

*The Problem of Expertise:
Epistocracy as a Potential Solution for
Emerging Technology Governance*

Daniel James Robertson

A thesis submitted in partial fulfilment of
the requirements for the degree of

Doctor of Philosophy in Political Science

at the University of Canterbury, Christchurch, New Zealand

30 June 2019

Abstract

This thesis investigates the epistemic performance of deliberative democracy regarding the governance of emerging science and technology policy. Central to this epistemic performance is the role of the expert in relation to that of the general public. This relationship has been described in terms of a problem of expertise, which is the tension between the epistemic quality of decision-making and ensuring a policy decision has the democratic legitimacy required to be authoritative; the more expertise that is required, the less democratic the policy process can be, and vice versa. In a time when emerging science and technology is becoming ever more complex, is it possible to create a governance system that has a strong tendency towards creating high-quality policy while still maintaining the democratic legitimacy that increasingly popular forms of public engagement are built upon?

To answer this question, an investigation of modern political conditions that drive the tension found in the problem of expertise is given, namely the wickedness of emerging science and technology policy problems and the legacy of a flawed history of technocratic governance. A theoretical analysis of the epistemic quality of deliberative democracy and its would be participants, relative to a hypothetical epistocracy, is then given to find which of these governance systems is most likely to produce high-quality emerging science and technology policy based on the criteria of sound reasoning with the best available information. Finding that a more epistocracy-oriented governance system is epistemically superior to a more participatory form of epistemic deliberative democracy, this thesis then proposes a moderate form of epistocracy called an “extended epistocracy” as a new governance regime for emerging science and technology policy where experts act as representatives for public groups. Lastly, given that the focus of the thesis has been on the epistemic quality of governance rather than

its legitimacy, potential future research is suggested to investigate whether an extended epistocracy can also meet the normative demands of democratic legitimacy.

While grounded in the philosophical field of political epistemology, this thesis presents an interdisciplinary study from a theoretical perspective that aims towards real world applications. The thesis integrates research from the fields of democratic theory, science and technology studies, and rational choice theory, but also draws on empirical data from the public understanding of science, public engagement exercises, management, and psychology. This thesis seeks to apply these interdisciplinary resources to the nascent but growing research on the epistemic quality of decision-making in governance systems.

Contents

Abstract.....	2
Acknowledgements	8
Chapter One: The Problem of Expertise as a Modern Problem	9
§1.1: Synopsis of the Thesis.....	9
§1.2: Development of the Understanding of the Problem of Expertise.....	14
§1.2.1: What Is the Problem of Expertise?.....	14
§1.2.2: Plato Proposes Philosopher Rulers.....	15
§1.2.3: The Lippmann–Dewey Debate on Easing Democracy’s Demands	18
§1.2.4: Truth to Power: Balancing Expertise and Democracy	23
§1.3: The Post-Normal Era Challenge for Democracy	25
§1.3.1: What Are Wicked Problems?	25
§1.3.2: What is Post-Normal Science?.....	27
§1.3.3: Society Is Also Complex.....	28
§1.3.4: Responding to a Post-Normal Era with an Extended Peer Community	31
§1.4: Conclusion	33
Chapter Two: Towards the Epistemic Turn of Deliberative Democracy.....	34
§2.1: Introduction	34
§2.2: Before the Deliberative Turn	36
§2.2.1: What Is Positivism?.....	37
§2.2.2: The Decline of Positivism.....	38
§2.2.3: Learning from Past Emerging Technology Governance	42
§2.3: The Deliberative Turn and the Emphasis on Procedure	44
§2.3.1: What Is Postpositivism?	44
§2.3.2: Procedural Emphasis within Deliberative Democracy.....	47
§2.3.3: Benefits of the Deliberative Procedure.....	49
§2.3.4: Issues Concerning Procedural Theories of Deliberative Democracy.....	51
§2.4: The Epistemic Turn.....	54
§2.4.1: Introducing Epistemic Democracy	54
§2.4.2: Challenges Facing the Epistemic Turn.....	57
§2.5: Potential Forms of Public Engagement to Overcome the Democratic Trilemma	61
§2.5.1: Nanotechnology as a Prime Opportunity for Public Engagement	61

§2.5.2: Different Forms of Public Engagement for Governance	63
§2.6: Conclusion	68
Chapter Three: Epistemic Democracy – Individual Citizen Competency.....	69
§3.1: Introduction	69
§3.2: Condorcet’s Jury Theorem Under Modern Political Conditions	71
§3.3: Citizen Competence Regarding Nanotechnology	74
§3.4: Will Competence Improve?	77
§3.4.1: Rational Ignorance and Rational Irrationality	78
§3.4.2: Fixing/Bypassing Rational Ignorance and Rationality	81
§3.4.3: Improving “Free” Information.....	82
§3.4.4: Shortcuts.....	85
§3.4.5: The Effect of Increasing Knowledge of Nanotechnology.....	91
§3.5: Conclusion	92
Chapter Four: Epistemic Democracy – Collective Intelligence	93
§4.1: Chapter Overview.....	93
§4.2: Performance of Nanotechnology Public Engagement Exercises.....	95
§4.2.1: Nanotechnology Consensus Conferences	96
§4.2.2: The Negative Slide Towards Nanotechnology Risks.....	100
§4.2.3: Implications for Nanotechnology Engagement	105
§4.3: The Effects of Cognitive Diversity on Deliberation	107
§4.4: Deliberator (In)dependence in Detail.....	109
§4.4.1: What Is Meant by (In)dependence.....	109
§4.4.2: The Effects of (In)dependence	111
§4.5: Does Diversity Trump Ability? (Theory).....	113
§4.5.1: What is the Diversity Trumps Ability Theorem?	114
§4.5.2: The Need to Carefully Select Diversity	117
§4.6: Does Diversity Trump Ability? (Computer Simulation).....	119
§4.6.1: Hong and Page’s Node Ring Deliberation Simulation	120
§4.6.2: The Node Ring Simulation Does Not Support Epistemic Democracy	121
§4.7: Summary	127
Chapter Five: Establishing an Extended Epistocracy	129
§5.1: Introduction	129
§5.2: Requirements for a Nanogovernance Regime	132

§5.2.1: Why Use Epistemic Deliberative Democracy?.....	133
§5.2.2: Upstream Engagement.....	134
§5.2.3: What Does Nanogovernance Need?	136
§5.3: What Is an Epistocracy?	138
§5.4: Who Are the Experts?.....	140
§5.4.1: General Ways to Identify Experts.....	141
§5.4.2: Relevant Expertise.....	142
§5.4.3: Lay Expertise	147
§5.5: Considering the Expert Actors as a Collective.....	148
§5.5.1: Finding Useful Categorizations of Diversity	150
§5.5.2: What is Task-Related Diversity?	151
§5.5.3: Managing Task-Related Diversity	152
§5.5.4: Effect of Bio-Demographic Diversity on Group Performance.....	158
§5.5.5: Managing Bio-Demographic Diversity	160
§5.6: What Role Should Experts Play in Policy-Making?	163
§5.6.1: The Role of Experts	163
§5.6.2: The Constitutional and Process Dimensions of an Extended Epistocracy.....	166
§5.7: The Selection Process.....	168
§5.8: Summary: What Does an Extended Epistocracy Look Like?.....	171
Chapter Six: Future Research — Legitimacy of an Extended Epistocracy	175
§6.1: Chapter Overview.....	175
§6.2: The Proposed General Structure for the Legitimacy Argument	178
§6.2.1: Using Epistemic Proceduralism as a Benchmark Governance System	178
§6.2.2: Argument Structure for Extended Epistocracy’s Legitimacy	179
§6.3: Overcoming the Legitimacy Barrier – The Expert/Boss Fallacy	181
§6.3.1: What Is the Expert/Boss Fallacy?	182
§6.3.2: Cannot Agree on Experts	183
§6.3.3: The Demographic Objection.....	184
§6.3.4: Epistocracy’s Extra Authority	186
§6.4: Comparing Democratic Legitimacy – Epistemic Injustice	188
§6.4.1: What Is Epistemic Injustice?	189
§6.4.2: Forms of Power	191
§6.4.3: Types of Injustice.....	192

§6.4.4: The Subject–Object Distinction.....	196
§6.4.5: Hermeneutical Injustice and Upstream Engagement	197
§6.5: Comparing Epistemic Legitimacy – Collective Ethics and Normative Consent.....	198
§6.5.1: What is Normative Consent?.....	199
§6.5.2: The Collective Ethics of Voting	200
§6.5.3: Responsibility for Collective Harm	203
§6.5.4: The Harm Caused by the Individual Voter.....	205
§6.5.5: The Responsibility of the Voter.....	206
§6.6: Conclusion	209
Chapter 7: Conclusion	210
Bibliography	213

Acknowledgements

First, I wish to thank my supervisors Assoc. Prof. Amy Fletcher, Prof. Diane Proudfoot, and Prof. Simon Brown for their guidance and thought-provoking discussions. Second, I wish to thank my mother, Jude Robertson, and her partner, Robert Greening, for their emotional and financial support, as well as my father, David Robertson, for his emotional support. Third, I wish to thank the University of Canterbury for funding through the UC Doctoral Scholarship. Lastly, I wish to thank my fellow political science colleagues for their support and interesting discussions.

Chapter One: The Problem of Expertise as a Modern Problem

§1.1: Synopsis of the Thesis

The purpose of this thesis is to evaluate the “problem of expertise” within the democratic governance of emerging science and technology and to consider whether and how an extended epistocratic model could better resolve the tension found in the problem of expertise between the need for both expertise and democratic legitimacy in public policy (Turner S. P., 2003, p. 5). To clarify terms, an epistocracy is rule by the knowers or the wise, i.e., experts, where an expert is someone with “an intangible but recognizable combination of education, talent, experience, and peer affirmation” (Estlund, 2003, p. 53) (Nichols, 2017, pp. 30–31). The investigation in this thesis is broken into three main parts. Part A discusses some of the epistemic challenges that governance faces, particular with regard to emerging science and technology policy, where Chapter 1 focuses on how the problem of expertise has been viewed over time and its relevance as a modern policy problem, while Chapter 2 focuses on the development of deliberative governance as a response to these challenges. However, while it is shown that deliberative democracy satisfies the democratic requirements towards solving the problem of expertise, it is still an open question whether it can satisfy the epistemic requirements as well. Therefore, Part B investigates the epistemic potential of deliberative democracy, with Chapter 3 analysing the epistemic performance of the general public as individuals, while Chapter 4 considers the epistemic performance of deliberative democracy as a collective exercise. This investigation of the epistemic potential of deliberative democracy strongly suggests that an epistocracy would better satisfy the epistemic demands of deliberative democracy, and as such, an epistocratic form of governance should be created. Therefore, Part C proposes an innovative, new mode of epistocratic governance for emerging science and technology policy that is intended to solve the problem of

expertise in Chapter 5, and demonstrates how its democratic legitimacy can potentially be defended in Chapter 6.

To speak of a rule by experts may initially conjure concerns that this research advocates a technocratic form of governance. However, the epistocratic model proposed in this thesis is *not* a technocracy, i.e., a rule by *technical*¹ experts alone, nor does this thesis advocate for the abandonment of democratic deliberation, broadly construed. Rather, the epistocracy that is proposed can be considered an epistemic form of deliberative democracy and is a governance model located in the middle ground between technocracy and participatory deliberative democracy, drawing on the advantages of both governance systems, while minimizing the disadvantages.

This project serves three scholarly purposes:

1. To investigate the epistemic properties of deliberative democracy from an interdisciplinary perspective that includes and synthesizes insights from rational choice theory, management, and psychology. For example, this thesis contributes to an emerging literature on epistemic democracy (see, for example, Scott Page (Page, 2007) and Ilya Somin (Somin, 2013)), which applies economic models to the study of the quality of democratic decisions.
2. Based on this investigation, propose an alternative form of governance called an *extended epistocracy* that accommodates the high level of expertise required in emerging technology governance, while still maintaining democratic legitimacy. The primary focus of an extended epistocracy in this study is to ensure that it improves epistemic outcomes relative to other models such as deliberative democracy. The question of extended epistocracy's democratic

¹ The term "technical experts" is used to denote those experts whose focus is on what are called "hard facts", discussed later in this chapter. Briefly, and somewhat coarsely, hard facts are quantitative, scientific facts, as opposed to "soft facts", which are more qualitative (Funtowicz & Ravetz, 1993, p. 750). Therefore, by way of distinction, an expert from the humanities would be considered "non-technical".

legitimacy and potential is considered in a final main section on future research.

3. Using the results from the first two aims, this thesis is to contribute to the emerging literature/theory on epistocracy and the epistemic properties of governance, as developed by, for example, researchers at the ARENA² Centre for European Studies in Oslo, Norway, involving scholars such as Cathrine Holst, H el ene Landemore, and Kasper Lippert-Rasmussen. As part of the EPISTO project, scholars examine and assess the legitimacy of expert rule in modern democracies with a particular interest in the European Union (ARENA Centre for European Studies, 2015) (Holst, 2014).

As noted above, the tension between expertise and democratic legitimacy is known as the problem of expertise (Turner S. P., 2003, p. 5). The proposed concern, and hence the motivation for this thesis, is that deliberative democracy and associated public engagement tools—the currently accepted approach to the problem of expertise—dominates the scholarly discussion of public policy governance, which is also played out in some practical examples, such as the Danish consensus conference model (Joss, 1998). There is little credence given to alternative governance systems, and there appears to be a general implicit assumption (explicit in some cases, see (Hardoř, 2014)) that a heavily participatory form of democracy should remain as the status quo decision-making process. This thesis takes a step back from the assumption that high levels of public engagement should remain the status quo and provides a critical analysis of the place of expertise in emerging technology policymaking. In doing so, it works within an emerging epistocratic turn in democratic theory (see, for example, (Landemore, 2017)) that considers new approaches to the definition and role of the expert and to the question of new models for democratic decision-making.

² Advanced Research on the Europeanisation of the Nation-State.

While the approach taken in this thesis towards a new solution to the problem of expertise is set within the relatively nascent field of political epistemology (American Political Science Association, 2014), this is an interdisciplinary project that provides a mixture of theoretical and empirical elements. Theoretically, this thesis draws upon science and technology studies, democratic theory, rational choice theory, and statistical formulations of voting behaviour. Empirically, this thesis draws upon research done in the social sciences (such as the public understanding of science and public engagement exercises), psychology, and management. Because of this mixture of scholarly fields and approaches, this thesis is able to fully recognise the epistemological shortcomings of both experts and the general public, and investigates the possibility of a governance regime for emerging technologies under modern political conditions as opposed to the ideal conditions assumed by deliberative democratic theory. By doing so, a practical solution to the problem of expertise can be found with the intention of being applied to real-world emerging science and technology policy problems. For this thesis, the emerging science and technology chosen as an example is nanoscience/nanotechnology³ (hereafter, simply “nanotechnology”) because it is seen to provide a great opportunity for the development of new forms of science and technology governance due to it being a newer technology (Priest & Greehalgh, 2011, p. 1521).

³ Nanotechnology can be described as:

the development and application of structures, materials, devices, and systems with fundamentally new properties and functions which derive from their size in the range of about 1 to 100 nanometres (Renn, Roco, & Litten, 2006, p. 19).

Nanotechnology is a highly interdisciplinary field that includes and extends work done in the more traditional fields of physics, chemistry, biology, and electrical engineering, among others. Nanotechnology is an enabling technology, meaning that it makes other technologies faster, lighter, stronger, smaller, etc., and as such, nanotechnology can be found in a vast range of applications. These applications include more efficient sunscreen, additives in food to provide extra vitamins or cosmetics to increase their glow, in car tires to increase their durability, and in paints and other types of coating to make a surface stain-resistant (Schmid et. al., 2006, pp. 241, 269, 340). However, future generations of nanotechnologies could include ever more complicated and influential technologies, where the progression of nanotechnology is thought to proceed from passive nanostructures, to active nanostructures, to systems of nanostructures, and finally to molecular nanosystems (Roco, Harthorn, Guston, & Shapira, 2011, p. 3568).

In terms of the chapter breakdown, the remainder of this introductory chapter forms the first half of Part A and answers the question: what is the problem of expertise and why is it important to study it today? This discussion begins by analysing how scholars' understanding of the causes of, and the solutions for, the problem of expertise have become more nuanced throughout political thought. It is shown that to solve the problem of expertise, an appropriate balance between expertise and democracy with experts and the general public each playing their part needs to be found (Wildavsky, 1979, p. 210). However, finding this balance is increasingly difficult under modern political conditions due to the rise of the post-normal era and its associated wicked problems, in which "facts are uncertain, values in dispute, stakes high and decisions urgent" (Funtowicz & Ravetz, 1993, p. 744). The problem of expertise is particularly salient at a time when emerging science and technology is creating significant "complexity, chaos, and contradictions", where it becomes ever more difficult to understand and predict the course of new technologies (Sardar, 2010, p. 436). The post-normal era can make it very difficult to find, understand, and appropriately connect the information required to make high-quality⁴, well-informed political decisions. To alleviate this difficulty, scholars have suggested extending the peer community, i.e., increasing the number of voices at the discussion table, to accommodate the use of "extended facts" (Funtowicz & Ravetz, 2003, p. 7). This chapter, therefore, finishes with a discussion regarding extending the peer community, which naturally leads to the discussion of deliberative democracy and the associated public engagement given in Chapter 2.

⁴ For the purpose of this thesis, "high-quality" means the use of the best accepted methods, theories, and information from both the natural and social sciences, and the policy created should be the result of logical (scientific, economic, moral, etc.) reasoning (RAND Corporation, 2015).

§1.2: Development of the Understanding of the Problem of Expertise

§1.2.1: What Is the Problem of Expertise?

The advancement of science and engineering disciplines in the early 21st century, along with the accelerating pace of technological change and diversification, means that there are increasing levels of specialization required at an ever-increasing rate (Berman & Dorrier, 2016)(Moschella, 2015). Furthermore, the commercialisation of the technologies produced using this advanced technical knowledge, and the increasing social impact that comes along with this commercialisation, means that the public want to play a role in developing policies regarding these technologies (Turner, 2003, p. 5). This process of public engagement in the governance process is subject to what is known as the “problem of expertise”, which asks how much weight should be given to expert judgment, especially where there is a potential tension between the expertise required to inform decision-making and the requirements of democratic methods of deliberation and decision-making. This tension is said to arise because an increase in the expertise required in the policy process, and hence an increase in policy quality⁵, means a decrease in the number of people that are able to contribute at that level of expertise. Conversely, increasing the number of people involved to democratic-sized constituencies, and hence increasing the democratic legitimacy of the resulting decision, often means the level of expertise involved must be lowered such that everyone can understand the discussion and effectively take part.

The idea of the problem of expertise is by no means new; indeed, it is arguably central to Western democratic theory and political science. Since ancient Greece and Plato’s proposal of philosopher rulers (Plato, 375BC [2007], pp. 191–192), the tension between the specialised knowledge of the expert and the democratic role of the public has constituted a major thread in the development of modern political science and

⁵ As shown in Chapters 3 and 4.

political theory over its approximately 2400-year history. The key figures discussed here, once the baseline has been established via a discussion of Plato's main ideas, are Lippman, Dewey, and Wildavsky, three contemporary figures who focused on the tension between experts and lay citizens from the emergence of modern industrial democracies onwards and proposed potential solutions relevant to the modern era. In analysing these key figures, an evolution in the understanding of both the causes of the problem of expertise, as well as potential solutions, can be seen. These four scholars are chosen because they provide key road markers in terms of the understanding of the problem of expertise, where the proposed causes of the problem of expertise have ranged from undisciplined citizens, to an overdemanding democracy, to experts overreaching their expertise, while the proposed solutions have ranged from dichotomous expert-only and democracy-only solutions, to a more nuanced balance where experts and citizens actively work together.

§1.2.2: Plato Proposes Philosopher Rulers

In *The Republic* (Plato, 375BC [2007])⁶, Plato provided a very early and sustained argument about the cause of what is today called the problem of expertise, which he attributed to the undisciplined nature of the average citizen. Plato believed that a democratic society is uncivilized at its core, and a democratic society is one in which there is an over-abundance of freedom and liberty, such that citizens may choose to ignore authority, become leaders themselves, and start wars as they wish. As such, in the short run, it seems like democracy is a pleasant form of anarchy (pp. 292–294). However, according to Plato, the democratic citizen is one who is strongly swayed by pleasure and is averse to reason and authority (pp. 295–298). This aversion to authority means the average citizen would reject the authority of the expert and consider them a “word-spinner” and a “star-gazer”, instead preferring a persuasive average citizen capable of controlling the system for the average citizen's selfish

⁶ All citations in this section come from this source.

benefit by whatever means necessary (pp. 210–211). This persuasion is effective because the average citizen has not heard enough free and fair discussion and is therefore liable to the subtle rhetorical tricks of those just trying to make a point for the sake of arguing rather than engaging in deliberation for the pursuit of knowledge (p. 222). Plato did, however, recognise that those pretending to be philosophers give the true philosophers a bad reputation, therefore making it difficult for the average citizen to respect true philosophers (pp. 210–211); however, Plato did not consider this an excuse for the democratic citizen's undisciplined behaviour (p. 223).

Plato's solution to the deficits of democracy was to propose a class of philosopher rulers. These philosopher rulers needed to have a love of knowledge and concentrate on the eternal and unchanging things in the field of knowledge, rather than on the field of opinion (pp. 65–66,203–204). Furthermore, Plato claimed that because the philosopher rulers understood the eternal and unchanging things, they were suitable guardians of the state; their expertise gave them natural legitimacy, which should be easily understood by the average citizen once it was explained clearly to them (pp. 204–205,210–211,222–223). Aside from the "love of knowledge" requirement, philosopher rulers were also meant to be elderly, dedicated to their community, hard-working, able to withstand pain and competitive trials, and not to be distracted by the whims of pleasure (pp. 113,114–115). Recognising that this requires complete dedication on the part of the philosopher ruler, philosopher rulers were not to take on other jobs, and they were to be trained from a young age (pp. 88–89). Given all these criteria, the group of those with the knowledge to rule would be far fewer in number than any other group with specialised knowledge; a governance system utilising philosopher rulers would definitely be a rule by the few, rather than a rule by the many (p. 132).

The mode of operation for philosopher rulers was to largely ignore the opinions of the average citizen while still maintaining an awareness of the events of average citizens.

If the philosopher rulers were to be concerned with public opinion, they would become beholden to the changing whims of the democratic citizen to keep them pleased, including those democratic citizens that would use the philosopher ruler's position for their own gain (pp. 215–216). That is to say, philosopher rulers were to primarily focus their minds on higher realities, rather than the affairs of the average citizen (pp. 223–225). However, just as Plato has argued that those who are entirely focused on the whims of the average citizen were unfit to rule due to their lack of understanding of the higher reality, those philosophers that only focussed on the higher realities were also unfit to rule due to their lack of understanding of the affairs of the average citizen. The philosopher rulers needed to bring themselves down from the upper world of philosophy into the political realm, doing so not because they hoped to gain glory or riches, but out of a sense of duty, since they were trained for the position from a young age. Furthermore, because the philosopher rulers sought pleasure in wisdom rather than in material riches as part of their duty to rule, they could not be bribed and were said to rule objectively (pp. 246–248).

In other words, Plato believed that the democratic citizen lacked the self-discipline and knowledge to rule, and so the philosopher ruler must take over to ensure high-quality policy is made. It could be said that Plato's formulation of the problem of expertise is coarse and his solution is extreme, though this is understandable given that it is one of the earliest attempts in Western philosophy to consider the relationship between governance, experts, and lay citizens. Plato's formulation thus provides a foundation for subsequent research into the tension between expertise and democracy, while his category of "philosopher rulers" is an essential precursor to the modern notion of the expert.

The question of how best to reconcile expertise with democracy resonates throughout the Western theoretical tradition (see, for example (Condorcet, 1995 (1785)), (Hamilton, Goldman, Jay, & Madison, 2008[1787–1788]), and (Mill, 1861)). For the purposes of this thesis, it is the development of modern technological societies of the

late 19th and early 20th centuries that puts this relationship at the centre of democratic governance and theory. Hence, another important development in the understanding of the problem of expertise that deserves analysis occurred when Lippmann and Dewey both in the 1920s argued that democracy is very epistemically demanding on citizens, i.e., the standards required by democracy were considered the issue regarding the citizens' "sub-standard" performance more so than the citizens themselves, in contrast with Plato arguing that the citizens were at fault. As a result, Lippmann and Dewey proposed solutions that act to make democracy less demanding on citizens, though they disagreed on how this was to be done. Lippmann proposed a representative scheme where experts were considered capable of withstanding the rigours of policymaking, while Dewey sought to modify democracy such that experts were able to help the general public through the democratic process.

§1.2.3: The Lippmann–Dewey Debate on Easing Democracy's Demands

Beginning with Lippmann's discontent at democracy's failing ability to manage an ever more complex society, Walter Lippmann and John Dewey were engaged in a debate during the 1920s regarding the nature of democracy (Lippmann, 1922, p. 16). This debate was largely centred around two books: Lippmann's *The Phantom Public* (1993[1925]) and Dewey's response in *The Public and its Problems* (1927). While this debate concerns more than just the problem of expertise, on the issue of the problem of expertise, there was a significant level of agreement on the cause of the problem of expertise, but they did disagree on the solution.

*Lippmann*⁷

Walter Lippmann believed that citizens had become distant from the political process and were simply carried along by the process, largely unaware of what was going on. Even if a citizen votes, "he knows that his sovereignty is a fiction", especially when

⁷ Within this section, all references come from *The Phantom Public* (1993[1925]).

compared with what is expected according to democratic theory (pp. 3–4). According to democratic theory, the citizen standing in the polling booth is “a highly intelligent and public-spirited voter” capable of making an informed and effective choice (pp. 30–31). Lippmann believed that, in reality, the citizen was bound by machinery of government that was created and maintained out of their control, and too numerous and complicated to keep track of to any significant degree (*Ibid.*). This dissonance between theory and reality was Lippmann’s primary concern with democracy: it set up an impossible ideal. The assumption was that if citizens only engaged more often, and with better quality information, they would be capable of self-government. Yet, overall, citizens have never come close to this ideal (p. 136); the citizen could not hope to have even the coarsest of opinions on all political problems and will at best catch glimpses of what policymakers do in place of the citizen’s judgment regarding these political problems (p. 115).

However, Lippmann did not fault the average citizen for this to anywhere near the same degree that Plato did; Lippmann was sympathetic to citizens since neither he, nor anyone he knew of, was capable of living up to the standards of the omniscient citizen that democracy demands, and he was a professional political scientist (pp. 10–11). Of the many public concerns the citizen is meant to keep track of, it is reasonable that “the citizen gives little of his time to public affairs, has but a casual interest in facts and but a poor appetite for theory” given that citizens have more pressing personal concerns to attend to first such as their work and family life (pp. 14–15). In short, Lippmann believed an individual does not, and simply cannot, have informed opinions on all public affairs, and further contrary to democratic theory, the aggregation of individual ignorance cannot provide a useful driving force in public affairs (p. 29). Therefore, the assumption that the popular will was wise if only one could get at it was wrong; as far as Lippmann was concerned, the cure for democracy was not more democracy (p. 25).

Instead, Lippmann's solution for governance when presented with low political competence recognised what he saw as the radical difference between the insider (the public official) and the outsider (the average citizen), which democracy missed by expecting those outside the immediate policy arena to be as competent as those within it (p. 137). Given democracy missed this distinction, the civic education it provided created a bewildered and insufficiently trained public, rather than the specialisation required that came from law and business schools (pp. 138–139): "[t]he outsider is necessarily ignorant, usually irrelevant and often meddlesome, because he is trying to navigate the ship from dry land" (p. 140). Surmising that a general civic education was not sufficient, Lippmann took advantage of the expertise that already existed; the insider was not only better trained but was also in a position to act according to that training (*Ibid.*).

Given the limited political competence of the general public, Lippmann believed that any theory of democracy must understate the potential of public action and that the average citizen should not meddle in affairs that were beyond them, especially on issues that were confused, subtly balanced, or hard to understand (pp. 53,130–131). More specifically, Lippmann recommended that the general public were not to take part in:

- executive action;
- discussing the intrinsic merits of a policy question;
- the anticipation, analysis, and solution of a question; or
- providing the specific, technical, or intimate criteria required in handling the question (p. 134).

Furthermore, Lippmann believed that the government should not be considered representative of the general will. Instead, he thought that it should be run by public officials, and only in the event of official failure, or when called upon, the public should intervene (pp. 62–63). The role for the public, under Lippmann's proposal, was

to determine whether insiders were behaving in a professional and objective manner, rather than to contribute based on a detailed understanding (p. 134).

*Dewey*⁸

John Dewey agreed with Lippmann that the public faced several problems when it came to self-governance, and in particular, that they were largely “ghosts” in a complex and technical government machine, and therefore when they did interact, they could cause serious harm to government action (p. 313). Society had become so nuanced and complex that citizens could no longer get by with a small set of political rules to vote with, and the general public did not have the time, ability, or desire to uncover this nuance fully (pp. 317,320,335–336). To add to this, the general public even found it difficult to know what process they should use to understand society’s complexity (pp. 334,337). Dewey also agreed that the disparity between what individual citizens perceived they were capable of and the potential effect they could have in the face of the overwhelming political machine caused apathy within citizens, noting that many citizens believed their votes made no difference in the battle between Lippmann’s “insiders” and “outsiders” (pp. 308–309,317). Dewey added that the increase in the forms of entertainment brought about by improvements in technology served as a great distraction for an apathetic public, mirroring Plato in this regard (pp. 321–322).

However, where Dewey disagreed with Lippmann was regarding the purpose of democracy, and hence Dewey proposed a different solution to solve the problems that the general public faced. Where Lippmann considered democracy to be a theoretical ideal given by democratic theorists, Dewey considered democracy to be a more practical product of need, and that any criticism of democracy should be done by comparing it with practical alternatives rather than an ideal theory (p. 304). For

⁸ Within this section, all references come from *The Public and its Problems* (1927).

Dewey, the purpose of democracy was to ensure that the actions of the governors were aligned with the interests of the governed rather than the governors (p. 293). Therefore, Dewey's solution to the problems of democracy was to maintain some form of democracy since the governed know their own interests and concerns the best, even if it was the experts who knew best how to fix the problems that arose. Furthermore, democracy was educative and forced a recognition of common interests and an elaboration of the true nature of political problems (p. 364). As a result, Dewey did not advocate for a general rule of experts, or "insiders", as Lippmann did. For Dewey, a rule by experts was necessarily ignorant of the concerns of the general public, and hence it was claimed that the governors would govern in their own interests (pp. 364–365). Here, it is important to note that Dewey has recognised that expertise is a broader concept than the one Plato and Lippmann focused on, which is to say that the average citizen can hold local expertise that the "expert" in the traditional sense was not privy to.

What Dewey saw as the main problem of the public was how to improve "methods and conditions of debate, discussion, and persuasion" (p. 365). While experts were a necessary part of this discussion, Dewey believed their role was to engage in technical inquiry and making facts known rather than in the framing and executing of policies; the role of the general public at large was to judge these facts against common concerns (*Ibid.*). Dewey did recognise that this was still a significant task for the general public given that important information could be hidden or misleading, and that the public would have to think outside their natural modes of thinking. Dewey, therefore, recommended that measures must be taken to make information transparent and objective, and that advancements in technology could be used to educate citizens on how to think from various points of view (pp. 365–366). Therefore, contrary to Lippmann, Dewey did agree that the cure for democracy was more democracy, as long as "more democracy" meant "more and better-informed public

deliberation” as opposed to “more people taking part in the same old democracy” (p. 325).

The Lippmann–Dewey debate is important for the development of understanding the problem of expertise for two main reasons. First, the quality of the democratic system, rather than just the participants, was recognised as being an important aspect regarding the interaction between expertise and democracy, and as such, could become a focus of methods to improve this interaction. Second, the way in which Lippmann and Dewey disagree perfectly highlight the tension found within the problem of expertise, with Lippmann choosing an expert-focussed solution and Dewey choosing a democracy-focussed solution. This left room of improvement for contemporary political scientists, such as Aaron Wildavsky, who went further to suggest that the democratic system be organised in such a way that a middle ground could be found where experts and citizens meet and work together to minimise the tension found in the problem of expertise.

§1.2.4: Truth to Power: Balancing Expertise and Democracy

Driven by his concern that the methods of policy analysis in his time may not be suitable for societal change, Aaron Wildavsky wrote his influential book, *Speaking Truth to Power: The Art and Craft of Policy Analysis* (Wildavsky, 1979, p. 1).⁹ In it, Wildavsky spoke of intelligence versus interaction, which directly corresponds to the expertise versus democracy tension found in the problem of expertise. In discussing intelligence versus interaction, he did not exclusively choose one side over the other. Instead, he sought to minimise this tension by creating a hybrid model between the two, but was drawn more towards the interaction aspect, with intelligence monitoring and strengthening the interaction (p. 12).

⁹ All citations in this section come from this source unless otherwise stated.

The reason why Wildavsky called for a hybrid form, in agreement with Dewey, was because even if there were what he called “Grand Planners” who acted on pure intelligence, perfectly understood the causal relationship between a policy and its effect, and acted in the common interest, the Grand Planners could not know the policy preferences of the public, nor the extent of the difference in preferences between groups without asking them (pp. 120–121). However, Wildavsky recognised that the policy preferences of the public could be unwise due to the effects of propaganda and indoctrination, among other effects. In this scenario, there would be an expert that knew better than the citizen with the unwise policy preference, but Wildavsky claimed that this only gave the expert the opportunity to discuss the issue with the citizen, not to overcome the citizen. This was meant to allow for the autonomy and mutual respect of the citizen to be maintained (pp. 121,393). When intelligence was infused with interaction, social forces were thought to guide intelligence, but this was at the risk of intelligence becoming the handmaiden of power of social preferences such that arguments were given for inexcusable behaviour. With that said, Wildavsky considered pure intelligence to be worse than pure interaction because it allowed for neither autonomy of the general public nor mutual respect between the general public and experts, and therefore he claimed that intelligence should be an aid to, rather than a substitute for, social interaction (p. 125). This is a more nuanced position than the one proposed by Dewey. While Dewey claimed that experts should act as a source of technical facts, Wildavsky proposed a deeper role for the expert where experts acted to propose solutions and genuinely engage in discussion with citizens. The role for citizens in Wildavsky’s model was to become issue specialists in such a way that they were prepared for policy issues in their specialist area as they arose (pp. 257,258); both experts and citizens each brought their own respective expertise and a basic trust in democracy and collective decision-making such that public participation could be meaningful, efficient, and fruitful (p. 255) (Forest, 2013, pp. 7,9).

§1.3: The Post-Normal Era Challenge for Democracy

The problem of expertise and how to best reconcile the demand for expert knowledge with the demand for public participation has been a major thread running through Western political theory for approximately 2400 years. Today, the question remains highly salient in the context of the modern challenge of wicked problems and the post-normal era. This section discusses how the post-normal era makes it difficult to make appropriate connections between cause(s) and effect in an increasingly complex and uncertain world, as well as reconcile differences in values regarding how to solve the problems that arise, but also presents the commonly suggested solution towards approaching such wicked problems, namely by utilising an “extended peer community”.

§1.3.1: What Are Wicked Problems?

Rittel and Webber were the first to characterise the essence of a wicked problem and the following is a summary of their description of a wicked problem (Rittel & Webber, 1973). First and foremost, there is no definitive formulation of a wicked problem. This lack of definitive formulation is due to the fact that in order to adequately describe a problem, one must do so in terms of potential solutions. For example, in a tame problem, such as solving a mathematics equation, the equation is recognised as part of a toolbox that also contains a process to solve that problem. With wicked problems, the policymaker begins with a toolbox that is incomplete, i.e., there is always a potential legitimate response to a wicked problem that no-one has thought of or will ever think of. Furthermore, with wicked problems, the systems involved are complex, which means it is difficult to identify the cause of the problem, especially because the cause of the problem is the result of another wicked problem; this puts the policy maker in a web of interlinked wicked problems. One may try to develop a resolution for the situation at hand (wicked problems are never solved, but rather iteratively resolved, i.e., wicked problems have no stopping rule), but it is only after this resolution has been tested and the results of the decision are making themselves

known over time that the problem can really begin to be understood. One may be able to contentiously say that progress is being made, i.e., a good decision was made by examining short- and long-term consequences of the decision, but one cannot say that the decision was correct since there are no well-defined criteria for doing so. Furthermore, since these consequences are substantial and not easily reversible, and every response to a wicked problem is a one-shot operation that is essentially unique, trial-and-error cannot be used. Also, due to the consequences of wicked policymaking being so substantial and irreversible, the policymaker has no right to be “wrong”; lives and livelihoods are significantly affected, and therefore the policy maker must make every conceivable effort to produce a good outcome since there are no safety nets.

In other words, wicked policy problems are policy problems characterised by high levels of chaos, complexity, and contradictions (Sardar, 2010, p. 436). A classic example is climate change, where small changes in the climate system can have significant unforeseen effects (chaos) due to the high number of climate interactions (complexity). Furthermore, there is great disagreement over how to proceed with potential solutions (contradictions). To put it another way, wicked problems are characterised by a combination of high levels of *complexity* regarding the elements, sub-systems, and interdependencies; *uncertainty* in relation to risks, consequences of action, and changing patterns; and *divergence* and fragmentation in viewpoints, values, and strategic intentions (Head B. W., 2008, p. 103). By subsuming uncertainty under complexity, a policy problem that has multiple stakeholders with conflicting values and interests (high diversity), as well as neither the problem nor the solution being well defined (high complexity), then this can be considered a very wicked problem. In direct contrast, a problem with a single stakeholder (low diversity) and both the problem and solution are known (low complexity) is a tame policy problem. Wicked policy problems, in general, are those with at least medium levels of diversity and complexity (Head & Alford, 2008, p. 10).

§1.3.2: What is Post-Normal Science?

When speaking of science and technology policy in particular, a very closely related idea to wicked problems is one known as “post-normal science”. Post-normal science is set up in contrast with Thomas Kuhn’s notion of normal science. According to Kuhn, when major scientific discoveries are made, such as Copernican astronomy or Newtonian mechanics, these discoveries form a paradigm within which normal science can take place, where these paradigms provide the foundation for further practice (Kuhn, 1970, p. 10). Scientists working within a paradigm are committed to the same rules and standards for scientific practice as governed by that paradigm, i.e., for any given “puzzle” facing a scientist within a paradigm, there are well-established methods that can be used to solve that puzzle, and it is possible to claim that the puzzle has been solved (*Ibid.*, pp. 11, 35–37). This is not to diminish the skill required to solve the “puzzle” since a great deal of theoretical and methodological understanding is required to obtain and interpret experimental results, but rather to emphasise that the process from puzzle to solution is better established under a normal science regime than it is in the case of post-normal science. As Funtowicz and Ravetz put it: “In this ‘normal’ state of science, uncertainties are managed automatically, values are unspoken, and foundational problems unheard of.” (Funtowicz & Ravetz, 1993, p. 740).

In contrast, post-normal science is the regime where “facts are uncertain, values in dispute, stakes high and decisions urgent” (*Ibid.*, p. 744). It is the realm in which systems uncertainties and/or decision stakes are high, where “systems uncertainties” relates to the comprehension and management of a complex reality rather than simply the collection of particular facts, and “decision stakes” are related to the costs, benefits, and value commitments of the various stakeholders (*Ibid.*). Because facts are uncertain and value commitments are of great importance, post-normal science breaks the status quo where scientific facts, i.e., “hard facts”, dominate. In a post-normal science regime, the status of social and ethical considerations, i.e., “soft facts”, hold at least the same

importance as hard facts, if not a higher importance due to the potential difficulty in obtaining hard facts (*Ibid.*, p. 750). However, this is not to say that soft facts are without uncertainties, or even that they are separated from hard facts (the collection of hard and soft facts, including the combined facts, is referred to as “extended facts” below) (Funtowicz & Ravetz, 1994b, p. 1884). In the post-normal regime, values are in dispute and they can influence the generation and interpretation of hard facts. This can be seen by noting that any statistical analysis requires categories to be set up. These categories and the lines between them are influenced by the researcher’s social, political, and ethical background. A classic example of this is the appearance of homosexuality in the *Diagnostic and Statistical Manual of Mental Disorders* up until 1986, where statistical analysis regarding homosexuality would be done by some researchers under the assumption that it was a mental disorder rather than simply a form of sexual preference (Herek, 2012).

§1.3.3: Society Is Also Complex

Complexity and uncertainty do not just relate to the natural world, they apply to society as well. As such, it is perfectly legitimate to extend the concept of post-normal science to the more general idea of a post-normal era. Social complexity and uncertainty are unavoidable because democracy cannot control society as an entirety by looking at it from the outside. Decisions must be made from within society as part of a particular context, and so there will be many factors outside this context that are uncontrollable as far as that particular policy case is concerned. Each of these factors interact in numerous ways with each other and with the case at hand (Bohman, 1996, p. 13). In this post-normal era, globalisation means that there is a vast array of interconnections and interdependencies in which changes, even small ones, can move rapidly through these connections and have the potential to create a large change as a result. In such a complex world, there are also contradictions, not only in terms of motives and values, but also in terms of the levels of change seen in the world and what is known. There are also other systemic social constraints that add to the complex

and uncertain nature of decision-making, such as asymmetries in competence, expertise, and availability of information, among other scarce deliberative resources (*Ibid.*). To add to this, these resources are dispersed in a complex manner, and can be found in many disjointed organisations—such as clubs, workplaces, and advocacy groups—and the information they are discussing comes from various resources with varying credentials—such as government institutions, mainstream media, and other agencies. Any plausible concept of governance must do justice to such dispersed discourse before it can consider a way of resolving moral disagreement (*Ibid.*, p. 177).

At this point, two aspects of extended facts that are particularly important for the purpose of this thesis that need to be highlighted. The first, and more commonly recognised aspect, is that scientific knowledge is complex. More importantly, it is becoming increasingly complex at a rapid rate and a significant level of dedication is required to keep up with the forefront of knowledge. This makes it difficult for those outside the related fields to keep up-to-date with relevant information to make decisions with regarding such issues and accentuates the problem of expertise. The second, and more controversial aspect, is that values (and soft facts in general) are becoming more complex. In a time where emerging technology is changing so rapidly, and futures assessment is required for governance just to keep up with technological development, the discussion of values can no longer rely on the ethical reasoning skills of the average citizen. It is very difficult for anyone, let alone the average citizen, to formulate a cogent ethical argument for or against, say, highly personalised nanomedicine or the rise of cyborgs. To have high-quality governance of these technologies, the values side of governance cannot just be left to the average citizen after scientists have passed on their factual knowledge of the case at hand, as seems to be the preferred *modus operandi* (see Chapter 2 for more detail). Moral reasoning is a highly technical form of expertise, and so is understanding the complexity of society. Therefore, this thesis argues that the soft facts discussed above will need to be evaluated by experts in the case of emerging science and technology, and these experts

will play a significant role in formulating emerging science and technology policy as part of an extended epistocracy. This is what makes the proposed extended epistocracy “extended”; it goes beyond the hard facts of the policy issue and the technical experts that know them and adds relevant experts that can provide the soft facts that are required.

These two aspects of extended facts are particularly salient considering the uncertainties found in the post-normal era, which makes emerging technology in particular more difficult to manage. Emerging technology, by definition, is still emerging, and as such, there are many uncertainties surrounding its ethical, social, political, legal, economic, and physical impacts. Ignorance abounds when it comes to emerging technology policy such there is an ignorance of potential risks, ignorance generated by information overload, and even an ignorance regarding the extent of the ignorance. Uncertainty, it seems, has taken centre stage (Sardar, 2010, pp. 437–440). Therefore, this uncertainty is much broader than simply just the health or environmental risks of a particular technology, i.e., the calculated probability that the technology would produce a negative outcome to health or the environment. In the post-normal era, it may not be known whether a negative event will occur, but on top of this, neither the probability nor the causes of a negative event may be known accurately, and it may even be difficult to ascertain what would constitute a negative effect due to uncertainties in soft facts. To add to this, not only are emerging technologies interdisciplinary in a technoscientific sense, but there are also many institutions vying for their place in emerging technology industries, along with other policy stakeholders, that need to work together to ensure the appropriate development of emerging technology. These institutions and other stakeholders, each with their own viewpoint, need to work together to manage uncertainty, not by pushing it to the side, but rather by embracing and managing it as an extended peer community. Given that extended facts have increased importance in the post-normal era, one way to make wicked problems more manageable, i.e., a wicked-ready

governance system, is to make effective use of this extended peer community, as discussed below.

§1.3.4: Responding to a Post-Normal Era with an Extended Peer Community

To develop a policy system capable of managing wicked problems, it is necessary to know what it means for a policy system to be “wicked-ready”. However, a difficulty arises when it is realised that there is no answer to the question: “what does wicked-readiness look like?”, at least not a definitive one. This is because, as noted above, when it comes to wicked problems, there is no stopping rule or set criteria to determine whether a wicked problem has been solved (Rittel & Webber, 1973, p. 162); this makes it impossible to know whether a governance system has achieved its goal. There are, however, some general guidelines that can be used to determine whether a policy process is going to produce a more successful outcome. One of these guidelines is to make sure an extended peer community is used, where an extended peer community is one that brings more voices to the discussion table over and above those that would come from the elite, i.e., technical sciences and politicians. Extended peer communities have the potential to bring extended facts, such as craft wisdom and community knowledge, as well as the expert knowledge found in science and economics into the discussion (Funtowicz & Ravetz, 2003, p. 7). This range of epistemic resources helps to provide a greater understanding of what issues should be considered regarding the policy case at hand in terms of both hard and soft facts, as well as in terms of the complexity, chaos, and contradictions involved.

With extended peer communities also comes the inclusion and elevation of the importance of soft facts in contrast to the positivist approach to policymaking where hard facts are revered.¹⁰ Soft facts come from value commitments (Funtowicz &

¹⁰ Positivism is discussed in more detail in Section §2.2.

