philosophy* Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/sum2023/entries/capability-approach/>>.

Rosenbaum, E. F. (2000). On Measuring Freedom. *Journal of Theoretical Politics*, 12(2), 205–227.

Rousseau, J. J. (1968). *The Social Contract* (originally published in 1762). London: Penguin Classics.

Rowlands, J. (1997). *Questioning Empowerment: Working with Women in Honduras*. Oxford: Oxfam GB.

Rozina, G., & Karapetjana, I. (2009). The Use of Language in Political Rhetoric: Linguistic Manipulation. *Süleyman Demirel Üniversitesi Fen-Edebiyat Fakültesi Sosyal Bilimler Dergisi*, 2009(19), 111-122.

Sabl, A. (2010). Exploiting the Clueless: Heresthetic, Overload, and Rational Ignorance. In Le Cheminant, W. & Parrish J. M. (eds.) *Manipulating Democracy*, NY: Routledge, pp. 247-262.

Saguy, T., Dovidio, J. F., & Pratto, F. (2008). Beyond Contact: Intergroup Contact in the Context of Power Relations. *Personality and Social Psychology Bulletin*, 34(3), 432-445.

## References

Samuelson, W., & Zeckhauser, R. (1988). Status Quo Bias in Decision Making. *Journal of Risk and Uncertainty*, 1(1), 7-59.

Scanlon, T. (1972). A Theory of Freedom of Expression. *Philosophy & Public Affairs*, 1(2), 204–226.

Schumpeter, J. (1942). *Capitalism, Socialism and Democracy*, NY: Harper & Brothers.

Schwartz, B. (2012). Choice, Freedom, and Autonomy. In Shaver P. R. & Mikulincer M. (eds.), *Meaning, Mortality, and Choice: The Social Psychology of Existential Concerns*, Washington: American Psychological Association, pp. 271–287.

Scoles, S. (2007). The Wisdom of Crowds—A Better Way to Think About the Markets. *Risk and Rewards*, 50, 10-13.

Sen, A. (1974). Informational Bases of Alternative Welfare Approaches: Aggregation and Income Distribution. *Journal of Public Economics*, 3(4), 387–403.

Sen, A. (1980). Equality of what? In Goodin, R. E. & Pettit, P. (eds.), *Contemporary Political Philosophy: An Anthology*, Hoboken: Wiley–Blackwell, pp.197-220.

Sen, A. (1988). Freedom of Choice: Concept and Content. *European Economic Review*, 32(2-3), 269-294.

Sen, A. (1991). Welfare, Preference, and Freedom. *Journal of Econometrics*, 50(1-2), 15–29.

Sen, A. (1993A). Capability and Well-Being, in Sen A. & Nussbaum M. (eds.) *The Quality of Life*, Oxford: Clarendon Press, pp. 30–53.

Sen, A. (1993B). Markets and Freedoms: Achievements and Limitations of the Market Mechanism in Promoting Individual Freedoms. *Oxford Economic Papers*, 45(4), 519–541.

Sen, A. (1995). *Inequality Re-examined*. Cambridge: Harvard university press.

Sen, A. (1999). *Commodities and Capabilities*. Oxford: Oxford University Press.

Sen, A. (2002). *Rationality and Freedom*. Massachusetts: Harvard University Press.

Sen, A. (2017A). Well-being, Agency and Freedom the Dewey Lectures. (1984). In Brooks, T. (ed.), *Justice and the capabilities approach*. NY: Routledge, pp. 3-55.

## References

Sen, A. (2017B). What do we Want From a Theory of Justice?. In Mancilla, A. & Campbell, T. (eds.) *Theories of Justice*. NY: Routledge, pp. 27-50.

Setälä, M., & Smith, G. (2018). Mini Publics and Deliberative Democracy. In Bächtiger, A., Dryzek, J. S., Mansbridge, J., Warren, M. E. (eds.), *The Oxford Handbook of Deliberative Democracy*, Oxford: OUP Oxford, pp. 300-314.

Sharma, S., Kar, A. K., Gupta, M. P., Dwivedi, Y. K., & Janssen, M. (2022). Digital Citizen Empowerment: A Systematic Literature Review of Theories and Development Models. *Information Technology for Development*, 28(4), 660-687

Somin, I., (2004). When Ignorance Isn't Bliss: How Political Ignorance Threatens Democracy. *Policy Analysis*, 525, 1-28.

Somin, I. (2013). *Democracy and Political Ignorance*. Stanford: Stanford University Press.

Spiekermann, K. and Goodin R.E. (2012). Courts of Many Minds, *British Journal of Political Science*, 42(3), 555-571.

Spiekermann, K. (2020). Epistemic Network Injustice. *Politics, Philosophy & Economics*, 19(1), 83-101.

Stoljar, N., Feminist Perspectives on Autonomy, *The Stanford Encyclopedia of Philosophy* (Summer 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <<https://plato.stanford.edu/archives/sum2024/entries/feminism-autonomy/>>.