Ravetz, 1994b, p. 1884), which include ethical concerns, political ideologies, and social and cultural points of importance. In post-normal science, the relative importance of persons and purposes becomes enhanced due to the technical uncertainties involved, and as such, it makes sense to bring in those most affected by a policy issue and its potential policy outcomes into the discussion (*Ibid.*, p. 1885). Each of these stakeholders will enter the discussion with their own perspective and soft facts, and conflicting interests and opinions will inevitably arise. A criterion of quality for the dialogue, therefore, will be one that presupposes that ethical principles are to be explicit as part of the dialogue. Furthermore, this plurality of perspectives does not deny the competence of those with special expertise, but it does mean that it is important to recognise that an expert is also a layperson in a range of areas, including the value commitments of others (Funtowicz & Ravetz, 1994a, p. 204).

In short, a wicked-ready policy system is one that embraces extended facts and uncertainties by way of extended peer communities. It would seem, then, that a form of democratic public engagement would be a natural extended peer community. However, this thesis argues that the extension of the peer community as part of the deliberative turn (as discussed in Section §2.3) has gone too far, giving preference to public participation at the expense of policy quality, including in the case of emerging science and technology policy. To temper this extension of the peer community, this thesis proposes a particular form of epistocracy, called an extended epistocracy, as a middle ground between technocracy and the more participatory public engagement. Given that “epistocracy” is a broad term, this thesis argues that it is well within the boundaries of an epistocracy to find a governance regime that respects both hard and soft facts, while still retaining a high level of expertise and hence obtaining high-quality policy. This is because the purpose of an extended peer community is to bring in diverse facts and ways of thinking into the deliberative process. An extended peer community does not require that everyone holding those diverse views *actually take part*, which is what would be required under a more democratic governance regime.

It is possible to create a deliberative scenario with expert representatives from various areas of relevant expertise covering the technical sciences, along with the social sciences and local knowledge. This, put simply, is the structure of an extended epistocracy, where a more detailed form is given in Chapter 5.

§1.4: Conclusion

The problem of expertise is defined as the tension between expertise and democracy, and scholars' understanding of this problem has evolved significantly over time, where it is now recognised that there needs to be a more nuanced interaction between experts and the average citizen. This sets up the goal of the thesis, which is to find an appropriate balance that minimises the tension between expertise and democracy. Doing so under modern political conditions is especially pertinent given the chaos, complexity, and contradictions of the modern post-normal era. The commonly accepted solution to this problem of expertise is the use of an extended peer community, particularly in the form of deliberative democracy in theory and public engagement for governance in practice. However, this thesis argues that deliberative democracy and the associated public engagement currently overemphasises public participation at the expense of policy quality. Examining the change in public participation in governance and the importance that has been assigned to it is the topic of the next chapter, which analyses the development of science and technology policymaking techniques since the mid-20th century, what lessons were learned, and how this influenced the creation of the extended peer community in use today. By doing so, important factors that need to be included in a suitable governance system can be established, and an understanding of why democratic methods have become the status quo is gained.

Chapter Two: Towards the Epistemic Turn of Deliberative Democracy

§2.1: Introduction

The discussion of wicked problems in Chapter 1 introduced the concept of the extended peer community as a way to create more effective public policy amidst the chaos, complexity, and contradictions of the modern post-normal world. This extension of the peer community is seen in theory in the form of deliberative democracy, and in practice in the form of public engagement for governance, for example, in the *A Nation of Curious Minds* report, which, among other recommendations, proposes methods for better interaction between the science sector and society in New Zealand (MBIE, 2014, p. 30). To understand the role that the problem of expertise has in this extension of the peer community, it is important to discuss the development of modern deliberative democracy, the role of expertise and epistemic standards within that development, and the consequences that this approach has for public engagement for governance. Sections §2.2 and §2.3 in this chapter, which are historical and philosophical in nature, thus aim to, respectively, answer the questions:

- What caused the increased emphasis on extending the peer community in the latter half of the 20th century?
- What effect did this extension of the peer community have on the epistemic nature of deliberative democracy?

After showing that initial forms of deliberative democracy focussed heavily on democratic participation and hence lacked an emphasis on the epistemic aspects of decision-making, Sections §2.4 and §2.5 focus on how to incorporate epistemic aspects more effectively into governance under modern political conditions. These two sections aim to, respectively, answer the questions:

- What challenges does an epistemic form of deliberative democracy or public engagement for governance face under modern political conditions?
- What forms of public engagement for governance could be used to satisfy these challenges?

Section §2.2 begins the discussion of the development of modern deliberative democracy by starting in the 1960s, a time before the deliberative turn when positivism was the dominant underlying philosophy for governance. In this section, positivism is described, and most importantly, so is its decline in popularity due to its ignorance of values and context, its inability to consider uncertainty properly, and its elitist nature. While this decline began in the 1960s, in practice, these faults were still present late into the 20th century, as seen in the deficit model. As a result of this ongoing positivistic influence, a deliberative turn took place around 1990, which had a distinctly postpositivistic approach, where postpositivism solves the problems of positivism by incorporating extended facts and uncertainty (Torgerson, 1986, p. 243). Section §2.3, then, discusses what postpositivism is, along with the very heavy focus on the procedural aspects of governance postpositivism inspired by elaborating on the principles of deliberative democracy. While these principles are not without their benefits, they were given priority over the epistemic aspects of policymaking, which tended to be ignored or even explicitly avoided. It is argued that this privileging of process over epistemic quality causes several problems regarding both the structure and output of deliberative democracy.

As a proposed method for countering these problems, the epistemic turn of deliberative democracy began in approximately 2010, as discussed in Section §2.4. While this epistemic turn allows for promising avenues of research through both the theoretical incorporation and empirical testing of epistemic features, it also faces challenges, such as the democratic trilemma (between political equality, deliberation, and mass participation (Fishkin, 2011, p. 248)), and the epistemic inequalities that arise

under modern political conditions. To evaluate whether an epistemic form of public engagement can succeed under modern political conditions, it is important to know the potential that practical deliberation has and what forms it can take. Given the focus of this thesis is the governance of emerging science and technology, nanotechnology governance is chosen as the example emerging technology due to it being a newer form of emerging technology that can be used to develop new forms of upstream engagement that avoid the problems associated with the governance of past emerging technologies (Priest & Greehalgh, 2011, p. 1521). As such, Section §2.5 presents the key features of different potential forms of public engagement that nanotechnology has utilised in terms of the participants involved, the information given, the interaction between participants, and the aims of the engagement process. However, these forms of public engagement face some challenges, notably the epistemic demands placed on participants while satisfying the mass participation horn of the democratic trilemma. One form of public engagement in particular, namely the consensus conference, is a comprehensive form of public engagement that has properties that are claimed to be able to help with these epistemic demands and provide a proxy solution to the democratic trilemma. Therefore, by discussing the consensus conference, this provides a paragon example of public engagement for governance that can be kept in mind for the discussions regarding epistemic democracy in Chapters 3 and 4.

§2.2: Before the Deliberative Turn

While not named as such, it could be argued that the concept of deliberative democracy goes back as far as the ancient Athenians, where deliberative democracy can be considered a process whereby participants “offer proposals for how best to solve problems or meet legitimate needs, and so on, and they present arguments through which they aim to persuade others to accept their proposals”¹¹ (Young, 2004, p. 227). In the Athenian case, a council of around 500 randomly chosen citizens would

¹¹ A more detailed description of deliberative democracy is given in §2.3.2.

hear appeals regarding legal and political cases (Fishkin, 2002, p. 222). However, the move towards modern forms of deliberative democracy arguably began with the move away from a positivistic approach to policymaking, beginning around the 1960s. This is an important starting point in the discussion towards the deliberative turn since it allows for an understanding of what the deliberative turn was seeking to avoid, namely positivism's technocratic influence on policymaking.

§2.2.1: What Is Positivism?

In general, the goal of positivistic policymaking is the abolition of politics, i.e., the abolition of the subjective aspects of decision-making, such as values and social context (Lynn, 1999, p. 418). Positivism aims to do this within social science by claiming that, just like the natural sciences, the social sciences are able to be studied using sensory experience in an inductive manner (Yadav, 1986, pp. 504–505). When this philosophical viewpoint is applied to policymaking, positivism uses the rigorous empirical methods found in the natural sciences to create causal generalisations that attempt to explain political behaviour across social and historical contexts (Fischer, 1998, p. 130). Positivism considers these empirical methods to be the only reliable approach to knowledge generation due to their emphasis on the separation of facts and values; to be considered scientific, positivism claims that the scientist must remain value-neutral and focus only on factual phenomena (*Ibid.*, pp. 130–131). By upholding this fact–value dichotomy, policymaking is intended to be more efficient by avoiding value conflicts; positivism dismisses pluralistic policymaking, going so far as to say that if politics and positivism do not match, then politics needs to adjust (*Ibid.*, p. 131). Since positivism considers the laws of social science to be general, positivism ignores the knowledge and beliefs of specific social groups, and without values and political context to complicate proceedings, positivism claims to be able to calculate precise and accurate solutions to policy problems (Torgerson, 1985, p. 241) (Fischer, 2000, p. 17) (Lynn, 1999, p. 418).

This overall goal of abolishing politics was prevalent in the Enlightenment period of the 18th century and was revived with the introduction of positivism in the 19th century (Torgerson, 1986, p. 34). As nations became technological states, ever more dependent upon emerging technology, there was an increasing emphasis on technical-scientific decision-making, and hence a requirement for technical expertise (Beck, 1996, p. 35). Spurred on by this industrialist, techno-scientific spirit and the persuasive nature of objective decision-making, positivistic methods in policymaking carried on through the early 20th century, reaching its peak in the 1960s. However, this dream that reason could escape politics is now seen as political naivete (*Ibid.*, pp. 35,38–39).

§2.2.2: The Decline of Positivism

Positivism is considered naïve because it is impossible for policymaking to abolish politics; no matter what methodology is used, policymaking is a form of political action with political goals in mind, at least implicitly (Lynn, 1999, p. 420). Once it was realised that the political process could no longer be ignored, along with the fact that positivism had limited empirical success, disillusionment with the positivistic method set in (Torgerson, 1986, p. 45) (Fischer, 1998, p. 131). More specifically, there were several significant flaws with positivism that accelerated this disillusionment, three of which will be discussed in turn: the ignorance of values/context, the poor ability to handle uncertainty, and the charge of elitism.

1) Positivism Ignores Values/Context

The first reason why the disillusionment with positivism set in is because positivism ignores values and social context. This ignorance is significant because social action is oriented towards notions of what is considered good and desirable by those affected by a policy decision, and therefore, it is necessary to include the publics' goals, values, and points of view in the decision-making process (Fischer, 2000, pp. 19,43,185–186). Furthermore, policy problems are framed differently by different actors based upon their value frameworks, hence discussion between actors is helpful not only for

solving the problem, but also for establishing what the problem is to begin with (de Leon, 1997, p. 87). However, positivism ignores all of this. This is because positivism holds to the fact–value dichotomy, where facts are considered cognitively meaningful and values are not.¹² By adhering to this dichotomy, positivism neglects:

distributive ends, procedural and historical principles, and the values ... associated with personal rights, public goods, and communitarian and ecological goals (Tribe, 1972, p. 105).

Furthermore, positivism is:

relatively blind to exactly the kind of disagreements, and conflicting interests, which need to be perceived in order to guide a search for solutions that, over the long run, do not harm significant values or groups. (Nelson, 1977, pp. 75–76)

By ignoring values and context, positivism is unable to explain why a policy problem is even a problem that needs solving and what would constitute an appropriate solution. Furthermore, positivism also ignores who the solution is made for. Even if positivism could develop laws that described political behaviour, it is unable to apply this to guide the development of public policy (Torgerson, 1986, p. 36); without a problem that is recognised as such, a socio-historical context, or a value-oriented goal to achieve, then a positivistic social science is simply a black box without adequate input or a reason for output. Even if an attempt was made to consider social values, as in a neo-positivist framework using cost–benefit analysis, the result is likely to involve a distorted guess regarding the values of the public that is biased in favour of the experts. This bias occurs because scientific methods, such as a cost–benefit analysis, are very poor at considering the values of different kinds of social relationships or the uncertainty a given individual is willing to endure to remain in their hometown, for example (Fischer, 2000, p. 17) (Wynne, 1996, p. 57) (Wagle, 2000,

¹² Only statements that could be measured or mathematically proven for their verity are meaningful according to (logical) positivism. Ethical statements cannot be measured or mathematically proven, and hence they are not considered meaningful statements by positivism (Creath, 2017).

p. 209). In short, problem setting, and solving, is non-technical in nature, which is not to say there are no technical aspects to political problems, but rather that policy decisions are first and foremost social and political decisions.

Aside from ignoring the values and social context inherent in policy problems, positivism also discounts the expertise that members of the public possess. For example, in the case of Cumbrian farmers in north-western England living near the Sellafield nuclear power plant, farmers were concerned about the effect the power plant was having on the soil and livestock. The farmers possessed local expertise regarding the environmental conditions and sheep behaviour, but this expertise was ignored and overruled by a generalised technocratic approach to the detriment of meaningful progress on the issue of food safety (Wynne, 1996, pp. 66,68). Positivism discounts the fact that the public possesses local expertise in the form of both scientific knowledge and value judgements because this local knowledge is not in the form of official technical expertise (Fischer, 2003, p. 128) (Fischer, 2000, p. 23).

2) Poor Ability to Consider Uncertainty Appropriately

The second reason why disillusionment with positivism set in is because positivism is unable to consider uncertainty appropriately. One of the aims of positivism is the creation of precise and accurate policy solutions. However, this ignores concerns regarding:

the strength and fit of evidence, the relevance and reliability of data, the intrinsic limitations of scientific tools ... and the assessment of incomplete and often contradictory evidence (Majone, 1989, p. 174).

By not allowing for the incompleteness of information, there is little room for disagreement or discussion in general. As such, positivism portrays scientists as having an “infallible messenger” quality when it comes to declaring factual knowledge (Hawkesworth, 1988, p. 52). If experts hold accurate and objective information, then there should be no need for disagreement among experts. However,

these disagreements do occur and the contradictory technical information cancels each other out in the minds of the public; if two experts, both with supposedly correct information, cannot agree, then the public has to fall back on value judgements to make a decision about the issue (Fischer, 2000, p. 137).

Another problem with not recognising uncertainties arises due to the complex and chaotic nature of both natural and social phenomena, as discussed in Section §1.4; positivists cannot claim the certainty and accuracy in their results that they would like to (Fischer, 1998, pp. 131–132,133). A good example of where experts' overconfidence in their results can give poor results was in the case of the energy crisis in the 1970s, where even with access to an abundance of information, the positivistic approach was unable to make any worthwhile predictions for such a complex technical problem (Fischer, 2003, p. 10).

3) *Elitism*

The third and final reason to be discussed regarding why the disillusionment of positivism set in is the elitist approach that experts took to policymaking. The first two reasons for disillusionment discussed above show signs of this elitism, where positivistic policy analysts were overconfident in their own abilities and neglected input from those directly affected by policies made; that is to say, positivistic policy analysts placed themselves above the general public as part of a hierarchical relationship. Positivistic policy analysts took on the role of a professional in a professional–client relationship, similar to that found in law and medicine, where the public took on the role of the client (Fischer, 1993, p. 168). This kind of professional–client relationship distorts communication since the client is meant to simply submit to the professional's authority, and this creates its own political universe that is biased towards the professional. In other words, the positivistic method is not as value-neutral as it appears; rather, it serves the interests of those in power, even if those in power are bound by codes of professionalism (*Ibid.*, p. 166).

Beyond the subordination of the public in this professional–client relationship, this hierarchical structure has been attributed to the death camps of World War II, the nuclear race, the ecological crisis, and “scientific” communism; as such, it is understandable why this charge of elitism is pressed against positivism and holds the potency that it does (Hoppe, 1999, p. 202). This positivistic hierarchical approach has been said to be “devastating to democracy”, with the positivistic policy analysts being “handmaidens of power”, and the positivistic methodology considered “tools of tyranny” (Lynn, 1999, pp. 419-420). Furthermore, the elitist, hierarchical approach has:

not only ‘turned off’ large numbers of people to the political process, but has also led to elitist policies that have benefits for only the few. The result has been serious economic and social imbalances that threaten the political equilibria of liberal democracies (Fischer, 1993, p. 166).

Given all of this, it is little wonder why when it comes to discussion regarding the problem of expertise, that expertise is considered to be the problem rather than democracy (even though the problem of expertise is a tension between expertise and democracy and so both could potentially be held accountable for this tension).¹³

§2.2.3: Learning from Past Emerging Technology Governance

While positivism started to decline from the 1960s onwards, continuing into the late 20th century, emerging science and technology policy was still utilising a largely positivistic mindset. Past attempts at biotechnology policy arguably provide the most important lessons regarding emerging science and technology governance in this regard. The development of biotechnologies involved heated debates that came as a surprise to many in the scientific and science policy communities, with objections to these biotechnologies raised on economic, environmental, animal welfare, and human health grounds (Priest & Greehalgh, 2011, p. 1522). These debates followed public

¹³ This thesis argues that both expertise and democracy are accountable, with democracy being more responsible than expertise. However, the point remains that at this stage in the historical discussion, it is easy to see why expertise is considered responsible for poor policy practice.

health controversies, such as the mad cow disease crisis, which helped to provide renewed interest in science and technology communication and the public understanding of science, albeit in a misguided manner (see the discussion on the deficit model shortly) (Kurath & Gisler, 2009, p. 563) (Pidgeon & Rogers-Hayden, 2007, pp. 193–194). In hindsight, many consider it a mistake to have only included the public at a late stage of development when products and applications were already on the market (Burri, 2009, p. 498). Policymakers have also been faulted for not engaging with a diverse enough group of stakeholders, and so calls were made to:

engag[e] the broad array of stakeholders outside government and industry-labour, health organizations, consumer advocates and environmental NGOs whose constituencies stand to be both beneficiaries of this new technology and those most likely to bear any risks that arise (Environmental Defence, 2005, p. 4).

Not only is an improvement in community engagement a key message to take from debates regarding biotechnology, but so is transparency and accountability. There was a heightened mistrust in the capability of governments to manage controversial technologies, such as genetically modified food, and so global opposition and consumer boycotts ensued. It is important to recognise, then, that public trust is going to be a key factor in the success or failure of any emerging technology and its related governance (Lyons & Whelan, 2010, pp. 55–56).

As noted above, even when the public was initially included in discussions, it was implemented poorly. This implementation was in the form of the deficit model, which is based on the assumption that the public's lack of acceptance of a technology is due to their lack of understanding of the technology and/or the management of its risks (Cormick & Hunter, 2014, p. 57). To fix this deficit, it was thought that once formal education ended, it was said to be the job of science media to educate the public about science and technology, where "communication" was in the form of transmission from scientists to the public rather than a dialogue between scientists and the public.

It was thought that simply giving the scientific facts to the general public would allow citizens to judge scientific issues like scientists, and controversies would disappear because facts were meant to speak for themselves. If this “communication” did not work, it was assumed that the deficit lay with the public rather than the experts, and if it continued, it was the fault of an irrational public and/or poor journalism rather than with the experts ignoring the socio-political factors that influenced the interpretation of scientific facts (Nisbet & Scheufele, 2009, p. 1767). As such, the deficit model continued the dismissal of socio-political factors and reinforced the distinct hierarchy between the expert and the layperson that was found in positivism, with the layperson being passive receivers of the information that was given to them independent of the socio-political context (Kurath & Gisler, 2009, p. 562). The deficit model treats the public as a monolithic whole, rather than recognising that public attitudes are complex and variable, and there is little evidence to suggest that this deficit model style of communication has been helpful (Cobb, 2011, p. 1534). Critics of the deficit model argue that laypeople are capable of meaningful engagement regarding emerging technologies, and that the general public holds a sceptical stance towards emerging technologies due to a mistrust of science and government officials, not due to ignorance and poor science reporting (Kurath & Gisler, 2009, p. 563).

§2.3: The Deliberative Turn and the Emphasis on Procedure

§2.3.1: What Is Postpositivism?

The overarching theme, then, from the above discussion regarding positivism is that up until the late-1900s, there was a distinct lack of effective public participation regarding policy issues, with public participation being discounted, ignored, and even shunned. In response to these attitudes towards public participation, there was a shift towards more participatory forms of governance, deliberative democracy in particular. This shift may have been 30 years in the making, but when it took place, it was done with “rapidity and enthusiasm”, so much so that in the early 1990s, an event

called the “deliberative turn” took place (Pateman, 2012)(Goodin & Dryzek, 2006, p. 219). Before this turn, democracy was more focussed on the aggregation of preferences and representation rather than the more direct form that is deliberation (Dryzek, 2000, p. v). This direct deliberation took a postpositivistic approach, where postpositivism solves the problems of positivism by incorporating extended facts and uncertainty (Torgerson, 1986, p. 243).

More specifically, postpositivism holds that not only is objective knowledge important, as positivism does, but also allows for the subjective dimensions of social life and the local expertise this produces to be incorporated, and for the possibility of uncertainty within objective knowledge (*Ibid.*). Postpositivism holds that it is important to analyse data through a social perspective, with its subjective dimensions of socio-political life, by way of a value-critical policy analysis *alongside* objective methodologies (de Leon, 1994, p. 83). Social reality is actively constructed by human actors by assigning meanings to events and actions, and therefore considers the context and processes that shape meaning structures to be important, emphasising that the same objects, situations, and actions can be viewed and described differently depending on one’s worldview (Fischer, 2003, p. 48) (Fischer, 2000, pp. 74-75). By considering the information held by citizens to be useful, postpositivism not only gives policymakers the opportunity to understand how citizens see the world around them, but it also gives the policymakers access to valuable local knowledge that is inaccessible to the more general and abstract methods of positivism (*Ibid.*, p. 2). Postpositivism, therefore, seeks input from relevant stakeholder groups regarding how they understand the policy problem at hand, what effects different policy alternatives will have on that society, and what values are most important to consider for that society. By adding the knowledge system of the society in question to that of the policymakers, this creates a more comprehensive knowledge system to deliberate with (Wagle, 2000, pp. 208–209).

Furthermore, given this emphasis on local knowledge by postpositivism, postpositivism fundamentally challenges the expert–client relationship found in positivism where postpositivism claims that the expert is to have a collaborative relationship with the public such that the expert acts as a facilitator of public learning and political empowerment (Fischer, 1993, p. 167) (Fischer, 1998, pp. 141–142). Public learning is achieved by giving citizens access to information, along with suitable explanations, with the intention that the general public can take part in informed discussions and make informed choices. By giving information in a form that can be understood by the general public, citizens are given the opportunity to understand the policy process, as well as influence decisions in areas that matter to them, and hence, the general public gains political empowerment (Fischer, 2003, pp. 15–16,215). Postpositivism holds that this kind of equal relationship between experts and the general public is necessary because those who are significantly affected by a policy decision should have their say in the policymaking process, and a policy science of democracy will only be meaningful if citizen participation takes place (Wagle, 2000, pp. 213–214). Therefore, the goal of participatory policymaking, as an ideal, is “a fair and impartial representation of all citizens’ values and preferences” such that there is a “mutual determination of goals, a shared ‘spirit of inquiry’, equity among everyone to influence others, and [the] freedom to end [the] practitioner–client relationship in [the] policymaking process” (Renn, Webler, Rakel, Dienel, & Johnson, 1993, p. 206) (Wagle, 2000, pp. 214–215). In postpositivism, those in authority should not be distinct from the socio-political context in the form of a neutral observer applying their knowledge *to* society, but rather the policy analyst must become part of a contextual, interdisciplinary, and problem-oriented inquiry *within* a particular society, even if this results in imprecise and ambivalent answers (Torgerson, 1986, pp. 40–41) (Lynn, 1999, pp. 418–419).

§2.3.2: Procedural Emphasis within Deliberative Democracy

The formulation of the principles of deliberative democracy as part of the deliberative turn took a postpositivistic approach, as can be seen by the procedural emphasis discussed in this section. Deliberative democracy, as the name suggests, has both deliberative and democratic features. The deliberative features are concerned with proposing and discussing policy options to identify the best available option with the aim of coming to a consensus on the option (Cohen, 1997, p. 75). The democratic features are concerned with treating citizens as equals and giving them the opportunity to engage effectively in the deliberative process. Different forms of deliberative democracy will assign varying levels of importance to the deliberative or democratic features, and hence varying levels of legitimacy will be generated for these features between theories, where legitimacy is the moral permissibility to enforce the policy created (*Ibid.*, p. 73) (Estlund, 2008, p. 2). While there are variations between different forms of deliberative democracy, the following is a detailed list of some common features found between them:

- To provide the essence of democratic legitimacy, deliberative democracy is comprehensive, where all potentially directly affected sections of the population are allowed to take part in the process, either in person or through a representative, both to have their views heard and to help make the decision (Cohen, 1997, p. 73) (Young, 2004, p. 228) (Baber & Bartlett, 2007, p. 6) (Luskin, Fishkin, & Hahn, 2007, p. 3) (Young, 2001, p. 672) (Fishkin, 2011, p. 251).
- As a guiding political procedure, these participants are to take part in balanced discussion, where they propose solutions and offer reasons for them to persuade the others present, while the other participants are given the opportunity to respond to the proposed solutions and reasons given with critiques of their own from their own perspective (Baber & Bartlett, 2007, p. 6) (Luskin, Fishkin, & Hahn, 2007, p. 3) (Young, 2001, p. 672) (Luskin & Fishkin, 2002, p. 1) (Fishkin, 2011, p. 251) (Young, 2004, pp. 227–228) (Cohen, 1997, p. 74).

- To satisfy these criteria, the reasons must be given publicly so they are available for open weighing, acceptance, or rejection by others. The reasons given are to be in terms that all other participants ought to accept, understand, and consider to be good reasons (if given the chance to reflect), such that they are related to the interests of other groups and can be expressed in more universal principles (Young, 2004, pp. 229–230) (Baber & Bartlett, 2007, p. 6) (Young, 2001, p. 672) (Goodin, 2003, pp. 25–26) (Benhabib, 1994, pp. 32–33) (Dryzek, 2009, p. 1381) (Dryzek, 2001, p. 651) (Bohman, 1996, p. 44).
- Furthermore, the reasons given are to be informed and informative using reasonably accurate factual claims based on relevant information given during the procedure and/or acquired as part of prior efforts towards information seeking and reflection. As such, deliberation should induce reflection among participants to increase the coherence of the beliefs the participants hold (Luskin, Fishkin, & Hahn, 2007, p. 3) (Fishkin, 2011, p. 251) (Luskin & Fishkin, 2002, p. 1) (Dryzek, 2009, p. 1381) (Benhabib, 1994, p. 32).
- As part of this reason-giving process, participants are to behave in a conscientious manner by displaying a reasonable disposition towards other participants. This disposition includes being willing to talk and listen with civility and respect, and take account of others' viewpoints, weighing them sincerely, and be willing to be criticised (Luskin, Fishkin, & Hahn, 2007, p. 3) (Young, 2001, p. 672) (Fishkin, 2011, p. 251) (Young, 2004, p. 228).
- All participants are to have an equal right and opportunity to engage meaningfully in all parts of the policy process. This equality means that all proposals, reasons, questions, and criticisms are considered, and decisions are made, purely on their merits rather than based on who presents them, i.e., there are no power differentials (coercive or otherwise) that can influence decisions. Equality also means that policies are selected for the benefit of the many rather than the few (Fishkin, 2011, p. 252) (Young, 2004, pp. 227,228) (Baber & Bartlett,

2007, p. 6) (Benhabib, 1994, p. 31) (Luskin, Fishkin, & Hahn, 2007, p. 3) (Young, 2001, pp. 671,672) (Dryzek, 2009, p. 1381) (Cohen, 1997, pp. 74–75).

§2.3.3: Benefits of the Deliberative Procedure

There are many democratic benefits to deliberative democracy that come from the procedures just stated, and this thesis assigns these benefits into three overlapping categories: quality of life, citizen interaction, and citizen attributes. Beginning with the quality of life improvements, for some, the ability to take part in political life is an essential part of what it means to live the good life and as such can be considered a good in and of itself independent of the outcome (Christiano, 1997, p. 245) (Fearon, 1998, p. 11). However, the most prominent of the quality of life improvements is that the deliberative process enhances the perceived legitimacy of the outcome through participation (Guido & Tesón, 2006, p. 1). Citizens feel they can at least accept the outcome of a deliberative process, even if they lose, since the process is impartial and respectful in that it gives those who are affected a chance to have their say or be represented (*Ibid.*, pp. 189,192). This legitimacy, in turn, helps to increase external political efficacy and political trust, where citizens believe their government to be more responsive and citizens have a greater appreciation for the work that governments do (Luskin & Fishkin, 2002, p. 3). Citizens are also given the opportunity to exercise their autonomy by taking part in self-government through this civic equality (Guido & Tesón, 2006, pp. 1–2). This is particularly important for those that would normally be given less attention in political discussions, such as minorities and/or those who would otherwise be effectively powerless in the political arena (Christiano, 1997, p. 259) (Gutmann & Thompson, 1996, p. 133). This legitimacy, in turn, helps to reduce social conflict, which is helped by the fact that deliberative democracy is an iterative process (Guido & Tesón, 2006, p. 4) (Gutmann & Thompson, 1996, pp. 142–143); those that have lost a deliberative process will have the opportunity to make their case again as part of future discussions and potentially reverse the decision based on new evidence, arguments, or social conditions.

Moving into the benefits regarding citizen interaction, the reduction in social conflict is aided by increased mutual respect between citizens outside of the deliberative process (Christiano, 1997, pp. 251–252). Within the deliberative process, mutual respect is accorded to citizens by allowing them to argue their case. This is carried on outside the deliberative process since the understanding gained allows citizens to appreciate the viewpoints of others and citizens can be more accommodating. This increased understanding of the viewpoints of others also allows citizens to appreciate the common ground they share between these viewpoints and so projects can be developed with a shared sense of community (Fearon, 1998, p. 60). In other words, citizens develop their political respect and empathy (Luskin & Fishkin, 2002, p. 3). Such a shared sense of community can allow for a large consensus on certain issues, and allows citizens to develop their sociotropism, where they tend to think more in terms of the greater good and the complex interactions between different worldviews contributing to this rather than just their personal interests (Elster, 1998a, p. 11) (Luskin & Fishkin, 2002, p. 4) (Ackerman & Fishkin, 2002, p. 145).

It was noted above that citizens gain a mutual respect and a better understanding of fellow citizens, i.e., citizens develop empathy and courtesy as part of the deliberative process (Fearon, 1998, p. 59). These are just two of the attributes that citizens are claimed to develop as part of the deliberative process. From an epistemic perspective, citizens develop their moral and intellectual qualities (Elster, 1998a, p. 11); they learn the nuances of ethical debate by being exposed to the reasoning of others and having their own reasoning tested, and they learn more about the policy issue at hand from the facts of the case to how the policy process works (Ackerman & Fishkin, *Deliberation Day*, 2002, pp. 145–146). From a capabilities perspective, citizens develop their eloquence, rhetorical skill, and imagination (Elster, 1998a, p. 59); citizens not only learn what and how to think from taking part the deliberative process, but also how to present their case and develop potential solutions. In other words, citizens develop their political sophistication and internal efficacy such that they have better epistemic

skills and the confidence to use them (Luskin & Fishkin, 2002, p. 3). The deliberation process also stokes political interest and participation within the citizens such that after the deliberation process is finished, citizens become more interested and engaged in the political process (*Ibid.*) (Ackerman & Fishkin, *Deliberation Day*, 2002, pp. 147–148).

§2.3.4: Issues Concerning Procedural Theories of Deliberative Democracy

One will notice from the deliberative turn description of deliberative democracy that it is very procedural, i.e., it falls under the category of procedural deliberative democracy. While such a process has many merits, as noted above, there are many reasons why an emphasis on procedure at the expense of substantive outcome is concerning. For procedural democrats, as the name suggests, the primary aim of democracy is to embody certain procedural virtues expressing who is to be involved and how the process should be run. There is contention regarding what these procedural virtues should be, but there is one thing they agree on that sets them apart from those who are more substantive oriented, where substantive-oriented theorists are more focussed on the external standards used to measure the quality of the policy output (Laurent, 2008, p. 2). For procedural democrats, the goodness or rightness of an outcome is not found in whether it tracks some independent truth of the matter, but is primarily derived from the nature of the procedure used to find the outcome (in the case of pure proceduralists, the quality of the procedure is all that matters for creating legitimacy) (Goodin, 2003, pp. 92–93). While there is a small emphasis on providing information and counter-argumentation, the major focus within the deliberative turn form of deliberative democracy described above is on how deliberation is done, namely in a highly democratic manner, rather than on the quality of what is being deliberated. This procedural focus is highlighted in relation to the writings of Dryzek after he notes that more discussion regarding the nature of deliberation was required given deliberation's central importance to democracy

(Dryzek, 2009, pp. 1379,1380). Regarding the procedural focus of deliberative democracy, Dryzek states:

Deliberative capacity may be defined as the extent to which a political system possesses structures to host deliberation that is authentic, inclusive, and consequential... deliberation must induce reflection noncoercively, connect claims to more general principles, and exhibit reciprocity¹⁴ (*Ibid.*, p. 1382).

Slightly further on, he also notes:

However, at the core of the idea of deliberation is a developed notion of democratic quality such that the greater the deliberative capacity of a system, the higher the quality of its democracy (*Ibid.*, p. 1390).

In other words, while providing a call to action to focus the discussion of deliberative democracy back on deliberation, Dryzek gives procedural principles as being the important aspects of deliberation, with no mention of epistemic reasons. Furthermore, as put by Landemore:

deliberative democrats have assumed that the value of democratic procedures in resolving disagreement lies essentially in the values they express (such as respect, equality, and reciprocity, among others), not explicitly or at all in their knowledge-aggregating, let alone truth-tracking, properties, or in their ability to produce “better” outcomes (Landemore, 2017, pp. 277–278).

Very little was said by early deliberative democrats regarding deliberative outcomes and their relation to any sort of procedural-independent standards; the epistemic properties of deliberation were essentially ignored at best, and actively avoided at worst (*Ibid.*, p. 280). This avoidance of epistemic properties was fuelled by what is known as the “fact of disagreement”, which states that citizens in free societies are

¹⁴ The principles of reciprocity states that the reasoning used as part of deliberation should be mutually justifiable such that the reasons used can be accepted by others who are similarly motivated to use reasons that can be accepted by others (Gutmann & Thompson, 1996, pp. 52–53).

committed to different and conflicting worldviews (Landemore, 2017, p. 277). Given this truism, scholars took for granted that the fact of disagreement entailed at least an agnosticism towards the truth value of moral and political claims, if not an epistemic abstinence, where this fact of disagreement functioned as a barrier to incorporating epistemic concerns into deliberative democratic theories (*Ibid.*, pp. 277,278). Citizens engaging in respectful deliberative debate were expected to not make appeals to the truth since it divisive and intolerant, or even exclusionary and sectarian (*Ibid.*, pp. 277,282).

This epistemic abstinence meant that while there was the opportunity for the citizens to exchange reasons, there was little emphasis on marking out any criteria to help make a policy decision with. In other words, a proceduralist form of deliberative democracy is epistemically indifferent with regard to the outcome and a coin toss between viewpoints would suffice while also treating citizens equally; to require otherwise would be to impose an external standard (Estlund, 1997, pp. 178–179). As part of a reason-giving deliberation, rather than simply a discussion, there needs to be external standards in mind such that the end goal of producing better policy can be aimed for; without external standards to measure policy proposals with, “better” has no meaning, and therefore there is no use for reason-giving deliberation (Richardson, 1997, p. 359) (Landemore, 2017, p. 284). As put by Benhabib (1994, p. 33):

Procedures can neither dictate outcomes, nor can they define the quality of the reasons advanced in argumentation nor can they control the quality of the reasoning and rules of logic and inference used by participants.

Procedural models of rationality are underdetermined.

In short, the incorporation of epistemic elements are essential for creating high-quality public policy.

§2.4: The Epistemic Turn

Given the issues just discussed regarding a procedural-heavy approach to deliberative democracy, it had become apparent that epistemic abstinence was unsustainable. This realisation brought about what has been labelled “the epistemic turn” by Landemore, which started around 2010 (Landemore, 2017, p. 278). Because of this epistemic turn, deliberative democracy is now able to not only potentially remove the problematic aspects of deliberative democracy that are caused by an overemphasis on procedural principles but can also allow scholars to test the epistemic capabilities of deliberative democracy, both in theory and in practice. Furthermore, it allows scholars to consider the implications that epistemic quality has on the legitimacy generated by deliberative democracy (*Ibid.*, p. 286). For the purpose of this thesis, the important questions to be investigated are:

- Where does deliberative democracy stand in terms of its epistemic capabilities?
- What effects does an increased epistemic emphasis have on the legitimacy generated by a governance system?

The first question is the topic of Chapters 3 and 4 and involves an analysis of the claim of epistemic democracy (introduced shortly), while the second question is considered in the form of potential future research in Chapter 6.

§2.4.1: Introducing Epistemic Democracy

By way of introducing epistemic democracy so that it can be analysed, it is important to place it in context with respect to deliberative governance systems. As noted earlier, there is variation in terms of the principles that are found in different forms of deliberative democracy. One way in which this variation can be thought of is along a democracy–expertise spectrum, as presented in Figure 1 below:

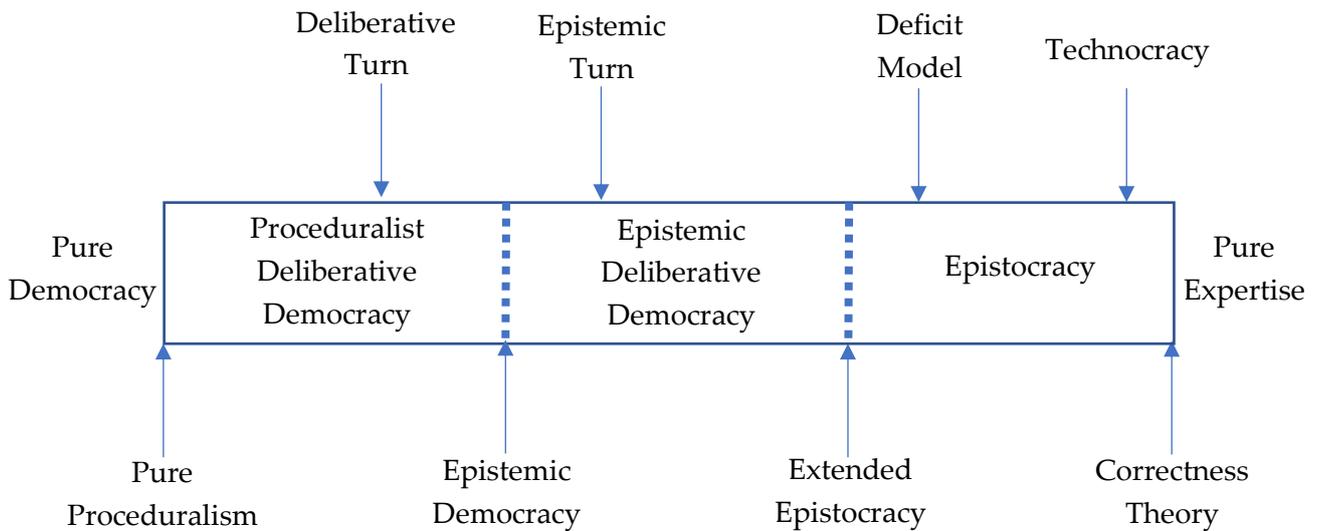


Figure 1: The Democracy–Expertise Spectrum

There are three important features of this spectrum that must be noted before placing epistemic democracy in context. First, the spectrum displays the relative qualitative levels of emphasis on expertise and democratic legitimacy; it is not intended to be to scale nor imply the complexities of governance can be represented along a single axis. Second, the top row of arrows represents the four stages of what could be considered a governance pendulum that has been discussed so far in this chapter regarding contemporary deliberative democracy; it starts with technocracy, sliding to the deficit model, making a large swing due to the deliberative turn, and returning as part of the epistemic turn deliberative democracy is currently undergoing. Lastly, the regions are separated by a dotted line to indicate that there is no sharp distinction between the three regions such that there can be overlapping mixtures between neighbouring regions.

Returning to placing epistemic democracy in context, it is important to recognise the positions that different forms of governance are found at on this spectrum depend on whether they consider the problem of expertise to be a significant issue, and what side of the problem of expertise they favour, i.e., democracy or expertise, as part of the

trade-offs made in solving it. For example, the range of principles and aims used in governance include:

- Whether deliberation in itself has value *or* is only instrumental as a means to achieving better results;
- Whether a deliberative theory should only define the procedural methods *or* it should also ensure that laws and policies created are imbued with certain epistemic standards;
- Whether deliberation should be participatory *or* representative (Gutmann & Thompson, 2004a, pp. 21,23,30).

Those forms of governance that favour the democratic side of the problem of expertise will tend towards the former option, while those that favour the epistemic side will tend towards the latter option. As discussed earlier, pure proceduralism exclusively values democratic aspects and does not allow for authority relations that arise due to levels of relevant expertise (Goodin, 2003, pp. 91–92). As a result, pure proceduralism is found on the very left side of this spectrum. At the opposite end of the spectrum is correctness theory, which is the claim that the only legitimacy a policy requires is the legitimacy that comes from the quality of the decision, and no legitimacy need come from the democratic nature of the process (Estlund, 2008, p. 99). In the central part of the spectrum, two mixture theories are presented. The first of these is found in the region where epistemic deliberative democracy and epistocracy overlap, which is the extended epistocracy proposed in Chapter 5. Extended epistocracy is an expert-focussed form of deliberative democracy that utilises representative experts as a way of incorporating democratic representation, and hence could be considered a combination of deliberative democracy and epistocracy that favours expertise over democracy with regard to the problem of expertise.

The second mixture theory is the epistemic democracy in question, which is loosely found in the overlap between procedural and epistemic forms of deliberative democracy. The word “loosely” is used because epistemic democracy is not so much

a form of governance, but rather a claim about democracy that says that democracy has truth-tracking capabilities because the democratic process is epistemically valuable such that it has a tendency to produce high-quality outcomes (Landemore, 2012a, pp. 44–45). As such, trade-offs do not need to be made between policy quality and democracy legitimacy, and hence epistemic democracy does not consider the problem of expertise to be a problem. Therefore, as part of the analysis of epistemic democracy in Chapters 3 and 4, it will also be determined whether the problem of expertise is a genuine problem; if the claims of epistemic democracy are true, the problem of expertise does not exist, and deliberative democracy can maintain its strong participatory emphasis. However, if the claim of epistemic democracy is false, then a new position on the democracy–expertise spectrum will need to be found that minimises the tension found in the problem of expertise.

§2.4.2: Challenges Facing the Epistemic Turn

This thesis agrees that the epistemic turn is a welcome development for deliberative democracy. However, the epistemic turn faces challenges if it is to satisfactorily include epistemic elements while also upholding the principles of deliberative democracy. The greatest challenge facing an epistemic form of deliberative democracy can be expressed in what is called the “democratic trilemma”. The democratic trilemma states that:

Under normal conditions, the three principles internal to the design of democratic institutions—political equality, mass participation, and deliberation—pose a trilemma. Serious efforts to realize any two will reliably run into clear roadblocks with respect to the third (Fishkin, 2011, p. 248).

where:

Political equality means that the people’s views are counted equally.
Deliberation means that the public has arrived at considered judgments

after it has weighted competing arguments. Participation is a kind of token of actual consent (*Ibid.*, p. 249).

Under ideal conditions, the trilemma is only present for logistical reasons since it is impossible to have a large group of people genuinely discussing an issue at once; this is known as the problem of scale (Parkinson, 2003, p. 181). However, if one is willing to accept a proportional representation, perhaps drawn by lot or by an enforced demographic representation, as a legitimate replacement for *mass* participation, then ideal deliberative democracy as described in detail above can provide fair, well-reasoned deliberation between free and equal participants from all stakeholder groups directly affected by the policy decision.

Under modern political conditions, though, recall that one of the demands of deliberative democracy is to maintain citizens' equal right and opportunity to engage meaningfully in deliberation and decision-making such that there are no power differentials that can influence the decision. This thesis agrees with deliberative democracy that there are many aspects in which citizens should be considered equals before the deliberation process, e.g., citizens should not be denied access because of their gender, ethnicity, sexual orientation, etc. However, there is one area particularly pertinent to epistemic forms of governance that could be a challenge for epistemic forms of deliberative democracy and that is differing levels of expertise among citizens. More specifically, under deliberative democracy:

- Citizens cannot be denied entry into, or respect within, deliberation for any reason, *including their level of expertise*;
- No-one is given greater authority for any reason, *including their level of expertise*.

In particular, for emerging science and technology policy, it is claimed that "to obtain the full mooted benefits of public engagement, it is necessary for scientists and the public to meet *directly* with one another, in an open and equal exchange of views" (Sturgis, 2014, p. 39)(emphasis in original). The emphasis on the level of expertise here is important because, from a specifically epistemic standpoint under modern political

conditions, epistemic inequalities play an important role in this trilemma, yet are not considered in the deliberative turn form of deliberation democracy described above (aside from rare exceptions, for example, (Knight & Johnson, 1997, pp. 289–290)). Epistemic inequalities are the differences in levels of expertise that arise due to the fact that natural abilities, time, background experiences, and education necessary to gain a significant level of expertise cannot be universal. Due to the division of labour, not everyone can become an expert in everything, especially since a great deal of expert knowledge is in the form of tacit knowledge, which is knowledge that cannot easily be codified and transmitted; instead, it must be experienced (Turner, 2003, pp. 102–103). Furthermore, some deliberators will be better orators than others, where they will not only be more imaginative and can come up with better potential solutions, but are much more capable of presenting them as publicly convincing (Gambetta, 1998, pp. 22–23) (Bohman, 1996, p. 127). Therefore, expert knowledge is a kind of possession that privileges its possessors with powers that, for a given policy area, most other citizens cannot successfully control, acquire, or share in (Turner, 2003, p. 19); this not only applies to technical, scientific expertise, but to moral expertise as well (Bell, 1999, p. 74).

What effect do epistemic inequalities have on the democratic trilemma? With epistemic inequalities in mind, mass participation is the weak link creating the trilemma from an epistemic standpoint. By including mass participation, epistemic inequalities are inevitable since not everyone will have the same level of expertise for a given policy area. Treating everyone's views equally as part of political equality will not naturally take place since there is a differential in power favouring those with higher deliberative capability; those with higher deliberative capability will have a better-reasoned grasp of their goals, how to achieve them, and how to defend them, while those of low deliberative capability will be open to confusion, arbitrary changes of opinion, and manipulation (Christiano, 1997, p. 257). That is to say, mass participation implies a lack of epistemic political equality. However, if epistemic

political equality is to be satisfied, then epistemic inequality implies that discussion must be done in the language of the non-expert in accordance with Turner's "paradigm plus authority" model. According to this model, there is a discursive asymmetry where experts have access to two paradigms, the expert and the common-sense paradigms (experts are also citizens in a world largely understood through common-sense reasoning), whereas the lay citizen only has access to the common-sense paradigm for that particular policy case and must rely on the cognitive authority of the expert with regard to matters found in the expert paradigm (Turner, 2003, p. 64). The expert must speak in the language of the lay citizen to be truly understood by lay citizens, but these are not the terms in which experts persuade one another, nor are they the terms that should necessarily be persuasive (Knight & Johnson, 1997, p. 139). A policy outcome, then, cannot be rationally justified through the force of the better argument to a lay citizen (*Ibid.*, p. 48). That is to say, the requirement that reasons given are well informed and in terms that others can accept and understand may not be able to be satisfied at the same time as part of a standard form of deliberative democracy. Giving reasons in such a way that other worldviews can accept and understand satisfies the need for political equality but leads to lower quality policy decisions because what is being discussed is only a small subset of what could be utilised if much larger and more technical comprehensive worldviews were allowed. Conversely, holding the discussion at the expert level satisfies the informed reasons criterion, but lay citizens are not being treated equally by not being able to effectively engage in deliberation since they have trouble understanding it. As a result, the deliberation requirement in the trilemma is not satisfied if mass participation is to be enforced.

In short, mass participation plus deliberation leads to a lack of epistemic political equality, and mass participation plus political equality leads to lower quality deliberation. To provide a democratic solution to the trilemma, epistemic democrats must argue that epistemic inequalities do not affect the epistemic quality of the policy created through deliberation. In other words, epistemic democrats need a way to

incorporate mass participation, or at least a satisfactory legitimacy equivalent, into their deliberative democracy governance system such that their claim of the truth-tracking capability of democracy can be upheld. The detailed analysis of this task is the focus of Chapters 3 and 4; Chapter 3 analyses the political competency of the general public as individuals, while Chapter 4 analyses the collective intelligence of the general public as group deliberators. However, before this analysis can begin, it is important to present ways in which epistemic democracy could utilise public engagement in practice under modern political conditions to bypass this trilemma. To do so, the next section discusses different forms of public engagement and their features.

§2.5: Potential Forms of Public Engagement to Overcome the Democratic Trilemma

§2.5.1: Nanotechnology as a Prime Opportunity for Public Engagement

Recall the criticisms, along with the poor results, discussed earlier that led many policymakers to the conclusion that science and technology governance needed to improve. This need for improvement had extra urgency given that nanotechnology development was gaining momentum, which helped lead to a turn towards governance approaches that were more participatory (Kurath & Gisler, 2009, p. 569). This urgency was amplified because there are many parallels between earlier technologies and nanotechnologies that have shaped the public's attitude towards nanotechnology, such as the level of initial hype and grand promises, as in nuclear technology, as well as concerns regarding what futures are desirable, who controls them, and what social and ethical issues are considered, if at all (Pidgeon & Rogers-Hayden, 2007, pp. 195–196) (Burri, 2009, pp. 505–506). Parallels have also been drawn by the general public between nanoparticles and nuclear waste due to their difficulty to clean up, between nanoparticles and asbestos due to their fibre-like characteristics, between nanotechnology and radiation from cell phones due to their invisible nature, between nanoparticles and dioxin due to potential harm to the ozone layer, and

between nanofood and genetically engineered food due to its “Frankenfood” nature (*Ibid.*, pp. 505–506).

Given these parallels, it is natural that nanotechnology governance was shaped from the beginning by the aforementioned lessons learned from previous emerging technologies, the realisation that the deficit model was not an appropriate way to create public acceptance of new technologies, and a desire for broader involvement of the public in the form of deliberative governance (Kurath & Gisler, 2009, p. 564). To counter these problems, a greater emphasis on state-funded initiatives aimed at encouraging wider public engagement and improving relations between science and society has arisen (Macnaghten, 2010, p. 23). This community engagement “master narrative” of public policy is intended to shift the development of nanotechnologies away from a source of distrust to one of responsible innovation (*Ibid.*)(Lyons & Whelan, 2010, p. 56). As such, nanotechnology governance has become part of a new deliberative governance movement, and unprecedented collaborations between those in nanotechnology and the social sciences have developed (Kurath & Gisler, 2009, p. 564). Over the space of two decades, there has been a dramatic shift towards an increased democratic involvement in emerging technology governance from the one-way deficit model, to an interest in the concerns of the lay public, to a genuine exchange of ideas and expertise (Matthew Kearnes, Nanotechnology, the Media and the Public conference, 2005, cited in (Kurath & Gisler, 2009, pp. 559–560)). Given this vigour surrounding the governance of nanotechnology, it is important to take a step back and consider whether the governance systems being introduced are appropriate, which is what this thesis seeks to do.

It should be noted that the above changes to the way in which nanotechnology is governed are experimental, and nanotechnology governance is said to provide a great opportunity for this, especially when upstream engagement is considered (Priest & Greehalgh, 2011, p. 1521). Upstream engagement is the use of public engagement

before particular technological trajectories and attitudes become set, allowing for key decisions to be influenced, and it is generally considered to be a way of providing more thoughtful and socially acceptable techno-scientific development (Delgado, Kjølberg, & Wickson, 2011, p. 834) (Pidgeon & Rogers-Hayden, 2007, p. 194). Given the issues that past technologies have raised, government authorities and other stakeholders are keen to avoid any repetition of them, particularly the issue of leaving public discussion too late (Burri, 2009, p. 498). Being a newer technology, nanotechnology poses a unique opportunity to put upstream engagement, and other science and technology studies recommendations discussed above, into practice (Delgado, Kjølberg, & Wickson, 2011, p. 836) (Paddock, 2010, p. 275). For some nanotechnologies (passive nanostructures), the opportunity to avoid dialogue being “too little, too late” is narrowing, but for later generations of nanotechnologies (active nanostructures, complex nanosystems, and molecular nanosystems), governance will need to evolve to reflect the increase in complexity of these technologies (Pidgeon & Rogers-Hayden, 2007, p. 195) (Roco, Harthorn, Guston, & Shapira, 2011, p. 3567).

§2.5.2: Different Forms of Public Engagement for Governance

What forms of public engagement are being used as a way of expressing the master narrative of nanotechnology governance? Examples of some of the many different forms of public engagement for nanotechnology policy include cultural festivals, card games, citizen juries, debates in schools, presentations in cafés, citizen schools, science shops, nanotrucks, consensus conferences, nanodialogues, forums, panels, formal inquiries, scenario workshops, deliberative polls, focus groups, staff workshops, and long-term engagements (Delgado, Kjølberg, & Wickson, 2011, p. 839) (Stebbing, 2009, p. 44) (Lyons & Whelan, 2010, p. 56) (Roco, Harthorn, Guston, & Shapira, 2011, p. 3563) (Fisher, Mahajan, & Mitcham, 2006, p. 488) (Priest & Greehalgh, 2011, pp. 1521–1522) (Powell & Kleinman, 2008, p. 329) (Powell & Colin, 2008, pp. 128,130–131) (Macnaghten & Guivant, 2011, p. 211) (Katz, Solomon, Mee, & Lovel, 2009, p. 537).

Of these forms of engagement, the consensus conference format is arguably the most famous form of public engagement and its paragon exemplar, having first been used in 1987, where it has been used in several countries since then. The consensus conference proceeds in several steps, beginning with the recruitment of around 12 to 15 participants, chosen to provide demographic representation¹⁵ but have no significant experience with the policy issue at hand, who are given a large amount of background reading about the conference topic. The participants then meet on three separate occasions:

- First, to discuss the background material with a facilitator.
- Second, to present their opinions and questions to a panel of experts, where the expert panel is sometimes chosen by the participants. This stage of the consensus conference is open to the general public.
- Third, the participants gather together to form a consensus regarding their viewpoints and to write a report summarising their findings and recommendations. This report is then distributed to the media, elected officials, and civil servants (Powell & Kleinman, 2008, pp. 331–332) (Cobb, 2011, p. 1536).

Highlighting the consensus conference is done for two main reasons. First, as the paragon exemplar, it provides the strongest form of public engagement for governance that should be kept in mind while considering the epistemic performance of potential participants, as discussed in Chapters 3 and 4. Second, regarding solving the mass participation horn of the democratic trilemma, it should be noted that consensus conferences (among some other forms of deliberation) choose their participants to be demographically representative. This is the same feature of deliberative polling that has been claimed to provide a proxy solution to the mass participation horn of the democratic trilemma; this is considered a proxy solution since, while mass participation is not feasible, demographic representation is said to

¹⁵ This means that participants are chosen such that the distribution of demographic features, e.g., gender, age, ethnicity, economic status, etc., is close to that of the surrounding area such that a range of viewpoints can be introduced into the discussion.

provide the second-best option that will provide as close as possible a representation of what the entire public would think if provided the conditions to deliberate (Fishkin, 2011, p. 251) (Fishkin, 2010, p. 74).

Beyond the consensus conference, other forms of public deliberation that have been used for discussing nanotechnology have permutations of the following attributes:

- **Participants:** Many forms of public engagement only choose citizens that are largely unfamiliar with the policy issue at hand to participate in deliberation, where participants are chosen to be representative of the population as a whole, or chosen to gain an understanding of the viewpoints of a specific section of society (Macnaghten & Guivant, 2011, pp. 211–212) (Macnaghten, 2010, pp. 26–27). Alternatively, a collection of only experts and dedicated stakeholders can be used (Kurath & Gisler, 2009, p. 565) (Katz, Solomon, Mee, & Lovel, 2009, p. 535). Sometimes there is a mixture of these two options, as in the case of a citizen’s panel (*Ibid.*, p. 536) (Pidgeon & Rogers-Hayden, 2007, pp. 197–198). The number of participants can range from 7–20 for a citizen’s jury, and up to 30–50 attending a forum (*Ibid.*, p. 197) (Kurath & Gisler, 2009, p. 565).
- **Information Given:** The amount of information given to participants, or expected from them, can also vary. In some forms of engagement, participants are given a brochure to read beforehand, and talks are usually given to participants during the engagement (Burri, 2009, p. 501)(Kurath & Gisler, 2009, p. 565). Even those forms of engagement involving only experts have talks such that everyone can begin to understand everyone else’s field of expertise (Katz, Solomon, Mee, & Lovel, 2009, p. 537). In some forms of engagement, participants are expected to do their own independent research and form their own discussion groups (Macnaghten, 2010, p. 26) (Powell & Colin, 2008, pp. 130–131).
- **Interaction:** The use of talks, as described above, tends to frame scientists as experts and the public as lay citizens. This framing means there is a hierarchical

relationship present akin to the deficit model; a lot of the time, participants were simply told information and then broke off into discussion groups (with a moderator) (Kurath & Gisler, 2009, pp. 565,567). However, there are forms of engagement that involve the participants questioning the experts and being treated more like equals (Pidgeon & Rogers-Hayden, 2007, pp. 198–199). There is also an example of long-term engagement where participants would work in conjunction with a group of experts over the course of more than three years, and have even garnered media coverage and influenced politicians (Powell & Colin, 2008, pp. 130–132).

- Aims: The aim of most public engagement seems to be directed towards being in-depth public opinion surveys, where participants are asked for their opinion on given nanotechnologies, and sometimes go further to discuss their hopes and fears regarding future nanotechnology (Kurath & Gisler, 2009, pp. 565–567) (Katz, Solomon, Mee, & Lovel, 2009, p. 536). However, for some forms of engagement, the aim (or at least the pretence) for engagement is to influence public policy, such as a citizen’s panel that was at least assured that their efforts would lead to greater social consideration of social issues in nanotechnology development; however, it is unclear whether this promise was followed up on (*Ibid.*). Similarly, citizen’s juries, are given a particular recommendation that they are to accept or reject (including giving suggested modifications where appropriate). While not legally binding, the views of the jury are given serious consideration (Pidgeon & Rogers-Hayden, 2007, pp. 197–198).

However, there is a great concern facing these forms of public engagement that an epistemic form of public engagement will need to address, which is that public engagement places a significant epistemic burden on the average citizen, especially when done in the form of a consensus conference. Given that citizens are chosen specifically due to their lack of knowledge in the policy area in question, the amount of information that they are expected to read and understand beforehand can be

overwhelming, even if they are given the opportunity to discuss with the facilitator in the first meeting. Noting that the information that policy experts have can take years, even decades to acquire, it seems unreasonable to believe the public involved in public engagement for governance has anything more than a surface understanding of the technical material and the issues involved. Granted, although giving (balanced and accurate) material to the participants provides a better background for policymaking than using general public understanding of the issue, it still does not provide the knowledge required to adequately predict cause and effect, to know what is (im)possible in terms of the scientific and policymaking capabilities, to identify many of the (potential) hazards, to form in-depth ethical/socio-political/economic arguments for and against their policies, etc. Even if the public engagement event was simply aimed at identifying the concerns of the public rather than creating policy, knowing what the legitimate concerns are still requires a significant amount of background knowledge. In other words, neither the type of participant chosen, nor the information given for public engagement are suitable for the aims the public engagement was created for (except in the case where only experts are invited for engagement); it would take too long to get average citizens up to an appropriate level of knowledge and make sure they understood it appropriately before it was suitable for them to be making critical policy decisions. This poses a significant challenge for an epistemic deliberative democracy, especially one that attempts to avoid the mass participation horn of the democratic trilemma by way of demographic representation. If it is possible that the average level of understanding of the public can rise such that a demographic sample chooses an informed public, or that an underinformed public can produce high-quality policy when they work together, then there remains the possibility that an epistemic democracy will be successful and the epistemic turn of deliberative democracy may continue. However, if not, then this opens the path for a more expert-focused governance system that does not need to satisfy the mass participation horn of the democratic trilemma, at least not in a manner traditionally ascribed to democracy, such as by way of demographic representation.

§2.6: Conclusion

The historical development of modern governance has gone from technocratic positivism, to the postpositivistic procedural forms of deliberative democracy after the deliberative turn, and has now found an epistemic middle ground after the epistemic turn. Each turn was caused by the recognition of flaws found in the dominant governance regime used at the time. In the case of positivism, it was the ignorance of values, the inability to handle uncertainty, and the elitist structure. In the case of procedural deliberative democracy, it was the lack of emphasis on providing epistemic checks on the quality of the policy output. However, now that research on deliberative democracy and public engagement for governance has begun to focus on the epistemic aspects of governance, it is important to consider how epistemic forms of deliberative democracy are going to achieve their goals, and whether these goals can be achieved at all. More specifically, two major challenges were presented to epistemic deliberative democracy: epistemic inequalities and their effect on the democratic trilemma, and the epistemic demands that deliberation places on participants. To overcome these challenges, epistemic democrats will need to show that, potentially through the various forms of public engagement presented, that the general public are capable of standing as epistemic equals with experts, individually or as a collective, such that the epistemic demand for the production of high-quality policy can be met. Whether this can be done is a discussion that takes place in Chapters 3 and 4, where the epistemic performance of citizens as individuals and as a collective, respectively, are analysed. By undertaking this analysis, it helps to establish what properties an epistemic deliberative democracy would need to have such that it was able to satisfy its goals, namely to satisfy the democratic trilemma as much as possible, while still producing high-quality policy, assuming it is possible to do so.

Chapter Three: Epistemic Democracy – Individual Citizen Competency

§3.1: Introduction

So far in the thesis, epistemic challenges posed by modern post-normal political conditions and the differing levels of expertise within deliberation have been presented. However, a claim has been made by epistemic democracy that these challenges can be overcome due to democracy's inherent ability to track the truth (Landemore, 2014, pp. 44–45) (Landemore, 2012a, pp. 8–9). If this claim is true, then the problem of expertise does not exist because there is no tension between the epistemic and democratic aspects of governance; no sacrifice of one aspect needs to be made to improve the other. Furthermore, if this claim of epistemic democracy were true, the methods of public engagement discussed in Chapter 2 are suitable for governing emerging science and technology policy. However, if this claim from epistemic democracy is false, then the tension between the epistemic and democratic aspects of governance does exist, and an alternative governance system will need to be proposed that is able to minimise this tension.

Therefore, Part B of this thesis investigates epistemic democracy, particularly epistemic deliberative democracy, to establish whether the claim of epistemic democracy is true. Chapter 3 investigates the political competence regarding emerging science and technology policy (henceforth, simply “political competence”) of *individual* voters, the causes of their competency, and the implications arising from those causes. Chapter 4 investigates the epistemic democracy in terms of the *collective* intelligence of the public, paying particular attention to the epistemic qualities of deliberative democracy, both in terms of actual emerging science and technology engagement exercises and in theory. Essentially, Chapters 3 and 4 investigate how much of the epistemic load of deliberation the average citizen can bear, and therefore,

how much the expert needs to take up to compensate. For example, if it were the case that the claim of epistemic democracy was true and that citizens have a high level of political competence, individually and/or collectively, then experts only need to become a guide to citizen decision-making on a close to epistemically level playing field. However, if the claim of epistemic democracy is false and the general public had low political competence, this would mean that experts should take a more authoritarian role in the decision-making process, at least from an epistemic perspective.

The scope for Chapters 3 and 4 is set out by considering Condorcet's jury theorem, which aims to show the epistemically beneficial nature of democratic voting (Berg, 1996, p. 229). Among the simplifying assumptions made by Condorcet, this chapter identifies three that stand out as being problematic for modern political conditions, particularly where emerging science and technology policy is concerned. These three assumptions are that: (1) citizens are competent, (2) citizens are unbiased, and (3) citizens vote independently. The first assumption is the primary focus of Chapter 3, the second assumption is considered at the end of Chapter 3 and beginning of Chapter 4, and the third assumption is the primary focus of Chapter 4.

Regarding the competence of individual citizens, this chapter investigates the evidence within the public understanding of nanotechnology literature, which strongly suggests that the average citizen does not have the nanotechnology understanding required to ensure that Condorcet's jury theorem supports the claim of epistemic democracy. However, much of this public understanding research is done in the form of snap polls without giving citizens the opportunity to do research and contemplate their answers. Therefore, to investigate the potential political competence of the general public, as well as the no prior bias assumption of Condorcet's Jury Theorem, this chapter draws upon literature on rational choice theory and heuristics to investigate whether true political competence will rise if given

the opportunity. This involves not only elaborating on the implications of rational choice theory, but also potential ways to increase the benefits or reduce the costs of political investigation, which include increasing the quality of media and public education and utilising what the average citizen already knows more effectively. If the average citizen can overcome this rational choice challenge, then the competence and no prior bias assumptions of Condorcet's jury theorem are satisfied, and epistemic democracy is supported. Conversely, if the rational choice challenge is not met, then the claim of epistemic democracy is less likely to be satisfied (but could still be saved by collective intelligence).

§3.2: Condorcet's Jury Theorem Under Modern Political Conditions

Condorcet's jury theorem, originally developed by Marquis de Condorcet in 1785 in his *Essay on the Application of Analysis to the Probability of Majority Decisions* (Condorcet, 1995 (1785)), has become a crucial feature of the epistemic defence of democracy (Ben-Yashar & Paroush, 2000, p. 197). In this section, Condorcet's jury theorem is introduced, not to critique the theorem itself since the conclusion logically follows from the assumptions made, but rather what can be drawn from it under modern political conditions to guide the discussion of citizen competence regarding decision-making. Although, as the name suggests, Condorcet's jury theorem was intended to be used regarding the performance of juries, many of the assumptions made by Condorcet can be extended to be applicable for modern political conditions, while other assumptions are still questionable. It is these questionable assumptions that guide the discussion for Chapters 3 and 4.

Condorcet's jury theorem shows that if a large, competent, and independent voting population is making a binary choice, then the majority will almost certainly select the

correct outcome. More specifically, the following assumptions are made to arrive at this result:

1. There are exactly two options, one of which is correct or better (and all voters share a preference for the correct/better option in light of full information).
2. The prior odds regarding which option is correct/better are even, i.e., before receiving any information signals, voters are equally likely to choose either option (or, put another way, they are not biased towards any particular option).
3. Voters have a homogenous level of competency, i.e., they are all equally likely to choose the correct/better option.
4. Voters' choices are independent of one another, i.e., the probability that a voter votes for the correct/better option is not linked to the probability that another voter does so.
5. The decision rule is a simple majority, i.e., the threshold proportion of votes required before an option can be enacted is 50% plus one vote. (Grofman, Owen, & Feld, 1983, p. 264) (Kanazawi, 1998, p. 69) (Austen-Smith & Banks, 1996, pp. 34,36)

By combining these assumptions, and if each voter is competent (i.e., each voter is more likely than not to vote for the correct/better option), then the following two conclusions can be drawn:

- (a) A group of voters will perform better, i.e., they are more likely to choose the correct/better option by majority rule, than any individual in the group is likely to do so alone.
- (b) The probability that the group chooses the correct option increases to certainty as the size of the group increases to infinity (Ben-Yashar & Paroush, 2000, p. 189).

The corollary of these conclusions is that, if the voters are incompetent (i.e., they are more likely to choose the incorrect/worse option), then a group of voters of homogenous competence will do worse than an individual voter, and the majority of a group of infinite size is guaranteed to vote incorrectly (Berg, 1996, p. 230).

Furthermore, the effect of these conclusions, whether positive or negative, is rapid. For example, if each voter were to have a 60% chance of selecting the correct/better option, it can be calculated that the chance of a mere 100 voters making the correct/better decision by way of simple majority rule is 98% (Ladha, 1995, p. 354). If this is extended to a common voting population of a few million voters, this voting population would reach near infallibility if each voter is competent. As a result, it is easy to see why Condorcet's Jury Theorem is used in support of democracy.

Condorcet's jury theorem can also be extended such that it is more applicable to modern political situations, rather than just to its initial application of a jury scenario.

More specifically:

- Assumption one can be extended to include any number of available voting options as long as the option voters are most likely to vote for is the correct/best one (List & Goodin, 2001, pp. 295–297).
- Assumption three can be extended to include varying levels of competence, as long as the average competence is just larger than $p = 0.5$ for a simple majority voting rule (Berend & Paroush, 1998, p. 482).
- Assumption five can be extended to other voting rules. This includes supermajorities, as long as the average voter competence is just larger than the proportion of the vote required to win, as the following decision rules: pairwise Condorcet¹⁶; and the Borda¹⁷, Hare¹⁸, and Coombs¹⁹ methods (Kanazawi, 1998, p. 72) (List & Goodin, 2001, p. 292).

¹⁶ Pairwise Condorcet: The candidate that beats or ties with every other candidate in pairwise elections is the winner.

¹⁷ Borda method: In an election with m candidates, each voter gives m points to their preferred candidate, $m-1$ to their next preferred candidate, and so on. The candidate with the most points is the winner.

¹⁸ Hare method: In each round, voters vote for their most preferred candidate and the candidate with the least votes is eliminated. This process continues until one candidate remains, who is then declared the winner.

¹⁹ Coombs method: This is similar to the Hare method, except in each round, voters vote for their least favourite candidate and the candidate with the most votes is eliminated. This process continues until one candidate remains, who is then declared the winner.

Condorcet's jury theorem, then, still applies to cases where there are independent, heterogeneous voters deciding amongst multiple candidates or policy options, using various decision methods. By applying these extensions, Condorcet's jury theorem becomes more realistic, and therefore, assumptions one, three, and five are applicable to modern political conditions. However, there are three aspects of Condorcet's jury theorem that are not so clearly satisfied under modern political conditions. These aspects are:

- The requirement that citizens need to be competent on average for collective intelligence be beneficial.
- The assumption that there is no prior bias regarding the options available (assumption two).
- The assumption that citizens are voting independently (assumption four).

These three assumptions will be considered in turn where Sections §3.3 and §3.4 will analyse the first aspect, Sections §3.4 and §4.2 will analyse the second aspect, and Sections §4.3–§4.6 will analyse the third aspect.

§3.3: Citizen Competence Regarding Nanotechnology

As noted above, Condorcet's jury theorem works well if citizens are competent on average. Conversely, Condorcet's jury theorem works if citizens are incompetent on average. For the purpose of this thesis, nanotechnology is chosen as the example emerging technology, and given nanotechnology's influence, which is likely to increase greatly, it is crucial that those taking part in deliberations are at least competent on average such that nanotechnology governance can create, or at least vote on, high-quality policy. This section investigates whether this is the case by asking: how well do the general public understand nanotechnology? To answer this question, this chapter draws upon social science research done on the public understanding of nanotechnology

It has often been stated that the general public have a very low understanding, or even awareness of nanotechnology (Kyle & Dodds, 2009, p. 82) (Powell & Kleinman, 2008, p. 332) (Burri, 2009, p. 499) (Cobb, 2011, p. 1534). This general lack of understanding of nanotechnology from the public has also been noted by nanoscientists that work in public engagement, where 35 of the 37 nanoscientists interviewed by Petersen et al. (2009, p. 517) considered the general public to be either uninformed or very uninformed. A survey done in 2004 showed that 83.6% of the United States general public had heard either “nothing” or “a little” about nanotechnology, and only 16.4% had heard “some” or “a lot” about nanotechnology. Furthermore, when these participants were asked three basic true/false questions regarding nanotechnology, only 3.1% of them knew the answer to all three of them, while a substantial 29.5% could not answer a single question correctly²⁰ (Cobb & Macoubrie, 2004, p. 397). In 2006, research found that amongst adults in the United States, there was a range of awareness regarding nanotechnology. The group that was most aware of nanotechnology were those aged 18–24, with 71% of them having heard of the term “nano”, and the same number having heard of “nanotechnology”. The least aware group were those aged 60+, with only 33% having heard of “nano”, and the same number having heard of “nanotechnology”. Of those aged 29 and older, 90% of them could not define the term “nanotechnology” (Waldron, Spencer, & Batt, 2006, pp. 572–573). In a quasi-experiment on citizen’s views regarding nanotechnology, also in 2006, 95% of the study’s participants had heard almost nothing or very little about nanotechnology (Macoubrie, 2006, p. 232). A more recent survey conducted in 2011 of the United States general public showed a slight increase in the self-reported familiarity with nanotechnology where 78.4% knew “nothing”, “very little”, or “a little” about nanotechnology, 16% had “some” knowledge, while 5.6% had “good” or “very good” knowledge of nanotechnology (Retzback, Marschall, Rahnke, Otto, & Meier, 2011, p. 6236). Another survey that same year showed similar results, with

²⁰ Note that a guess was marked down as an incorrect answer by the researchers (Cobb & Macoubrie, 2004, p. 397).

81.5% of the United States general public being “unfamiliar” with nanotechnology (“no knowledge” = 57.6% or “low knowledge” = 23.9%), while 18.5% were “familiar” with nanotechnology (“moderate knowledge” = 13.4% or “high knowledge” = 5.1%) (Vandermoere, Blanchemanche, Bieberstein, Marette, & Roosen, 2011, p. 200). Yet another survey has shown around 77% of Americans had heard “just a little” or “nothing at all” about nanotechnology (Liang, et al., 2015, p. 584). Furthermore, Batt, Waldron, and Broadwater have shown that the general public’s lack of understanding of nanotechnology is a symptom of their general lack of scientific knowledge, where the public have specific problems with understanding the size, scale, and symbolism of nanotechnology (Batt, Waldron, & Broadwater, 2008, p. 1141). This supports the earlier findings of Cobb and Macoubrie that show that the United States’ general public pay little attention to science in general, and to nanotechnology in particular (Cobb & Macoubrie, 2004, p. 403).

From the above results, it can be said that, aside from the quasi-experiment outlier, roughly 80% of the United States public are unfamiliar with nanotechnology, while only 20% are familiar enough to take part in a somewhat informed discussion. Furthermore, only about 5% of the public could be said to have any significant expertise in nanotechnology. It is unclear how this exactly translates into a precise average competency level for the general public that can be used in Condorcet’s Jury Theorem, but it certainly cannot be higher than the $p = 0.5$ that is required for aggregation to be beneficial. As Cobb notes, this low level of political competence makes for a perilous foundation to build public engagement on (Cobb, 2011, p. 1534).

One will notice that the above discussion of public understanding of nanotechnology has focused on the United States. This was done to be able to make a comparison of the changes in the public understanding of nanotechnology over time where the country of interest needs to be held constant, and since most surveys in this area were conducted in the United States, this makes the United States a logical choice for this

comparison. To this end, it can be seen that the level of understanding of nanotechnology among the public has remained essentially constant over the time period in question, which is in line with previous research done in the United States (Kleinman, Delborne, & Anderson, 2011, p. 230). This lack of change of understanding regarding nanotechnology is also present in the Netherlands, where intensive media attention over 2009/2010 resulted in very little change in the public's understanding of nanotechnology, and there was a general lack of change in understanding even up to 2014 (van Giesen, Fischer, & van Trijp, 2018, pp. 170,178). Furthermore, by way of comparing the United States to other regions, the 2005 Eurobarometer study covering the 25 members of the European Union at the time showed that only 8% of citizens showed an interest in nanotechnology (European Commission, 2005, p. 13). In France in 2011, 81.5% reported having little or no knowledge, 13.4% with moderate knowledge, and only 5.1% with high knowledge about nanotechnology (Vandermoere, Blanchemanche, Bieberstein, Marette, & Roosen, 2011, p. 199). In Iran, 22% of the public knew somewhat about nanotechnology, and only 5% had specific information about nanotechnology (Rahimpour, et al., 2012, p. 121). In Australia, only 8% of the Australian public know what nanotechnology is and how it works (Lyons & Whelan, 2010, p. 61). All of these results are consistent with those found in the United States.

§3.4: Will Competence Improve?

So far, this chapter has explained that a passing level of competence regarding nanotechnology from the average citizen is important if Condorcet's jury theorem is to work in favour of epistemic democracy regarding nanotechnology governance. This chapter has also shown that this nanotechnology competence is currently not high enough, which has been noted as being a "troubling finding", especially given the increased calls for public engagement and the public still willing to give their views despite their lack of knowledge (Cacciatore, Scheufele, & Corley, 2011, p. 387).

Even so, it could be the case that the above public understanding of nanotechnology research is an unfair representation of the (potential) competence level of average citizens. Given that policy decisions are made with the ability to take some time to consult relevant resources, rather than by way of an on-the-spot poll method that polls generally use, it could be the case that the potential competence level of citizens is high enough to satisfy the requirements of Condorcet's jury theorem. Furthermore, there are positive signs of willingness from nanoscientists to engage with the low knowledge public in an educational capacity (Petersen, Anderson, Allan, & Wilkinson, 2009, p. 517). Therefore, this section investigates the following question: if citizens were given the opportunity to find out more about nanotechnology in their day-to-day lives, is it likely that the competence levels of individuals would increase enough for Condorcet's jury theorem to work in epistemic democracy's favour regarding nanotechnology governance? To answer this question, two closely related concepts from rational choice theory are used to explain why it might be unlikely for most citizens to improve their current political competence significantly, namely rational ignorance and rational irrationality. However, methods have been proposed to be able to alleviate the effects of rational ignorance/irrationality, such as decreasing the cost of information and increasing the effectiveness of a small amount of information. These methods are analysed in detail to establish whether it is likely they can alleviate the effects of rational ignorance/irrationality.

§3.4.1: Rational Ignorance and Rational Irrationality

Anthony Downs gave the first application of rational choice theory to the behaviour of governments and voters²¹ (Downs, 1957, p. 14). During this application, he analysed the process of becoming informed and argued that "it is individually irrational to be well-informed", where being (economically) rational means choosing an available

²¹ As a reminder, it should be noted that rational ignorance/irrationality can apply to deliberation as well, rather than simply aggregative voting. This section speaks specifically of voting because that is how it was presented by Downs.

action that has the highest net benefit (*Ibid.*, pp. 215,246). The corollary to this is that it is (economically) rational to be ignorant, hence the term “rational ignorance”. Downs argued that making appropriate political decisions requires time, effort, and money, which the voter could spend on other aspects of their life, i.e., there is an opportunity cost to becoming well-informed on political matters. To make an informed decision, a voter must select appropriate sources and gather information from those sources, analyse that information, and evaluate how that information relates to their political goals (*Ibid.*, p. 210). These are the costs of becoming well-informed. The voter will also consider the benefits that come from making an informed decision, where the benefits of voting could include the expected increase in policy quality from that citizen’s vote, the chance that the information will change the voter’s mind to the better option, or the inherent happiness gained by the political process. If the costs of becoming well-informed are greater than the expected benefits, then an economically rational voter will not continue to search for information and either will not vote, or will not put much effort into voting and free-ride on the information-finding efforts of others (who are also likely to be free-riding)(*Ibid.*, pp. 244–246).²² The presence of rational ignorance is supported by the fact that political knowledge levels have been roughly stable at low levels for decades despite the increases in education level and availability of information. Newer information sources, such as cable television and the Internet, have diverted the attention of many citizens away from politics towards other types of entertainment from these sources (just as Plato was concerned about, albeit with different distractions) (Somin, 2013, p. 20). For a rationally ignorant voter, it is not

²² This is because a rational decision maker, in the economic sense of “rational”, is defined as one who will invest resources into this information-seeking process up until the marginal returns equals the marginal costs (assuming decreasing marginal returns and/or increasing marginal costs). At this point, the rational decision maker will make their decision (Downs, 1957, p. 215). That is to say, a voter will continue to invest in information up until the utility gained from the last bit of information equals the utility spent on obtaining it. If a piece of information is unlikely to change their mind (the voter perceives one candidate to be far superior to the rest), or the piece of information is likely to make little difference (candidates are proposing similar policies), then the voter is unlikely to search for new information (*Ibid.*, pp. 242–243). This narrows down the range of scenarios in which a voter will make a significant information-seeking effort.

education or information availability that is the main constraint, but rather the time taken to sort through and learn the information (*Ibid.*, p. 66).

Rational irrationality is a similar concept to rational ignorance. While rational ignorance is concerned with whether it is rational for the voter to search for information, rational irrationality is concerned with whether it is rational for the voter to reason well;²³ it may be the case that it is (economically) rational to be (epistemically) irrational (Caplan, 2007, p. 141). Rational irrationality is possible because citizens have preference over beliefs, where someone will prefer to hold on to a belief to maintain a “private happiness” rather than sincerely investigate the matter for fear of being proved wrong (*Ibid.*, pp. 80,206).²⁴ To put rational irrationality into rational choice theory terms, maintaining a political belief costs very little in most cases (assuming the citizen is part of a social group that accepts that ideology). As a result, any benefit gained from ideological loyalty (such as maintaining friendships or the satisfaction of finding confirming information) is almost certainly going to outweigh the costs of investigating alternative viewpoints. It is important to realise that the presence of rational irrationality means that the second assumption of Condorcet’s Jury theorem regarding no prior bias is incorrect. Due to rational irrationality, reasoners will at least initially tend to approach new information in a manner that is skewed towards the

²³ Or, to put the contrast between rational ignorance and rational irrationality another way:

Both [rational ignorance and rational irrationality] treat cognitive inadequacy as a choice, responsive to incentives. The difference is that rational ignorance assumes that people tire of the search for truth, while rational irrationality says that people actively avoid the truth (Caplan B. , 2007, p. 123).

This is not to say that the effects of rational ignorance and rational irrationality are independent of each other; one needs to be able to reason appropriately to decide whether or not they should accept a certain piece of information into their decision-making process (Somin, 2013, p. 63).

²⁴ Rather than assuming that citizens always process information to the best of their abilities, it is important to consider that citizens can hold beliefs for non-epistemic reasons, such as greed or conformism, or because the belief is more comforting, flattering, or exciting (*Ibid.*, pp. 2,115–116). These effects, when combined with rational ignorance, means that citizens are susceptible to misinformation and deception because they are more likely to accept information that conforms to their bias, and less capable of checking it due to a lack of prior knowledge (Somin, 2013, pp. 84,88).

easiest method for processing it, which is to confirm what they already believe, i.e., there is a prior bias in their decision-making towards what they already believe.

§3.4.2: Fixing/Bypassing Rational Ignorance and Rationality

From the above discussion, it is important to realise that political ignorance and irrationality, then, cannot simply be thought of as the result of stupidity or selfishness; they are largely the result of rational behaviour, and can even be displayed by those who are intelligent and genuinely concerned for the wellbeing of society. Because it is rational to behave this way, it makes it difficult to stop; to expect a citizen to lower their political ignorance or irrationality (i.e., increase their political competence) would be to ask them to engage in (economically) irrational behaviour, which is not enticing (Somin, 2013, pp. 3–4). It seems, then, that political ignorance and irrationality are a permanent feature, at least for the population at large. However, it may be possible for the public to reduce the impact of political ignorance and irrationality, and therefore, improve the epistemic quality of democratic decision-making (*Ibid.*, p. 4). In turn, this improvement in decision-making quality would make it more likely that democracy is able to track the truth, and hence satisfy the claim of epistemic democracy. To investigate whether the effects of rational ignorance and irrationality can be reduced, several methods are presented that either increase the incentives or reduce the costs within the rational choice model. In the interests of space, the focus here will be on rational ignorance since the methods and analysis used can also be applied to rational irrationality. The methods to be discussed come in three forms:

1. Decreasing the cost of information by having good quality “free” information.
2. Increasing the effectiveness of a small amount of information gathering.
3. Concentrating power in the hands of those who are already well-informed (Somin, 2013, p. 170).

The first two of these will be critically evaluated in turn in this section. The fourth option, which is the use of an epistocracy, will be developed in detail in Chapter 5.

§3.4.3: Improving “Free” Information

Improving Education

The first method to be discussed to potentially create a better-informed citizenry is to improve “free” information. In general, there are two main suggested ways this could happen: by improving education and improving the media. Both would work by allowing citizens to receive better quality information at low cost, and hence provide a net increase in the rational choice benefit. On the face of it, this idea seems plausible. Regarding education, those with higher education levels tend to have greater political knowledge, even when controlling for other factors. For nanotechnology specifically, education level has a correlation with the knowledge level of nanotechnology in the megacities of Iran (Rahimpour, et al., 2012, p. 119). Among those without any higher education, none of those surveyed had professional level knowledge about nanotechnology compared with 7% of those that had undergone higher education. Also, among those with a university degree, 22% did not know anything about nanotechnology compared with 58% of those without a university degree (*Ibid.*, p. 125). However, for the purposes of emerging technology governance, such as for nanotechnology, the effect of improved formal education is limited. This is because even if better quality political information were introduced in formal education, there is a significant time constraint involved regarding and presenting and learning it. Just introducing enough information into the schooling system to give citizens a suitable general basis for making political decisions regarding nanotechnology is a difficult task in itself, assuming such courses were to be taken by a significant portion of the population. Furthermore, even if the schooling system could accommodate all of this knowledge, it is difficult for citizens to keep up-to-date with everchanging political issues after they have finished their schooling (Somin, 2013, pp. 176–177); citizens only spend a certain number of years in formal education, and as such, too keep up-to-date with emerging science and technology, others methods need to be used, such as media sources.

Improving Nanotechnology Media

Overall, it has been shown that the general public largely pay little attention to nanotechnology in the media (van Giesen, Fischer, & van Trijp, 2018, p. 170). This thesis proposes that there are two main reasons for this: the availability of information and the lack of intention when searching for information. Regarding the availability of information, in the early years of nanotechnology, such as around the year 2000, there was very little coverage of nanotechnology (Brossard, Scheufele, Kim, & Lewenstein, 2009, p. 548). While there have been some spikes in nanotechnology coverage, such as after controversies, the level of coverage has remained low since then, but is slowly increasing (Kjaergaard, 2010, pp. 86–87) (Liang, et al., 2015, p. 585). Regarding the lack of intention when searching for information, it is important to note that the forms of media that have the most effect on the representation of nanotechnology are newspapers and television, with science news on television having the greatest effect, newspapers having a mixed effect, and online information not achieving significance (Brossard, Scheufele, Kim, & Lewenstein, 2009, pp. 554,555) (Cacciatore, Scheufele, & Corley, 2011, pp. 389,393,398). This is particularly interesting due to the abundance of nanotechnology information online (*Ibid.*, p. 399), but accessing this information is more intentional than passively receiving information from television or browsing a newspaper, supporting the claims of rational ignorance (Liang, et al., 2015, p. 585).

As such, when the public does pay attention to nanotechnology in the media, television and newspapers have the potential to play a key role in shaping public perception of nanotechnology, since for many, this is how they form their perception of the reality of science, and as a result, journalists and others in positions of power exercise significant control over the perceived reality that is created (Cacciatore, Scheufele, & Corley, 2011, p. 388) (Kjaergaard, 2010, pp. 80,92) (Habermas, 2006, p. 419). As such, this type of media coverage can act to provide a simple, albeit controlled, information to help make nanotechnology policy decisions (Brossard,

Scheufele, Kim, & Lewenstein, 2009, p. 548). In general, the information provided by media portrays nanotechnology positively, but media in Europe and more likely to recognise the risks of nanotechnology than the media in the United States (Gaskell, Ten Eyck, Jackson, & Veltri, 2005, p. 86); even so, two thirds of nanotechnology newspaper articles from 2000 to 2007 in Norway were either unequivocally positive or mostly positive towards nanotechnology (Kjølberg, 2009, p. 66). In more recent years in the U.S., the dominant positive framing has remained (Liang, et al., 2015, p. 585). Furthermore, the discussion of nanotechnology in early years tended to be about nanotechnology's use of taxpayer funds and nanotechnology education, but shifted in later years towards an oversensationalised revolutionary and deterministic science-fiction vision (Kjaergaard, 2010, pp. 86–87) (Kjølberg, 2009, pp. 67–68) (Petersen, Anderson, Allan, & Wilkinson, 2009, p. 521).

While part of the reason why this oversensationalisation occurs is to bring in media views and draw attention to the issue, it is also caused by some journalists not having the scientific background to be able to critically scientific sources adequately, especially those that come from rogue scientists that may wish to cause controversy (*Ibid.*, pp. 521,524) (Kjølberg, 2009, pp. 68,70). While this is not true of all journalists, it is important to note that there is a great deal of variability in the quality of the journalists, ranging from those that are very prepared and have science PhDs to those with “astonishingly poor” scientific knowledge that give “unpredictable” (relative to what was discussed in the interview), “inaccurate”, and “fictional” portrayals of nanotechnology that are detrimental to science (Petersen, Anderson, Allan, & Wilkinson, 2009, pp. 516–520). Not only did scientists prefer the better prepared and qualified journalists, who usually worked for broadsheet newspapers (as opposed to tabloid newspapers), but they also preferred, and even encouraged, journalists that could provide a realistic discussion of the potential societal effects of nanotechnology, both positive and negative (*Ibid.*, pp. 523,525). Furthermore, it is possibly due to the variability in the quality of nanotechnology media that there have been mixed results

regarding the link between nanotechnology literacy and science media use; Vandemoere et al. (2011, p. 197) claim that there is no link, while Yi-Fan Su et al. (2014, p. 363) claim there is a positive link, i.e., the more attention citizens paid to science content in newspapers and television, the higher their nanotechnology literacy tended to be.

In short, as it currently stands, nanotechnology receives little coverage in the media, and when it does, there is a great variance in its quality but tends to be unbalanced, oversensationalised, and inaccurate. Therefore, if the public understanding of nanotechnology is to improve, the public need to overcome rational ignorance and intentionally search out and identify the high-quality sources of information and consume this media regularly. While this is possible, it will require and concerted effort on the part of both media producers and the general public, which is difficult enough to achieve on its own without also considering the other policy areas citizens may also be interested in that will also take up their time.

§3.4.4: Shortcuts

Introduction to Shortcuts

The second way to reduce the impact of political ignorance to be discussed is by using epistemic shortcuts, which are a way for citizens to increase the efficiency of their information searching, i.e., their benefits per unit cost increases. For the purpose of this thesis regarding nanotechnology deliberative democracy, the general public will need to develop their own viewpoints that can be defended at the deliberating table, rather than simply trusting in the opinions of experts (see, for example, (Downs, 1957, pp. 230–231)), and so shortcuts will be the focus of this section. Individuals use shortcuts to reach quick judgements based on information that is most easily available to them, especially when the issue is not directly related to their everyday lives (Brossard, Scheufele, Kim, & Lewenstein, 2009, p. 550). This is because, as indicated by the discussion on rational ignorance/irrationality, people are generally cognitive

misers such that they minimise the effort used in making decisions (Liang, et al., 2015, p. 585). Over the long term, citizens are said to form reasonable attitudes towards policy issues through what they take in from casually received, low-cost information sources, such as watching the news or casual discussions with friends. These interactions build up an ideological framework that can be matched to political parties and/or opinion leaders that share a similar worldview (Caplan, 2007, p. 153)(Somin, 2013, p. 90). This method drastically reduces the amount of resources that the citizen needs to spend on specifically finding and processing information because it has already been done by having it casually assimilated into, or rejected from, the citizen's worldview.