Sugden, R. & Jones, P. (1982). Evaluating Choice, *International Review of Law and Economics*, 2(1), 47-65.

Sunstein, C. R. (1999). The Law of Group Polarization. *Journal of Political Philosophy*, 10, 175–195.

Sunstein, C. R. (2009). *Going to Extremes. How Like Minds Unite and Divide*. Oxford: Oxford University Press.

Sunstein, C. R. (2019). *On Freedom*. NJ: Princeton University Press.

Surowiecki, J. (2012). *The Wisdom of Crowds*. London: Abacus.

Sweller, J., Ayres P. & Kalyuga S., (2011). *Cognitive Load Theory*. NY: Springer.

Székely, M., & Michael, J. (2021). The Sense of Effort: A Cost-Benefit Theory of the Phenomenology of Mental Effort. *Review of Philosophy and Psychology*, 12(4), 889–904.

The Electoral Commission, Public Attitudes, (2025), Available at: <<https://www.electoralcommission.org.uk/research-reports-and-data/public-attitudes/public-attitudes-2025>>[accessed September 2025].

Thaler, R. H. & Sunstein C. R. (2009). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. London: Penguin.

Thaler, R. H. (2016). Behavioural Economics: Past, Present, and Future. *American Economic Review*, 106(7), 1577-1600.

Thompson, D. F. (2008). Deliberative Democratic Theory and Empirical Political Science. *Annual Review of Political Science*. 11, 497-520.

Tolbert, C. J., McNeal, R. S., & Smith, D. A. (2003). Enhancing Civic Engagement: The Effect of Direct Democracy on Political Participation and Knowledge. *State Politics & Policy Quarterly*, 3(1), 23-41.

Tormey, S. (2015). *The End of Representative Politics*. NY: John Wiley & Sons.

UK Parliament. (2020) *About Citizens' Assemblies*. Available at: <<https://www.parliament.uk/get-involved/committees/climate-assembly-uk/about-citizens-assemblies/>> [Accessed 9 October 2022].

University of Southampton, (19 February 2025) Democracy in Crisis: Trust in Democratic Institutions Declining Around the World. Available at: <https://www.southampton.ac.uk/news/2025/02/democracy-in-crisis-trust-in-democratic-institutions-declining-around-the-world.page> [accessed September 2025].

Urbinati, N. (2000). Representation as Advocacy: A Study of Democratic Deliberation. *Political Theory*, 28(6), 758-786.

Velikanov, C., & Prosser, A. (2017). Mass Online Deliberation in Participatory Policy-Making. In Alois A. P. & Leonidas G. A. & Reddick, C. G. (eds.), *Beyond Bureaucracy*. NJ: Springer, pp., 209-234.

## References

Vile, John R. (January 23, 2025) Trump Executive Order on Restoring Freedom of Speech, Available at: <https://firstamendment.mtsu.edu/article/trump-executive-order-on-restoring-freedom-of-speech/> [accessed September 2025].

Warren, M. E., & Gastil, J. (2015). Can Deliberative Minipublics Address the Cognitive Challenges of Democratic Citizenship? *The Journal of Politics*, 77(2), 562-574.

Webster, R. K. , Brooks, S. K. , Smith, L. E., Woodland, L. Wessely, S. and Rubin, G. J. (2020). How to Improve Adherence with Quarantine: Rapid Review of the Evidence. *Public Health*, 182, 163-169.

West Midlands Recovery Coordination Group. (WMRCG, 2020). *Full Debrief of Findings July 2020 Citizens' Panel*. Available at: [www.wmca.org.uk/media/4053/final-report-wmrcg-citizens-panel-on-post-covid-recovery.pdf](http://www.wmca.org.uk/media/4053/final-report-wmrcg-citizens-panel-on-post-covid-recovery.pdf) [accessed October, 2022].

Wolff, J. (1994). Democratic Voting and the Mixed-Motivation Problem. *Analysis*, 54(4), 193-196.

World Health Organization. (2021, May 31). *Voice, Agency, Empowerment: Handbook on Social Participation for Universal Health Coverage*. Available at: <https://www.who.int/publications/i/item/9789240027794> [accessed October 2022].

World Bank. (2023). *Open and Inclusive: Fair Processes for Financing Universal Health Coverage*. Washington DC: World Bank.

Xu, Y. (2004). On Ranking Linear Budget Sets in Terms of Freedom of Choice. *Social Choice and Welfare*, 22(1), 281-289.

Young, I. M. (1997). Punishment, Treatment and Empowerment: Three Approaches to Policy for Pregnant Addicts. *Feminist Studies*, 20(1), 33-57.

Zenon, A., Solopchuk, O., & Pezzulo, G. (2019). An Information-Theoretic Perspective on the Costs of Cognition. *Neuropsychologia*, 123, 5-18.

Zhao, Y., Wang, Y., Liu, Y., Cheng, X., Aggarwal, C. C., & Derr, T. (2025). Fairness and Diversity in Recommender Systems: a Survey. *ACM Transactions on Intelligent Systems and Technology*, 16(1), 1-28.