When it comes to shortcuts that can be used for nanotechnology policy, public attitudes towards nanotechnology comes from limited knowledge and media attention, as noted above, as well as non-cognitive attitudes and the associated quick calculations regarding potential consequences that go beyond using only scientific knowledge (van Giesen, Fischer, & van Trijp, 2018, p. 170) (Brossard, Scheufele, Kim, & Lewenstein, 2009, p. 549). It has been found that hopes and expectations have a stronger effect than knowledge when it comes to nanotechnology and this is particularly true for those who know little about nanotechnology (Dijkstra & Critchley, 2016, p. 73) (van Giesen, Fischer, & van Trijp, 2018, p. 171). These attitudes act as perceptual filters that are used as conceptual frameworks to organise and interpret the little information they have and reach judgements regarding controversial technologies (Liang, et al., 2015, p. 586). The mental associations that the public have regarding nanotechnology are comprised of a complicated set of factors, and these mental associations are potentially the strongest influence on the public's opinions regarding nanotechnology and its potential dissemination, where potential beneficiaries missing out due to misconceptions regarding nanotechnology (Cacciatore, Scheufele, & Corley, 2011, pp. 387–388,399) (Hosseini & Rezaei, 2011, p. 514). Aside from the general perceptions of nanotechnology that are portrayed

through the media, there are two other general forms of shortcuts to be considered regarding nanotechnology, namely analogies and predispositions, which are discussed in turn

Forms of Shortcuts Used Regarding Nanotechnology

As discussed above, the media is an important way in which nanotechnologies are framed and presented to the general public. Based on the discussion above, the average citizen will tend to see nanotechnologies presented in the media in a generally positive manner, though with a sense that an unavoidable science fiction future awaits them. Furthermore, the information presented was either inaccurate and sensationalised or reasonably accurate and balanced depending on where one looked. However, the media cannot be the entire source of the public's shortcuts since the public's attitude to previous technology can cause perceptions of nanotechnology that are in contrast with the generally positive outlook presented by the media²⁵ (Vandermoere, Blanchemanche, Bieberstein, Marette, & Roosen, 2011, p. 197).

Regarding analogies, as noted earlier in §2.2.3, previous technologies, such as genetic modification, biotechnology in general, and nuclear power, have met with resistance and rejection from the public at large, and nanotechnology faces similar issues (van Giesen, Fischer, & van Trijp, 2018, p. 169). The effect of this analogy with previous technology as a type of shortcut is particularly strong since citizens are naturally more familiar with previous technology, and have little knowledge of nanotechnology, so the views on previous technology remain prominent and quite stable (*Ibid.*, p. 171). Furthermore, a downside to using analogies is that if they become central in shaping

²⁵ However, this effect does depend on the region the public live in. It was shown that those living in the United States require a lot less optimism about previous technologies before they consider nanotechnology in a positive light than was the case in Europe (considering two previous technologies as being positive was enough in the United States, whereas seven positive previous technologies were required in Europe to reach the same level of optimism found in the US) (Gaskell, Ten Eyck, Jackson, & Veltri, 2005, p. 84).

discourse and are used unreflectively, they carry with them implicit assumptions and expectations that may not be appropriate with regard to deliberative quality (Schwarz-Plaschg, 2018, p. 156). Also, some have argued that the ubiquity of the analogy between nanotechnology and a given past technology is unfortunate since the analogy is not particularly strong or helpful due to nanotechnology being a term for a range of technologies, while the past technology in question is a specific form of technology (Kearnes, Macnaghten, & Wilsdon, 2006, p. 15). In contrast, others have argued that it is the broad nature as an enabling technology that makes the analogy with previous technologies all the more reasonable; as an enabling technology, nanotechnology affects other types of technology by making them faster, stronger, lighter, etc. and so it makes sense to conceive of nanotechnology in terms of other technologies (Cacciatore, Scheufele, & Corley, 2011, p. 388) (van Giesen, Fischer, & van Trijp, 2018, p. 168). The important aspect to note here regarding analogies is that the reasoning a citizen uses based off of an analogy is only as good as the analogy itself. While it may be epistemically easier for the citizen to transfer the opinions of past emerging technology onto nanotechnology through the broad analogy of “emerging technology”, a lot of information becomes lost and the risk of the transferred opinion being misguided is high. To counter this, the more appropriate and specific the analogy one wishes to use, the more knowledge the citizen requires about past technology as well as current nanotechnology. It may be possible for citizens to be able to make broad statements regarding not wanting to repeat the health risks of asbestos or the miscommunication with the public regarding genetically modified food, but to use analogies between those past technologies and a given nanotechnology respectively requires:

- knowing information regarding both the science of the risks and the logistics behind regulation for both asbestos and the given nanotechnology in question to find a solution.

- understanding both the mechanisms for communication and the surrounding context for both genetically modified food and the chosen nanotechnology to find a solution.

However, to require such detailed information is to defeat the purpose of using shortcuts in the first place since shortcuts are meant to be an epistemically easy replacement for detailed information searching. Using shortcuts may be more appropriate if there was, for example, a referendum regarding the use of a particular nanotechnology where the discrete nature of the vote and the already determined option is more accommodating to loose analogies. However, in deliberative democracy, policy positions must be developed and defended with appropriate reasoning by those taking part, where unjustified analogies should not be convincing. That is to say, in deliberative democracy, analogies can be used but citizens cannot rely on analogies as a shortcut.

Beyond analogies with previous technologies, predispositions can play an important role when it comes to shortcuts regarding nanotechnology. A major predisposition is religious affiliation; positively, this may be achieved by recognising that technology can be used to spread a religious message, and negatively, nanotechnology may be seen as “playing God” and against the natural order (Brossard, Scheufele, Kim, & Lewenstein, 2009, p. 547). Accordingly, the results regarding the effects of religiosity on perception of nanotechnology are mixed. In one study, it was found that religiosity negatively affected perceptions of usefulness, but it had no effect on support for nanotechnology (Cacciatore, Scheufele, & Corley, 2011, pp. 388,393). In contrast, another study found that strength of religiosity had a significant negative relationship with support for nanotechnology. Furthermore, the views towards nanotechnology of those with high religiosity did not change with level of knowledge about nanotechnology, but an increase in knowledge increased support for nanotechnology for those with low religiosity (Brossard, Scheufele, Kim, & Lewenstein, 2009, p. 554).

Other predispositions that affect support for nanotechnology are political affiliation,²⁶ deference to experts,²⁷ the specific type of nanotechnology,²⁸ and a general approach to risk evaluation.²⁹

The use of such predispositions not only violate the no prior bias requirement of Condorcet's jury theorem, but it also faces the same problems as the use of broad analogies discussed earlier, namely a great deal of information is left unconsidered, the predispositions can be misguided, and it would take a great deal of knowledge and resources to confirm whether the predisposition was appropriate, which would defeat the purpose of using the predisposition as a shortcut; without such an epistemic investment, the predispositions or opinion leaders chosen using these predispositions are as likely to be as misleading as they are informative (Somin, 2013, p. 99). Another problem with the use of predispositions as a shortcut is that they might not be used for reasons that are related to epistemic value, e.g., they could be used due to an emotional attachment to a particular ideology, and hence the reasoning based on these predispositions are likely to become biased away from tracking the truth (*Ibid.*, pp. 90–91).

²⁶ Regarding political affiliation, liberals were more likely to be supportive of nanotechnology than conservatives were, but political affiliation had no effect on perceived usefulness (Cacciatore, Scheufele, & Corley, 2011, p. 393).

²⁷ Regarding deference to experts, higher deference is related to a greater support for technology in general and nanotechnology in particular, though this effect is stronger in some countries than others (in Singapore over the U.S., for example) (Liang, et al., 2015, pp. 594–595).

²⁸ Regarding the specific type of nanotechnology, it has been shown, for example, that the attitude towards nanotechnology used in food and its packaging is less positive than for other types of nanotechnology (van Giesen, Fischer, & van Trijp, 2018, p. 170) (Vandermoere, Blanchemanche, Bieberstein, Marette, & Roosen, 2011, p. 201).

²⁹ Regarding risk evaluation, in general, most people perceive the benefits of nanotechnology to outweigh the risks; however, this depends on the location where U.S. citizens are generally more positive than scientists, who in turn are more positive than European citizens (who are more likely to give “non-attitude” answers) (Kearnes, Macnaghten, & Wilsdon, 2006, p. 42) (Dijkstra & Critchley, 2016, p. 72). More specifically, risk perception had a negative effect on support for nanotechnology, while benefit perception had a stronger positive effect on support for nanotechnology (Cacciatore, Scheufele, & Corley, 2011, p. 396) (Brossard, Scheufele, Kim, & Lewenstein, 2009, p. 554).

§3.4.5: The Effect of Increasing Knowledge of Nanotechnology

Having a knowledge base in nanotechnology allows for informed decisions to be made. When citizens become better informed about nanotechnology, they are better able to recall evidence and are affected more by the knowledge rather than by the non-cognitive attitudes in the form of the predispositions just discussed. That is to say, when better informed, citizens make decisions using reasoning based on value structures and cognitive beliefs rather than on predispositions, and their beliefs become more consistent and connected in a larger technology cognitive framework (van Giesen, Fischer, & van Trijp, 2018, p. 171). However, because nanotechnology is more unfamiliar to citizens relative to technology in general, the initial balance between affect and cognition lies closer to affect for nanotechnology than for other technologies; therefore, a greater shift is required for nanotechnology, especially since the average citizens' general technology cognitive framework is not fully crystallized, and so it is difficult for nanotechnology to find a place within it (*Ibid.*, pp. 178,179). This lack of knowledge and established opinion means that the general public are overly sensitive to change, particularly negative breaking stories (*Ibid.*, p. 179).

Having citizens that were motivated enough to inform themselves regarding nanotechnology would potentially be able to overcome the force of affect, though this is unlikely given citizens are not known to search for information on nanotechnology themselves (van Giesen, Fischer, & van Trijp, 2018, pp. 179–180). It has been shown that increases in basic knowledge and familiarity regarding nanotechnology has only a small positive effect on changing support for nanotechnology, and does not have a significant effect on the perceptions of the usefulness of nanotechnology; to be able to truly influence the public's opinions of nanotechnology in such a way that it overcomes their affective responses, in depth engagement rather than simple risk communication gathered through casual interactions is required (Cacciatore, Scheufele, & Corley, 2011, pp. 393,396) (Vandermoere, Blanchemanche, Bieberstein, Murette, & Roosen, 2011).

In other words, not only are the general public likely to have received an inaccurate representation of nanotechnology through casual interactions with the media and their shortcuts are likely to be underdeveloped and dominated by affective responses, it will take a great deal of information to overcome these issues, which is unlikely to be searched for (or at least searched for appropriately) due to rational ignorance and rational irrationality. This is not to say that the general public cannot become appropriately informed to engage in nanotechnology governance, but rather that the cognitive and resource barriers in place are high and a concerted effort will be required to overcome them.

§3.5: Conclusion

The aim for Chapters 3 and 4 is to establish whether epistemic democracy's claim that democracy can track the truth to a suitable degree is true. While this question cannot be fully answered until the end of Chapter 4 after collective intelligence is considered, so far it looks like the claim of epistemic democracy regarding nanotechnology governance is false. This negative result is because, as an individual, the average citizen displays low understanding of nanotechnology. Furthermore, it is unlikely that this understanding is going to improve, even if methods are used such that rational choice theory favours epistemic democracy. However, democratic decision-making is done as a group, and this chapter only considered the competence of the average citizen as an individual. As such, leaving the discussion of epistemic democracy here would give an unfair representation of its capabilities. It may be the case that the diversity provided by a democratic public means that the collective competence of general public is greater than what an epistocracy can provide. Evaluating this claim is the topic for the next chapter.

Chapter Four: Epistemic Democracy – Collective Intelligence

§4.1: Chapter Overview

Chapter 3 demonstrated, via analysis of the literature on the public understanding of science and rational choice theory, that the political competence of the average voter as an individual is low and it is likely to remain low because it is (economically) rational to be politically ignorant/irrational. What does this mean for voting in general? One could say that if each voter is, on average, politically incompetent, then putting them together will produce an incompetent electorate that will vote poorly as a collective. However, one must be careful that this line of reasoning does not commit a fallacy of composition (Page, 2015, p. 375); just because members of a group have a certain property does not mean that the group has that property. It could be the case that the voting population at large displays a collective wisdom that allows democracy to be epistemically viable or even superior to all other forms of governance, where collective wisdom is an emergent property that depends not only on the intelligence of individual citizens, but also collective properties not found in the citizens themselves (Landemore, 2012a, p. 18). If it is the case that collective wisdom allows the group to be smarter than any individual within that group, or at least adequately competent, then political scientists can bypass the ignorance and irrationality of the average citizen and focus on the emergent epistemic properties of the democratic public (*Ibid.*, p. 2). Therefore, Chapter 4 continues the task that was initially set out in Chapter 3, namely to analyse the epistemic quality of democracy to determine whether the claim of epistemic democracy is true, except this chapter considers the collective rather than the individual.

This chapter is split into two parts, namely voting (Sections §4.2 and §4.3) and deliberation (Sections §4.4–§4.7). More specifically, Section §4.2 investigates the

miracle of aggregation, which has commonly been used as a defence of the epistemic benefits of democracy. It states that as long as some citizens are well-informed, and the rest effectively vote close to randomly (i.e., their errors cancel out enough such that these errors do not overcome the influence of the well-informed voters), then the citizens as a collective votes as if they were well-informed (Erikson, 2007, p. 26) (Caplan, 2009, pp. 198-199). Section §4.2 demonstrates that, while the miracle of aggregation has the potential to work in theory, evidence from past elections show that in practice there is too much of a bias towards voting incorrectly for the miracle of aggregation to be a reliable defence of epistemic democracy.

However, the miracle of aggregation is a simple representation of the effects of votes cancelling out. To provide a more realistic representation, it must be realised that the public do not vote randomly, but instead they have a diversity of viewpoints and vote accordingly. Therefore, what needs to be introduced is not randomness, but rather the correlations between votes. Section §4.3, then, draws on statistical formulations of voting behaviour to analyse the effects of voting correlation to establish whether they support the claim of epistemic democracy. It is demonstrated that, while epistemic democracy is supported under certain conditions, those conditions are not modern political conditions. Rather, an investigation of the effects of voting correlation in relation to voter competence and the size of the voting population suggest that a relatively small group of competent, diverse voters (i.e., an epistocracy) produces the voting group with the highest collective competence.

However, voting is only one part of the policy process. Where democracy really can show its epistemic qualities is when it comes to deliberation. This epistemic strength of deliberation is due to cognitive diversity, where cognitive diversity is the various ways the members of the group can think about the policy issue at hand (Page S. E., 2007, p. 7). Speaking as an epistemic democrat, Hélène Landemore (2012a, pp. 8-9) makes the probabilistic claim that because of the immense cognitive diversity of the

general public, democracy is the smartest method for making group decisions in the long run. This claim of Landemore's is split into two steps: first to show that deliberation is beneficial, and second, to show that these benefits remain as the number of deliberators increases (Landemore, 2012a, pp. 96-97). The first step is effectively granted for the sake of argument, though several reasons are briefly given to explain why it is most likely true in Section §4.4. The second step of Landemore's claim is critically analysed in Sections §4.5 and §4.6. During these sections, the Diversity Trumps Ability Theorem (a mathematical formulation of voting performance) is analysed in both its theoretical and computer-simulated forms, where the Diversity Trumps Ability Theorem states that n randomly selected people will tend to outperform n people with the highest individual political competence on a problem-solving task (Page S. E., 2007, p. 162). Throughout these two sections, it is argued that Landemore's claim is correct in the case of an ideal democracy. However, for modern political conditions, the Diversity Trumps Ability Theorem has not been shown to favour epistemic democracy. Furthermore, it is argued that to make the best use of cognitive diversity, the cognitive diversity must be carefully chosen, and that it seems best to choose this cognitive diversity from a collection of experts.

§4.2: Performance of Nanotechnology Public Engagement Exercises

As noted by Powell and Colin (2008, p. 129), there is an “unspoken premise that ‘average’ citizens *can* have meaningful impacts on scientific and technological developments—and that limited citizen engagement ‘exercises’ will help facilitate this—[that] seems to underlie many engagement projects”. The recurring question in this thesis is whether average citizens can, in fact, have a quality, meaningful impact on scientific and technological developments. Towards addressing this question, this section focusses on analysing the performance of public engagement regarding nanotechnology as an example of emerging technology governance. Given the consensus conference format has been presented in this thesis as the exemplar of

public engagement for governance to be used to represent epistemic deliberative democracy, this discussion will centre around a consensus conference known as the 2008 National Citizen's Technology Forum (NCTF). There has also been another prominent nanotechnology consensus conference, namely the Madison Citizens' Consensus Conference, but in the interests of space, the NCTF is chosen as the focus because it is the more rigorous conference, both in terms of the information given to participants and the interaction that participants had with each other and experts. Beyond the analysis of a consensus conference, another form of engagement, namely focus groups in the United Kingdom and Switzerland, will also be considered. These focus groups are analysed to understand how the perception of nanotechnology changed throughout the focus groups, and in particular, how a strongly negative attitude towards nanotechnology developed that was unable to be critically evaluated due to the lack of suitable expertise present.

§4.2.1: Nanotechnology Consensus Conferences

Participants in the NCTF were chosen to maximise the demographic representation as much as possible, with 74 participants³⁰ fully taking part over six deliberation sites. Of these participants, most participants had heard "nothing" or "just a little" about nanotechnology prior to deliberating, and those with self-interested agendas were excluded since they were presumed to be unwilling to engage appropriately in deliberation (Hamlett, Cobb, & Guston, 2008, p. 1) (Cobb, 2011, p. 1536) (Kleinman, Delborne, & Anderson, 2011, pp. 227–228). Selecting participants in this manner is based on what is called the "blank slate" approach where participants should have no vested interest in the particular policy such that they can be open-minded during deliberation. While this was technically true of the NCTF participants, the participants

³⁰ The number of participants chosen varies slightly between accounts. One account states that there were 74 participants (Hamlett, Cobb, & Guston, 2008, p. 1). Another account states that 89 participants were chosen, of which, four did not attend the first meeting and 12 did not complete the post-test, leaving 73 full participants (Cobb, 2011, p. 1537).

did come in with other biases such as being optimistic about the learning experience and about the prospects of nanotechnology (Kleinman, Delborne, & Anderson, 2011, p. 234). In other words, while there were no vested interests, there were still prior biases, violating the no prior bias assumption of Condorcet's jury theorem. Furthermore, it has been suggested that, given the recommendations for both consensus conferences were similar given the differences in prior bias, using a "blank slate" method is unnecessary, and even harmful since it seems to needlessly restrict the viewpoints available within deliberation (*Ibid.*, p. 235).

In terms of the information given, the NCTF participants were given a 61-page document covering technical information about nanotechnologies, as well as debates about their anticipated social impact. The document had been vetted by experts and written in language suitable for non-experts, as well as further information obtained upon request (Cobb, 2011, p. 1536) (Hamlett, Cobb, & Guston, 2008, p. 5). Regarding the process, the NCTF had three phases:

- The first phase involved the participants discussing the information they were given in groups with moderators and suggesting questions to ask a panel of experts.
- The second phase involved nine sessions over a month where the participants could discuss with experts including, among others, a chemist, a nanoscientist, a public policy expert, and an ethicist. This was done electronically due to the geographic range of the participants (Cobb, 2011, pp. 1536–1537) (Hamlett, Cobb, & Guston, 2008, pp. 5–6).
- The third phase involved returning to the discussion groups with moderators to formulate policy recommendations (*Ibid.*, pp. 334–335).

Overall, given that deliberation is taking place with participants that were unfamiliar with nanotechnology coming into the process, the NCTF can arguably be considered the strongest candidate as a demographically representative form of epistemic deliberative democracy.

However, despite the fact that the aims of a consensus conference is the publication of policy recommendations to be incorporated into real world applications, the recommendations given in the final report from the NCTF were not particularly in depth, nor did they solve any policy problems, and some already had solutions. The recommendations were essentially a list of concerns³¹ that others were expected to solve with no meaningful suggestion as to how this was to be achieved, even though the topic was restricted to a single topic (human enhancement), and hence there was the opportunity to be more specific. The recommendations include (Hamlett, Cobb, & Guston, 2008, pp. 7–8):

- Concerns over regulatory adequacy and safety standards. There is the recommendation that they should be strengthened and carefully monitored, but there seems to be no recommendation regarding how this is to be done (other than recommending possibly using a new agency).
- Similarly, with concerns regarding funding and equitable access priorities, there are no suggestions regarding the process or criteria that are to be used for making these decisions other than to treat disease before using enhancements, nor are there any explanations regarding who is going to pay for treatment or enhancements other than a stated concern regarding health insurance.
- Concerns regarding privacy and maintaining civil liberties already have well-established solutions within biomedical ethics.
- Concerns regarding the public being informed and involved in the decision-making process were raised, though no justification seems to have been given as to why citizens with no familiarity with nanotechnology should be involved or how this is to be done on a regular basis, and the suggestion to include ethicists in the process seems to negate the need for lay citizen involvement.

³¹ The same problem occurred with both the Madison Citizens' Consensus Conference and the Melbourne Citizen's Panel where the outcomes seem to only consist of a list of issues regarded as important without any solutions to those issues (Powell & Kleinman, 2008, pp. 335–336) (Katz, Solomon, Mee, & Lovel, 2009, p. 536).

- Finally, concerns regarding potential military uses and environmental damage appear to be simply vaguely stated concerns without any attempt as to how those concerns are to be addressed.

One minor counterexample was found in the Nanojury UK that not only identified issues, but also gave the beginnings of how these issues were to be resolved; however, they did not seem to go further than this (Rogers-Hayden & Pidgeon, 2006, p. 174). Even so, the important aspect to recognise regarding this lack of in-depth recommendations is that this appears to be a sign that the NCTF did not satisfy the requirements for deliberative democracy given that reason giving in a deliberate democracy should involve presenting and defending one's policy solution and receiving counterarguments in return; a list of concerns without presenting solutions does not satisfy these criteria. This is particularly notable since the NCTF is the strongest example of a nanotechnology consensus conference found, and consensus conferences are designed in such a way that they should, at least in principle, be able to satisfy the requirements for deliberative democracy by aiming towards policy recommendations; while the NCTF, as a *consensus* conference, may have satisfied its aim in terms of coming to a consensus regarding what concerns were to be considered important, deliberative democracy requires more than this. Regarding the prospects of epistemic deliberative democracy at this stage, one may say that, at best, consensus conferences can genuinely satisfy deliberative democracy in theory, but any future consensus conferences on nanotechnology will need to emphasise finding policy solutions to (potential) situations presented by nanotechnology rather simply highlighting concerns. As such, no comment can be made regarding the quality of an epistemic deliberative democracy for nanotechnology governance in practice since it is yet to occur. However, at worst, it could be the case that the short-term nature of a consensus conference and the types of participants chosen to attend, namely those that have no significant prior experience with nanotechnology, are the reasons why the consensus conference did not live up to its potential. As will be argued regarding the

negative slide towards nanotechnology risks presented in the next section, it is this latter worst-case scenario that appears to be the most likely.

§4.2.2: The Negative Slide Towards Nanotechnology Risks

In nanotechnology public engagement, participants correctly recognise that there are concerns with nanotechnology, as there are with any technology. However, there is a common theme regarding these concerns in that these concerns commonly appear to move from “nanotechnology *can* go wrong” to “nanotechnology *will* go wrong” (labelled the “negative slide”), with no attempt to propose solutions to stop this negative slide from occurring; sometimes it is even explicitly stated that the concerns suggested are too complicated to be governed effectively. Three different instances of focus groups where this occurred is discussed below to show that this effect is not an isolated incident: two of them were from the United Kingdom, and one was from Switzerland. At this point, one may argue that this represents an unfair sample since Europeans tend to be more pessimistic regarding nanotechnology than other regions, and hence, are more likely to take part in the negative slide (Kearnes, Macnaghten, & Wilsdon, 2006, p. 42). In response, while this is true, the purpose of this section is not to criticise the focus group participants for having a negative attitude towards nanotechnology since risks are associated with all technology and a critical approach to technology governance will always involve an investigation of the risks involved. Rather, the main point here is to identify that no attempts were made to identify potential solutions to alleviate these concerns. At this point, yet another potential argument against this selection of focus groups arises because it was not the purpose of the focus groups to produce policy solutions, but rather to discuss reactions regarding the policy issue in question (Kurath & Gisler, 2009, pp. 565–567). In response, while it may not have been the purpose of the focus groups to develop policy solutions, this does not justify what appears to be a consensus across the groups that nanotechnology is far too difficult to govern and that disaster is inevitable; there should at least be the recognition that measures can be put in place to alleviate harms

and take advantage of benefits, even if those participants themselves are not the ones to devise those measures.

United Kingdom (versus Brazil) Focus Groups

The first example to be discussed of the negative slide taking place was in focus group discussions in the United Kingdom³² between groups of 6–8 non-scientists without specialist knowledge of nanotechnology, though groups were organised around a common lifeworld experience, such as environmentalists, technophiles, business leaders, and religious followers, among others (Macnaghten & Guivant, 2011, pp. 211–212). Over the course of 3 hours, each group discussed technology in general; were given information regarding nanotechnology current and future uses, as well as funding levels; and then a concept board was used to prompt discussion about the social and ethical issues surrounding nanotechnology (*Ibid.*, p. 212). To see the power of the negative slide that took place in these focus groups, “without exception, all our UK groups discussions ended in tragedy” because it was concluded that nanotechnology was too profound and complex to be predicted and controlled (*Ibid.*, p. 214). Environmental, social, and moral disasters were thought to be inevitable because nanotechnology was thought to mess with nature, humanity, and the moral order of things to create a continual acceleration towards a science fiction dystopia (*Ibid.*, p. 215–216). From the account of these exchanges given, there is no mention of policy suggestions that could be used to avert this tragedy or ways in which potential benefits might outweigh the potential risks, which further emphasises how inevitable the U.K. participants considered this outcome to be.

³² There were also similar group discussions that took place in Brazil, though the opposite slide effectively took place from “nanotechnology *can* go right” to “nanotechnology *will* go right” with statements such as: “I believe that nanotechnology will help solve everything”, “I’m totally optimistic”, and “I believe it will be a positive thing!” (Macnaghten & Guivant, 2011, pp. 211–212). While this may be optimistic for nanoscientists to hear, it still presents the same problem as the negative slide in that there was no critical and detailed consideration of policy solutions that balance risks and benefits.

Swiss Publifocus

The second example to be discussed is a Swiss Publifocus between laypersons chosen based on their demographic variables that were given an informative brochure beforehand that was well-balanced and written for laypeople. Over the course of a half-day seminar, participants witnessed a balanced set of talks from experts about nanotechnology and then discussed their fears and hopes regarding nanotechnology in the presence of a moderator (Burri, 2009, pp. 500–501). Many of the participants heard about nanotechnology for the first time when reading the brochure and their instinctive reaction was one of fear, or at least rejection, particularly regarding breathing in nanoparticles or applying nanoparticles to the skin where the concern was that nanoparticles could cross the blood–brain barrier or cause other illnesses. Other concerns included silver nanoparticles killing animals and other nanoparticle build-up, as well as privacy and job loss concerns (*Ibid.*, p. 505). Some example statements from participants include:

- “At the first moment, when I read in the brochure about the suntan lotion [containing nanoparticles], I was scared.” (*Ibid.*, p. 503)
- “At home, I have some soya powder in an aluminium bag. I will immediately get that clarified, and will not use it anymore.” (*Ibid.*, p. 505)
- “Just like I don’t want any [genetically modified] food, I don’t want any nano food.” (*Ibid.*, p. 506)

Presenting such concerns is indeed beneficial for deliberation on nanotechnology governance and may even be considered insightful (it is unclear how many of these concerns were brought up by the participants themselves or whether they were mentioned in the brochure). However, the problem here regarding the potential quality of deliberation for governance is that the initial reaction was one of fear and there is no discussion in the report given that could be considered an analysis of the concerns suitable for governance deliberation. To the credit of the participants, they did consider some of the benefits of nanotechnologies to counterbalance some of the risks, but these benefits were for different technologies to those that were feared.

Furthermore, it was stated that participants appeared positive overall regarding nanotechnology and no one opted against nanotechnology entirely, but this appears to be more resignation than acceptance as there was a stated inevitability because the path of nanotechnology could not be stopped (*Ibid.*, p. 504).

DEMOS Focus Groups

The third example to be discussed are DEMOS focus groups, which involved three sessions taking place in the United Kingdom, but what is intriguing this time is that half of the participants in the third session were scientists rather than all of the participants being lay citizens. In the first session, lay citizens with no prior involvement or exposure to nanotechnology were placed in five groups based on a shared lifeworld (i.e., professional men, professional women, politically active, technophiles, and mothers). This first session began with a general discussion of emerging technology, and then nanotechnology and some of its applications were introduced, and further discussion ensued (nothing is mentioned regarding the balance or accuracy of this information) (Kearnes, Macnaghten, & Wilsdon, 2006, pp. 43–44). When nanotechnologies were introduced, participants went from a state of ignorance, to surprise and enthusiasm for the level of funding and possibilities, to unease and anxiety (*Ibid.*, p. 48). When presented with mainstream and optimistic scenarios, the mainstream scenario was treated with scepticism due to the influence of commercial interests, and the optimistic scenario regarding human enhancement was considered to be dangerous hubris leading to a science fiction future by most participants, using the controversies of genetically modified food as a reference point to call for a slowing down of development (*Ibid.*, pp. 50–51).

After a week of investigating nanotechnology for themselves, the second session involved discussion of how their viewpoints may have changed (*Ibid.*, p. 44). In this session, the uncertainty surrounding nanotechnology and its long-term effects caused alarm. In particular, the effects of nanoparticles were considered to be scary, with one

initially optimistic participant stating: “I wouldn’t touch it now with a barge pole if you paid me to put that stuff on my face. It’s so frightening.” (*Ibid.*, pp. 52–53). Again, in this session, the loss of control or having control in the wrong hands was seen as dangerous, whether it was the technology itself being uncontrollable, or terrorists or the socially privileged making decisions to the disadvantage of everyone else (*Ibid.*, pp. 53–54). There was the consideration of the governance of nanotechnologies, but it was concluded that it would be too difficult to do so effectively, especially due to the economic pressures the government faced to develop nanotechnology. Overall, participants felt nervous, apprehensive and unsure about future nanotechnology developments, which was a feeling that became stronger throughout the discussions (*Ibid.*, pp. 54–55).

In the third session, 12 of the participants from the first two sessions were chosen to meet with 12 nanoscientists to discuss the development of nanotechnology and its potential social impact in groups (*Ibid.*, p. 59). During this session, there was a lot of discussion regarding regulation, control, accountability, and governance, but it is important to note that no solutions regarding these areas were proposed, other than a request for more labelling and a call for more upstream engagement; instead, there was simply a general sense of helplessness (*Ibid.*, pp. 60–62). What is particularly notable is that the scientists felt they were just as helpless as the citizens, which reassured the public due to the sense of unity, but this still left open the question of who was responsible and to be held accountable during the developmental process and beyond as there were concerns that scientists, industry, and marketing departments were all in it for their own gain and that no one was present to take care of the social and ethical aspects of the larger picture (*Ibid.*, pp. 62–65). Overall, the scientists appeared impressed by the contributions by the public, but did note that these participants were already more informed and attentive than the average citizen and that this was not a process that could simply be scaled up to the population at large (*Ibid.*, p. 67).

What is interesting to note about this third session where the scientists agreed with the concerns of the citizens and the “sense of helplessness” in establishing ways of alleviating those concerns is the absence of social scientists, policymakers, or ethicists present in the discussion of this third session (and all other focus groups). While this general reluctance or inability to find potential solutions has not been actively investigated to this author’s knowledge, the suspected cause for the negative slide and the defeatist attitude towards governance is that neither the lay citizens, nor the nanoscientists, had the specific expertise required to develop policy solutions to the concerns raised. As such, this suggests that ethical, social, and political expertise are, in fact, relevant forms of expertise that need to be included in policymaking as opposed to relying on lay citizens in this regard. An example of an engagement exercise that helps to support this point is an Australian Office of Nanotechnology workshop on Social Inclusion and Engagement in Nanotechnology where members of the general public found the anti-nanotechnology groups to be too alarmist and wanted to hear a more balanced view in terms of both benefits and risks. However, participants at the workshop included those from government, community groups, and activists, among others, who were already well versed in considering policy solutions and would have contributed the possibility of policy solutions as part of the discussions (Cormick, 2009, p. 441). In other words, by adding the relevant ethical/social/political expertise, the negative slide did not appear.

§4.2.3: Implications for Nanotechnology Engagement

The aim of this section was to consider the question regarding whether citizens can have quality impacts on science and technology development. As will be noticed from the forms of engagement discussed that the exercises were very top-down in that these events were initiated, organised, and facilitated by experts (Powell & Colin, 2008, p. 129). Furthermore, there was a distinct epistemic gap where experts, present in person or through information brochures, were distinguished as such, with lay citizens being labelled as such. This outcome is inevitable given both the lack of resources and

knowledge of the average citizens, and so it is important to ask whether these kinds of engagements are capable of creating quality, impactful policy from lay citizens, even assuming political, socioeconomic, and cultural barriers were not present (*Ibid.*). The answer to this question appears to be: no.

Considering the engagement exercises discussed above, which are all short-term engagements, recall from Section §3.4.5 that it takes a significant amount of increased knowledge to change the reasoning of an engagement participant from affective to cognitive reasoning, especially for a technology they are unfamiliar with, such as nanotechnology. For focus groups in particular, even the DEMOS focus groups that occurred over three sessions, it appears that affective reasoning is still the dominant mode of reasoning. Furthermore, even after an intensive month of discussion, the NCTF still produced a list of only concerns, which may indicate that affective reasoning, as opposed to cognitive reasoning, was still in effect. However, even if cognitive reasoning had taken over, a month is still a very short timeframe to develop the relevant expertise from a complete unfamiliarity of nanotechnology; it would be unlikely that the citizens had both the technical and ethical/political knowledge to make quality decisions after only a month.

There have been many calls “across government, business, academe, and public interest and advocacy groups” for *informed* citizen input (Hamlett, Cobb, & Guston, 2008, p. 3). As such, it would seem best that any citizen input would come from those citizens that were already proficient in nanotechnology policy before they entered the policymaking discussion. To this end, focus groups may be useful in providing such public education to allow citizens to build up their proficiency, but it is recommended that this is done in the presence of those with ethical, social, and political expertise (as well as technical scientific expertise) such that citizens also learn the skills associated with these domains so that citizens are able to take part in “real-world” public engagement, which is long-term, seldom controlled, and often contentious;

furthermore, citizens would also need to have the power and resources to be able to engage in such situations (Powell & Colin, 2008, pp. 129,130).

An example of such long-term, real-world policy engagement was organised at the University of Wisconsin Madison's Nanoscale Science and Engineering Center. As part of this engagement, the citizens met weekly to co-design and co-organise engagement exercises and produce policy recommendations and media releases (*Ibid.*, p. 131). However, while the initially lay citizens did develop their capabilities regarding knowledge of nanotechnology and the policy process in ways that short-term engagement cannot, even after three years, this process still needed to be a top-down endeavour with citizens requiring "*intensive*"³³ investment from researchers in terms of time, information, and event organisation (*Ibid.*, p. 132). The citizens did not feel ready enough to become independent, and the researchers agreed they were not ready, which confirmed the scientists' suspicion that even an empowered general public are unlikely to be capable of independent political engagement in a meaningful manner due to their lack of resources, access, energy, and power (*Ibid.*). As such, given the amount of resources required for the empowered citizens to achieve a supported competence, and the fact that this level of resources could not be spent to create competence on a mass democratic scale, it would be more efficient and effective to spend these resources on those that already have a significant skill base, or take these empowered citizens in and give them professional policymaking roles. Either way, from these practical examples given that the lay citizen does not have the capability required to satisfy the requirements of an epistemic deliberative democracy.

§4.3: The Effects of Cognitive Diversity on Deliberation

However, it could be the case that the apparent failure of the above examples to meet the requirements of an adequately informed epistemic deliberative democracy is the

³³ Emphasis in the original.

result of poor execution rather than being an inherent property of a demographically representative epistemic deliberative democracy. As such, the rest of this chapter seeks to investigate the theoretical potential of epistemic deliberative democracy to determine whether the results discussed in Section §4.2 can be improved upon. It could be the case that by having voters interact, rather than just be aggregated, the claim of epistemic democracy may be true. After all, H el ene Landemore makes the claim that:

democracy is overall the smartest method for making group decisions, [but] does not exclude the possibility that some democratic decisions will be mistaken, nor does it exclude the possibility that a particular democracy will do worse than a particular oligarchy. ... On average, however, and in the long run, the claim is that democracy is a safer bet than a dictatorship or even an aristocracy (Landemore, 2012a, pp. 8–9).

The reason why Landemore makes such a bold claim is due to the “cognitive diversity” held within the general public. This cognitive diversity is an emergent property that causes *collective* wisdom, where cognitive diversity refers to the different ways in which a group can see, interpret, and predict the world around them. More specifically, cognitive diversity involves the:

- “diversity of perspectives (the way of representing situations and problems),
- diversity of interpretations (the way of categorizing or partitioning perspectives),
- diversity of heuristics (the way of generating solutions to problems), and
- diversity of predictive models (the way of inferring cause and effect).” (Page S. E., 2007, p. 7)

For example, cognitive diversity can arise due to differences in:

- training (an ecologist, an economist, and a sociologist are probably going to approach the same problem differently).

- experience (someone that has delivered healthcare to a poor African village will probably approach medicine differently to someone practicing in an expensive, private hospital).
- identity (a white, Christian man is probably going to approach problems differently than a Middle Eastern, Muslim woman) (*Ibid.*, pp. 302–309).

Higher cognitive diversity in the group is said to be useful because, under the right circumstances, their interaction can produce new solutions, and this effect can increase when there are more perspectives and heuristics to interact with each other (Page S. E., 2007, pp. 216,217,221–222). However, this thesis disagrees that cognitive diversity can be used as a means to justify democracy epistemically to the extent that it is in the literature. This chapter argues that even though cognitive diversity is the greatest epistemic strength of democracy, it does not have the ability to compensate for the low political competency of the average voter, particularly regarding emerging science and technology policy. To make this argument, this chapter considers what effect cognitive diversity has as the number of deliberators changes, along with the quality of the deliberators and the dependence they have on each other. If it is the case that high-quality policy can be achieved with a high number of diverse, average citizens, then epistemic deliberative democracy can be recommended as an epistemically adequate form of governance. However, if a small number of experts is required to create high-quality policy, then epistemic deliberative democracy will need shift to a more expert-oriented form of governance.

§4.4: Deliberator (In)dependence in Detail

§4.4.1: What Is Meant by (In)dependence

Recall that three assumptions from Condorcet's jury theorem were identified that were potentially problematic under modern political conditions. The first two, namely individual citizen competence and the assumption of no prior bias, have already been

considered. The third assumption to be considered, and the focus of this chapter, is Condorcet's assumption of voter independence. Therefore, to begin the investigation into the effects of cognitive diversity, a detailed statistical analysis of the effects of voter (in)dependence will be given. It should be noted that while the sources used originally focussed their investigation on the effects of voter (in)dependence, given that the probably p of each voter voting correctly is a dummy variable representing performance, the same mathematics also applies to the epistemic quality of deliberators, and hence to deliberative democracy.

What is meant by independence? Independence means that the probability that a voter votes for the correct/better option is not connected to the probability that another voter does so (Austen-Smith & Banks, 1996, p. 34). However, this is a statistical interpretation of what is meant by "independence". How is "independence" to be interpreted in practice? Berg claims that:

individual votes in an assembly cannot possibly be independent of each other if the voters are members of the same society; shared information, common beliefs, communication, and opinion leaders are obvious sources of dependence (Berg, 1993, p. 94).

Berg (1996, pp. 231,235) also adds party discipline and the reinforcement of opinion, i.e., peer pressure, to the list of ways in which votes can potentially be dependent, while Ladha (1992, p. 621) (1995, p. 354) adds common training and experience (such as culture or religion). In other words, there will always be some level of dependency between voters in practice. The forms of dependency suggested so far are all ways in which votes can be positively correlated, i.e., voters are more likely to vote for the same option. Votes can also be negatively correlated. For example, two jurors that interpret evidence in different ways, or two opposing voting factions within an assembly, are likely to vote in opposition to one another (*Ibid.*, p. 233). This will be elaborated in more detail regarding cognitive diversity later in this chapter.

§4.4.2: The Effects of (In)dependence

Given that the independence assumption in Condorcet's Jury theorem is untenable under regular political conditions, many scholars have tried to incorporate voting correlations into a Condorcet style analysis to understand the effects of voting correlations. Kaniovski (2010) has analysed the effects of correlation for a homogenous jury, i.e., one in which the probability of voting correctly is the same for each juror, and each pair of votes has the same coefficient of correlation (*Ibid.*, p. 454). He shows that the likelihood the jury will make the correct decision using majority rule:

- increases (decreases) for a competent jury with decreasing negative (increasing positive) correlation, and vice versa for an incompetent jury.
- increases (decreases) with increasing (decreasing) voter competence for positive correlation, and vice versa for negative correlation.
- increases with increasing voting population after a threshold size for positive correlation (*Ibid.*, p. 462).

Furthermore, Kaniovski derives the following formula regarding the effects of a change in size of the voting population:

$$M_n^B(p, c) - M_{n-2}^B(p, c) = C_{n-2}^{\frac{n-1}{2}} p^{\frac{n-1}{2}} q^{\frac{n-1}{2}} \left[(p - q) + c(n - 1)(0.5 - p) \left(n - \frac{n - 3}{4pq} \right) \right]$$

where $M_n^B(p, c) - M_{n-2}^B(p, c)$ gives the change in the probability the population will vote correctly when increasing the voting population by two voters, p is the probability of a juror voting correctly, $q = 1 - p$, c is the level of correlation (a highly negative c refers to a lot of cognitive diversity and a highly positive c refers to a lot of cognitive agreement), and:

$$C_{n-2}^{\frac{n-1}{2}} p^{\frac{n-1}{2}} q^{\frac{n-1}{2}} > 0$$

Based on this formula, this section draws the following conclusions regarding adding more voters by varying correlation (positive or negative), competency ($p > 0.5$ or $p < 0.5$), and population levels (high or low):

<i>Correlation</i>	<i>Competence</i>	<i>Population</i>	<i>Effect</i>	<i>Situation</i>
Positive	Competent	Low	Worsen	A few non-diverse experts
Positive	Competent	High	Improve	Following a good opinion leader
Positive	Incompetent	Low	Ambiguous	
Positive	Incompetent	High	Worsen	Following misleading broadcasts
Negative	Competent	Low	Improve	Epistocracy
Negative	Competent	High	Worsen	Competing expert factions
Negative	Incompetent	Low	Ambiguous	
Negative	Incompetent	High	Improve	Homogenous democracy

Table 1: Effects of Voter Correlation, Competence, and Population on Collective Performance

For clarification:

- “Competent” means that jurors are more likely that not to vote for the best of two options, and vice versa for “incompetent”.
- To determine whether the voting population is high or low, the critical population size, n_c is:³⁴

$$n_c = \frac{3}{1 - 4pq}$$

where less than n_c constitutes a “low” population, more than n_c is a “high” population.³⁵

- An “improve” effect means that that group is more likely to vote for the correct option after adding more voters, and vice versa for “worsen”.

³⁴ This critical population size is found by setting the $n - \frac{n-3}{4pq}$ term in Kaniovski’s equation equal to zero; any n higher than n_c will make the term positive, and hence help improve the likelihood that the group selects the better option, and vice versa for n lower than n_c .

³⁵ It is also worth noting that increasing the number of jurors produces a diminishing return. In other words, adding another juror for the “improve” scenarios creates a better jury (assuming the critical number of jurors has not been crossed), but the last juror added does not have as much of an effect as adding the juror before them, and similarly for the “worsen” scenarios. As team size increases, it adds extra resources and skills to the team up until an optimal size, at which point there are significant process losses and lower team integration causing dissatisfaction and distrust; these losses overcome the marginal benefits from increasing diversity at larger sizes. The optimal size is usually around three to five members, but this depends on many factors (Horwitz, 2005, pp. 233–234) (Horwitz & Horwitz, 2007, p. 997).

- “Homogenous” in the last row refers to a population that has a homogenous competence, i.e., all voters have the same probability of voting correctly.

Therefore, to get the best outcomes (the “improve” scenarios), it is recommended to have either:

- As many people as possible following a high-quality opinion leader.
- An epistocratic council where the number of people on the council is just less than the critical number.
- A homogenous democracy with a large turnout.

The first option is untenable for the purpose of deliberative democracy since following a high-quality opinion leader does not satisfy the reason giving requirements of deliberative democracy. Furthermore, under the model Kaniovski based his work on, correlations must decrease as the number of jurors increases for the distribution to make sense, so there cannot be too many people relying on the same opinion leader (*Ibid.*, p. 458). Regarding a large democratic turnout, this is also untenable for deliberative democracy due to the problem of scale; it would simply be impractical to have a large amount of deliberators. Therefore, since the two democratic options listed are not suitable under modern political conditions, this section concludes, by process of elimination, that an epistocracy is the most epistemically successful option of the three (or at least leaves democracy’s epistemic benefits further in doubt) based on voting dependence considerations. Ladha also comes to the same conclusion as the table above for high populations (Ladha K. K., 1995, pp. 361–362). Furthermore, Ladha also agrees that an epistocracy is the most epistemically successful form of decision-making, where the experts are chosen from different, even opposing, schools of thought (Ladha K. K., 1995, p. 360).

§4.5: Does Diversity Trump Ability? (Theory)

The result just stated is in contradiction with that of Landemore, who makes the claim that “it is epistemically better to have a larger group of average but cognitively diverse

people than a smaller group of very smart but homogeneously thinking individuals” since “cognitive diversity actually matters more to the production of smart collective solutions than individual ability does” (Landemore, 2012a, p. 90). Landemore bases this claim on the Diversity Trumps Ability Theorem, which comes from the work of Hong and Page (*Ibid.*, p. 102), which is commonly used in the defence of epistemic democracy. Therefore, another theoretical method used in this thesis to establish the interaction between cognitive diversity, deliberator quality, and number of deliberators is an analysis of the Diversity Trumps Ability Theorem. This is done both in terms of elaborating upon the theory (Section §4.5) and through the analysis of computer simulations (Section §4.6).

In terms of elaborating upon the Diversity Trumps Ability Theorem in theory, it is essential to understand exactly what the Diversity Trumps Ability Theorem states, and under what conditions it works, to understand to what extent it can be used to justify epistemic democracy, if at all. As this section will argue, some of the conditions required for the Diversity Trumps Ability Theorem to be satisfied are not found under modern political conditions. Furthermore, this section argues that to make the best use of cognitive diversity, deliberators must be chosen carefully to create an appropriate cognitive diversity.

§4.5.1: What is the Diversity Trumps Ability Theorem?

Between Hong and Page’s initial work (2004), and Page’s later book (2007), five sufficient conditions were assumed for the Diversity Trumps Ability Theorem to hold.

These five conditions are:

1. *Intelligence*: If given an initial potential solution, each problem solver will find (or stay with) a solution that is at least as good (Hong & Page, 2004, p. 16387).³⁶
Intelligent problem solvers have access to better solutions, it is easier for them

³⁶ The number of potential solutions a problem solver could settle on, known as “local optima”, is finite (Hong & Page, 2004, p. 16387).

to settle on their better solutions than their worse ones, and they can intentionally discern a good solution from a bad one (Page S. E., 2007, pp. 146,149). The intelligence condition also requires that the cognitive skills an agent brings to a problem must be relevant (*Ibid.*, pp. 159–160).

2. *Difficulty*: The problem in question is assumed to be difficult enough such that no single problem solver is guaranteed to find the optimal solution (Hong & Page, 2004, p. 16387) (Page S. E., 2007, p. 159). This criterion is important because if the problem was simple enough for someone to be guaranteed to find the optimal solution, there would be no need for the group.³⁷
3. *Diversity*: For every potential solution that a problem solver could settle on, other than the optimal solution, there is at least one other problem solver in the group that will not settle on this solution (Hong & Page, 2004, p. 16387) (Page S. E., 2007, pp. 160–161). This is important because it means that the group is less likely to get stuck at sub-optimal solutions if there is diversity in the potential solutions.³⁸
4. *Uniqueness*: There is a unique, best-performing individual problem solver (Hong & Page, 2004, p. 16388). This is simply a mathematical technicality rather than a criterion that must be upheld when selecting groups for deliberation; it simply gives Hong and Page a reference point when claiming that the group performs better than the best individual.
5. *Size*: The total number of potential problem solvers that could be selected is large, and there are more than a few actual problem solvers. The sizes of the potential and actual problem solver sets depend on the difficulty of the problem and the diversity found within the problem solvers; more difficult

³⁷ Note that this criterion requires that the problem is objectively difficult, e.g., a wicked problem; making the problem subjectively difficult for deliberators by choosing low competence deliberators is not the intention here (Page S. E., 2007, p. 345).

³⁸ It should be noted that, as stated, this is a sufficient condition, but it can be relaxed; as long as the problem solvers do not all have the same low local optima, then the Diversity Trumps Ability Theorem can still follow (Page S. E., 2007, p. 165).

problems need more problem solvers and a diverse group can be smaller (Page S. E., 2007, p. 162).

With these five criteria in mind, if there is a collection of potential problem solvers and one compares the performance of the N best performing individuals³⁹ from this collection as a group, and N randomly selected individuals from this collection as a group, then Hong and Page prove:

The Diversity Trumps Ability Theorem: Given [the above conditions], a randomly selected collection of problem solvers outperforms a collection of the best individual problem solvers (*Ibid.*).

That is to say, *if the above conditions hold*, then a public engagement involving N random citizens will outperform an epistocracy with N members. This thesis does not contest the Diversity Trumps Ability Theorem as a theorem; given the above conditions hold, then the theorem is a mathematical truth (Page S. E., 2007, p. 162). However, this thesis does investigate its applicability as a defence for epistemic democracy under modern political conditions, particularly in the case of emerging science and technology policy.

Landemore makes the claim that the above conditions “for this theorem to apply are not utterly demanding” (Landemore, 2012a, p. 102). This claim is correct for some of the conditions, but not others. The difficulty condition is satisfied since emerging science and technology policy are wicked problems, and hence are difficult for a democratic population given they are even difficult for experts. A democratic population also has the size to satisfy the fifth condition (and the fourth condition is a mathematical technicality rather than a criterion to be applied). What remains are the first condition, intelligence, and the third condition, diversity, as potential problems for the application of the Diversity Trumps Ability Theorem to emerging science and technology policy under modern democratic conditions. Chapter 3 gave evidence showing that the public understanding of nanotechnology is poor (and is unlikely to

³⁹ For the purpose of this thesis, “best performing individuals” means those individuals with the highest political competence.

improve), and hence the intelligence criterion is not met. Furthermore, given that most citizens are getting their information, if any, from the same source, namely television, then the diversity criterion is also not likely to be met (Brennan, 2014, p. 47). This is backed up by Althaus' finding that poorly informed opinions, i.e., the majority of opinions, have a tendency to be one-sided (Althaus, 2003, pp. 59–60); this can also be seen in the negative slide that was discussed in Section §4.2.

§4.5.2: The Need to Carefully Select Diversity

Having detailed the actual claims that the Diversity Trumps Ability Theorem claims, there is another important factor that needs to be realised before the Theorem can be analysed in earnest. It is important to realise that cognitive diversity does not come for free; in general, more deliberators cannot simply be added to the group in the without paying an epistemic price for it. This price can be seen when other theorems are considered. Some of these other theorems include the Diversity Prediction Theorem, and its corollary, the Crowd Beats the Average Law (Page S. E., 2007, pp. 208,209):

Diversity Prediction Theorem: Given a crowd of predictive models, each contributing to a larger cognitive diversity, then:

$$\text{Collective Error} = \text{Average Individual Error} - \text{Prediction Diversity}$$

Here, a predictive model is a combination of the worldview and reasoning processes of a deliberator that produces a prediction and/or policy decision.

Because prediction diversity must be positive, then the following is true:

Crowd Beats the Average Law: Given any collection of diverse models, the collective prediction is more accurate than the average individual predictions, or:

$$\text{Collective Prediction Error} < \text{Average Individual Error}$$

The Diversity Prediction Theorem says that prediction diversity and individual skill are equally important (contrary to Landemore's claim that diversity is more important). This is because there needs to be a balance struck with equal weights

between ability and diversity. Having highly skilled but very homogenous problems solvers will not have enough diversity to make up for the individual error, as low as it may be. Conversely, increasing diversity will tend to lower the average skill level and increases the average individual error, and this error may be too high for the diversity to be able to compensate for. As Page puts it: "Being different is *as important as being good*" (*Ibid.*, p. 208)(emphasis added). In other words, one cannot simply add democratic citizens to deliberation and expect that the resulting policy will improve due to added cognitive diversity; one must be careful about who is added to the deliberation to become more confident that adding that deliberator improves the policy created.

By way of adding to this conclusion, one may contend that this section has been vague regarding the notion of combining predictive models, and rightfully so. In theory there are an infinite number of ways that models can be combined. Factors that need to be considered are:

- How should everyone's models be weighted? Models could be combined democratically such that everyone's models are given equal weight, or experts could be weighed more heavily, or those models with more emphasis on the environment, for example, are given more weight.
- What would it even mean to combine predictive models together? If there are discrete policy options, such as voting for parties in an election, then combining voter models as a weighted average would likely mean a policy solution that does not exist "wins".
- Who, if anyone, checks that the combined model is a high-quality one? Even if a weighted average would make sense as a policy solution, such a weighted average may be a poor solution since it could be a trough between two policy quality peaks.
- What properties do the problem solvers need for the combined model to be of a good quality? High prediction diversity is an important quality the group

needs to reduce the collective error, so if the group has a significant systematic bias regarding the effect of a policy aspect and/or its interaction with other aspects, then prediction diversity will be low.

This section does not ask these questions to give definite answers since the answers will depend on the type of policy problem under consideration. Rather, the purpose of asking these questions is to realise that citizens as a collective will not automatically be smart; choosing the collective must be an intentional process. Furthermore, this guidance must be intelligent to find the correct weighting of individual models and policy selection procedure, as well as to be able to discern a high-quality collective solution from a poor one. This guidance must also be able to create a collective with an appropriate diversity. As shown in Chapter 3, neither of the intelligence or appropriate diversity conditions is found in the democratic population at large, and as such, the group of policy deliberators must be carefully guided. It is the task of Chapter 5 to provide that guidance.

§4.6: Does Diversity Trump Ability? (Computer Simulation)

Having elaborated upon the theory behind the Diversity Trumps Ability Theorem, this section now elaborates upon the implications of this Theorem. To investigate these implications, this section draws on the computer simulation forms of the Diversity Trumps Ability Theorem. If it is the case that a broad range of deliberators chosen from the general public can be shown to produce high quality policy, then it is more plausible that an epistemic democracy can solve the problem of expertise for emerging science and technology policy (and probably any other policy field). However, if a high level of expertise become necessary for emerging science and technology policy, perhaps despite a potential lack of cognitive diversity, then it is more plausible that an expert-focussed form of governance, like an epistocracy, can solve the problem of expertise for emerging science and technology policy.

§4.6.1: Hong and Page's Node Ring Deliberation Simulation

The Diversity Trumps Ability Theorem is not just the result of a mathematical proof; it can also be shown using computer-simulated models of group discussion. One such model set up by Hong and Page (2004) models a group of cognitively diverse problem-solving agents. To understand how their model works, imagine a circle with n nodes on it, and each node is given a random value Q between 1 and 100 inclusive, where n represents the number of potential solutions to a problem, and Q is the quality of the solution. Each problem solver is given the same diversity k out of a potential group diversity ℓ , where k and ℓ are the number of potential ways the problem solver and group can think about a problem, respectively. From this, a diversity heuristic is assigned for each problem solver, where k numbers are chosen randomly from between 1 and ℓ inclusive. A possible example would be $k = 3$ and $\ell = 12$ such that a diversity heuristic might be [1,4,9]. This diversity heuristic means that, given a randomly assigned starting node, the problem solver will search the 1st, 4th, and 9th nodes away clockwise from their starting point. If at least one of these nodes has a higher assigned quality value than the starting point, the problem solver makes note of this highest value and sets it as their new local optimum. The problem solver then continues the process again, starting ℓ nodes away from their starting point. This process is repeated until the problem solver does not land on a higher valued node during each search of k nodes, at which point, the next problem solver comes in to search, starting at the highest local optimum found by the previous problem solver. This process is repeated until no more problem solvers can land on a higher-quality node. The node associated with this highest collective local optimum is the solution decided on by the group. The assigned performance of a problem solver or group is the average highest local optimum found across all starting nodes. A measure of the diversity between heuristics, i.e., how differently two problem solvers think, is also generated by considering whether the nodes searched at each step of the cycle are

different for different problem solvers⁴⁰ (*Ibid.*, p. 16386). In a representative sample run for the ten agents and $\ell = 12$ case, “[t]he best agent scored 87.3; the worst agent scored 84.3; the average score of the 10 best agents was 87.1, and the average score of the 10 randomly selected agents was 85.6” (*Ibid.*, p. 16387); in other words, the absolute difference between the quality of the best agent and the worst agent was a mere 3%. Over multiple sets of 50 runs of this simulation, with varying numbers of agents and potential levels of diversity, although the randomly chosen agents had significantly higher levels of diversity relative to the best agents, the performance of the random agents was no more than 2.5% better in absolute terms than that of the best agents (*Ibid.*).

§4.6.2: The Node Ring Simulation Does Not Support Epistemic Democracy

An important question to consider now is whether the above described model can defend epistemic democracy as a potential solution to the problem of expertise. This section identifies many reasons why epistemic democrats should not rely on such a model, especially if they are trying to show that an epistemic democracy is the best of all decision-making systems. Since one of the aims of this thesis is to provide methods that can be used to investigate the epistemic performance of deliberative democracy, solutions to some of the identified problems are given such that a more accurate simulation of the deliberative process can be achieved. The reasons identified are as follows:

⁴⁰ Qualitatively speaking, the measure works fine since it gives larger diversity values where one would roughly expect them to be. Quantitatively though, the measure used is flawed since it would say that problem solvers with heuristics [1,4,9] and [4,9,1] are completely diverse since each problem solver would search a different node first, a different node second, and a different node third. However, overall, they are both searching the same three nodes, so they have no diversity at all.

Each Problem Solver Has Access to the Same Circle of Nodes

Having access to the same circle of nodes means that both the experts and the randomly selected citizens have access to the same quality nodes. In other words, each problem solver can access the same high-quality ideas and recognise them as high-quality ideas without pluralism, uncertainty, or an expertise gap obscuring their judgement. To say that each problem solver has access to high quality ideas and can recognise them as such means that, by default, the model has either already excluded average citizens from the process (the average citizen is not going to know certain potential solutions to an advanced emerging science and technology problem), or the experts do not get to make full use of their expertise (which puts the problem at the average citizen level by default). By shifting the competence levels like this, the possibility of describing an epistemic democracy appropriately with this model disappears. Furthermore, in saying that the problem solvers can see the quality of the solution on each node without pluralism or uncertainty being an issue, this means that citizens are able to see the problem and potential solutions perfectly objectively, which is highly unreasonable. As Moore and Brennan note, a group of problem solvers with the same circle of nodes implies that all problem solvers are seeing the same problem and have a common goal, and that there are no incentive or communication problems. However, democracies are much more complicated than this such that even if problem solvers all aim towards a vague notion of “the common good”, they will not all agree on what the common good is (Moore, 2014, pp. 104,106) (Brennan, 2014, p. 39).

To rectify these problems, this section suggests that the model needs to somehow weaken⁴¹ nodes for average citizens and introduce a random element on each node for

⁴¹ This weakening effect could be dampened if the effect of giving information to the problem solvers is modelled, akin to a consensus conference. Adding this into the simulation would mean that the problem solvers have a higher average node value than they otherwise would have, but it would also lower the diversity each citizen brings to the discussion since they are now statistically dependent with respect to the information given. Whether these two effects provide a net benefit is a matter to be determined by the simulation itself.

each citizen. Each node could have some sort of visibility filter such that average citizens might be blocked from seeing very high-quality solutions, or maybe add some sort of random fidelity check representing the trust they place in an expert “presenting” this solution. Furthermore, a given node could have different random values assigned to it to represent both the worldview of each citizen and the uncertainty each problem solver has regarding the harms and benefits of each solution. Therefore, the number each problem solver would see on each node, i.e., for each given policy solution, will be a function of their potential skill level, their fidelity to experts, their worldview, and their subjective uncertainty. Admittedly, this makes the model much more complicated, but also much more realistic.

The Model Does Not Capture Genuine Deliberation

In making the claim that the model does not capture genuine deliberation, it is recognised that this is a harsh criticism since capturing genuine deliberation within a computer model is very difficult. There would need to be some sort of iterative machine learning procedure that filled in nodes as solutions are discovered during the deliberative “discussion” rather than the policy solutions being predetermined by the problem solvers before searching. There would also need to be a machine learning procedure that allowed the problem solvers to learn from each other during the “discussion” process, which would allow for the quality of solutions to change as its merits are analysed. However, as it stands, some improvement can be made to the model relatively easily. Currently, the model effectively has each problem solver sitting in a room by themselves searching for a locally optimal solution. The problem solver then leaves a note for the next problem solver that only tells the new problem solver about the recently found solution, forgetting all other solutions that were searched for. Even if each problem solver could agree on the quality of each solution, and hence the other solutions searched do not matter, this is still not a very efficient problem-solving strategy.

This section suggests that a more accurate way of representing deliberation, though by no means perfect, is to model all the problem solvers in the room at the same time searching for solutions. For example, if there were three problem solvers with diversity heuristics of [1,4,9], [3,7,9], and [4,8,11], then these problem solvers could collectively search through seven different nodes each cycle, namely [1,3,4,7,8,9,11], rather than the three they would each be searching individually. This method is much more efficient as a search procedure, but it does have diminishing returns associated with it due to the overlap in nodes being searched. If problem solvers can search up to 12 nodes ahead, and each has a diversity heuristic of size three, then it will not take many problem solvers to cover all twelve nodes in the search cycle, especially if problem solvers are chosen specifically for their given diversity. Because of these diminishing returns, this more efficient form of decision-making would lend itself towards an epistocracy rather than an epistemic democracy. Having an overlap in these heuristics adds no significant benefit, and substantial opportunity costs, at least as far as the model is concerned (Quirk, 2014, p. 142). In real deliberation situations, however, having some overlap in heuristics would be beneficial since it would allow problem solvers to confer their solutions with each other without having to explain the heuristic itself, and it would allow those with a heuristic that leads to a minority dissenting opinion to team up rather than be overwhelmed by the deliberation process. However, even with the benefits gained by having these overlaps, this still does not support the claim that the deliberating group should be particularly large because the diminishing returns are still present.

Quirk also suggests two other ways in which the diversity heuristics can be improved. The first one is to give some problem solvers a larger k value, i.e., they can search more nodes per cycle. This would represent the real-world intuition that those of higher ability have access to a greater number of relevant problem-solving techniques. The second improvement would be to give problem solvers the ability to identify which heuristics are the most effective (*Ibid.*, p. 143); while the time taken to deliberate is not

a concern during computer simulations, decisions need to be made within an appropriate timeframe in the real world, so if problem solvers were able to identify better heuristics during the process, this could speed up the decision-making process.

*The Model Compares Experts with Experts*⁴²

In the representative sample run mentioned above, the individual performance of the best 10 problem solvers was only 1.5% better in absolute terms than the 10 random problem solvers (87.1 compared with 85.6) (Hong & Page, 2004, p. 16387). For all practical purposes, these are equally and highly qualified problem solvers; none of these are likely to be average citizens. Therefore, what Hong and Page's model shows, if anything, is that a diverse group of experts will tend to narrowly beat a less diverse group of experts. Quirk also notes that given the artificial nature of the computer simulation and the lack of attention to issues, such as costs of communication, emotional or attitudinal barriers to cooperation, and incentives for investment of effort, the small difference provides flimsy evidence for the "finding" that diversity "trumps" ability (Quirk, 2014, pp. 141,145). With that said, even if it is true that communication and cooperation issues do erode these benefits, however small, then this does at least show that cognitive diversity can be beneficial if managed properly (Page S. E., 2007, p. 314). The main point here, however, is that this result says nothing about the performance of average citizens, and therefore cannot be used as a defence

⁴² Quirk makes a similar argument that the computer model is self-fulfilling. He argues that because of the random nature of the computer model and the large number of nodes compared to the possible search heuristics, then by virtue of the law of large numbers and individual ability being defined as an average over all starting nodes, individual ability is bound to converge and the differences in ability between individuals will be small. Any slight advantage in terms of individual ability will come from having just the right small subset of heuristics, and as such, all the best problem solvers will be cognitively homogenous. Therefore, because the effect of ability is so small, and the best problem solvers are, by design, cognitively homogenous, the entire model simply becomes a diversity competition. Multiple high-ability problem solvers are therefore destined to be redundant to the group's success (Quirk, 2014, pp. 141–142). However, for Quirk, this is counterintuitive to what it means for a problem solver to have high ability; to have high ability means making contributions that are much better and more varied than those with low ability as opposed to only slightly better and less varied (*Ibid.*, p. 143).

of epistemic democracy. Rather, it only shows the benefits of cognitive diversity at relatively even levels of ability. Brennan says something similar. He reads the Diversity Trumps Ability Theorem as saying that the top 3% of problem solvers should be used rather than just the top 1%, and that it is misleading to extend the theorem to a general rule of the many rather than a rule of the few; even epistocrats can agree to rule of the many, as long as it is an intelligent many (Brennan, 2014, pp. 38–39).

The Model is a Computer Simulation

Adding to the concern just discussed, Hong and Page’s model only provides conditions under which the Diversity Trumps Ability Theorem holds; there is no testing of the model as part of real deliberations, nor is there adequate explanation of how real world conditions (in terms of what citizens know, how they get that knowledge, and the processes of deliberation), are capable of satisfying the model (Quirk, 2014, p. 135)(Gunn, 2014, pp. 59,62)(Brennan, 2014, pp. 34-35). This follows a common theme in the literature; defences of epistemic democracy tend to be theoretical (if conditions X, Y, and Z hold, then epistemic democracy works), while critiques of epistemic democracy tend to be empirical (conditions X, Y, and Z do not hold) (*Ibid.*, p. 36). As such, epistemic democrats need to show that actual democratic conditions are modelled by these theorems and simulations or something similar. That is to say:

- Find or adapt a computer model that has a similar intelligence and diversity profile to the general public and still verifies the Diversity Trumps Ability Theorem (it might be the case that a diverse population with heterogeneous intelligence still satisfies the Diversity Trumps Ability Theorem as long as the average intelligence is high enough), or
- Show that the public has an intelligence and diversity profile that does satisfy the conditions that Hong and Page have set out as sufficient conditions for the Diversity Trumps Ability Theorem to be fulfilled.

Without being able to do so, the Diversity Trumps Ability Theorem and the associated simulations are at best curiosities for the epistemic democrat, and at worst, they doom the cause of epistemic democrats.

At best for epistemic democracy, this section has shown that the claim of epistemic democracy is as yet unsupported, but potential future research has been presented. At worst, it appears that an epistocracy can at least satisfy the Diversity Trumps Ability Theorem and appears to be the most epistemically beneficial governance system, though this last claim is yet to be proven outright.

§4.7: Summary

The purpose of this chapter was to continue the analysis of the claims of epistemic democrats, who claim that democracy can at least track the truth, and is even the epistemically best form of governance on average in the long run, by considering democracy as a collective endeavour. Section §4.2 analysed the performance of instances of actual nanotechnology public engagement and found that the results were unsatisfactory, largely because those taking part did not have the social and political expertise to provide adequate solutions to policy problems. However, to analyse the potential of epistemic deliberative democracy, the remainder of the chapter provided a theoretical analysis of deliberative democracy. Section §4.3 provided an in-depth analysis of the correlations associated with deliberation and showed that a small collection of diverse and competent deliberators, i.e., an epistocracy, is most likely to produce the epistemically best outcomes. Furthermore, Sections §4.4–§4.6 considered what effect the cognitive diversity of a democratic public would have on deliberation. It was demonstrated that Landemore’s claim that the general public’s cognitive diversity gave it an epistemic advantage over an epistocracy was, at best, unfounded. Instead of defending epistemic democracy, these sections demonstrated that to make the best use of cognitive diversity, a deliberating group should be guided, and that

having an epistocracy as this guided group satisfies the Diversity Trumps Ability Theorem. Overall, throughout Chapters 3 and 4, the claims of epistemic democrats have been shown to be (as yet) unfounded based on the arguments considered, which means that there is strong reason to believe that the problem of expertise does exist. Furthermore, this chapter showed that an epistocracy containing cognitively diverse members is likely to be the most epistemically beneficial form of governance. Given this positive result for epistocracy, the next chapter seeks to create a form of epistocracy that can best utilise cognitive diversity in practice, which this thesis calls an “extended epistocracy”.

Chapter Five: Establishing an Extended Epistocracy

§5.1: Introduction

In Chapter 4, it was demonstrated that an epistocracy containing cognitively diverse members is likely to be the most epistemically beneficial form of governance. Because of this result, Part C of this thesis proposes a new epistocratic governance system for emerging science and technology policy called an *extended epistocracy* in Chapter 5, that has the greatest tendency towards producing epistemically high-quality policy decisions. However, a focus on maximising epistemic quality is in tension with democratic legitimacy, as stated in the problem of expertise. Therefore, Chapter 6 presents various areas of future research that could be used to defend the legitimacy of an extended epistocracy. This is done to answer the call from Jasanoff, who states that:

Participation alone, then, does not answer the problem of how to democratize technological societies. Opening the doors to previously closed expert forums is a necessary step – indeed, it should be seen by now as a standard operating procedure. But the formal mechanisms adopted by national governments are not enough to engage the public in the management of global science and technology. What has to change is the *culture* of governance, within nations as well as internationally; and for this we need to address not only the mechanics, but also the substance of participatory politics. The issue, in other words, is no longer *whether* the public should have a say in technical decisions, but *how* to promote more meaningful interaction among policy-makers, scientific experts, corporate producers, and the public. (Jasanoff, 2003, p. 238)(emphasis in original)

To begin developing an extended epistocracy, Section §5.2 identifies the key features from lessons learned from past emerging technology governance, useful substantive properties of epistemic deliberative democracy, and challenges that upstream engagement faces. By analysing these features, it becomes even clearer that nanotechnology governance should take the form of an extended epistocracy. To provide a framework for the development of an extended epistocracy, inspiration is taken from Cathrine Holst's classification of an epistocracy, as discussed in Section §5.3 (Holst, 2012, p. 42). From the eight dimensions that make up Holst's classification, six of them are used in this chapter since they directly apply to the epistemic quality of the policy created. In brief, these six interrelated components are *substance* (what the experts should know), *cognitive* (the incentives experts face), *actor* (who the experts are), *social organisational* (the role of the expert), *constitutional* (the rights and authority of experts), and *process* (how much the experts do). This chapter adds two extra dimensions beyond the six taken from Holst's classification. The first component added is a *selection* component, which explains the process by which the experts are selected for the epistocracy. The second component added is a *structure* component, which explains how experts are to interact with one another given the need for guidance of cognitive diversity, as discussed in Chapter 4. Therefore, there will be a total of eight dimensions relevant to the epistemic quality of an epistocracy that are considered in this chapter using the following questions answered in Sections §5.4, §5.5, §5.6, and §5.7, respectively:

- Who can be considered an expert?
- How are the experts to interact with one another?
- How are the experts to interact with the public?
- How are the experts chosen?

To establish who can be considered an expert, Section §5.3 considers the expertise required for emerging technology policy and those that are to provide it. To establish who is eligible for entry into an extended epistocracy, several general methods are

given that allow one to identify experts, which are decision-making quality, social construction, and professionalism. Four main areas of expertise are also defined to understand who is eligible to take part in policy discussion due to their relevant expertise, or lack thereof. Relevant expertise is defined to be a broad concept comprised of not just technical, scientific expertise, but also ethical and socio-political expertise, and it is held by trained experts, advocates, and local experts. A final group of citizens who do not have relevant expertise, called lay citizens, are also considered.

Once it is known what set of criteria would qualify someone as an expert, Section §5.4 discusses how the chosen experts are to act together as a group. Given the importance of diversity on group performance, as discussed in Chapter 4, the task-related (cognitive) and bio-demographic (identity) diversity of groups are analysed based on the psychology and management literature. Analysing the task-related diversity is especially important for emerging science and technology policy because the correct balance of task-related diversity helps to improve decision-making for uncertain, complex, and high stakes problems, i.e., wicked problems (Olsen, Parayitam, & Bao, 2007, pp. 200,203–204) (Jehn, Northcraft, & Neale, 1999, pp. 746–747). Analysing bio-demographic diversity is important because it is important to analyse recent and increasingly louder claims that increasing bio-demographic diversity is able to increase group performance (Bourke, Garr, van Berkel, & Wong, 2017). It is shown that on one hand, task-related diversity must be carefully balanced to reach its greatest potential, while on the other hand, bio-demographic diversity has little effect on group performance and so no specific recommendations are made for it.

Once it is known which experts have been selected for the extended epistocracy, Section §5.5 discusses what the role of experts is and the expectations placed upon them. Drawing upon Pielke's ideal expert types (Pielke, 2007, p. 14), the role of experts in an extended epistocracy governing emerging science and technology policy is placed on a spectrum such that an extended epistocracy can be considered a

democratic epistocracy (an epistocracy with democratic elements). Following on from assigning the roles of experts is assigning the level of authority the experts are to have and how much of the policy process the experts are involved with. It is argued that experts in an extended epistocracy are to be given authority as expert policymakers in their given area of expertise, but given the same rights as the average citizen in all other respects. The experts are also capable of creating policy at a local, national, or international level, and they are to take an active role through the entire policy process as decision-makers rather than as simply expert guides.

Finally, knowing all the features regarding who is involved in an epistocracy and how the epistocracy is to operate, it is possible to describe how the members of the epistocracy are chosen. Therefore, Section §5.6 proposes a selection procedure for an extended epistocracy. This chapter recommends that a database is developed and maintained of those that can fill specific expertise roles within the epistocracy. In terms of the cognitive dimension, the incentive for epistocrats to perform at their epistemic best is maintained by ensuring that the hiring process is open, fair, and balanced, and that there is an explicit ban on bribery or bargaining that may sway policy decisions.

§5.2: Requirements for a Nanogovernance Regime

Given the purpose of this chapter is to develop a governance regime that could be used for emerging technology governance, with nanotechnology being the chosen example of an emerging technology, it is important to ask: what requirements should a nanogovernance regime fulfil? Based on the discussion regarding the flaws of positivism in Section §2.2, it should be clear, first and foremost, that a technocracy is to be avoided. Therefore, to reiterate the disclaimer given in the opening section of this thesis: the extended epistocracy proposed in this chapter is not a technocracy. As the name suggests, an *extended* epistocracy makes use of an extended peer community beyond the one used for a technocracy. Furthermore, this engagement with the

extended peer community should be done in such a way that it avoids the use of the deficit model, also discussed in Section §2.2. This section also discusses important properties of epistemic deliberative democracy that can be incorporated into an extended epistocracy, as well as challenges regarding upstream engagement that need to be taken into account.

§5.2.1: Why Use Epistemic Deliberative Democracy?

In the process of developing a new governance system, it is useful to be able to identify what works well in governance systems that are already in place. To this end, this section aims to identify the benefits of using public engagement and what can be carried over into an extended epistocracy. Overall, it has been said that public engagement, and by extension deliberative democracy, has three overlapping types of justification: 1) normative (justifies the process in terms of ethical values), 2) instrumental (provides a means to an end), and 3) substantive (produces better end products)⁴³ (Laurent, 2008, p. 2) (Delgado, Kjølberg, & Wickson, 2011, pp. 830–831). While some normative and instrumental forms of justification have been given in Section §2.3.3, for the purpose of this thesis, the main concern is the substantive justification of epistemic deliberative democracy, primarily in terms of its epistemic output, but also in terms of its ability to produce democratic legitimacy. The substantive properties of epistemic deliberative democracy that should be incorporated into a governance system are as follows. First, it creates better decisions due to cognitive diversity enhancing the critical exchange of arguments and counterarguments with supporting information, and as a result, the overall discussion becomes better than the sum of the parts (Cobb, 2011, p. 1535) (Kyle & Dodds, 2009, p. 83). Second, this discussion allows for the policy created to be considered legitimate

⁴³ To clarify the difference between instrumental and substantive justifications since they are both concerned with ends, instrumental reasons aims towards predefined ends where what could be considered success can be defined before the procedure; on the other hand, the success of substantive justification cannot be defined before the procedure since the procedure itself is meant to determine the successful substantive result (Delgado, Kjølberg, & Wickson, 2011, pp. 830–831).

because potentially affected stakeholders are able to be heard and taken seriously in terms of their societal values and concerns regarding uncertainty as opposed to purely objective, scientized considerations (*Ibid.*, pp. 83,84) (Rogers-Hayden & Pidgeon, 2006, pp. 169,170)(Kjølberg, 2009, p. 62). Third, the actions of policymakers are subject to accountability and transparency, especially those regarding decisions made using public funds, through the use of public deliberation (Kyle & Dodds, 2009, p. 84) (Katz, Solomon, Mee, & Lovel, 2009, p. 534). Lastly, it satisfies the considerable political will regarding the democratisation of science and technology (Kjølberg, 2009, p. 62).

§5.2.2: Upstream Engagement

Another important consideration regarding emerging technology governance is when in the technology development cycle the governance process should occur (Kyle & Dodds, 2009, p. 87). As noted in Section §2.5.1, a strong consensus has developed among both governmental and nongovernmental stakeholders that informed citizen input is required early in the process of developing technologies, i.e., “upstream” engagement is the best time to engage the public regarding nanotechnologies, where “upstream” refers to the time before the development path and attitudes are determined (*Ibid.*, p. 88) (Hamlett, Cobb, & Guston, 2008, p. 1) (Burri, 2009, p. 498). Nanotechnology, as a relatively recent developing type of technology, is an important avenue for experimenting with forms of upstream engagement since it allows for a critical examination of “the driving purposes, expectations, imaginations, and social ends of upstream knowledge”, as well as “ownership, control, and responsibility”, and goes beyond contemplating only the scientific management of objective risks, impacts, and consequences (Macnaghten & Guivant, 2011, pp. 208–209) (Macnaghten, Kearnes, & Wynne, 2005, pp. 278,281,282).

While upstream engagement can promote meaningful engagement in the research and development process, it also poses challenges in terms of:

- trusting the public to understand and contribute meaningfully to discussions, and being able to provide good quality information such that they are able to do so.
- finding ways to incorporate the public directly and, more generally, finding those willing and capable of managing and taking part in future-oriented engagement.
- coping with and exploring the epistemic uncertainty regarding the future development and potential impacts of nanotechnology.
- whether the risk, ethics, and regulation frameworks are adequate and its interaction with supposed independence of science and economic dynamism of the applications.
- to separate hype from what is actually achievable.
- ensuring that genuine engagement towards impactful policymaking is taking place rather than being an opportunity for manipulating dialogue and perceptions of nanotechnology (Kyle & Dodds, 2009, p. 88) (Rogers-Hayden & Pidgeon, 2006, p. 176) (Katz, Solomon, Mee, & Lovel, 2009, pp. 539,540–541)(Burri, 2009, p. 499) (Kearnes, Macnaghten, & Wilsdon, 2006, p. 12) (Petersen, Anderson, Allan, & Wilkinson, 2009, p. 513) (Delgado, Kjølberg, & Wickson, 2011, p. 835).

Or, as summarised by Kyle and Dodds (2009, p. 90):

What needs to occur ‘upstream’ is the social and political framework for the conditions for such engagement. That is, the conditions of an educated, engaged public and governments’ and researchers’ commitment to open, accountable, contestable evidence for technological claims. Then, as specific risks emerge or applications become realisable, concrete public engagement with those factors is possible. To this end, we suggest that there needs to be synchronous development of ethics and nanotechnologies within a framework of public awareness and information (*Ibid.*, p. 90).

Similarly, Schwarz-Plaschg (2018, p. 162) states that engagement should not just be upstream but should also be long-term. In other words, for the synchronous development within the social and political framework to be enacted, those taking part need to remain up-to-date with recent developments in nanotechnology, or at least be able to be brought up-to-date easily.

§5.2.3: What Does Nanogovernance Need?

From the above discussion, the requirements for nanogovernance can be summarised as follows. Nanogovernance should accommodate the public's will for the democratisation of science by allowing a range of views into the discussion. This range of views is also used to enhance cognitive diversity, where the cognitive diversity should include not only technical aspects of nanotechnology, but also the social, ethical, legal, and economic implications, as well as uncertainty as part of future-oriented engagement. Those participating should be able to contribute meaningfully and in a way that does not involve manipulation or coercion of any kind, and they must have the ability to keep up-to-date with relevant policy issues. The process should be open (as far as possible) to promote transparency, and the outcome should be politically influential. Lastly, incentives should be provided to encourage this process to occur over the long-term.

These requirements completely cover the suggestions that are provided by Roco et al. (2011, p. 3560) when they suggest that nanogovernance should be transformative (i.e., politically influential), responsible (i.e., include social values), inclusive (i.e., allow all stakeholders to participate), and visionary (i.e., incorporate future-oriented, long-term planning). The above requirements also cover those suggested by Groves (2011, p. 789) and Powell and Colin (2008, p. 134), who recommend transparency and, again, that deliberation be politically influential. However, Powell and Colin go further by focussing more on the logistics and planning stage of the engagement exercise, where they suggest that both short- and long-term exercises take place, and that citizens,

scientists, and academic and government organisers are trained in terms of how to organise and engage within the deliberation process (*Ibid.*).

While it is the purpose of this chapter to provide a detailed governance system that is able to accommodate all of these recommendations, a quick overview is given here to highlight an important feature. Given the discussion regarding political competence and understanding of nanotechnology in Chapters 3 and 4, even just the epistemic recommendations for nanogovernance suggest that lay citizens are not going to be able to satisfy these requirements due to the high level of knowledge and maintenance of that knowledge required.⁴⁴ Therefore, this chapter proposes that nanogovernance is to be undertaken by an epistocracy, i.e., a rule by experts (Estlund, 2003, p. 53), where those taking part have the knowledge and resources to more effectively engage in the governance process. However, the requirements for nanogovernance require that the political will for democratic engagement be upheld by incorporating the viewpoints of relevant stakeholders. If “epistocracy” is taken to be synonymous with “technocracy”, then there appears to be a dilemma. This dilemma is resolved once it is realised that “epistocracy” is a broader concept than “technocracy” and that all forms of relevant expertise, not just technical expertise, can be incorporated into an epistocracy; it is precisely this extension of the relevant forms of expertise that makes the extended epistocracy created “extended”. As such, each relevant stakeholder and their viewpoint can be represented by someone with expertise relevant to that stakeholder, and it is these relevant experts that make up a dedicated extended epistocratic council. In doing so, all relevant expertise and the associated cognitive diversity is democratically represented among epistemic equals who are politically effective and have the incentives to remain engaged in a long-term open process. This forms the broad outline of the extended epistocracy that is proposed in this chapter.

⁴⁴ Training could be given to lay citizens, as suggested by Powell and Colin (2008, p. 134), but an impractical amount of training would be required, and those citizens would no longer be lay citizens at the end of it, hence defeating the purpose of deliberating with lay citizens.

§5.3: What Is an Epistocracy?

In short, an epistocracy is a rule by the knowers or the wise, i.e., it is a rule by experts (Estlund, 2003, p. 53). However, to give a more specific structure to the concept of an epistocracy, Cathrine Holst has provided eight dimensions that can be used to classify different forms of epistocracy: historical, organisational, constitutional, process, substance, actor, cognitive, and normative. The organisational dimension can be further broken down into cultural and social subdimensions (Holst, 2012, p. 42). These dimensions are defined shortly. This chapter also adds a selection dimension to clarify what process is used to select the participants of an epistocracy, known as epistocrats, and a structural dimension since the eight dimensions listed above only consider an epistocracy as a collection of individual experts without considering how they are to work together as a team. For this thesis, the historical dimension⁴⁵ is considered only inasmuch as the form of epistocracy developed throughout the thesis is for emerging science and technology policy, and therefore it can be considered a modern form of epistocracy. Furthermore, the cultural organisational,⁴⁶ and normative⁴⁷ dimensions are concerned with democratic legitimacy rather than the epistemic quality of the policy created, and hence they are not included within this chapter but instead are the topic of suggested future research.

This chapter will use the six remaining (sub)dimensions from Holst's framework, plus the two created for this chapter, to propose an extended epistocracy to be used as a

⁴⁵ The historical dimension can be separated into pre-modern and modern epistocracies, though this distinction is a vague matter of degree rather than a sharp distinction. To provide a rough guide to identifying a modern epistocracy, the more post-normal the society is, the more modern it is (Holst, 2012, pp. 42-43).

⁴⁶ "*Cultural epistocracy* refers to societies where respect for knowledge and knowers is considerable and many subscribe to the idea that decision-making must be knowledge-based and knowers must play a significant role in decision-making" (Holst, 2012, p. 44).

⁴⁷ The normative dimension concerns the moral justification for the level of epistocracy used in a governance system. For example, one may believe that democratic input is the most important aspect, and therefore they will leave little room for epistocratic features. Conversely, one may believe that policy quality is the greatest value, and therefore will prefer a high level of epistocratic features (Holst, 2012, pp. 50-52).

potential form of emerging science and technology governance. As a brief description, these eight dimensions are:

- The *substance* is concerned with the range of expertise required from the epistocrats, i.e., the factual/technical and/or ethical dimensions of policy decisions (*Ibid.*, p. 47).
- The *actor* dimension considers who the knowers are and hence who is invited to become part of the epistocracy (*Ibid.*, pp. 48–49).
- The *structure* dimension finds a configuration of the properties of epistocrats that most effectively utilises the cognitive diversity of the epistocrats as they work together.
- The *social organisational* dimension is the level of respect given to experts within the policy-making community,⁴⁸ and hence the range of influence experts have when making policy decisions (*Ibid.*, p. 44).⁴⁹
- The *constitutional* dimension outlines the level of social/political rights given to experts and non-experts, and identifies where along a democracy–epistocracy spectrum the epistocracy falls (*Ibid.*, p. 45). The constitutional dimension is also concerned with the level at which the epistocracy acts, e.g., local, national, or international (*Ibid.*).
- The *process* dimension describes how much of the process the experts are a part of. For example, the experts might only act as resources of scientific knowledge, or they might be fully responsible for decision-making and implementation (*Ibid.*).
- The *selection* dimension proposes the methods by which experts are chosen to be part of an epistocracy, e.g., democratically voted, nominated, job interview, etc.

⁴⁸ The social organisational dimension is in contrast with the cultural organisational dimension, which is the level of respect and authority the *public* gives to experts (Holst, 2012, p. 44).

⁴⁹ For the purpose of this chapter, the social organisational dimension is so closely linked to the next two dimensions (constitutional and process) that this chapter considers them as being subdimensions of the social organisational dimension.

- The *cognitive* dimension is related to the incentives the epistocrats face: do the experts have an incentive to behave as deliberating knowers, or are experts driven by monetary or other non-epistemic incentives? (*Ibid.*, pp. 49–50).

Each of these aspects of an extended epistocracy will be developed in the order just listed.

§5.4: Who Are the Experts?

The defining feature of an epistocracy is that it is comprised of experts, so the core question for this section is: who can be considered an expert? To answer this question, the substance (i.e., what expertise is included?) and actor (i.e., who can take part?) dimensions of an extended epistocracy are developed. After giving a brief description of the general ways in which experts can be identified, this section develops the substance dimension of an extended epistocracy by considering citizens as being members of one of four epistemic groups for the policy issue at hand. These four groups are trained experts, advocates, local experts, and lay citizens, where the first three of these groups are considered to hold relevant expertise. It will be explained, per postpositivism and the extension of the peer community, that there is a large range of relevant knowledge to be considered in the policymaking process, and so the substance dimension will contain not just technical scientific knowledge, but also ethical and local expertise. In contrast, lay citizens are those that do not have significant relevant expertise for the policy issue at hand, and are therefore not considered eligible for selection into an extended epistocracy for that particular policy issue.

It is important to emphasise that these four expertise types just mentioned are case specific, i.e., a citizen could be in any one of these four main groups depending on the policy issue at hand. A particular set of trained experts will not be in charge of all policy, nor will average citizens never have the opportunity to take part in the political

process. The opportunity for participation for the average citizen is still open as long as they can provide relevant expertise. For the vast majority of policy areas, though, any given citizen will be in the “lay citizen” category.

§5.4.1: General Ways to Identify Experts

In general, an expert is someone with “an intangible but recognizable combination of education, talent, experience, and peer affirmation” (Nichols, 2017, pp. 30–31). More specifically, some methods that can be used for the identification of someone with expertise are decision-making quality, social construction, and professionalism. Each of these three methods of identification will be discussed in turn.

Decision-Making Quality

The first way someone might be identified as an expert is through the quality of their decisions. That is to say, an expert is someone “whose level of performance exceeds that of most others” (Cianciolo, Matthew, Sternberg, & Wagner, 2006, p. 614). By considering performance over time, experts can still be considered experts if they are wrong on occasion (and sometimes even justifiably wrong⁵⁰ given the uncertainty involved), but also emphasises the fact that they are more likely to be correct than non-experts (Yates & Tschirhart, 2006, p. 424). Put another way, experts are those that know more than most about their given field, along with having a natural aptitude for it, and as a result, their expertise is more likely to be authoritative than most other people (Nichols, 2017, pp. 29–30,32–33,36).

Social Construction

A second way to represent expertise is through “social construction”, where experts are those who are claimed to be experts by other established experts. This does have

⁵⁰ By “justifiably wrong”, it is meant that the expert makes a reasonable decision based on the best information and theories reasonably available to the expert, and takes adequate account of the uncertainty that the expert faces.

the potential flaw that it is possible for someone to “be an expert” by knowing the right people, or by putting on an appropriate act (Yates & Tschirhart, 2006, p. 426). However, in general, this “social construction” is through the rigorous use of accreditation, such as peer review, board certification, and membership to professional associations (*Ibid.*, p. 35).

Professionalism

Finally, a third way that an expert can be identified, and that includes the identification methods discussed above, is to consider their profession. An expert can be identified relatively easily based on their profession due to the profession’s institutionalised nature. Furthermore, this institutionalisation is a good indicator of quality because it generally requires training, accreditation, and experience to enter the profession and rise to positions of seniority (Evetts, Mieg, & Felt, 2006, pp. 105–106). While not perfect, this kind of “credentialism” is at least a quick and tangible method of recognising one’s expertise in the form of degrees, awards, publications, and track record (Nichols, 2017, pp. 31,33).⁵¹

§5.4.2: Relevant Expertise

Now that general methods for identifying experts have been identified, one might object that these three general methods just mentioned for identifying general expertise appear unreasonably elitist since it appears that one must have an officially documented track record that is recognised by others that have a close connection to the profession (e.g., published research articles that have gone through a peer review process). However, one of the aims in developing an extended epistocracy is to move

⁵¹ Science is a particularly strong example of the institutionalisation of expertise with its own system of “checks, validation procedures, recognition and authority processes, and hence claims to legitimacy” (Evetts, Mieg, & Felt, 2006, pp. 105–106). This system of mechanisms is also intended to act as a legal restriction to competitors, such as pseudoscientists (*Ibid.*, p. 118). While scientific professions may be an example of knowledge-based professions, these forms of identification credentialism can be extended to practical and artisanal professions as well (Stehr & Grundmann, 2011, p. 16).

away from such an exclusively elitist form of selecting experts such that the peer community can be extended. It is therefore necessary to identify how far this extension reaches in terms of what this thesis refers to as “relevant expertise”. This extension of what expertise is in accordance with the fact that expertise itself is a broad notion. While it may be easy to conceive of the scientist in their lab coat as an expert, a farmer is also an expert in cultivating plants and a travel guide is an expert in the tourist trade, since they have appropriate knowledge and experience, and have gained trust and social respect for it (Stehr & Grundmann, 2011, p. x); such examples satisfy all three qualifiers for expertise just discussed in Section §5.4.1. Put another way, one does not need to be an intellectual to be considered an expert (*Ibid.*, p. 28). Therefore, in relation to the decision quality, social construction, and professionalism criteria for recognising experts described above, an example of one way to extend the peer community of experts is that a member of the public who is well-respected in their community for their high-quality community work can be considered as having relevant expertise. While it may be more difficult for community work to be related to emerging science and technology policy than it is for other policy areas, examples might include a farmer concerned about nanomaterials in waterways, or a nurse that is concerned about nanodevices to alleviate disabilities.

However, even if expertise is a broad notion, an extended epistocracy still requires that expertise is held to a high standard. Given that an extended epistocracy is designed for solving wicked policy problems, which requires a great deal of complex knowledge under uncertain circumstances, this high standard would involve satisfying the skill set for naturalistic decision-making⁵² is also important, where the skills required for naturalistic decision-making are:

⁵² The key contextual features of naturalistic decision-making “are: 1. Ill-structured problems (not artificial, well-structured problems). 2. Uncertain, dynamic environments (not static, simulated situations). 3. Shifting, ill-defined or competing goals (not clear and stable goals). 4. Action/feedback loops (not one-shot decisions). 5. Time stress (as opposed to ample time for tasks). 6. High stakes (not situations devoid of true consequences for the decision maker). 7. Multiple players (as opposed to

- The ability to make fine discriminations and identify complex patterns.
- Significant levels of declarative knowledge (including tacit knowledge) and spend more trying to understand situations with that knowledge.
- Being able to recognise uncertainty, as well as their own strengths and weaknesses.
- The ability to have rich mental models and run mental simulations, and be able to implement strategies using a repertoire of tactics and leverage points (Ross, Shafer, & Klein, 2006, pp. 405–406).

All of those chosen to take part will need to be able to satisfy these four skills, regardless of which relevant expertise source group they come from.

There is one more general criterion that places another restriction on who is suitable to take part in an epistocracy, though this restriction may be loosened depending on the availability of appropriate experts. Those taking part within the extended epistocracy are to be considered representatives from their respective areas of expertise. For example, a nanoscientist on an extended epistocracy will be considered the representative for nanoscience and its community of nanoscientists. However, it is important to note the distinction here between the representatives and other members of the expertise group since the representatives should, if available, hold an additional type of knowledge, i.e., they should have interdisciplinary knowledge. In an extended epistocracy, the representatives need to be fluent not only in the area of expertise they are representing, i.e., *contributory expertise*, but also in the general policy issue at hand and the policy processes involved, i.e., *interdisciplinary interactional expertise*, to allow for a greater fluency of understanding between members of the extended epistocracy,⁵³ where:

individual decision making). 8. Organizational goals and norms (as opposed to decision making in a vacuum).” (Zsombok, 1997, p. 5)

⁵³ For more detail on why this is the case, see Section §5.4.3 later in chapter regarding cognitive diversity versus cognitive consensus.

- *Contributory expertise* is the expertise required to take part in the activity with competence and contribute to the domain of expertise (Gallopín & Vessuri, 2006, pp. 41–42) (Collins & Evans, 2007, pp. 14,24)
- *Interactional expertise* is expertise of the language of a specialist subject without having practical experience within that subject (*Ibid.*, p. 28). To achieve the required level of fluency involves interacting with the contributory experts themselves rather than simply reading the primary literature (*Ibid.*, pp. 31–33).

Having contributory expertise on its own is very demanding, so to also aim for interactional expertise is especially demanding, which is another reason why policy decision-making should be left to experts rather than to the average citizen wherever possible.

In summary, the concept of relevant expertise as used for the purpose of this thesis is broader than the technical expertise that would be used in a positivistic technocracy, but not so broad as to allow everyone to take part in decision-making for any given policy issue. The rest of this section elaborates on the three main overlapping types of expertise that comprise relevant expertise, namely trained expertise, advocacy, and local expertise, and they will be discussed in turn.

Trained Experts

The first type of relevant expertise is that held by trained experts, who are those that have been formally trained in their particular area of expertise. A key distinction between technocracy and the broader notion of epistocracy can be found here, even with this “elite” type of expertise. In a technocracy, the set of relevant experts only contains those who are trained experts from technical fields, such as the natural sciences, engineering, mathematics, computing, economics, etc. Under an epistocracy, the concept of trained expertise can be expanded to include all forms of trained expertise that are relevant to the policy case at hand, which can include social scientists, ethicists, and others from the humanities. It is precisely this extension to the

relevant fields of expertise beyond the technical fields that is the purpose of an extended epistocracy, hence the name.

Advocacy

To continue the extension of a governance system beyond that of a technocracy, the second type of relevant expertise is advocacy expertise. For the purpose of this thesis, advocates are defined as those that are actively politically engaged in relation to a specific cause outside of their work as part of an extended epistocracy (as discussed later in Section §5.6, all members of an extended epistocracy are to act as advocates), and these advocates hold expertise specific to their area of advocacy. This advocacy form of expertise is not mutually exclusive with trained expertise in that there are trained experts who act as advocates for civic groups, such the movement scientists (Fischer, 2000, pp. 45,74,110,151–154). Aside from those advocates that have trained expertise, there will also be advocates that have self-taught, primary source expertise, which is expertise gained by reading the primary or quasi-primary literature.

Local Knowledge

The third group of relevant expertise is local expertise, which is a type of contributory expertise. Local expertise is expertise that focuses on understanding the immediate situation and environment (Gallopín & Vessuri, 2006, p. 43). Examples of local expertise may include cyclists who know which roads are the most dangerous for them to cycle along, parents who know what facilities in the area are missing that could help their child's development, or citizens who have anecdotal evidence of an increase in a disease's occurrence and suspect the cause may be a tainted water supply. Another form of local knowledge is traditional expertise, which is based on experience and adapted to a local culture and environment over time and continues to evolve, i.e., traditional expertise is a long-term version of local knowledge (*Ibid.*, pp. 42–43).

§5.4.3: Lay Expertise

The fourth group of citizens to be considered are lay citizens, who are those citizens that, by definition, do not possess relevant expertise. While it may be the case that the relevant expertise just discussed is broad, it is not all-encompassing. It may seem trivial to consider lay citizens in such a way where lay citizen is simply someone without relevant expertise, but this distinction is contrary to the core of democratic theory. In democratic theory, everyone that wishes to have a say about the issues that affect them can do so, as discussed in Chapter 2, and as such, democratic theory considers that all (at least adult) citizens have expertise that is relevant. However, as shown in Chapter 3, there are those citizens that do not meet the criteria for political competence regarding nanotechnology policy. It is important to recall that the distinction between relevant and lay expertise is relative to the policy issue being discussed. For a given policy issue, the vast majority of citizens will be in the lay citizen group, even those that are world-renowned in their field for other policy areas, and so the term “lay citizen” should not be considered a pejorative. Lay citizens do have expertise, it is just not high-quality or relevant enough for high-quality emerging science and technology policymaking. The most advanced form of expertise that lay citizens, relative to the policy area being discussed by the extended epistocracy in question, is known as popular understanding, which is understanding of a field gained through mass media and popular books. Popular understanding knowledge is usually unproblematic in the case of settled knowledge, but is problematic in the case of unsettled knowledge, such as with the uncertainty surrounding emerging and future nanotechnology. The problems arise because the distance between lay citizens and knowledge creation, and the narrow portrayal of emerging technology knowledge by the media, means that the popular understanding is a long way from a deep understanding of the concepts involved (Collins & Evans, 2007, pp. 19–22).

§5.5: Considering the Expert Actors as a Collective

The discussion of the actor dimension in Section §5.4 considered which individuals were eligible for entry into the extended epistocracy. However, as shown in Section §4.5.2, to create the best governance council, it is not enough to simply collect the best experts into a single team. It is important to consider how these experts are to work together and that the cognitive diversity is guided appropriately. To this end, the literature on diversity within work-groups is analysed to find which features of epistocrats (actor dimension) and of the extended epistocracy itself (structure dimension) can create optimal emerging technology policymaking conditions when the epistocrats are combined. Part of the challenge regarding this analysis is that if one were to consider “diversity” as a broad category and its effects on group performance, the results are mixed (Crisp & Turner, 2011, p. 256) (van Knippenberg & Schippers, 2007, p. 517). In this broad sense, diversity is what Milliken and Martins call a double-edged sword; diversity allows for a greater range of perspectives and the generation of higher-quality policy, but it can also cause group members to become more dissatisfied as they fail to integrate with the group (Milliken & Martins, 1996, p. 403). Because of this double-edged nature of diversity, Milliken and Martins ask if there are ways in which organisations can balance the positive and negative aspects of diversity (*Ibid.*, p. 421); “[i]f managed properly, team heterogeneity can create a significant operational synergy, whereas mismanaged team diversity can become a major impediment to optimal functioning” (Horwitz, 2005, p. 219). The general purpose of this section, then, is to find a way to manage diversity properly to achieve high-quality emerging technology policy.

More specifically, this section provides a greater understanding of the effects of diversity beyond what was given in Chapter 4 by answering the question: How are the experts to interact with each other? In Chapter 4, it was shown that “all things equal otherwise (that is, controlling for a number of interfering factors like

communication costs)”, cognitive diversity is beneficial for the quality of the policy created (Landemore, 2012a, p. 2). However, controlling for these interfering factors is a significant idealisation. Furthermore, the understanding of cognitive diversity in Chapter 4 was implemented in a vague way such that even though a brief description of cognitive diversity was given, cognitive diversity was reduced to a set of numbers during its application to represent problem solver diversity, group diversity, node quality, and competence. By reducing these factors to a single number each, critical information was not emphasised, such as what aspects of diversity are useful, how cognitive diversity overlaps rather than creates a difference, and the effects of diversity over time. In contrast, this section takes a more nuanced approach by breaking diversity into two useful categories, namely task-related diversity and bio-demographic diversity, to be defined shortly. Each form of diversity is elaborated on in turn to give a deeper understanding of what each form entails, its effects on deliberation within a policy-making group, and the ways in which they can be managed to achieve a better understanding of the actor and structure dimensions of an extended epistocracy.

For task-related diversity and its associated task-related conflict, analysis will be given of the psychology and management literature regarding the interaction of these task-related effects with experience, competence-based trust, and to what extent there needs to be an overlap in the expertise provided by group members to create the most epistemically capable group. It will be shown that the greater the experience of the group, the more effective task-related conflict is, that competence-based trust effectively focuses task-related conflict, and that there needs to be an overlap in expertise regarding a general understanding of the policy issue at hand, but each member also contributes their own unique expertise the group’s task-related diversity. Regarding the bio-demographic diversity, analysis of the psychology and management literature shows that bio-demographic diversity has an ambiguous to

non-existent effect on group performance, and therefore no specific recommendations are made regarding a bio-demographic diversity quota.

§5.5.1: Finding Useful Categorizations of Diversity

The main distinction that is used in the literature to differentiate between diversity types is between a person's (visible) identity and a person's cognitive skill. A particularly clear example is from Horwitz and Horwitz:

“bio-demographic diversity represents innate member characteristics that are immediately observable and categorized (e.g., age, gender, and race/ethnicity) whereas task-related diversity is acquired individual attributes (e.g., functional expertise, education, and organizational tenure) that have been postulated to be more germane to accomplishing tasks than bio-demographic diversity”⁵⁴ (Horwitz & Horwitz, 2007, p. 990)

Note that “task-related diversity” from the above quote and “cognitive diversity” from Section §4.3 can be used interchangeably. Bio-demographic and task-related types of diversity are related to emotional⁵⁵ and task-related conflict,⁵⁶ respectively, where the more visible the diversity trait is, the stronger the emotional conflict, and the more task-related the diverse trait is, the stronger the task-related conflict (Pelled, 1996, p. 617). A middle ground between strongly bio-demographic diversity and strongly task-related diversity would be those forms of diversity that are not visible, but not necessarily task-related, such as sexual orientation, personality, attitudes, and

⁵⁴ At first, it may seem somewhat self-fulfilling to specifically place all task-related forms of diversity into one category in a thesis whose aim is to maximise task-related performance. However, this would only be the case if the other category were just a catch-all for “other” forms of diversity. It is important to note for later in this chapter that bio-demographic diversity is the type of diversity that is not useful for task-related performance. It will be shown that it is one's skill set, rather than one's bio-demographic identity, that truly matters when considering the group's performance. These effects of diversity naturally lead to the formation of an epistocracy rather than a more representative style of public engagement based on identity representation.

⁵⁵ Emotional conflict “is individual-oriented disagreement arising from personal disaffection” (Amason & Sapienza, 1997, p. 495)

⁵⁶ Task-related conflict is disagreement about the content and/or process of the task the group is facing (Jehn, Northcraft, & Neale, 1999, p. 743)

values (van Knippenberg & Schippers, 2007, p. 519) (Jehn, Northcraft, & Neale, 1999, p. 745).

Making this distinction between bio-demographic and task-related diversity helps to clarify the “double-edged” nature of diversity by identifying the competing effects of diversity on group performance, and hence the mixed results from studies on diversity (Horwitz, 2005, p. 225). In some studies, bio-demographically heterogeneous groups could outperform bio-demographically homogeneous groups, and vice versa for other studies. This seeming inconsistency can be explained by realising that a bio-demographically heterogeneous group also just happened to contain task-related diversity; it was the task-related diversity that was of primary importance, not the bio-demographic diversity, and the task-related diversity was able to outweigh the emotional conflict caused by bio-demographic diversity. The converse can also be true in cases where bio-demographically homogeneous groups outperform heterogeneous groups; the emotional conflict caused by bio-demographic diversity in the homogeneous group was low enough that it did not outweigh the task-related diversity that was present (Jehn, Northcraft, & Neale, 1999, pp. 741–742). Due to the generally competing effects of bio-demographic and task-related diversity, diversity must be managed carefully rather than letting the group more naturally form as would be the case in deliberative democracy (Kurtzberg, 2005, pp. 62–63).

§5.5.2: What is Task-Related Diversity?

Cognitive ability is the “capacity to understand complex ideas, learn from experience, reason, problem solve, and adapt” and is “one of the best predictors of individual job performance (Devine & Philips, 2001, p. 507). However, the relationship between individual performance and group performance is not simple, as discussed in Chapter 4. It may be the case that a group of individuals with high cognitive ability fail to utilise it effectively as a group, so task-related diversity is also important if a group is to make high-quality decisions (*Ibid.*, pp. 507–508). Groups that have higher task-

related diversity, have access to a greater range of knowledge, skills, and abilities (KSAs) that comes from a range of educational, occupational, and life experiences. These greater cognitive resources allow for the consideration of multiple and better-quality alternatives, leading to better, more creative, problem solving; better solution implementation; and the ability to cope with non-routine, complex, and uncertain problems (Horwitz, 2005, pp. 224–225,234)(Evans & Carson, 2005, p. 303) (Milliken & Martins, 1996, p. 416) (van Knippenberg & Schippers, 2007, p. 518) (Olsen, Parayitam, & Bao, 2007, pp. 200,203–204) (Wegge, Roth, Neubach, Schmidt, & Kanfar, 2008, p. 1302). An increased diversity of KSAs can also be associated with a different range of social networks coming from a range of employment or education background, for example, and access to these social networks can help a group to be more innovative (Milliken & Martins, 1996, pp. 410–411). Communication with networks outside the group is beneficial for the performance of the group as it helps the group member to draw upon extra expertise if needed (Evans & Carson, 2005, pp. 304,308,309–310).

§5.5.3: Managing Task-Related Diversity

While task-related diversity has the potential to produce high-quality outcomes, as noted above, careful management is required to achieve these outcomes; a balance needs to be struck between the positive and negative aspects of task-related diversity. By definition, a range of KSAs within a task-related diverse group means that there are going to be a range of different perspectives, interpretations, and assumptions regarding the issue at hand (Mohammed & Dumville, 2001, p. 98). These differences need to be reconciled as part of what is known as task-related conflict, which is concerned with disagreements about the content and/or process of the task; the greater the task-related diversity, the greater the task-related conflict (Jehn, Northcraft, & Neale, 1999, p. 743)(Amason & Sapienza, 1997, p. 495). Despite being termed a form of “conflict”, task-related conflict can be very beneficial to group performance because by taking part in the conflict, a group is forced to thoroughly investigate information

and potential solutions, and they are less likely to settle on a sub-optimal solution⁵⁷ (*Ibid.*); in this way, disagreement itself becomes a basic resource (Miller, Burke, & Glick, 1998, p. 41). To utilise this resource effectively, this section considers various effects of task-related diversity, and as a result, show that long-term interaction between those with high levels of interdisciplinary relevant expertise are able to make the best use of task-related diversity.

Solving Initial Disagreements

In the initial stages of a decision-making group's lifetime, task-related diversity can tend to have a negative effect on performance due to disagreements regarding how the group should proceed (Milliken & Martins, 1996, p. 411).⁵⁸ These disagreements lead to an increase in the turnover rate and a decrease in performance due to members being less comfortable communicating ideas (Milliken & Martins, 1996, pp. 411–412). If these initial disagreements can be overcome, then in general, task-related conflict leads to higher group performance (Olsen, Parayitam, & Bao, 2007, pp. 213–214) (van Knippenberg, Carsten, & Homan, 2004, p. 1009) (Pelled, 1996, pp. 624–625) (Horwitz & Horwitz, 2007, pp. 1005–1006). This thesis therefore recommends that the members

⁵⁷ This settling can take place in the form of groupthink, which occurs when members of the group would rather retain unanimity and be seen as part of the group rather than create controversy by criticizing others' views, particularly those of the senior members of the group (Pelled, 1996, p. 624). Overcoming groupthink does require a measure of open-mindedness/cognitive motivation (the willingness to listen to, learn, and engage with varying viewpoints), but if this is present, particularly in complex tasks, then it will increase the likelihood of creative new knowledge emerging (Mitchell & Nicholas, 2006, p. 69) (Kearney, Gebert, & Voelpel, 2009, pp. 584–585); challenging, deliberating, and discussing different perspectives in a comprehensive manner is crucial for generating innovative solutions (Mitchell, Nicholas, & Boyle, 2009, pp. 539–540). It is therefore essential that open-mindedness/cognitive motivation and the norm of free expression is present so the utilisation of diverse viewpoints through brainstorming can take place (Mitchell & Nicholas, 2006, pp. 71–72). This norm of free expression requires psychological safety such that members of the group are not afraid to express different opinions, and therefore miss out on making use of their diversity (Roberge & van Dick, 2010, pp. 303–304).

⁵⁸ This difficulty in proceeding is a part of task-related conflict known as process conflict (Jehn & Mannix, 2001, p. 239). Process conflict hinges on (work) value diversity, e.g., those who value quality will have conflict with those that value efficiency or cost-effectiveness (Pelled, 1996, p. 622)(Jehn, Northcraft, & Neale, 1999, p. 745). Higher process conflict means that group members are spending more time arguing about how to solve the problem and spend less time actually solving the problem, and therefore making good use of task-related diversity (*Ibid.*, p. 746).

chosen to form an extended epistocracy are to remain within the epistocracy for a significant amount of time and work together on multiple policy cases. This long timeframe gives members of an epistocracy the opportunity to work through disagreements about how the group should proceed for their first policy case together, and future policy cases are more efficient since this process will not need to be done in as much detail for further policy cases. There are going to be situations where group members will need to be replaced due to life circumstances or bad practice (a position on an extended epistocracy should not be tenured and is open for review), but the intention is that the group stays together for a significant period of time.

Further Benefits of Experience

Task-related conflict may also lead to emotional conflict as personality clashes become triggered during heated task-related debate. However, this becomes less important with more experienced members since these members are more used to negative criticism and can separate the meaning from the message and continue with the task at hand (Olsen, Parayitam, & Bao, 2007, p. 202). It may even be the case that task-related conflict reduces emotional conflict since people are given the opportunity to voice their opinions and they become more committed to the outcome as a result because they have had a chance to make a contribution and they enjoy the learning process that comes with hearing the viewpoints of others (*Ibid.*, p. 203). Therefore, this thesis further recommends that members of an extended epistocracy should have experience in a policy-making role, whether it is in the same extended epistocracy, or with another group, such that the cognitive diversity can be used appropriately. Because policy experience is valued so highly in an extended epistocracy, this thesis recommends the use of specialised policymakers wherever possible, who have interactive expertise regarding the stakeholder group they are representing, but also have policymaking expertise. Such a policymaker would not be all that different from a standard representative politician, except the requirement of high-quality contributory and interactive expertise is strictly enforced in that the epistocrats

become representatives by virtue of their expertise rather than by a popular vote from the general public.

Effectively Focussing Task-Related Conflict

One way in which ineffective task-related disagreements can be overcome is by having competence-based trust, which is where members choose who to trust based on believing that their reasons and evidence are of a high quality (i.e., evidence and reasons that are competent, responsible, and reliable). As such, those trusted members are believed to be able to carry out their epistemic role within the group (Olsen, Parayitam, & Bao, 2007, pp. 200–201). Higher levels of competence-based trust will mean that diversity can be utilised better because members are more accepting of the expertise that others bring to the discussion (*Ibid.*, p. 200). Having competence-based trust does not mean that there will be no task-related conflict, but rather that the information given is taken more seriously and task-related conflict can be utilised effectively (*Ibid.*, p. 201). The need for competence-based trust is yet another reason why it is important to have a governance group comprised of those, and only those, with a high level of relevant expertise rather than lay citizens. Lay citizens are not going to be able to achieve the competence-based trust that is required to establish themselves as being an epistemically useful member of the group, and so it will be more difficult to overcome initial disagreements with lay citizens in the group.⁵⁹

Cognitive Diversity versus Cognitive Consensus

Another potential problem that can arise with task-related conflict is that there may be increased difficulty in finding a consensus and members are unwilling to move from their fundamental beliefs (Horwitz & Horwitz, 2007, pp. 991–992). To alleviate this problem, there needs to be a balance between cognitive diversity and cognitive consensus, where cognitive consensus is agreement regarding how key issues are

⁵⁹ The potential for this distrust in lay citizens to cause what is known as a testimonial epistemic injustice (“a deflated level of credibility to a speaker’s word” (Fricker, 2007, p. 1)) is discussed in section §6.4.

defined and conceptualized. Increasing cognitive diversity increases problem solving resources, but without an appropriate level of cognitive consensus, there will be more disagreements regarding how issues are to be understood⁶⁰ (Mohammed & Dumville, 2001, pp. 312–313). For complex tasks, such as emerging science and technology policy, the appropriate balance lies closer to the “more cognitive diversity” end of this balance to gain access to a greater range of perspectives and potential solutions (*Ibid.*, p. 313).

A concept that is closely related to the balance between cognitive diversity and cognitive consensus is task interdependence. Task interdependence is the degree to which a member of a group is reliant on other members of the group doing their job for the initial member to do their own job (van der Vegt & Janssen, 2003, p. 731). An example of a highly task-interdependent group is one concerned with emerging science and technology governance because there is a distinct amount of division of labour present. This division of labour is due to the complexity of the problem, which requires each member of an extended epistocracy to represent their own area of expertise, and each member relies heavily on other members of the group to represent their own expertise effectively to the discussion table. However, too much unique specialization can create misunderstandings about the policy issue at hand, so one does need to be careful to make sure that there is a significant level of overlap regarding a general understanding of the policy issues at hand to avoid co-ordination losses (Milliken & Martins, 1996, p. 415)(Evans & Carson, 2005, pp. 307,311).⁶¹

⁶⁰ The variation in the way that an issue is understood is due to selective perception, “a mental process in which information is interpreted after being filtered through a cognitive base” (Pelled, 1996, p. 622). In other words, the differences in groups members’ functional background or bio-demographic traits, for example, will make some information more important than others, and therefore will be used to interpret the issue at hand (though functional background will have a greater effect) (*Ibid.*, pp. 622-623).

⁶¹ This is technically referred to as transactive memory, which is the amount of expertise that is able to be effectively passed between members of the group. Groups that have high transactive memory are more easily able to accurately understand the expertise of others, which is to say there is enough effective overlap to allow for information transferral, co-ordination, and creative new knowledge emerging (Mitchell & Nicholas, 2006, p. 69). There is, however, a balancing act here. Those groups with very high overlap between expertise are not maximising the potential for cognitive diversity since

Furthermore, task interdependence has a U-shaped relationship with respect to performance, where higher task interdependence creates better performance than moderately task interdependent groups. High task interdependence forces group members to work together, while moderate task interdependence does not develop strong shared norms, information flow, nor agreement on the members roles within the group (Horwitz, 2005, p. 235). This is accentuated by the fact that for highly task interdependent tasks, members of the group have a greater opportunity to help or hinder the work of others in the group, so there needs to be high levels of perceived goal interdependence so the group works together (van der Vegt & Janssen, 2003, p. 732).

This greater performance with higher task-interdependence means that for an extended epistocracy, as opposed to general deliberative democracy, the specialisation of experts with respect to a given policy area allows for shared norms, information flow, and agreement on their roles. That is to say, as part of having a specialised nanomedicine epistocracy, for example, each member has expertise regarding their own respective fields they are representing and they also all have a high level of understanding regarding nanomedicine policy in general. In terms of the types of expertise given above, each member of the epistocracy will have their own contributory expertise regarding the field they are representing, but (at least) interactive expertise regarding policymaking within the particular policy area in question.

everyone in the group knows the same expertise, and hence there is very low transactive memory since no-one is learning anything new. While this makes discussion easier to understand, it is not very creative and there is a significant risk of groupthink occurring (Mohammed & Dumville, 2001, pp. 95-96,103-104). Conversely, those groups with very low overlap technically have high cognitive diversity, but it is much more difficult to access since transactive memory is low. Transactive memory also makes discussion more efficient since not a lot of time is used trying to explain concepts to members of the group, and there is a smoother synchronization of joint actions (*Ibid.*, p. 95) (DeChurch & Mesmer-Magnus, 2010, p. 47).

Summarising Task-Related Diversity Recommendations

In this section it was shown that having task-related diversity within a group can produce many benefits, such as managing the uncertainty and complexity of wicked problems like emerging science and technology policy by introducing a range of knowledge, skills, and abilities. However, with task-related diversity comes task-related conflict. This thesis recommends that for an extended epistocracy, the effects of initial disagreements and emotional conflict during heated discussion are to be alleviated by having epistocrats that are experienced in policy-making. This thesis further recommends that task-related conflict is focussed effectively by utilising the competence-based trust that comes from having members of the group with a high level of political competence. Lastly, to find a balance between cognitive diversity and cognitive consensus, as well as manage task interdependence, this thesis recommends that members of an extended epistocracy at least have interactional expertise with regard to a general understanding of the policy issue at hand to provide a common basic understanding between members, while also providing their own expertise from the field they represent so they can contribute to the task-related diversity of the group.

§5.5.4: Effect of Bio-Demographic Diversity on Group Performance

Having considered the effects of task-related diversity on group performance, this section now moves on to consider the other type of diversity mentioned by Horwitz and Horwitz, which is bio-demographic diversity. With the greater influence of globalisation in the workforce, along with the opening up of the workforce to a wider array of bio-demographic groups, the issue of bio-demographic diversity is becoming ever more important (Bourke, Garr, van Berkel, & Wong, 2017)(Horwitz & Horwitz, 2007, p. 988) (Pelled, 1996, pp. 615–616). Based on this greater importance placed on bio-demographic diversity, scholars have posited that this increase in diversity could be used to improve group performance. However, evidence supporting this improvement is limited (Roberge & van Dick, 2010, p. 295). This support is lacking

because, in general, the results are inconclusive. Some scholars have found that the effects of (visible) bio-demographic diversity are said to be mixed but leaning towards being negative for group performance by some (Milliken & Martins, 1996, pp. 408–409). Yet others have found that bio-demographic diversity has no meaningful overall effect on group performance (Horwitz & Horwitz, 2007, p. 1006) (Jehn, Northcraft, & Neale, 1999, p. 757).⁶²

If anything, the strongest claim one could make regarding the salience of (visible) bio-demographic diversity is its impact on emotional conflict, where emotional conflict “is individual-oriented disagreement arising from personal disaffection” (Amason & Sapienza, 1997, p. 495). Emotional conflict can be caused when bio-demographic diversity is salient, such that group members apply social categorization, i.e., they classify themselves and others into groups, and then form a similarity attraction, i.e., they favour those that are part of their own subgroup (and feel hostile and anxious around those in other subgroups) (Roberge & van Dick, 2010, p. 298) (Jehn, Northcraft, & Neale, 1999, p. 745) (Wegge, Roth, Neubach, Schmidt, & Kanfar, 2008, p. 1302) (Pelled, 1996, p. 622). Emotional conflict can also be caused by differences in educational level due to the frustration felt working with those of lesser ability (*Ibid.*). Once in a state of emotional conflict, group members are in a psychological state that

⁶² A few examples of the mixed effects of bio-demographic diversity are:

- Age diversity can improve performance in complex tasks, but decrease performance in interdependent tasks (Wegge, Roth, Neubach, Schmidt, & Kanfar, 2008, p. 1301) (Timmerman, 2000, p. 601). Age diversity also provides less emotional conflict, but more relationship conflict (*Ibid.*, 593–594). However, the general trend for age diversity is said by some to be negative (Wegge, Roth, Neubach, Schmidt, & Kanfar, 2008, p. 1302).
- Gender diversity has been demonstrated to create high levels of conflict, low cohesion, and increased turnover, and decreased performance, but has also been shown to promote innovation, creativity, and productivity, with no evidence for increased conflict (Wegge, Roth, Neubach, Schmidt, & Kanfar, 2008, pp. 1301,1303) (Lee & Farh, 2004, p. 139)
- Race diversity results are said by some to be very mixed with a tendency towards negative group performance (Timmerman, 2000, pp. 594–595) (Horwitz, 2005, p. 229). It could also be said that there is an overall insignificant to positive effect on innovation, along with increased cooperation, and better-quality brainstorming ideas, but not necessarily more ideas (Milliken & Martins, 1996, p. 406)(Cady & Valentine, 1999, pp. 734,745).

makes it difficult for them to process complex information, they become less willing to engage with the viewpoints of others, and valuable time is used up on unproductive arguments (*Ibid.*, p. 625). However, there is a very consistent tendency for the salience of bio-demographical features, and hence the negative influence of emotional conflict, to decay over time as members of the group get to know each other better; this decay in bio-demographic diversity salience means that task-related diversity becomes more salient, and the group can take better advantage of its task-related diversity (Milliken & Martins, 1996, pp. 407,409) (Cady & Valentine, 1999, pp. 744–745). As a corollary, (visible) bio-demographical features are going to be salient and negative for an ad-hoc group, such as those that would normally be constructed for public engagement (Lee & Farh, 2004, p. 140). As such, this is further evidence that a long-term council should be used for making policy decisions rather than the use of a short-term public engagement exercise.

§5.5.5: Managing Bio-Demographic Diversity

Given the above discussion of the effects of bio-demographic diversity, or the lack thereof, what role does bio-demographic diversity play within an extended epistocracy? Even though there may be an overall negative effect of bio-demographic diversity, if any, on group performance, this thesis makes no specific recommendations regarding the bio-demographic make-up of the members of an extended epistocracy. That is to say, there is no recommendation regarding the age/gender/race/etc. distribution within an extended epistocracy. This recommendation is made because an extended epistocracy is intended to stay together as a relatively cohesive unit, working together on many policy cases as much as possible. Furthermore, by not having any recommendation regarding bio-demographic features, choosing members for an extended epistocracy can instead be based on their high, relevant expertise. In doing so, bio-demographic features are not salient, and therefore have even less effect on the group's performance. By focussing on selecting members of an epistocracy based on their high, relevant expertise, there

is unlikely to be emotional conflict when the extended epistocracy is newly formed, and it is emotional conflict is less likely to be a problem as the members of the group stay together.⁶³

There are three overarching caveats that need to be mentioned here though. First, to say that no bio-demographic recommendation is made is not to say that important stakeholders should be ignored. If it is the case that the relevancy of a stakeholder is due to a bio-demographic property, then this bio-demographic property should be adequately represented in the epistocratic council. For example, if the policy case at hand concerns Māori health, then Māori should be represented in the extended epistocracy. However, as noted above, their representation is due to the relevancy of the expertise provided, i.e., their contribution to the task-related diversity, rather than simply because a bio-demographic quota must be filled. This is also in accordance with commitments to the Treaty of Waitangi that require meaningful consultation in areas that directly affect Māori (Munshi, Kurian, Morrison, & Morrison, 2016, p. 289).

This last point relates to the second caveat, which is that task-related diversity should take priority over bio-demographic diversity. The goal of an extended epistocracy is the creation of high-quality policy and task-related diversity has a strong, positive relationship to group performance, while bio-demographic diversity has an ambiguous, possibly negative, relationship. Therefore, task-related diversity should be prioritised to create a group focussed on creating high-quality policy, rather than focussing on bio-demographic diversity, which cannot be used to specifically aim for high-quality policy. If an organisation wishes to create an extended epistocracy with

⁶³ In contrast, public engagement involves experts and lay citizens meeting for a single policy case, as discussed in Chapter 2. For public engagement, lay citizens are chosen either randomly or because of their biodemographic properties, i.e., there is likely to be biodemographic diversity and/or biodemographic properties are salient, respectively. There is also a significant expertise gap between experts and lay citizens, by definition. Therefore, in the case of public engagement, both bio-demographic saliency and an expertise gap will accentuate affective conflict and the group is not together long enough to temper it.

a bio-demographic-based quota in mind, they are free to, as long as they meet the requirement that the task-related diversity of the group is capable of creating high-quality policy.

The third caveat is that a lack of recommendation regarding bio-demographic diversity does not mean that bio-demographic diversity should be intentionally avoided. Given that the purpose of an extended epistocracy is the creation of high-quality policy, purposefully limiting bio-demographic diversity also unnecessarily limits the pool of people from which task-related diversity can be drawn. Therefore, it is important to ensure that any bio-demographic diversity that is introduced provides a net increase in effective⁶⁴ task-related diversity so there is an overall expected net gain in policy quality by introducing that person into the governance council.

The lack of a specific recommendation regarding the bio-demographic make-up of an extended epistocracy, along with the second caveat in particular, is in line with suggestions from Horwitz and Horwitz (2007, pp. 1006–1007) and Olsen, Parayitam, and Bao (2007, p. 199). Both sets of authors recommend that the focus when creating a governance group is task-related diversity, and that one should not place too much importance on bio-demographic diversity. This recommendation is backed up by the finding that those diverse traits that are low in visibility and high in job-relatedness have favourable effects on cognitive task performance, and the effect is more prominent for cognitively complex tasks. The reverse is also true: diverse traits that are highly visible and low in job-relatedness are associated with lower cognitive task performance (Pelled, 1996, pp. 625–627). In other words, the insistence of public

⁶⁴ The use of the word “effective” here is a reminder that task-related diversity must be added in such a way that the KSAs provided are relevant, and that there needs to be an overlap in the general understanding of the policy issue.

engagement for governance in using demographic representation as a method for choosing participants is not recommended.

§5.6: What Role Should Experts Play in Policy-Making?

So far, this chapter has answered the following questions: who can be considered an expert, and how are the epistocrats to interact with each other? The next question to be answered concerns how the experts are to interact with the public, which will establish the social organisational dimension of an extended epistocracy, and its two subdimensions, the constitutional and process dimensions. These dimensions are established by understanding the role that experts are to fulfil using Pielke's four idealised types of expert. Out of these four types of expert, two of them (the issue advocate and the honest broker, to be defined shortly) are most beneficial for emerging science and technology policy. These two types of expert can also be thought of as ends of an epistocracy–democracy spectrum. Based on this spectrum, this section argues that epistocrats within an extended epistocracy should behave in a way that is closer to the epistocracy end of this spectrum by being advocates for the expertise groups they represent, and by being involved in the entire policy process rather than simply acting as policy guides as recommended by postpositivism. Furthermore, it is also recommended that an extended epistocracy can act at local, national, and international levels as is appropriate.

§5.6.1: The Role of Experts

To identify the role of experts within an extended epistocracy, Pielke's four different types of expert is used. While these expert types are idealised, they do give a good range of the possible roles of experts within policymaking. These four different types of experts can be organised based on their views of science and democracy, where the views of science are the linear model and the stakeholder model, and the views of democracy are the Madison view and the Schattschneider view (Pielke, 2007, p. 14).

The first important distinction to make is between the views of science. The linear model of science, which involves the pure scientist⁶⁵ and the science arbiter⁶⁶, is the view that there is a simple, linear path from basic science to applied science to policymaking. Along this path, scientific consensus acts as the necessary basis for political consensus and decision-making (*Ibid.*, pp. 12–13). The pure scientist and science arbiter, therefore, work best in cases with high value consensus and low uncertainty, which describes a tame policy problem, not a wicked problem (*Ibid.*, p. 19). They are therefore unsuitable for an extended epistocracy because an extended epistocracy is designed for wicked problem policymaking. In contrast, the stakeholder model (involving the issue advocate and the honest broker) is more appropriate for policymaking for wicked problems because it states there are feedback loops regarding the use of science and its products that should be taken seriously such that the public has a chance to respond to emerging science and technology and its associated chaos, complexity, and contradictions by giving the public the opportunity to deliberate. In doing so, the public bring their diverse worldviews into the discussion, which results in lower value consensus.

An extended epistocracy is therefore a type of stakeholder model with the issue advocate and honest broker as ideal types. To make the distinction between the two, the views of democracy are considered. The Madisonian view of democracy, held by the issue advocate, states that society would be best served if experts/elites aligned themselves with their preferred faction or interest group and make decisions accordingly. In contrast, the Schattschneider view of democracy, held by the honest

⁶⁵ The *pure scientist* is involved with the policy process as little as possible by giving policymakers relevant information and then allowing the policy makers to do what they will with that information; the focus of the pure scientist is their own research with little to no concern for how that research is used by others (Pielke, 2007, pp. 1-2,15).

⁶⁶ The *science arbiter* acts as a resource to answer factual questions the policy makers may have but will not tell the decision-makers which decision to make. This means the science arbiter will likely be on an advisory panel answering positive question that science can answer (in principle), but will refrain from engaging in normative questions and the political fray (Pielke, 2007, pp. 2,16).

broker, states that policymaking should be a competition between democratic interest groups with experts acting as resources to allow the public to make their decision (*Ibid.*, pp. 11–12). By combining this understanding of the Madisonian and Schattschneider views of democracy, along with the understanding of the stakeholder model given above, the two contending expert roles can be clarified as follows:

- An *issue advocate* is an expert who acts to convince policymakers to choose a narrow range of options, or even a particular option, and openly aligns themselves with group(s) proposing these option(s). The issue advocate is very invested in how scientific knowledge is used and makes sure science is actively engaged in the decision-making process (*Ibid.*, pp. 2,16–17).
- An *honest broker* is an expert who will act to expand (or at least clarify) the options that are available to the policy makers and act as a resource to help the policy makers make a decision. The honest broker acts to open up options for decision-makers to choose between, contrary to the issue advocate who seeks to narrow down options (*Ibid.*, pp. 3,17–18).

This thesis, then, presents the issue advocate and the honest broker as ends of an epistocracy–democracy spectrum, where the issue advocate is at the epistocratic end, and the honest broker is at the democratic end.⁶⁷ With this spectrum as a theoretical apparatus, an extended epistocracy can be placed somewhere along this spectrum by considering the constitution and process dimensions, the two subdimensions of the social organisational dimension, which are the subject of the next section.

⁶⁷ The epistocratic end could, in theory, be extended further to include a technocracy, but this option is ruled out based on the discussion of positivism in Chapter 2. The democratic end could also be extended further to a governance system that does not recognise the epistemic qualities of experts, such as pure proceduralism; however, this option was ruled out by the discussion regarding the epistemic turn in Chapter 2.

§5.6.2: The Constitutional and Process Dimensions of an Extended Epistocracy

The Constitutional Dimension

The first of these dimensions to be considered is the constitutional dimension, i.e., where along an epistocracy–democracy spectrum an extended epistocracy lies. This section proposes that the role of the expert in an extended epistocracy should be somewhere between that of an ideal issue advocate and an ideal honest broker but erring on the side of issue advocate. That is to say, an extended epistocracy is an example of a democratic epistocracy, which is primarily an epistocracy with democratic components. The reason why this middle ground is taken is, given that the political competence of the public is low, as discussed in Chapters 3 and 4, it is not desirable to leave emerging technology decision-making up to non-experts, even if guided, which is what an ideal honest broker would do. Furthermore, a pure issue advocate could act in a stubborn manner, arguing only for their given option without being open to deliberation from any other experts and/or they may take a technocratic approach to policymaking. In striking a balance between the issue advocate and the honest broker, the extended epistocracy, as a group, would act in an honest-broker-like manner in the sense that members would expand the range of policy options through discussion, but individually, the experts would act as issue advocates in that they would openly argue in accordance with the groups they were representing, and also make decisions on their behalf.

As a democratic epistocracy, the democratic component allows for members of the public that do not have trained expertise to still be able to take part in epistocratic deliberations. Furthermore, epistocrats can be held accountable by other experts that did not take part in the deliberative process, by giving the non-deliberating experts an opportunity to review the decision that was made and the reasoning for it and proposing changes wherever possible. Meanwhile, the epistocratic part of a democratic epistocracy reinforces the idea that those taking part will still need to have

an expert level understanding of the policy issue at hand. Therefore, in terms of the rights of experts and the general public, an extended epistocracy would still allow for the standard democratic set of social/political rights with the only difference being that when it comes to emerging science and technology policymaking, only those with a high-level of relevant expertise are eligible to take part in the decision-making process. Even in the case of emerging science and technology policy, the social/political rights of those outside the extended epistocracy are represented by their respective representative expert. Furthermore, restricting emerging technology decision-making to an epistocratic council does not mean that someone who is currently an average citizen is permanently cut off from the policymaking process. Everyone has the right to earn the ability to take part in emerging technology policymaking; however, for epistemic and logistical reasons (only a few people can genuinely deliberate at a time), not all those that may wish to take part are able to.

The other aspect of the constitutional dimension is the level at which the experts are to act, namely the local, national, or international level. Given that the relevance of the expertise is a crucial component when it comes to choosing participants, an extended epistocracy is organised for a particular policy area, rather than a level of action. That is to say, an extended epistocracy can be created at any of the three levels of action as long as those taking part are chosen appropriately for the policy issue at hand.

The Process Dimension

The other dimension to be considered in this section is the process dimension, which is how much of the process the experts are a part of. Given that an extended epistocracy is an epistocracy, then by definition, experts are involved in the entire policy process. More specifically, the experts in an extended epistocracy do not act as advisors and leave the decision-making to the general public as an honest broker would do, or even leave decision-making up to politicians. Instead, the experts in an extended epistocracy not only provide the advice and potential options, but they are

also the decision-makers. This is not to say that the general public cannot participate at some point since they are able to have discussions with their representative expert to make their viewpoints known. However, the general public's views are filtered through the expertise of their expert representative. The expert representative uses their best judgement to decide what views of the general public are to be brought to the deliberating table, whether in a complete or modified form, or not at all.

§5.7: The Selection Process

Now that it is known who is eligible to be chosen for an extended epistocracy, and how those chosen are to interact, both with each other and with the public at large, it is possible to present a method for selecting the members of an extended epistocracy. That is to say, this section answers the final of the four main questions proposed for this chapter, which is: how are the experts chosen? In answering this question, the selection and cognitive dimensions of the extended epistocracy are established. This section will argue that in terms of the selection method used, members of an extended epistocracy will be chosen by selecting them from a database of known experts in the field, where the database is built up and maintained for quality over time. In terms of the cognitive dimension of an extended epistocracy, i.e., the incentives the epistocrats have to perform at their epistemic best, this thesis recommends that the selection process is open and transparent to the public so the public know the selection process is fair and balanced rather than biased by monetary incentives.

Selection Dimension

When deciding on a method for selecting members of an epistocracy, it must be remembered that members must have a high level of relevant knowledge, and the group's individual knowledge, skills, and abilities must fit together such that there is significant overlap of the general understanding of the policy area, but there is enough task-related diversity to produce high quality results. Due to these restrictions, there

are selection methods that cannot be used. For example, an extended epistocracy cannot be selected by allowing members of the public to put their names forward for selection as an epistocratic council member and then have the general public vote for the epistocrats. Even if there were a screening process that ensured that those who were on the voting ballot had a high level of relevant expertise, a democratic vote for the council at large would still not guarantee that the task-related diversity was managed appropriately; it may be the case that if this selection method is used, then there is little control over whether a field of expertise becomes over- or under-represented.

To help alleviate this problem of over/under-representation of a particular area of expertise, members of the extended epistocracy could be selected to fill specific roles within the extended epistocracy. For example, if an extended epistocracy was considering the implementation of a nanodevice in the brain, one position on the council could be set aside for a nanoscientist, another one for a bioethicist, another for a representative of those with disabilities, and so on. Selecting members for an extended epistocracy in this manner means that an extended epistocracy is largely going to be an invited form of engagement since highly technical and hypothetical future-oriented applications are unlikely to create uninvited forms of engagement due to a lack of awareness from the general public and pressing problems to form a movement around (Delgado, Kjølberg, & Wickson, 2011, pp. 834,839). Lay citizens also tend not to have access to the information, resources, or power that are required to hold engagement exercises for emerging science and technology issues, even if they were aware of them (Powell & Colin, 2008, p. 129).

However, two important issues that need to be considered when designing an invited form of engagement are knowing who has what expertise and making the process trustworthy. In response to the first of these issues, a database can be established that contains information on the expertise of academics and other researchers willing to

take part in deliberation, related research institutes and advocacy groups, and highly engaged local citizens that have presented themselves as willing to take part, among others. Initially, this database will be small, and the time taken to assess the quality of expertise of those that are in the database will be significant, but over time and with many calls for applicants, the database will grow to include those that will be selected for roles on an extended epistocracy, and those that are available as peer reviewers for the policy decisions that are made.

Cognitive Dimension

The second issue, which is making the selection process trustworthy, falls under the cognitive dimension of an extended epistocracy, which relates to the incentives the members of the extended epistocracy have to perform at their epistemic best. The call for applicants mentioned above is done in the same way that job vacancy positions are advertised. These advertisements will primarily be sent to those on the database of potential epistocrats, but they will also be open to the public such that members of the public are able to learn what the selection criteria are and therefore know that the process is fair and balanced rather than unduly swayed by monetary interests. Furthermore, because membership on an extended epistocracy is a full-time, salaried position, members are compensated for their efforts by way of the salary and are forbidden from receiving any bribes and engaging in any bargaining that would disincentivise them from performing at their epistemic best. By having this open process and an enforcement of incentives, there is the potential for building public trust (Paddock, 2010, p. 277). Furthermore, social activities, such as public engagement and policymaking, is often seen as a social add-on by experts rather than an important component of their work (Katz, Solomon, Mee, & Lovel, 2009, pp. 540–541). By making a position on an extended epistocracy a salaried position, this provides an incentive to take part in social activities as opposed to focussing solely on their research.

§5.8: Summary: What Does an Extended Epistocracy Look Like?

The purpose of this chapter was the creation of an extended epistocracy that could be used as a governance system for emerging technology policy. This extended epistocracy builds upon Holst's description of an epistocracy by adding two extra dimensions, which were the structure dimension that established how the epistocrats were to work together to effectively utilise diversity, and a selection dimension that outlines the process by which members are selected. Given the amount of detail and the number of recommendations that were made regarding the establishment of an extended epistocracy, it is useful to provide a reference summary of all the recommendations that were made in one place, which is the aim of this section.

First and foremost, an extended epistocracy is made up of experts who can potentially be chosen from those citizens that cover a broad range of relevant, high-quality expertise. More specifically, potential epistocrats must be able to display high-quality decision-making skills, be socially recognised as an expert, and display a high level of professionalism. The breadth of expertise can come from trained experts, advocates, and local experts such that the range of expertise is extended beyond the purely technical expertise that was used under technocratic governance. There are also citizens, called lay citizens, that do not have relevant expertise, but this is not to say that lay citizens will never have to opportunity to take part in politics. All citizens will find themselves in the "lay citizen" category for the vast majority of policy cases, but if lay citizens wish to take part in a particular extended epistocracy, they are required to gain relevant expertise for that particular policy area.

Once this pool of potential epistocratic candidates is identified, they need to be placed together in a team that best optimises their diversity. Diversity was broken into task-related diversity and bio-demographic diversity. Task-related diversity allows for a

greater range of knowledge, skills, and abilities, which is crucial for complex decision-making, e.g., emerging technology policymaking, but task-related diversity can also cause inefficient task-related conflict. To ensure that task-related conflict is effective, an extended epistocracy:

- Is intended to stay together for a long period of time so that initial disagreements are less disruptive.
- Preferably uses experienced epistocrats because it is more likely that task-related conflict will be taken as constructive criticism rather than as a personal attack.
- Has deliberation between experts because the greater competence-based trust means that task-related diversity can be more efficiently utilised.

An extended epistocracy also has a significant level of common understanding regarding the general issues of the policy area in question between members, which means that information flows more smoothly. However, each member also needs to provide their own unique specialised knowledge, skills, and abilities such that the group can take advantage of task-related diversity.

Contrary to the more definite empirical findings regarding task-related diversity, the effects of bio-demographic diversity are ambiguous at best, and so no specific recommendation is made regarding any bio-demographic quota. However, general recommendations were made such that:

- Again, the group should stay together for a long time, but this time it is to reduce affective conflict.
- Again, experts are used, but this time it is to decrease frustration from working with low expertise members.
- Experts are chosen specifically for their relevant expertise rather than to fill a bio-demographic quota. Emphasising relevant expertise means that bio-demographic features become less salient and hence cause less affective conflict.

Given the wicked nature of emerging technology policy problems, the role of experts is placed on an issue advocate–honest broker spectrum, i.e., an epistocracy–democracy spectrum. An extended epistocracy falls on the democratic epistocracy side of this spectrum. The epistocrats are to act closer to honest brokers as a group such that a range of policy options are presented, but closer to issue advocates as individuals to openly argue for a position and to represent their respective expertise groups. Epistocrats are also intended to be part of the entire policymaking process, from possibly meeting those in the public they represent, to proposing options, to making the policy decisions, and overseeing implementation and monitoring. That is to say, the epistocrats are not just advisors, as would be the case in public engagement for governance; rather, they are in charge of the entire process.

Lastly, members of an epistocracy are chosen to fill specific expertise roles, where knowledge of those with experience in each expertise role is taken from a database that is developed by collecting trained experts, advocates, and local experts who are willing to take part in the deliberative and/or peer review process. This collection is done by advertising open vacancies such that members of the public are aware of the criteria used for selection, and that the general public can be more assured that the epistocrats are performing at their epistemic best rather than being bribed or engaging in bargaining.

With these aforementioned recommendations in mind, it is important to note that an extended epistocracy is designed for the purpose of creating epistemically high-quality policy, i.e., it does not explicitly focus on achieving democratic legitimacy (though democratic features have been included with democratic legitimacy in mind). However, democratic legitimacy is crucial to the success of a governance system, as discussed in Chapter 2, so the next chapter is dedicated to proposing potential future

research that could be conducted to establish whether an extended epistocracy has the legitimacy required to be an effective governance system.

Chapter Six: Future Research – Legitimacy of an Extended Epistocracy

§6.1: Chapter Overview

One of the primary aims for this thesis has been to investigate the epistemic quality of governance by considering both an epistemic deliberative democracy and an epistocracy. To this end, Chapters 3 and 4 strongly suggest that an epistocracy is likely to be the best performing governance system from an epistemic perspective. However, there is more to good governance than being epistemically successful. As discussed in Chapter 2, a governance regime must also be considered legitimate, where legitimacy is taken to mean “ethically permitted to use acts and threats of coercive enforcement” (Estlund, 2008, p. 41). Regarding democratic legitimacy, any proponent of an extended epistocracy (or epistocracy in general) have a serious objection to answer coming from deliberative democrats, among others. The purpose of this chapter, then, is to investigate how, if at all, this legitimacy objection may be answered. It is beyond the scope of this chapter to consider the legitimacy of an extended epistocracy in detail such that a conclusion can be formed regarding any legitimacy objections. This chapter does, however, continue the overall interdisciplinary approach of the thesis by proposing three areas of future research, along with suggested questions to guide that research, that would help in considering any legitimacy objections. These three areas of research are the literature considering the expert/boss fallacy, the epistemic injustice literature, and the collective ethics literature as applied to normative consent, which are to be explained shortly.

These three areas of research are presented in the form of an argument structure, which is laid out in detail in Section §6.2. The general outline for this chapter’s argument structure is established by comparing an extended epistocracy with a competing governance system, namely epistemic proceduralism, which aims for a

procedure that “is better than random [regarding decision-making capabilities] and is epistemically the best among those that are generally acceptable in the way that political legitimacy allows” (Estlund, 2008, p. 8). Extended epistocracy and epistemic proceduralism are compared along two legitimacy dimensions, namely democratic and epistemic, where democratic legitimacy comes from the democratic procedure as described in Chapter 2. In contrast to democratic legitimacy, what this thesis calls “epistemic legitimacy” is the legitimacy generated by a governance system’s tendency towards producing high quality policy; the stronger the tendency towards high quality policy, the higher the epistemic legitimacy generated. This classification of legitimacy along these two dimensions is important because extended epistocracy and epistemic proceduralism are effectively mirrors of one another; an extended epistocracy emphasises the epistemic dimensions while sacrificing the democratic dimension, and vice versa for epistemic proceduralism. Therefore, epistemic proceduralism will generate more democratic legitimacy than an extended epistocracy, and an extended epistocracy will generate more epistemic legitimacy than epistemic proceduralism. The question then becomes: who generates more legitimacy overall?

However, before the comparison between the two competing governance systems can truly begin, a legitimacy barrier within epistemic proceduralism must be overcome, which is the task of Section §6.3. This legitimacy barrier occurs because governance systems that are epistemically superior to epistemic proceduralism are allegedly too controversial to be implemented (Estlund, 2008, p. 42). Epistemically superior governance systems are considered too controversial because they commit what Estlund calls the expert/boss fallacy, which states that it is wrong to give someone political authority simply on the basis that they know better (*Ibid.*, p. 40); to give the better knowers political authority by allowing them to rule would not be acceptable to all reasonable points of view (*Ibid.*, p. 33). More specifically, three areas in which an epistocracy would be unacceptable are considered: citizens cannot agree on experts,

experts are demographically biased, and the extra authority an epistocracy places over citizens. These three claims of violated acceptability are evaluated in turn and potential future research is suggested.

On the assumption that the above legitimacy barrier can be overcome, this chapter includes two more major sections to discuss the democratic and epistemic legitimacy dimensions, respectively. Section §6.4 discusses how an extended epistocracy might compare to epistemic proceduralism along the democratic legitimacy dimension by considering epistemic injustice, which occurs when someone is not given the respect they deserve in their capacity as a knower (Fricker, 2007, p. 20). Is it the case that a democratic form of governance is required to avoid epistemic injustice, or is emerging science and technology policy an exception that means it is acceptable to not treat someone as a knower? In response to this question, four areas are considered that identify potential ways an extended epistocracy might commit an epistemic injustice, which are the forms of epistemic power, the types of injustice they can commit, the treating of knowers as subjects/objects, and upstream engagement as a potential form of epistemic injustice.

In contrast, Section §6.5 seeks to compare the epistemic legitimacy of both epistemic proceduralism and an extended epistocracy. This section does this by applying collective ethics to the concept of normative consent. Estlund claims that normative consent is said to be given to an authority figure by default if it would be immoral not to give consent to that authority figure if asked (Estlund, 2008, p. 10). For the purpose of this section, the authority in question is an extended epistocracy, and the immorality in question is the collective harm caused by bad voting. That is to say, does the higher epistemic legitimacy of an extended epistocracy, and hence the lower immorality caused by collective harm justify the authority of an extended epistocracy to the individual voter, or is the minimal epistemic standard of epistemic proceduralism enough to provide the epistemic legitimacy it needs? To answer this

question, three aspects of the individual voter's response to the collective harm are evaluated by considering the collective ethics literature, namely: what makes them responsible, what have they done wrong, and what is their responsibility?

§6.2: The Proposed General Structure for the Legitimacy Argument

§6.2.1: Using Epistemic Proceduralism as a Benchmark Governance System

The method this thesis proposes to evaluate the legitimacy of an extended epistocracy is to compare it with a competing governance system. To this end, Estlund's "epistemic proceduralism" is chosen because it is the mirror governance system to an extended epistocracy along a problem of expertise spectrum. To explain what is meant by the phrase "mirror governance system", recall that the problem of expertise states that there is a tension between expertise and democratic legitimacy. In Chapters 3 to 5, it was demonstrated that appropriate use of expertise leads to a greater tendency towards higher policy quality. This tendency towards higher policy quality generates what this thesis calls "epistemic legitimacy"; the stronger the tendency towards high quality policy, the higher the epistemic legitimacy generated. Therefore, when the problem of expertise claims there is a tension between expertise and democratic legitimacy, this can be interpreted as a trade-off between epistemic legitimacy and democratic legitimacy. Regarding this trade-off, an extended epistocracy aims for a high level of epistemic legitimacy while also aiming for at least a minimum acceptable level of democratic legitimacy. On the other hand, epistemic proceduralism uses a democratic process because it "is better than random [regarding decision-making capabilities] and is epistemically the best among those that are generally acceptable in the way that political legitimacy allows" (Estlund, 2008, p. 8). In other words, epistemic proceduralism is the mirror of an extended epistocracy in that epistemic proceduralism aims for a high level of democratic legitimacy while also aiming for at least an acceptable minimum level of epistemic legitimacy. By considering a "mirror"

governance system, this creates what could be considered two ends of a reasonable spectrum in that both systems seek to maximise one of either democratic or epistemic legitimacy while maintaining the other form of legitimacy at a reasonable level. It may be the case that the most appropriate form of governance lies somewhere between epistemic proceduralism and extended epistocracy; finding this balance is the purpose of further research.

There are two more features of epistemic proceduralism that are important to note. First, epistemic proceduralism is “generally acceptable in the way that political legitimacy allows”, which means that epistemic proceduralism is said to generate democratic legitimacy in the way required by deliberative democracy, by definition.⁶⁸ Second, the “epistemic” part of “epistemic proceduralism” is to signify that the decisions made by epistemic proceduralism generate further legitimacy because the decisions are said to be made by a process with a tendency to produce correct decisions (*Ibid.*).⁶⁹

§6.2.2: Argument Structure for Extended Epistocracy’s Legitimacy

Three important features that can be drawn from the above brief overview of epistemic proceduralism are as follows: the concept of a trade-off between democratic

⁶⁸ An important question concerns whether epistemic proceduralism can be considered legitimate by default, but for the purpose of having a reference point in the upcoming general argument in this chapter, the democratic legitimacy of epistemic proceduralism can be taken for granted.

⁶⁹ This focus on the epistemic quality of the process is in contrast with requiring that the decisions themselves are correct for them to be considered legitimate as correctness theories would suggest (Estlund, 2008, pp. 98–99). Extended epistocracy is also a proceduralist governance system in this epistemic sense since it aims to be a governance system that strongly tends towards high-quality policy decisions through an epistemically superior policy process; there is no claim that decisions must be correct to be legitimate within an extended epistocracy. Because an extended epistocracy is proceduralist in this epistemic sense, it avoids the deference challenge posed by Estlund. Under both an extended epistocracy and epistemic proceduralism, citizens do not have to defer their judgement to experts if the citizen disagrees with the expert. It is possible for the citizen to disagree with the policy result, while still maintaining that the policy process was reasonable, and it is the reasonableness of the policy process that is required for epistemic legitimacy under both governance systems (*Ibid.*, pp. 102–103,104).

legitimacy and epistemic legitimacy of the procedure, epistemic proceduralism is legitimate by definition, and epistemic legitimacy can be generated by the epistemic quality of the procedure. By considering these three features, a general argument can be made that can compare the relative legitimacies of an extended epistocracy and epistemic proceduralism. The general argument runs as follows. Given its improved decision-making capabilities, an extended epistocracy generates more epistemic legitimacy than epistemic proceduralism. However, epistemic proceduralism generates more democratic legitimacy by satisfying the tenets of deliberative democracy. However, if an extended epistocracy gains at least as much epistemic legitimacy as it loses in democratic legitimacy, relative to epistemic proceduralism, and epistemic proceduralism is already legitimate by definition, then an extended epistocracy is also legitimate.⁷⁰ In other words, a sufficient condition for extended epistocracy's legitimacy is if an extended epistocracy's total legitimacy (epistemic plus democratic legitimacy) is at least as much as epistemic proceduralism's total legitimacy, then an extended epistocracy is also legitimate, since epistemic proceduralism is already considered legitimate by default.⁷¹ Such an argument, then, requires a comparison with epistemic proceduralism with respect to two legitimacy dimensions, namely the democratic dimension and the epistemic dimension.⁷²

⁷⁰ Such a structure for an argument could be considered problematic because one cannot measure legitimacy. Legitimacy comparisons that are made are qualitative, and therefore vague by necessity, but provides an interesting avenue for future research nonetheless.

⁷¹ The converse may not necessarily be true since it might be the case that epistemic proceduralism is more legitimate than necessary, and so an extended epistocracy could have a lower total legitimacy than epistemic proceduralism and still have a satisfactory level of legitimacy. It might also be the case that both types of legitimacy need to be weighted differently such that any drop in democratic legitimacy will need to be compensated by an even greater increase in epistemic legitimacy to be considered a "fair trade". As far as the structure of the general argument of this chapter goes, asking for a breakeven sufficient condition with equal weights can be considered a middle ground between these two potential modifications to the general argument, though the modifications can be finessed in future research.

⁷² Another way of framing this dimensional comparison comes from Brennan, who recognises that both democracy and epistocracy may be unjust, but argues that a moderate epistocracy would be a moral improvement over a democracy (Brennan, 2011b, p. 701). This alleged moral improvement is due to a moderate epistocracy's unequal distribution of authority being less unjust than democracy's decisions being made in an incompetent and morally unreasonable way (*Ibid.*, p. 704); the better policy quality under a moderate epistocracy can compensate for the injustice the unequal authority distribution (*Ibid.*,

To evaluate these two dimensions of legitimacy, three potential areas of future research are considered. These areas are the expert/boss fallacy, epistemic injustice, and collective ethics as applied to normative consent. The expert/boss fallacy is considered to examine whether what this chapter terms a “legitimacy barrier” (defined shortly) can be overcome. On the assumption that this legitimacy barrier can be overcome, epistemic injustice and collective ethics are considered to examine the relative levels of democratic and epistemic legitimacy, respectively. Each of these areas will be discussed in turn.

§6.3: Overcoming the Legitimacy Barrier – The Expert/Boss Fallacy

Recall that epistemic proceduralism is a governance system that “is epistemically the best *among those that are generally acceptable* in the way that political legitimacy allows” (Estlund, 2008, p. 8). As part of this definition, general acceptability (to be defined shortly) creates an epistemic legitimacy barrier. This legitimacy barrier is said to exist because if there are methods for creating policy that are epistemically better than democracy, these methods are too controversial to be generally acceptable (*Ibid.*, p. 42). If this barrier remains, then there is no need to compare the democratic and epistemic legitimacy dimensions of an extended epistocracy because the barrier rules an extended epistocracy as being illegitimate by default. Therefore, the first step in the proposed argument for this chapter is to investigate whether this legitimacy barrier can be overcome, which requires an analysis of the expert/boss fallacy to establish if it is indeed a fallacy.

p. 717). In either case, i.e., in both Brennan’s framing and the framing used for this thesis, it is the combined legitimacy that decides a governance system’s true legitimacy.

§6.3.1: What Is the Expert/Boss Fallacy?

Estlund argues that the argument for epistocracy⁷³ fails because it commits the expert/boss fallacy. The expert/boss fallacy occurs when making the allegedly illicit inference from “S would rule better” to “S is a legitimate or authoritative ruler” (*Ibid.*, p. 40). Estlund claims that the expert/boss fallacy is a fallacy because the just-stated inference violates the *qualified acceptability requirement* (QAR):

No one has authority or legitimate coercive power over another without a justification that could be accepted by all qualified points of view (*Ibid.*, p. 33).

Estlund does not give a general account of what “qualified” means, though he does use “qualified” as if it were a synonym for “reasonable”. He does not give a general account because he claims that his argument does not require one, though he does recognise it would be good to have one for some purposes (*Ibid.*, pp. 63–64). As a result, this chapter recommends that finding a general account of “qualified” would be worthwhile future research. Without further research in this area, it seems that the QAR succeeds in blocking anything the arguer wishes to be blocked, but only because “qualified” is such a flexible concept as it currently stands. Nevertheless, the QAR is a core component of the expert/boss fallacy, and as such, the QAR needs to be

⁷³ For Estlund, the argument for epistocracy is in three steps:

1. *The Truth Tenet*: there are true (at least in the minimal sense) procedure-independent normative standards by which political decisions ought to be judged.
2. *The Knowledge Tenet*: some (relatively few) people know those normative standards better than others.
3. *The Authority Tenet*: The normative political knowledge of those who know better is a warrant for their having political authority over others (Estlund, 2008, p. 30).

In other words, if there are standards by which to judge better political decisions, then those who know those standards best should make those decisions (if policy quality is the aim of deliberation), and the decision-makers have authority (*Ibid.*, p. 206) (Quong, 2010, pp. 36–37). Estlund is willing to grant the truth and knowledge tenets, and even assume that experts in an epistocracy act in good faith, and so the focus of Estlund’s resistance against epistocracy is the authority tenet (*Ibid.*, pp. 31–33, 40, 206–207). Because Estlund grants the truth and knowledge tenets, this chapter also takes them for granted and focuses on the authority tenet as well.

examined to establish whether the legitimacy barrier can be overcome.⁷⁴ As such, this section considers three ways in which an epistocracy is said to violate the QAR, namely that citizens cannot agree on experts, the demographic objection, and that an epistocracy involves unjustified extra authority. These three factors are discussed in turn and potential future research is presented accordingly.

§6.3.2: Cannot Agree on Experts

The first way in which epistocracy is claimed to violate the QAR is due to there being qualified points of view that disagree on who count as experts, particularly political experts; that is to say, anybody who is put forward as an expert could potentially be the subject of qualified controversy (Estlund, 2008, p. 36) (Quong, 2010, p. 37). Furthermore, an epistocracy gives some citizens rather than others extra political authority based on their level of political wisdom, which is a form of what Estlund calls an “invidious comparison”. Estlund argues that no invidious comparisons are able to satisfy the QAR, and therefore, the case for epistocracy is said to be defeated (Estlund, 2008, pp. 35–36). To analyse this line of reasoning from Estlund as part of future research, some questions that could be considered are:

1. Why should “acceptability” matter in the QAR, and what exactly is it about expertise that needs to be accepted? Do qualified citizens need to agree on the level of expertise, the relevance of the expertise, or something else?
2. Related to (1), if someone is unable to recognise true and/or relevant expertise, why are they able to be the crucial objection to an epistocracy? Estlund does note that crazy or vicious objections do not count as being qualified (*Ibid.*, p. 4), but what about well-meaning, yet ignorant, objections? Could it not be the actual quality of the policy process as evaluated by experts that is important, not the perceived quality as seen by an average citizen? This last question potentially begs the question in favour of an

⁷⁴ There may be other ways in which the expert/boss fallacy could still be considered a fallacy, but given that the QAR is such an important tenet in at least liberal political philosophy, the strength of democracy’s legitimacy depends strongly on the strength of the QAR (Brennan, 2011b, p. 714).

epistocracy in that it appears to assume that only those with high levels of relevant expertise are considered qualified enough to be objectors, but the reverse potentially begs the question in favour of democracy.

3. If, as Quong (2010, p. 37) notes, there are no criteria for political expertise (as opposed to scientific expertise), can a set of criteria be developed capable of identifying political or moral expertise? Would it be possible to test this set of criteria in the form of an exam to establish who can take part?
4. Is it possible to separate the qualified acceptability requirement and invidious comparisons, as Lippert-Rasmussen (2012, p. 245) suggests? In other words, could it be the case that political competence is the acceptable exception to equal universal suffrage? If so, what level of political competence would be considered an acceptable exception?⁷⁵

§6.3.3: The Demographic Objection

The second way in which an epistocracy is claimed to violate the QAR relates to Estlund's argument against a form of epistocracy presented by John Stuart Mill known as a scholocracy. In a scholocracy, every adult receives at least one vote, while the educated receive multiple votes to counteract the fact the educated were likely to be outnumbered by the lesser educated voters (Estlund, 2008, p. 207). Estlund argues against this scholocracy with what he calls the demographic objection, which states that:

⁷⁵ For example, different levels of political competence to be considered could be basic literacy (Estlund, 2008, pp. 217–218), or a minimum voting age (e.g., see (Cook, 2013) and (Lecce, 2009)), or a policy area specific test that tests both factual knowledge and reasoning ability, such as a bioethics test found in (Schivone, Mameli, & Boniolo, 2015, pp. 278–279). By considering and comparing increasingly demanding criteria for political competence, it may be possible to determine at what level and why political competence becomes too unjust as a method for selecting voters/deliberators. Alternatively, one could consider which levels of political competence are not suitable for political decision making and deny authority to those that hold such low levels of political competence. It may be easier to identify those that are definitely not competent (e.g., those that hold extremely bigoted views) and excluding them from the political process rather than trying to find those that are experts and only including the experts. A benefit of this methods is that it does not rely on the authority tent of the epistocracy argument, and so does not commit any expert/boss fallacy (Brennan, 2011b, p. 713).

The educated portion of the populace may disproportionately have epistemically damaging features that countervail the admitted epistemic benefits of education⁷⁶ (*Ibid.*, p. 215).

Given that greater education is biased towards privileged groups in society, the demographic objection states that giving these privileged groups more voting power acts to compound the disadvantaged status of the lesser educated groups (*Ibid.*). Estlund suggests that even if those citizens with multiple votes were arranged to be statistically representative of empirically measurable features of the public as a whole, there might always be unmeasurable properties of the public that harm the quality of the vote. Therefore, a rule by scholocracy should not take place (*Ibid.*, pp. 215,216–217). Estlund suggests that given the history of ruling arrangements that privilege some over others, such as voting tests in the United States to disqualify African-American voters, the just stated objection to scholocracy seems to be a qualified objection.⁷⁷

An important question that arises regarding the demographic objection is that it appears to make many moves from possible to actual. For example, the arguments go from “possibly bigoted educated voters” to “actually bigoted educated voters”, or “possibly epistemically harmful” to “actually epistemically harmful”. While this is consistent with the form of the QAR that Estlund proposes where “any *possible* qualified objection is a justification defeater” (*Ibid.*, p. 47)(emphasis mine), this method of argumentation seems to allow for too many objections. It would seem odd to deny a form of governance because a possible reasonable objection existed, while no-one

⁷⁶ Estlund contends that a scholocracy is the most formidable form of epistocracy since it is very reasonable to believe that a population that has received a good political education will tend to rule more wisely (Estlund, 2008, pp. 207,211,212). Furthermore, marks of education come in the form of degrees and are therefore at least easily identifiable (if potentially crude). Both of these factors (good political education and the marks of education) help towards satisfying the qualified acceptability requirement (*Ibid.*, p. 210).

⁷⁷ Estlund does, however, leave the question of whether the demographic objection is a qualified objection open for debate. In general, though, Estlund believes that any educational criterion for extra voting power is open to qualified objection; this even includes basic literacy because this would harm the epistemic outcome of voting because the viewpoints of illiterate people would be left out of discussion (Estlund, 2008, p. 219).

actually held that objection.⁷⁸ Furthermore, the structure of the QAR means that even if Estlund's argument against a scholocracy is not valid, democrats can simply look for another argument. For an epistocracy to be considered legitimate, epistocrats would need to show that every possible objection to an epistocracy is unreasonable, which is a much more difficult task than finding one objection that seems reasonable. Applied in this manner, the QAR stacks the argument heavily in favour of equal universal suffrage.

In contrast, and concurring with Lippert-Rasmussen, any balance between bigotry and epistemic quality is an open empirical question that needs to be investigated, focussing on actual properties of citizens rather than simply considering possible objections (Lippert-Rasmussen, 2012, p. 251). Therefore, if there were to be a fair comparison of the reasonableness of both an extended epistocracy and epistemic proceduralism, this chapter recommends that empirical evidence is required that shows whether either governance system is epistemically harmful due to bigotry, or at least one has a better overall balance between bigotry and epistemic harm than its competitors.⁷⁹

§6.3.4: Epistocracy's Extra Authority

Lastly, the third way in which an epistocracy is claimed to violate the QAR is due to the extra authority an epistocracy is claimed to have. Estlund believes that anarchy

⁷⁸ This is especially odd if the argument can be applied against itself. An interesting property of the demographic objection is that the argument can simply be turned around to say that democracy suffers from a qualified objection: it is possible that equal universal suffrage carries a net bigoted bias or causes epistemic harm. As a result, democracy fails the qualified acceptability requirement, and hence is not legitimate. This is a specific case of a more general argument: could there not be qualified disagreement all around? (Estlund, 2008, pp. 36–37) Furthermore, Brennan (2011b, pp. 720–721) suggests that his arguments against democracy count as a qualified objection against democracy. Similarly, this thesis, even if wrong, could still be considered a qualified objection to democracy. The point being that democracy should not be exempt from the qualified acceptability requirement, and that democracy does not seem to satisfy the qualified acceptability requirement.

⁷⁹ It may even be the case, as Lippert-Rasmussen suggests, that higher education also leads to lower levels of bigotry such that a scholocracy actually alleviates both the epistemic and legitimacy issues that Estlund is concerned about (Lippert-Rasmussen, 2012, p. 249).

should be considered the default position of “authority”, and any addition of authority or coercive power to a governance system must satisfy the QAR (Estlund, 2008, p. 37). According to Estlund, the move from democracy to scholocracy is problematic because there is an extra, unjustified, element of authority. In a democracy, those that lose an election may be under the authority of those that win, but this ruling relationship only exists for that election. Citizens are otherwise considered equals. Under a scholocracy, the authority relationship between the educated and less educated is such that there is said to be a permanent inequality of citizens (*Ibid.*, p. 37).

For the purpose of establishing future research avenues, it could be argued that the authority relations within an extended epistocracy are more equal than a scholocracy. There is a lower difference in authority on average (henceforth, “differential authority”) in an extended epistocracy than in a scholocracy in the sense that everyone is considered a lay citizen for the vast majority of policy cases with a spike in authority only appearing in a small number of policy cases; a citizen only leaves lay citizen status if they have high-quality relevant expertise for a certain policy issue. On the other hand, it could be argued that it is more difficult to change one’s lay citizen status, i.e., citizens have little of what this thesis terms “authority mobility”, in an extended epistocracy for a given policy issue than in a scholocracy since the required standard for expertise is higher in an extended epistocracy. Therefore, an important future research question is: Is the reduction in differential authority for an extended epistocracy relative to a scholocracy enough to make an extended epistocracy justifiable? Or does the lower authority mobility found in an extended epistocracy still make an extended epistocracy unjustifiable? One way to proceed regarding these questions is to consider Lippert-Rasmussen’s suggestions that issue-sensitive (akin to differential authority) and time-sensitive (akin to authority mobility) invidious comparisons may not be as controversial as other forms of epistocracy (Lippert-Rasmussen, 2012, pp. 247–248).

Another potentially fruitful area of future research comes from considering the defence of equal voting, or something close to it. Estlund does not officially attempt a defence of equal voting since he sees it to be a questionable ideal (Estlund, 2008, p. 222). Quong, however, argues in favour of equal voting by saying that even though democracy gives a reduction in the freedom from authority for citizens, the large advantages democracy has over anarchy outweigh this reduction in freedom from authority, and this fact is claimed not to be subject to qualified disagreement (Quong, 2010, p. 46). However, if the reduction in the freedom from authority moving from an anarchy to a democracy can be justified due to the increased benefits democracy has over an anarchy, this thesis asks: Can the same argument be run for an epistocracy over anarchy since an epistocracy provides significant benefits over an anarchy? Or, as a more challenging task, can the increase in policy quality going from democracy to extended epistocracy argued for in Chapters 3 and 4 outweigh the decrease in freedom from authority of the lay citizen? (This last question is essentially the main question being answered in this chapter). Either way, attempting to find arguments for or against equal voting, or something close to it is, in itself, an important area of future research. Such research will, at the very least, open the legitimacy of democracy for critical examination, and make it seem more plausible for non-democratic forms of governance to be considered legitimate.

§6.4: Comparing Democratic Legitimacy – Epistemic Injustice

Section §6.3 proposed research that could be conducted regarding the possibility of overcoming what this chapter called the epistemic legitimacy barrier by drawing on literature concerning the expert/boss fallacy. On the assumption that this barrier can be overcome, the comparison between levels of democratic and epistemic legitimacy generated by an extended epistocracy and epistemic proceduralism can begin in earnest. Section §6.4, then, considers the democratic legitimacy side of this comparison, which is the strength of epistemic proceduralism since it is already

assumed to be democratically legitimate by default. Therefore, the focus of this section will be to propose ways to analyse the democratic legitimacy of an extended epistocracy. To do so, this section adds the literature surrounding epistemic injustice (to be defined shortly) to the interdisciplinary approach of the chapter, and thesis at large. This epistemic injustice literature is examined to present future research that could determine whether an extended epistocracy does present a significant injustice through epistemic injustice, or whether emerging science and technology governance provides an appropriate exception that exempts an extended epistocracy from the charge of epistemic injustice.

§6.4.1: What Is Epistemic Injustice?

A central feature of an extended epistocracy, being an epistocracy, is that members of the population are treated differently based on their level of relevant expertise. However, in the interests of establishing whether an extended epistocracy is legitimate, it is important to ask: does discrimination based on levels of relevant expertise create what is known as an epistemic injustice? The term “epistemic injustice” comes from the work of Miranda Fricker, who defines two different forms of epistemic injustice, namely *testimonial* and *hermeneutic*, as follows:

Testimonial injustice occurs when prejudice causes a hearer to give a deflated level of credibility to a speaker’s word; hermeneutical injustice occurs at a prior stage, when a gap in collective interpretive resources puts someone at an unfair disadvantage (Fricker, 2007, p. 1).

An example of testimonial injustice might be discounting someone’s level of credibility based only on their gender or race, while an example of hermeneutical injustice might be trying to explain an occurrence of sexual harassment in a culture that lacks a critical understanding of that concept (*Ibid.*).

Fricker considers two types of primary harm that are caused by epistemic injustice: *epistemic* and *ethical*. The epistemic harm occurs when knowledge is not effectively

passed from the speaker to the listener, causing a dysfunction in the epistemic system (*Ibid.*, p. 43). The ethical harm is that the speaker is wronged in their capacity as a knower, which is said to be a capacity essential to human value. Following Kant, the speaker has been dishonoured regarding their humanity by undermining their capacity for reason (*Ibid.*, pp. 44,46).⁸⁰ For the purpose of considering epistemic injustice against lay citizens under political situations, it is the ethical harm that is most important.⁸¹ Regarding ethical harm, epistemic injustice could be thought of as the converse version of the expert/boss fallacy, at least in political scenarios. While the expert/boss fallacy claims that it is illegitimate to grant decision-making authority based on one's higher political competence, epistemic injustice raises an important question regarding whether it is illegitimate to deny decision-making authority based on one's lower political competence.

This chapter considers four different areas in which an extended epistocracy might commit an epistemic injustice, three of which are concerned with testimonial injustice, and one is concerned with hermeneutical injustice. The three areas of testimonial injustice to be considered are the different types of social power that are active, the types of injustice these social powers create, and whether the public are treated as epistemic subjects or objects. In terms of hermeneutical injustice, this section considers whether upstream engagement presents a hermeneutical injustice due to society's general lack of understanding of future technologies. In considering each of these types of potential injustice in turn, potential research questions are presented that can

⁸⁰ Secondary harm from epistemic injustice may include not being believed at crucial times, such as during a trial, or losing confidence in one's ability as a knower and their future education thwarted as a result (Fricker, 2007, pp. 46,46–47,58).

⁸¹ While epistemic harm could technically exist, by not granting decision-making authority to lay citizens, it is not clear that this brings about dysfunction to the epistemic system. As argued in Chapters 3 and 4, *including* lay citizens in decision-making is likely to bring dysfunction to the epistemic system, and as such, it may be an injustice to include them. However, this form of injustice is a topic for Section §6.5 where normative consent and collective ethics are discussed.

be investigated to understand whether an extended epistocracy is truly committing these epistemic injustices and to what extent.

§6.4.2: Forms of Power

Testimonial injustice can be thought of as a form of agential power, where an agent (individual, group, or institution) actively discounts the level of a speaker's credibility (*Ibid.*, p. 10). An extended epistocracy, as a governance system, may be an example of agential power because it is actively denying citizens without relevant expertise from entering into the decision-making component of the policy process, though they are given the opportunity to express their concerns with their expert representative. However, any information given by the average citizens is discounted in that this knowledge is filtered through the expertise of the expert representative before entering deliberation. The average citizens are only speaking for themselves if the expert representative deems the information to be of a high enough quality and relevant; the information is otherwise filtered and modified at best, and ignored at worst.

Agential power is in contrast to structural power, which occurs when no particular agent is actively disenfranchising members of a group, yet through social circumstance, the members of that group effectively disenfranchise themselves (*Ibid.*, p. 11). An extended epistocracy may be an example of structural power since even though an extended epistocracy does not actively remove anyone's right to participate in deliberations, it is difficult for many citizens to exercise their right to deliberate in a given policy area due to the requirement that they have high-quality relevant expertise. As such, it might be the case that an extended epistocracy causes citizens to disenfranchise themselves because they do not feel capable of meeting the expertise requirements for active participation.

An extended epistocracy, then, could be committing an epistemic injustice through agential and/or structural power, and these will form major threads through the

subsequent discussion of epistemic injustice. However, whether the epistemic injustice is in the form of agential or structural power, the central focus of Fricker's work is concerned with identity power, where the epistemic injustice occurs against a citizen or group because of their social identity, i.e., their race, gender, sexual orientation, socio-economic status, etc. (*Ibid.*, pp. 14,16). If identity power were the only way to create epistemic injustice, then it could be said that implementing an extended epistocracy does not create an epistemic injustice. The active process for choosing extended epistocrats is identity blind (or bio-demographically blind, to use the terminology of Chapter 5). As a form of structural power, it might, however, be the case that those who disenfranchise themselves from an extended epistocracy are more likely to fall within particular identity groups. For example, some members within the working class may be too busy to fulfil the expertise requirements of an extended epistocracy, and are therefore effectively disenfranchised due to their lack of time to devote to political matters.

§6.4.3: Types of Injustice

Once it is known what types of social power are used to commit potential epistemic injustices, it is also important to know what types of injustices that power can commit. Two relevant types of potential epistemic injustice, or lack thereof, identified by Fricker are:

- *Innocent Error*: An innocent error in judging a speaker's credibility occurs due to human fallibility rather than immoral hatefulness or epistemic carelessness. Someone committing such an error is neither ethically nor epistemically culpable (Fricker, 2007, p. 21);⁸² and

⁸² An example of an innocent error might be an electrician explaining the problem with the wiring in the house in layman's terms when, unknown to the electrician, the house owner is an electrical engineer. By explaining the problem in layman's terms, the electrician does not intend to insult the house owner, but rather is making a safe assumption that since not many people have a full understanding of electronics, the house owner probably does not have a full understanding of electronics either.

- *Incidental Injustice*: An incidental epistemic injustice is one that affects a speaker in a limited area of their life.⁸³ Being an incidental injustice is not to say that it is a minor injustice, especially if it is persistent, but rather to say that the injustice is localised (*Ibid.*, pp. 27,29).⁸⁴

Regarding an extended epistocracy, any potential epistemic injustice caused could be found in either of these categories to at least some extent. Therefore, both types of potential epistemic injustice will be considered in turn.

Innocent Error

An innocent error may occur in an extended epistocracy by accidentally denying a genuine expert entry into the selection database for entry an extended epistocracy, or including someone in the selection database that is not a genuine expert. Given that it is difficult to give a truly accurate test for political competence, and there may be many citizens wishing to take part in the deliberative process, a relatively quick proxy test must be used to decide who are suitable experts with relevant expertise. For example, this test will, in part, be based on generalisations regarding professionalism, as discussed in Section §5.4.1 (*Ibid.*, p. 32). However, such a test may cause false positives: a “professional” may be mistakenly granted entry into the extended epistocracy, at the expense of someone more qualified, even though the one granted entry underperforms at their job. A professionalism test may also produce false negatives: a “non-professional” may be denied entry to an extended epistocracy even though they do possess high-quality relevant knowledge. Even though a professionalism test is not

⁸³ An example of an incidental injustice could be a student who has a lecturer that discredits the student’s contributions to class discussion based on a bio-demographic feature the student has. However, all the student’s other acquaintances, including their other lecturers, all respect that student’s opinions. While it is unfortunate that this student is not able to make the most of their learning experience in this class, this is a localised epistemic injustice that is not experienced elsewhere in this student’s daily life.

⁸⁴ In contrast to incidental epistemic injustice, systematic epistemic injustice affects many areas of a speaker’s life and reflects a larger social prejudice (*Ibid.*, p. 27). Given that, in its current form, an extended epistocracy is only intended to be used for emerging technology policy, and hence leaves other areas of a citizen’s political engagement untouched, an extended epistocracy does not constitute a systematic epistemic injustice.

the only test that could be used, the point remains that due to pragmatic necessity regarding the selection process, some citizens may be mistakenly granted or denied entry into deliberation based on a generalisation. However, this mistaken judgement is not due to immoral hatefulness or a lack of honest attempt at establishing suitable experts, so the selection process is not guilty of epistemic injustice. While this is technically not an epistemic injustice because it at least appears to be an innocent error, important questions this chapter proposes to ask are:

- Are the mistaken judgements simply unfortunate but expected elements of a selection process that still provides a net benefit over democratic methods, or do they truly constitute a more general injustice, such as an unfair bias against certain professions?
- Given the distrust of experts discussed in Chapter 2, could the mistaken judgements be perceived as a lack of honest attempt, or even an attempt to act out an ulterior motive by way of expert selection? What methods could be used to tell the public that an honest attempt was made when experts were selected?

Furthermore, as Fricker notes, if there is a particular person (e.g., a perpetual liar), or even a group of knowers (e.g., used car salesmen), that has been unreliable in the past, then it is entirely possible that no epistemic injustice is committed against them in a knowledge transaction (*Ibid.*, p. 42–43). The question then becomes: does this lack of epistemic injustice translate to the political sphere? If the description given of the political competence of the average citizens in Chapters 3 and 4 is true, then there is no epistemic injustice committed against them as a group since they are unreliable knowers. However, if the description given is false, then an extended epistocracy may be committing an epistemic injustice because the average citizen is being prejudged in their competence regarding emerging technology before being given a true opportunity to display or improve their competence. Furthermore, it is important to ask: are there overriding factors such that even if the average citizen is an unreliable

knower, should they still be given the opportunity to take part in decision-making as unreliable knowers?

Incidental Injustice

Even if an epistemic injustice were committed by denying someone entry into the decision-making process for emerging science and technology policy, this would be a case of incidental epistemic injustice, since the citizen could otherwise function well as a knower in their everyday life despite being denied entry to an emerging science and technology extended epistocracy. However, denying someone in such a manner could create a persistent feedback effect where people of a certain group are not taken seriously regarding emerging science and technology policy, and therefore, their views on the policy issue are not asked for. Because the views of this citizen group are not heard, the views cannot be taken seriously, which perpetuates the cycle. This cycle of injustice may begin by having, say, professionalism or education tests required for entry into deliberation. Those groups that are already disadvantaged are less likely to have achieved high levels of professional or educational status, and so they are left out of discussion (Anderson, 2012, p. 169). This cycle could lead to a disadvantaged group remaining disadvantaged due to their ineffective testimony caused by the “listener” committing an epistemic injustice (Fricker, 2007, p. 130).⁸⁵

For the purpose of searching for participants to take part in an extended epistocracy, it may be the case that the expertise provided by these disadvantaged groups is preemptively judged to be irrelevant to the policy issue at hand, when in fact they are crucial to the creation of high-quality policy. These groups are being silenced in the area of emerging science and technology policy due to not even being allowed to speak

⁸⁵ This is akin to an example of non-identity based social power, which is a student who is too shy to take part in discussion, and therefore does not take part. The teacher then learns to never expect the student to take part in discussions and the student is therefore never asked for their opinion (Hookway, 2010, p. 156).

(*Ibid.*, p. 131). It is therefore important that care is taken to avoid the silencing becoming systematic in nature and potentially extend further than emerging science and technology policy without appropriate justification, and that attempts to amend professional or educational disadvantages are made in the future (Bohman, 2012, p. 183) (Anderson, 2012, pp. 169,170). As such, potential questions for future research proposed by this chapter are:

- What methods can be introduced into the selection process for an extended epistocracy such that it is efficient enough to produce policy in a timely manner, but broad enough to include all the relevant required expertise? Would it be necessary to explicitly seek out citizens from disadvantaged groups that could be added to the database of potential epistocrats?
- Can the selection process potentially be linked to public education programs such that those citizens who are interested in emerging science and technology policy are able to learn about, and give more direct feedback towards, policy issues?

§6.4.4: The Subject–Object Distinction

The third and final area of testimonial injustice to be considered is the subject–object distinction. As part of the process for an extended epistocracy, average citizens can discuss their concerns with an expert representative, but they are not invited to take part in making the policy decision. This distinction between “decision-maker” and “concerned citizen” could be considered the same distinction Fricker makes between epistemic subjects and objects, respectively, where an epistemic subject is an agent who conveys information, whereas an object is merely a source of information that an inquirer can glean information from (Fricker, 2007, p. 132). Put another way, an epistemic subject is a trusted informant able to contribute a worthwhile and reasonable viewpoint, whereas an epistemic object has their capacity as a giver of knowledge undermined and is treated, for example, as a data point in a social science

survey (*Ibid.*, pp. 132–133). Due to this possible injustice, there are several important questions to be considered by this chapter for potential future research:

- Following from a suggestion by Medina (2012, p. 203), where exactly does the distinction lie between an epistemic subject and object? Must the general public take part in decision-making if they are to be considered an epistemic subject, or is it enough to allow the general public to simply voice their concerns?
- If the general public were considered only as data points in a “users of nanoproductions” survey, for example, but were never asked for their opinion on nanoproductions, are they being treated as *mere* epistemic objects (à la Kant (Fricker, 2007, p. 134)), even though they are respected as informants in their everyday lives?
- As Fricker (*Ibid.*, pp. 135–136) suggests, is it ethical, or at least beneficial, to deny someone being an epistemic subject based on what the speaker has to say being untrustworthy or unreasonable? Does this denial of someone as an epistemic subject extend to political situations where deliberative democracy says that all (non-crazy and non-vicious) viewpoints are to be treated equally?

§6.4.5: Hermeneutical Injustice and Upstream Engagement

Aside from the three forms of testimonial injustice just presented, an extended epistocracy could cause an epistemic injustice through a hermeneutical injustice. Recall that a hermeneutical injustice occurs “when a gap in collective interpretive resources puts someone at an unfair disadvantage” (Fricker, 2007, p. 1). Put another way, hermeneutical injustice is:

The injustice of having some significant area of one’s social experience obscured from collective understanding owing to hermeneutical marginalization (*Ibid.*, p. 158).

Though potentially stretching the definition of hermeneutical injustice, this chapter offers the suggestion that upstream engagement poses a hermeneutical injustice. Given the increasing specialisation of technoscientific fields, and the increasing speed

of their development, upstream engagement is required to use ever more specialised knowledge and look further into the future. As a result, larger sections the general public are simply unable to maintain the expertise required to take part in deliberation under an extended epistocracy governance regime. This lack of expertise on the part of the general public creates what is known as a hermeneutical inequality where the public are “rendered unable to make communicatively intelligible something which it is particularly in his or her interests to be able to render intelligible” (*Ibid.*, 162). Members of the general public are then excluded from the upstream engagement, either effectively due to their inability to fully understand the required futures concepts, or actually due to the selection process for an extended epistocracy. While it may be unfortunate for the members of the general public to be excluded in such a manner, one must ask if it is epistemically unjust to do so. The general public’s concerns regarding emerging technology are still concerns that will affect the uptake of that technology, even if the concerns may not be reasonable. To what extent, then, should the concerns of the general public be taken seriously? Is it enough that the public’s concerns are heard by their expert representative in an extended epistocracy, only to be dismissed before these concerns reach the deliberating table? Or must all (well-meaning) concerns be taken seriously, regardless of how reasonable they are, to avoid epistemic injustice?

§6.5: Comparing Epistemic Legitimacy – Collective Ethics and Normative Consent

Recall that the purpose of this chapter is to compare the democratic and epistemic legitimacies of both epistemic proceduralism and an extended epistocracy. To this end, Section §6.4 considered the democratic legitimacy of an extended epistocracy such that it could be compared to the democratic legitimacy of epistemic proceduralism. Section §6.5, then, compares an extended epistocracy and epistemic proceduralism along the epistemic legitimacy dimension. In doing so, the total legitimacy from the combined effects of Sections §6.4 and §6.5 can be compared to

determine whether an extended epistocracy or epistemic proceduralism (or a governance system in between) is most appropriate to be used for emerging science and technology governance. Given that an extended epistocracy is designed to be an epistemically superior form of governance, one could simply say that an extended epistocracy generates more epistemic legitimacy simply due to the tendency to produce higher quality policy than epistemic proceduralism such that this epistemic legitimacy gap can be compared against the democratic legitimacy gap between the two governance systems. However, while useful, this method would not provide a complete legitimacy picture of the epistemic legitimacy gap since there are other research areas that can be investigated. In particular, this section combines the relatively nascent concepts of normative consent (to be defined shortly) and collective ethics (the application of ethics to collective actions, such as voting) as a potential way of comparing both governance systems along the epistemic legitimacy dimension.

§6.5.1: What is Normative Consent?

One potential way to consider the epistemic legitimacy of a governance system is to ask whether there is a moral obligation to use the most epistemically beneficial governance system even though it is less democratically legitimate, or can it be morally acceptable to use a less epistemically successful governance system, even if it does produce worse results? Put another way, would it be morally wrong to refuse to consent to the authority of the epistemically superior governance system? If the refusal to consent to this authority is immoral, a moral obligation is said to be created and an agent under the moral obligation is said to give what Estlund calls “normative consent” (Estlund, 2008, p. 10). As an example of normative consent, imagine an experienced police officer giving instructions in the event of an emergency. Among other things, it is the police officer’s experience in emergency situations that gives them the knowledge of the appropriate procedures and the ability to lead accordingly. In turn, this experience and the knowledge gained gives the police officer authority in the emergency. By not accepting the authority of the police officer, there is the risk of

potentially significant harm to oneself and/or others, and as such, normative consent is given to the authority of the police office by bystanders. In a similar manner, experts in an epistocracy might play the same role as a police officer in an emergency. The experience and expertise of experts may give them the authority to lead such that the risk of potentially significant harm is reduced. As a result, important questions this chapter proposes for the purpose of future research are:

- Is the analogy between the emergency and the epistocratic governance situation strong enough to support the legitimacy of an epistocracy? Or is the short timeframe of an emergency compared with the relatively long policy decision process enough to disqualify the analogy? Or, as a middle ground, is the expertise of experts enough to grant them authority as policy guides, but not as leaders?
- If authority is based to some extent on expertise, as Estlund suggests, when does authority fade based on poor judgements, if at all? (*Ibid.*, p. 125)

These questions aside, the main line of further research this section proposes to regarding the epistemic legitimacy of governance systems is to examine the collective ethics of voting/deliberating. The overarching question when it comes to considering the collective ethics of voting/deliberating is as follows: if democratic decision-making, and hence epistemic proceduralism, does tend to produce worse outcomes than an extended epistocracy, then is it possible to go from the collective harm⁸⁶ caused by democracy in this way to argue that an individual voter with low political competency should give normative consent to an extended epistocracy?

§6.5.2: The Collective Ethics of Voting

In *The Ethics of Voting* (Brennan, 2011a), Jason Brennan lays out some moral obligations for voters based on their epistemic capabilities. Voting, and policymaking in general,

⁸⁶ "A collective harmful activity is a harmful activity caused by a group or collective, where individual inputs into the harmful action are negligible" (Brennan, 2011a, p. 71).

is morally significant, at least as a collective, since it can affect the health, wealth, and general wellbeing of citizens. If it is done well, a nation can prosper, but if done poorly, citizens must live under harmful and unjust rule. Therefore, there are moral obligations regarding how voting should be done (*Ibid.*, p. 1). However, it is the collective outcome of a voting decision rather than a certain individual vote that determines policy quality, assuming the government acts in accordance with the voting outcome. Therefore, what are the moral implications regarding the epistemic quality of votes for the *individual* voter? Brennan believes that citizens have no duty to vote, but if they do vote, then they must vote in an epistemically justified⁸⁷ manner for the common good⁸⁸ based on the best reasoning and evidence available. Those who are not going to vote in such a manner should abstain from voting (*Ibid.*, pp. 4,70). Put in another way, Brennan argues that citizens have a moral obligation not to vote in an epistemically unjustified⁸⁹ manner (*Ibid.*, p. 68).

There are several areas of study within collective ethics that can help to illuminate whether Brennan's conclusion is correct. While some study has been done regarding the collective ethics of voting, the main area of collective ethics can be drawn from environmental ethics and their work in determining individual responsibility towards anthropogenic climate change. The collective ethics discussed in environmental ethics

⁸⁷ What is considered to be epistemically justified is left to epistemologists; Brennan does not argue for or use any particular theory of epistemic justification (Brennan, 2011a, p. 70).

⁸⁸ Voting for the "common good" means that the voter should vote for the collectively best option, which might not be for the best candidate, i.e., the voter should use strategic voting where necessary. It may be the case that the best candidate does not have much of a chance of becoming elected, in which case, the voter may vote for the best candidate likely to get elected (assuming this candidate is of a reasonable quality and this kind of compromise does not harm the prospects for good outcomes in future electoral outcomes) (Brennan, 2011a, pp. 131–133).

⁸⁹ Epistemically unjustified voting can take two forms, namely unexcused harmful voting and fortuitous voting:

- *Unexcused harmful voting* when a person votes, without epistemic justification, for harmful policies or for candidates likely to enact harmful policies.
- *Fortuitous voting* occurs when citizens vote for what are in fact beneficial policies or candidates likely to enact beneficial policies, but they lack sufficient justification to believe that these policies or candidates are good. (Brennan, 2011a, p. 68)

provides a useful parallel to the collective ethics of voting since almost all the concerns found in collective environmental ethics hinge on how the ineffectiveness of individual actions should be considered; if individual actions are ineffective, then while they may not produce much good, they also cannot do much harm either (*Ibid.*, p. 71). The same can be applied to voting; while each individual vote does not do much good, each vote does not do much harm either. The challenge then becomes: if necessary, how can an individual voter be convinced they are doing something immoral if they are causing, at worst, very little harm?

To place the just stated question in context, imagine a consensus conference has taken place discussing the potential introduction of a new nanodevice. Experts in the field have had their chance to speak on the issue and 15 citizens have had the opportunity to question the experts and propose their own recommendations, which have been laid out in a policy document. The policy document is then voted on by a larger democratic body, which vote in favour of implementing the recommendations. As it turns out, implementing the recommended policies was a poor idea; a collective harm has been done. Let us assume that the policy implemented was exactly the policy that was voted on, i.e., the problem is the policy itself rather than its implementation. What does this collective harm mean for the individual voters? To answer this question, this section breaks it down into three general component questions that can be used to propose future research:

- What, if anything, makes an individual voter responsible for the collective harm?
- What, if anything, has the individual voter done wrong?
- What, if anything, does collective ethics require of the voter?

These three questions will be discussed in turn. Many further questions are posed throughout the discussion, and while answers from the literature are presented, it would also be value for there to be further contributions to these discussion by way of future research.

§6.5.3: Responsibility for Collective Harm

Regarding the first question proposed, collective ethics proposes several ways the voter might be considered responsible, namely in terms of “contribution”, “causation”, and “groups”, which are discussed in turn. Contribution and causation are considered first and together since they both suffer from the same problem when it comes to arguing the case for an extended epistocracy. Groups are discussed last since this section contends that it is the most crucial method for determining the case for or against an extended epistocracy. Considering contribution, then, the *contribution principle* states that “we have a moral obligation not to make problems worse” (Sinnott-Armstrong, 2005, p. 301).⁹⁰ Considering causation, the *harm principle* states that “we have a moral obligation not to perform an act that causes harm to others” (Sinnott-Armstrong, 2005, p. 297). While both the contribution principle and the harm principle seem reasonable at face value, they both face the problem noted earlier, which is what Hiller calls *individual causal inefficacy*: given the scale of an election, an individual voter is causally ineffective (Hiller, 2011, p. 349). These problems are exacerbated due to elections being decided by thresholds, especially in the current example with a majority pass/fail outcome. It is difficult, then, to establish what harm a single poor vote has done, or what the single vote has caused given the outcome would be exactly the same without that vote (i.e., the policy recommendations still would have passed), unless it is the extremely rare deciding vote (*Ibid.*, p. 301) (Pettersen, 2004, pp. 292–293) (Brennan, 2011a, p. 72). If a single vote causes no (significant) harm and cannot change a harmful outcome, then Sinnott-Armstrong argues it would be wrong to consider the voter morally blameworthy (Sinnott-

⁹⁰ Brennan’s *Clean Hands Principle* is very similar. He argues that even though we are not obligated to solve the problem on our own, the individual should not contribute to it with the caveat that there must be no significant personal costs imposed on the individual if they were to refrain for the principle to apply (Brennan, 2011a, pp. 72–73).

Armstrong, 2005, pp. 306–307).⁹¹ The ineffectiveness of individual votes causes a paradox: an election result is caused, yet no-one caused it. A blameworthy event has occurred, but no-one is blameworthy (Vanderheiden, 2006, p. 86).

In response to the ineffectiveness of a single vote, Parfit would claim that Sinnott-Armstrong has made the second mistake of moral mathematics:

The Second Mistake: If some act is right or wrong *because of its effects*, the only relevant effects are the effects of this particular act (Parfit, 1984, p. 70).

According to Parfit, “[e]ven if an act harms no one, this act may be wrong because it is one of a *set* of acts that *together* harm other people” (and vice versa for beneficial acts) (*Ibid.*). This idea of a set of acts leads to the third way a voter may be considered responsible for the collective harm, namely the group principle:

Group Principle: We have a moral obligation not to perform an action if this action makes us a member of a group whose actions together cause harm (Sinnott-Armstrong, 2005, p. 306).

Considering the set of acts helps to solve the paradox above; the group, as a whole, becomes effective and a poor vote can be considered part of the causal genesis of the poor policy outcome (Pettersen, 2004, p. 293).⁹² The question remains: how is the group to be identified? To answer this question, it is necessary to consider the second of the three questions proposed above regarding what, if anything, the individual voter has done wrong.

⁹¹ Sinnott-Armstrong’s argument is related to the pollution caused by a single drive in a gas-guzzling vehicle, but the point remains the same.

⁹² Many poor voters are likely to understand that their individual vote is ineffective, and will vote with the belief that at least they cannot, individually, do any harm. If enough of them do it, though, these voters become harmful as a group; that is to say the voter is harmless, but the mindset is not. Therefore, given that all voters are in a similar situation regarding the effectiveness of their votes, and therefore will probably think in a similar manner regarding that effectiveness, it is safest to assume that a poor voter’s action will effectively be multiplied.

§6.5.4: The Harm Caused by the Individual Voter

As noted above, each individual voter does not have a crucial effect on the outcome of the election as an individual. However, are there other ways in which harm can be attributed to voters in such a way that it avoids the aforementioned paradox, where a blameworthy event has occurred, yet no one is to blame? According to Brennan, voting based on poor reasoning is wrong because it imposes an unacceptable risk (Brennan, 2011a, pp. 79–80). Just as it would be unacceptable for a government to use epistemically poor methods in the process of governing, it is also unacceptable for citizens to choose their government (and expose the population at large to potential harm) with those same methods (*Ibid.*, pp. 80–81). In the case of the individual voter, the risk is the increased probability that the voting threshold is achieved for a negative outcome (Almassi, 2012, p. 13).

The question then becomes: what counts as an unacceptable risk? Sinnott-Armstrong's *risk principle* states that "we have a moral obligation not to increase the risk of harms to other people" (Sinnott-Armstrong, 2005, p. 302). However, the risk principle is not specific enough for the purpose of this section; any political action, or inaction, could potentially increase the risk of a negative outcome. Searching for relevant information or engaging in political discussion may bias the information the voter has and negatively sway their decision. Trying to change policy may risk making the situation worse, but keeping the status quo may risk further harming those already negatively affected. To make any sort of political progress, there must be some level of *acceptable* risk, but how is it defined? If, as noted above, the harm caused is the increased risk due to poor political reasoning, an investigation into what would be considered poor political reasoning relative to what could be reasonably expected by voters is proposed as future research. Or perhaps the goal should be to reasonably minimise risk rather than find an acceptable level of risk. The difference between acceptable and minimised risk can be seen in the difference between an extended epistocracy and epistemic proceduralism. Using an extended epistocracy may reasonably minimise risk in a

pragmatic sense since it is a governance system that is intended to be least likely to make mistakes. On the other hand, having a more democratic method like epistemic proceduralism may still be within the realm of acceptable risk, and it has the added bonus of including more democratic legitimacy than an extended epistocracy. Furthermore, another potential research question to consider is: how is the chosen level of risk implemented with regard to choosing who can partake in decision-making? That is to say, will there still need to be some epistemic restrictions, albeit less strict than for an extended epistocracy, such that the acceptable level of risk is not exceeded, or is a fully democratic selection method able to create policy within the realm of acceptable risk?

§6.5.5: The Responsibility of the Voter

If the harm caused by the poor voter is that the risk of surpassing a negative threshold has increased, what is the responsibility of the poor voter? The natural answer would be to at least aim to decrease the probability that a negative threshold is reached⁹³ (Almassi, 2012, p. 13). To do so, a poor voter could simply abstain. What is interesting about this potential solution is that the reason why a voter may be a poor voter in the first place (the high cost to vote well, as discussed in Chapter 3), is the same reason why this solution is more palatable, i.e., it may be a lot less costly for the individual to vote (Brennan, 2011a, pp. 80–81). This line of reasoning is consistent with Hiller's claims that "it is *prima facie* wrong to perform an act which has an expected amount of harm greater than another easily available alternative" (Hiller, 2011, p. 352). When applied to collective voting, Hiller's *prima facie* duty states that it is *prima facie* wrong to vote poorly rather than take the easier option and not vote. As a *prima facie* statement, this is fine. However, the potential lack of psychological benefits for the

⁹³ It should be noted that an "all things considered" clause needs to be applied to allow for countervailing moral considerations (Almassi, 2012, p. 13); murdering the worst candidate may guarantee they do not win the election, but countervailing moral considerations would maintain it is still wrong to do so.

individual that come along with voting, or a lack of social cohesion and drive for social action at the collective level may be potential costs that need to be investigated (*Ibid.*, pp. 75–76).

A further potential question to be considered as future research is: Is Brennan's suggestion of decreasing the risk of a negative outcome by having poor voters abstain enough? Should a voter be actively aiming to increase the probability of a positive outcome instead?⁹⁴ What is meant here by actively increasing the probability of a positive outcome is not simply having a poor voter indirectly increase the probability of a positive outcome by abstaining since good voters now form a larger proportion of voters. Here, increasing the probability of a positive outcome is the more demanding task of actively voting well. However, as Vanderheiden notes, an *ought implies can* principle should be put in place (Vanderheiden, 2006, p. 90).⁹⁵ Asking all eligible voters to vote well is too demanding, as discussed throughout this thesis. Is it possible, then, to argue for a middle ground where if one can vote well, they are obligated to, but if it is unreasonable to expect one to vote well, they are obligated to abstain?

Or perhaps, future research may want to investigate whether the standard for what is considered to be good voting should be lowered. Serota and Leib claim that when a voter is taking part in direct democracy, the voter is acting as a representative for the public at large. As such, voters have the obligation to vote in accordance with a

⁹⁴ Or is there a middle ground that Sinnott-Armstrong suggests where it may be morally good to refrain from committing an action like voting poorly, but it is not obligatory (Sinnott-Armstrong, 2005, pp. 306–307).

⁹⁵ Along the same lines, it may also be necessary to enforce a reasonable expectation that enough voters are likely to contribute to the goal of voting well, or at least not voting poorly. This is because no voter on their own is going to make a difference with their vote, so there is little incentive in putting in a lot of effort to vote well if no one else is (Sandberg, 2011, p. 242). One way to enforce cooperation is to invoke satisfying fairness as a moral value to incentivise voting/abstaining appropriately. If someone votes badly, usually because they have not put in the effort to research their options, then they increase the burden on those voting well in their attempt to achieve a good collective outcome. Not only does a good voter have to be more certain, i.e., do more research, to vote well to cancel the bad vote, but they must also convince someone else to vote well to make progress (Brennan, 2011, p. 75).

credible and good faith conception of the public interest (Serota & Leib, 2013, p. 1598). With that said, Serota and Leib argue that if a voter makes a poor vote, they are not “morally culpable so long as [their] mistaken belief was credible and based on meaningful evidence when [their] vote was cast” (*Ibid.*, p. 1617). In a similar manner, Smith argues that people have the duty to inform themselves before acting, which when applied to voting, means that voters should make a good faith effort to find credible evidence and reason appropriately (Smith, 2014, p. 11). However, for the purpose of future research, it is important to qualify what is meant by a “credible mistaken belief” or “meaningful evidence” and how much voters need to inform themselves before voting. Should “meaningful evidence” mean “best available evidence” such that “credible mistaken belief” means that even if a decision was wrong, it was still the best decision based on current understandings? Or should the standard be lowered to make “meaningful evidence” commensurate with the epistemic resources each citizen has available to them, or could reasonably be expected to gain within the timeframe available?⁹⁶

Overall, if further investigation finds that the individual voter within epistemic proceduralism is unnecessarily increasing the risk of a negative voting outcome past what is determined to be an appropriate level by not voting in a justifiable manner, then it can be said that normative consent should be given to an extended epistocracy. This normative consent at least increases the epistemic legitimacy gap between and extended epistocracy and epistemic proceduralism in an extended epistocracy’s favour, if not automatically qualify it as a legitimate governance system. However, if

⁹⁶ Serota and Leib’s answer to this line of questioning claims that voters do not need to take part in deep individual empirical and normative analysis, but instead can use heuristics such as endorsements from political newspapers and public interest organisations. Furthermore, they also claim that given each voter has little influence, their duty of care can be calibrated relative to the risk of harm each one creates (Serota & Leib, 2013, p. 1618). On the other hand, Smith recognises that although moral theories can be epistemically demanding, and therefore agents can only reasonably act on their subjective knowledge to maximise expected utility, agents still have a *prima facie* duty to inform themselves such that they maximise their ability to fulfil other duties (Smith, 2014, p. 11).

it is found that an individual voter is not required to give normative consent to an extended epistocracy, then the epistemic legitimacy gap between an extended epistocracy and epistemic proceduralism is simply the difference between the average policy quality of the two governance systems (unless other ways of contemplating epistemic legitimacy are found), and a comparison of total legitimacy is necessary to determine which of the two governance systems should be used for emerging science and technology policymaking.

§6.6: Conclusion

While the focus of this thesis has been on the epistemic quality of governance systems, a governance system must also be considered legitimate by the public to be effective. To this end, this chapter proposed several areas of future research that could potentially be used to investigate and compare the total legitimacy of epistemic proceduralism and an extended epistocracy to determine which of the two are best suited for emerging science and technology policy governance, and maybe whether an intermediate governance system is optimal. After suggesting ways in which the legitimacy barrier blocking this comparison can be overcome by analysing the expert/boss fallacy, two further research areas were used to compare the democratic and epistemic legitimacy of an extended epistocracy with epistemic proceduralism. Evaluating the potential epistemic injustices of an extended epistocracy was suggested as a method for comparing the democratic legitimacies, while the collective ethics of voting as applied to normative consent was the suggested method for comparing epistemic legitimacies. If it can be shown that the total legitimacy of an extended epistocracy is at least as high as the total legitimacy of epistemic proceduralism, then it is acceptable to implement an extended epistocracy as a governance regime for emerging technology policy since epistemic proceduralism is already assumed to be legitimate by default.

Chapter 7: Conclusion

Recall that this thesis aimed to serve three scholarly purposes, which were to use interdisciplinary methods to investigate the epistemic properties of deliberative democracy; to propose a form of governance that is intended to be able to solve the problem of expertise that arose as part of the first aim; and overall, to contribute to the growing literature on epistocracy and epistemic properties of governance. These aims are to be discussed in turn.

The first aim was satisfied throughout Chapters 1–4 in Parts A and B of the thesis. The role of Part A was to establish what the problem of expertise was and the challenges that it presents under modern political conditions. By utilising standard forms of democratic theory and science and technology studies, Chapter 1 presented the problem of expertise as it has appeared throughout political thought, as well as the challenges faced when trying to solve the problem of expertise for wicked policy problems in a post-normal era. Knowing that an extended peer community is the accepted governance method for such policymaking, Chapter 2 continued the use of democratic theory and science and technology studies, alongside literature on public engagement for governance, to show how the emphasis on epistemic properties for modern governance has shifted over time to utilise an extended peer community and then back to what forms of governance could be used to make best use of the epistemic properties of deliberative democracy.

The purpose of Part B of the thesis, then, was to establish whether epistemic deliberative democracy has the epistemic properties to overcome the problem of expertise. To this end, Chapter 3 used resources from political epistemology such as statistical evaluations of voting and public understanding of science, as well as rational choice theory from economics, to establish the actual and potential individual citizen competency of the average citizen regarding nanotechnology governance.

Having shown that the individual competency of the average citizen regarding nanotechnology governance was low and likely to stay low, it could have been the case that a collective democratic effort could overcome the problem of expertise. To investigate this possibility, Chapter 4 examined the actual performance of public engagement for governance exercises regarding nanotechnology, as well as its potential performance through statistical representations of deliberation. It was shown that not only did the practical engagement exercise not satisfy the requirements for deliberative democracy, it appears that an epistocracy would be a theoretically more suitable epistemic alternative.

With this result regarding epistocracy in mind, the second aim of this thesis was satisfied through the development of an extended epistocracy for maximal epistemic performance and potential future research that could be used to defend its democratic legitimacy. The extended epistocracy was proposed in Chapter 5 by using literature from science and technology studies, psychology, and management studies. It is recommended that members of an extended epistocracy are to be those with relevant expertise and are carefully chosen such that their largely distinct cognitive diversity slightly overlaps. Furthermore, the extended epistocracy is ideally to remain as a long-term, cohesive unit that is in control of all parts of the policymaking process. However, such a selective group of deliberators opens up concerns regarding the democratic legitimacy of the policy process. As a way of potentially alleviating these concerns, Chapter 6 took a very broad interdisciplinary approach by including further scholarship in political epistemology, as well as from collective ethics. The suggested method for establishing the legitimacy of an extended epistocracy was to compare it to another form of governance, namely epistemic proceduralism, along both the epistemic and democratic dimensions to establish which procedure had the greater total legitimacy.

With the possible argument structure that was laid out in Chapter 6, this helps to satisfy the third aim of the thesis by providing a method for defending the legitimacy of an extended epistocracy or even providing a way in which a modified version could be found that better minimises the tension between expertise and democracy found in the problem of expertise. The broad, though not exclusive, interdisciplinary range of potential resources provided throughout the thesis will aid future researchers in searching for such an epistocracy. Furthermore, the results of this thesis and the interdisciplinary range of resources used highlight the need for further investigation into the epistemic properties of democracy that goes beyond the status quo assumption of democratic theory. While the turn towards epistemic forms of deliberative democracy is indeed encouraging, the field of political epistemology is still young and underdeveloped. Sustained research into the epistemic properties of deliberative democracy involving a great deal of interdisciplinary collaboration is still required, especially for wicked problems found in these modern post-normal times, to help to create a governance system that is capable of undertaking the difficult task of creating high-quality policy under such conditions. To fail to do so could be devastating.

Bibliography

- Ackerman, B., & Fishkin, J. (2002). Deliberation Day. *The Journal of Political Philosophy*, 10(2), 129–152.
- Ackerman, B., & Fishkin, J. S. (2002). Deliberation Day. *The Journal of Political Philosophy*, 10(2), 129–152.
- Almassi, B. (2012). Climate Change and the Ethics of Individual Emissions. *Perspectives: International Postgraduate Journal of Philosophy*, 4(1), 4–21.
- Althaus, S. L. (2003). *Collective Preferences in Democratic Politics: Opinion Surveys and Will of the People*. Cambridge: Cambridge University Press.
- Amason, A. C., & Sapienza, H. J. (1997). The Effects of Top Management Team Size and Interaction Norms on Cognitive and Affective Conflict. *Journal of Management*, 23(4), 495–516.
- American Political Science Association. (2014). *Ideas, Knowledge, and Politics (Section 46)*. Retrieved from American Political Science Association: <http://www.apsanet.org/section46>
- Anderson, E. (2012). Epistemic Justice as a Virtue of Social Institutions. *Social Epistemology*, 26(2), 163–173.
- ARENA Centre for European Studies. (2015, August 21). *University of Oslo*. Retrieved from Why Not Epistocracy? Political Legitimacy and the "Fact of Expertise" (EPISTO): <http://www.sv.uio.no/arena/english/research/projects/episto/index.html>
- Austen-Smith, D., & Banks, J. S. (1996). Information Aggregation, Rationality, and the Condorcet Jury Theorem. *American Political Science Review*, 90(1), 34–45.
- Baber, W. F., & Bartlett, R. V. (2007). Problematic Participants in Deliberative Democracy: Experts, Social Movements, and Environmental Justice. *International Journal of Public Administration*, 30(1), 5–22.
- Bartels, L. M. (2008). The Irrational Electorate: Many of Our Worst Fears about America's Voters are True. *Wilson Quarterly*, 32(4), 44–50.
- Batt, C. A., Waldron, A. M., & Broadwater, N. (2008). Numbers, Scale and Symbols: The public understanding of nanotechnology. *Journal of Nanoparticle Research*, 10(7), 1141–1148.
- Beck, U. (1996). Risk Society and the Provident State. In S. Lash, B. Szerszynski, & B. Wynne (Eds.), *Risk, Environment and Modernity: Towards a New Ecology* (pp. 27–43). London: SAGE Publications.
- Bell, D. A. (1999). Democratic Deliberation: The Problem of Implementation. In S. Macedo, *Deliberative Politics: Essays on Democracy and Disagreement* (pp. 140–173). New York: Oxford University Press.
- Benhabib, S. (1994). Deliberative Rationality and Models of Democratic Legitimacy. *Constellations*, 1(1), 26–52.
- Ben-Yashar, R., & Paroush, J. (2000). A Non-Asymptotic Condorcet Jury Theorem. *Social Choice and Welfare*, 17(2), 189–199.

- Berend, D., & Paroush, J. (1998). When Is Condorcet's Jury Theorem Valid? *Social Choice and Welfare*, 15(4), 481–488.
- Berg, S. (1993). Condorcet's Jury Theorem, Dependency among Jurors. *Social Choice and Welfare*, 10(1), 87–95.
- Berg, S. (1996). Condorcet's Jury Theorem and the Reliability of Majority Voting. *Group Decision and Negotiation*, 5(3), 229–238.
- Berg, S. (1996). Condorcet's Jury Theorem and the Reliability of Majority Voting. *Group Decision and Negotiation*, 5(3), 229–238.
- Berman, A. E., & Dorrier, J. (2016, March 22). *Technology Feels Like its Accelerating - Because It Actually Is*. Retrieved from Singularity Hub: <https://singularityhub.com/2016/03/22/technology-feels-like-its-accelerating-because-it-actually-is/>
- Bohman, J. (1996). *Public Deliberation: Pluralism, Complexity, and Democracy*. Cambridge: MIT Press.
- Bohman, J. (2012). Domination, Epistemic Injustice and Republican Epistemology. *Social Epistemology*, 26(2), 175–187.
- Bourke, J., Garr, S., van Berkel, A., & Wong, J. (2017, February 28). *Diversity and Inclusion: The Reality Gap*. Retrieved from Deloitte Insights: <https://www2.deloitte.com/insights/us/en/focus/human-capital-trends/2017/diversity-and-inclusion-at-the-workplace.html>
- Brennan, J. (2011a). *The Ethics of Voting*. Princeton: Princeton University Press.
- Brennan, J. (2011b). The Right to a Competent Electorate. *The Philosophical Quarterly*, 61(245), 700–724.
- Brennan, J. (2014). How Smart is Democracy? You Can't Answer that Question A Priori. *Critical Review: A Journal of Politics and Society*, 26(1–2), 33–58.
- Brossard, D., Scheufele, D. A., Kim, E., & Lewenstein, B. V. (2009). Religiosity as a Perceptual Filter: Examining Processes of Opinion Formation about Nanotechnology. *Public Understanding of Science*, 18(5), 546–558.
- Burri, R. V. (2009). Coping with Uncertainty: Assessing Nanotechnologies in a Citizen Panel in Switzerland. *Public Understanding of Science*, 18(5), 498–511.
- Cacciatore, M. A., Scheufele, D. A., & Corley, E. A. (2011). From Enabling Technology to Applications: The Evolution of Risk Perceptions About Nanotechnology. *Public Understanding of Science*, 20(3), 385–404.
- Cady, S. H., & Valentine, J. (1999). Team Innovation and Perceptions of Consideration: What Difference Does Diversity Make? *Small Group Research*, 30(6), 730–750.
- Caplan, B. (2007). *The Myth of the Rational Voter: Why Democracies Choose Bad Policies*. Princeton: Princeton University Press.
- Caplan, B. (2009). Majorities Against Utility: Implications of the Failure of the Miracle of Aggregation. *Social Philosophy and Policy*, 26(1), 198–211.
- Christiano, T. (1997). The Significance of Public Deliberation. In J. Bohman, & W. Rehg, *Deliberative Democracy: Essay on Reason and Politics* (pp. 243–277). Cambridge: MIT Press.

- Cianciolo, A. T., Matthew, C., Sternberg, R. J., & Wagner, R. K. (2006). Practical Intelligence and Expertise. In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 613–632). Cambridge: Cambridge University Press.
- Cobb, M. D. (2011). Creating Informed Public Opinion: Citizen Deliberation about Nanotechnologies for Human Enhancements. *Journal of Nanoparticle Research*, 13(4), 1533–1548.
- Cobb, M. D., & Macoubrie, J. (2004). Public Perceptions about Nanotechnology: Risks, benefits and trust. *Journal of Nanoparticle Research*, 6(4), 395–405.
- Cohen, J. (1997). Deliberation and Democratic Legitimacy. In J. Bohman, & W. Rehg, *Deliberative Democracy: Essays on Reason and Politics* (pp. 67–92). Cambridge: MIT Press.
- Cohen, J. (1998). Democracy and Liberty. In J. Elster, *Deliberative Democracy* (pp. 185–231). Cambridge: Cambridge University Press.
- Collins, H., & Evans, R. (2007). *Rethinking Expertise*. Chicago: Chicago Press.
- Condorcet, J.-A.-N. d. (1995 (1785)). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Retrieved from Histoire de la pensée économique, Département d'économie, Université Laval: <http://www.ecn.ulaval.ca/~pgon/hpe/documents/adamsmith/Condorcet%20Probabilites.pdf>
- Cook, P. (2013). Against a Minimum Voting Age. *Critical Review of International Social and Political Philosophy*, 16(3), 439–458.
- Cormick, C. (2009). Piecing Together the Elephant: Public Engagement on Nanotechnology Challenges. *Science and Engineering Ethics*, 15(4), 439–442.
- Cormick, C., & Hunter, S. (2014). Valuing Values: Better Public Engagement on Nanotechnology Demands a Better Understanding of the Diversity of the Publics. *Nanoethics*, 8(1), 57–71.
- Creath, R. (2017). *Logical Empiricism*. Retrieved from Stanford Encyclopedia of Philosophy: <https://plato.stanford.edu/entries/logical-empiricism/#EmpVerAntMet>
- Crisp, R. J., & Turner, R. N. (2011). Cognitive Adaptation to the Experience of Social and Cultural Diversity. *Psychological Bulletin*, 137(2), 242–266.
- de Leon, P. (1997). *Democracy and the Policy Sciences*. Albany: University of New York Press.
- DeChurch, L. A., & Mesmer-Magnus, J. R. (2010). The Cognitive Underpinnings of Effective Teamwork: A Meta-Analysis. *Journal of Applied Psychology*, 95(1), 32–53.
- Delgado, A., Kjølberg, K. L., & Wickson, F. (2011). Public Engagement Coming of Age: From Theory to Practice in STS Encounters with Nanotechnology. *Public Understanding of Science*, 20(6), 826–845.
- Devine, D. J., & Philips, J. L. (2001). Do Smarter Teams Do Better? A Meta-Analysis of Cognitive Ability and Team Performance. *Small Group Research*, 32(5), 507–532.

- Dewey, J. (1927). The Public and its Problems. In *The Later Works of John Dewey 1925–1953* (Vols. 2 (1925–1927), pp. 235–372). Carbondale: Southern Illinois University Press.
- Dijkstra, A. M., & Critchley, C. R. (2016). Nanotechnology in Dutch Science Cafes: Public Risk Perceptions Contextualised. *Public Understanding of Science*, 25(1), 71–87.
- Doubleday, R. (2007). Risk, Public Engagement and Reflexivity: Alternative Framings of the Public Dimensions of Nanotechnology. *Health, Risk & Society*, 9(2), 211–227.
- Downs, A. (1957). *An Economic Theory of Democracy*. New York: Harper and Row Publishers.
- Dryzek, J. S. (2000). *Deliberative Democracy and Beyond: Liberals, Critics, Contestations*. Oxford: Oxford University Press.
- Dryzek, J. S. (2001). Legitimacy and Economy in Deliberative Democracy. *Political Theory*, 29(5), 651–669.
- Dryzek, J. S. (2009). Democratization as Deliberative Capacity Building. *Comparative Political Studies*, 42(11), 1379–1402.
- Elster, J. (1998a). *Deliberative Democracy*. Cambridge: Cambridge University Press.
- Environmental Defence. (2005). Retrieved from Getting Nanotechnology Right the First Time: http://www.edf.org/sites/default/files/4533_EDnanopersp_0.pdf
- Erikson, R. S. (2007). Does Public Ignorance Matter? *Critical Review*, 19(1), 23–34.
- Estlund, D. (1997). Beyond Fairness and Deliberation: The Epistemic Dimension of Democratic Authority. In J. Bohman, & W. Rehg, *Deliberative Democracy: Essays on Reason and Politics* (pp. 173–204). Cambridge: MIT Press.
- Estlund, D. (2003). Why Not Epistocracy? In N. Reshotko (Ed.), *Desire, Identity and Existence: Essays in Honour of T.M. Penner* (pp. 53–69). Academic Printing and Publishing.
- Estlund, D. (2008). *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press.
- European Commission. (2005). *Europeans, Science and Technology*. Brussels: Directorate General Press and Communication.
- Evans, W. R., & Carson, C. M. (2005). A Social Capital Explanation of the Relationship between Functional Diversity and Group Performance. *Team Performance Management: An International Journal*, 11(7/8), 302–315.
- Evetts, J., Mieg, H. A., & Felt, U. (2006). Professionalism, Scientific Expertise, and Elitism: A Sociological Perspective. In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 105–124). Cambridge: Cambridge University Press.
- Fearon, J. D. (1998). Deliberation as Discussion. In J. Elster, *Deliberative Democracy* (pp. 44–68). Cambridge: Cambridge University Press.
- Fearon, J. D. (1998). Deliberation as Discussion. In J. Elster, *Deliberative Democracy* (pp. 44–68). Cambridge: Cambridge University Press.

- Fischer, F. (1993). Citizen Participation and the Democratization of Policy Expertise: From Theoretical Inquiry to Practical Cases. *Policy Sciences*, 26(3), 165–187.
- Fischer, F. (1998). Beyond Empiricism: Policy Inquiry in Postpositivist Perspective. *Policy Studies Journal*, 26(1), 129–146.
- Fischer, F. (2000). *Citizens, Experts, and the Environment: The Politics of Local Knowledge*. Durham: Duke University Press.
- Fischer, F. (2003). *Reframing Public Policy: Discursive Politics and Deliberative Practices*. New York: Oxford University Press.
- Fisher, E., Mahajan, R. L., & Mitcham, C. (2006). Midstream Modulation of Technology: Governance from within. *Bulletin of Science, Technology, and Society*, 26(6), 485–496.
- Fishkin, J. S. (2002). Deliberative Democracy. In R. L. Simon (Ed.), *The Blackwell Guide to Social and Political Philosophy* (pp. 221–238). Malden: Blackwell Publishers.
- Fishkin, J. S. (2003). Consulting the Public Through Deliberative Polling. *Journal of Policy Analysis and Management*, 22(1), 128–133.
- Fishkin, J. S. (2010). Response to Critics of When People Speak: The Deliberative Deficit and What To Do About It. *The Good Society*, 19(1), 68–76.
- Fishkin, J. S. (2011). Deliberative Democracy and Constitutions. *Social Philosophy and Policy*, 28(1), 242–260.
- Fishkin, J. S., & Luskin, R. C. (2005). Experimenting with a Democratic Ideal: Deliberative Polling and Public Opinion. *Acta Politica*, 40(3), 284–298.
- Forest, P. (2013). "Citizens As Analysts" Redux: Revisiting Aaron Wildavsky on Public Participation. *Journal of Public Deliberation*, 9(1:7), 1–13.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Friedman, J. (2014). Political Epistemology. *Critical Review: A Journal of Political and Society*, 26(1–2), i–xiv.
- Funtowicz, S. (2006). Why Knowledge Assessment? In Â. G. Pereira, S. G. Vaz, & S. Tognetti, *Interfaces Between Science and Society* (pp. 138–145). Sheffield: Greenleaf Publishing Ltd.
- Funtowicz, S. O., & Ravetz, J. R. (1993). Science for the Post-normal Age. *Futures*, 25(7), 739–755.
- Funtowicz, S. O., & Ravetz, J. R. (1994a). The Worth of a Songbird: Ecological Economics as a Post-Normal Science. *Ecological Economics*, 10(3), 197–207.
- Funtowicz, S. O., & Ravetz, J. R. (1994b). Uncertainty, Complexity and Post-normal Science. *Environmental Toxicology and Chemistry*, 13(12), 1881–1885.
- Funtowicz, S., & Ravetz, J. (2003). *Post-Normal Science*. Retrieved April 2014, from International Society for Ecological Economics: <http://isecoeco.org/pdf/pstnormsc.pdf>
- Gallopín, G., & Vessuri, H. (2006). Science for Sustainable Development: Articulating Knowledges. In A. G. Pereira, S. G. Vaz, & S. Tognetti, *Interfaces Between Science and Society* (pp. 35–52). Sheffield: Greenleaf Publishing Ltd.

- Gambetta, D. (1998). "Claro!": An Essay on Discursive Machismo. In J. Elster, *Deliberative Democracy* (pp. 19–43). Cambridge: Cambridge University Press.
- Gaskell, G., Ten Eyck, T., Jackson, J., & Veltri, G. (2005). Imagining Nanotechnology: Cultural Support for Technological Innovation in Europe and the United States. *Public Understanding of Science, 14*(1), 81–90.
- Goodin, R. E. (2003). *Reflective Democracy*. New York: Oxford University Press.
- Goodin, R. E., & Dryzek, J. S. (2006). Deliberative Impacts: The Macro-Political Uptake of Mini-Publics. *Politics and Society, 34*(2), 219–244.
- Grofman, B., Owen, G., & Feld, S. L. (1983). Thirteen Theorems in Search of the Truth. *Theory and Decision, 15*(3), 261–278.
- Groves, C. (2011). Public Engagement and Nanotechnology in the UK: Restoring Trust of Building Robustness? *Science and Public Policy, 38*(10), 783–793.
- Guido, P., & Tesón, F. R. (2006). *Rational Choice and Democratic Deliberation: A Theory of Discourse Failure*. New York: Cambridge University Press.
- Gunn, P. (2014). Democracy and Epistocracy. *Critical Review: A Journal of Politics and Society, 26*(1–2), 59–79.
- Gutmann, A., & Thompson, D. (1996). *Democracy and Disagreement*. Cambridge: Harvard University Press.
- Gutmann, A., & Thompson, D. (2004a). *Why Deliberative Democracy?* Princeton: Princeton University Press.
- Gutmann, A., & Thompson, D. (2004b). Deliberative Democracy Beyond Process. In C. Farrelly (Ed.), *Introduction to Contemporary Political Theory* (pp. 232–241). London: Sage Publications.
- Habermas, J. (2006). Political Communication in Media Society: Does Democracy Still Enjoy an Epistemic Dimension? The Impact of Normative Theory on Empirical Research. *Communication Theory, 16*(4), 411–426.
- Hamilton, A., Goldman, L., Jay, J., & Madison, J. (2008[1787–1788]). *Federalist Papers*. New York: Oxford University Press.
- Hamlett, P., Cobb, M. D., & Guston, D. H. (2008). *National Citizens' Technology Forum: Nanotechnologies and Human Enhancement*. Arizona State University. Tempe: Center for Nanotechnology in Society. Retrieved from https://cns.asu.edu/sites/default/files/library_files/lib_hamlettcobb_0.pdf
- Hardoš, P. (2014). *"Expert Review": On Reconciling Politics of Expertise and Democracy*. Budapest: Central European University (PhD Thesis).
- Hawkesworth, M. E. (1988). *Theoretical Issues in Policy Analysis*. Albany: University of New York.
- Head, B. W. (2008). Wicked Problems in Public Policy. *Public Policy, 101*–118.
- Head, B., & Alford, J. (2008). Wicked Problems: The Implications for Public Management. *International Research Society for Public Management 12th Annual Conference*, (pp. 1–26). Brisbane.
- Herek, G. M. (2012). *Facts About Homosexuality and Mental Health*. Retrieved from UC Davis Psychology: http://psychology.ucdavis.edu/rainbow/html/facts_mental_health.html

- Higgins, C. (2013, Summer). *Post-Truth Pluralism: The Unlikely Political Wisdom of Nietzsche*. Retrieved from The Breakthrough: <http://thebreakthrough.org/index.php/journal/past-issues/issue-3/post-truth-pluralism>
- Hiller, A. (2011). Climate Change and Individual Responsibility. *The Monist*, 94(3), 349–368.
- Holst, C. (2012). What is epistocracy? Dimensions of Knowledge-Based Rule. In S. A. Øyen, T. Lund-Olsen, & N. S. Vaage (Eds.), *Sacred Science?: On Science and Its Interrelations with Religious Worldviews* (pp. 41–54). Wageningen: Wageningen Academic Publishers.
- Holst, C. (2014). *Expertise and Democracy*. Oslo: ARENA Centre for European Studies. Retrieved from <http://www.sv.uio.no/arena/english/research/publications/arena-reports/2014/report-01-14.pdf>
- Hong, L., & Page, S. E. (2004). Groups of Diverse Problem Solvers Can Outperform Groups of High-Ability Problem Solvers. *Proceedings of the National Academy of Sciences*, 101(46), 16385–16389.
- Hookway, C. (2010). Some Varieties of Epistemic Injustice: Reflections on Fricker. *Episteme*, 7(2), 151–163.
- Hoppe, R. (1999). Policy Analysis, Science and Politics: from "Speaking Truth to Power" to "Making Sense Together". *Science and Public Policy*, 26(3), 201–210.
- Horwitz, S. K. (2005). The Compositional Impact of Team Diversity on Performance: Theoretical Considerations. *Human Resources Development Review*, 4(2), 219–245.
- Horwitz, S. K., & Horwitz, I. B. (2007). The Effects of Team Diversity on Team Outcomes: A Meta-Analytic Review of Team Demography. *Journal of Management*, 33(6), 987–1015.
- Hosseini, S. M., & Rezaei, R. (2011). Factors Affecting the Perceptions of Iranian Agricultural Researchers towards Nanotechnology. *Public Understanding of Science*, 20(4), 513–524.
- Iyengar, S., Luskin, R. C., & Fishkin, J. S. (2003). *Facilitating Informed Public Opinion: Evidence from Face-to-Face and Online Deliberative Polls*. Retrieved from Semantic Scholar: <https://pdfs.semanticscholar.org/7b78/8da09a51e1367a803b6a502ce52e6cea5985.pdf>
- Jasanoff, S. (2003). Technologies of Humility: Citizen Participation in Governing Science. *Minerva*, 41(3), 223–244.
- Jehn, K. A., & Mannix, E. A. (2001). The Dynamic Nature of Conflict: A Longitudinal Study of Intragroup Conflict and Group Performance. *The Academy of Management Journal*, 44(2), 238–251.
- Jehn, K. A., Northcraft, G. B., & Neale, M. A. (1999). Why Differences Make a Difference: A Field Study of Diversity, Conflict, and Performance in Workgroups. *Administrative Science Quarterly*, 44(4), 741–763.

- Johnson, J. (1998). Arguing for Deliberation: Some Sceptical Considerations. In J. Elster, *Deliberative Democracy* (pp. 161-184). Cambridge: Cambridge University Press.
- Joss, S. (1998). Danish Consensus Conferences as a Model of Participatory Technology Assessment: An Impact Study of Consensus Conferences on Danish Parliament and Danish Public Debate. *Science and Public Policy*, 25(1), 2–22.
- Kanazawi, S. (1998). A Brief Note on a Further Refinement of the Condorcet Jury Theorem. *Mathematical Social Sciences*, 35(1), 69–73.
- Kaniovski, S. (2010). Aggregation of Correlated Votes and Condorcet's Jury Theorem. *Theory and Decision*, 69(3), 453–468.
- Katz, E., Solomon, F., Mee, W., & Lovel, R. (2009). Evolving Scientific Research Governance in Australia: A Case Study of Engaging Interested Publics in Nanotechnology. *Public Understanding of Science*, 18(5), 531–545.
- Kearnes, M., Macnaghten, P., & Wilsdon, J. (2006). *Governing at the Nanoscale: People, Policies, and Emerging Technologies*. London: Demos.
- Kearney, E., Gebert, D., & Voelpel, S. C. (2009). When and How Diversity Benefits Teams: The Importance for Team Members' Need for Cognition. *Academy of Management Journal*, 52(3), 581–598.
- Kjaergaard, R. S. (2010). Making a Small Country Count: Nanotechnology in Danish Newspapers from 1996 to 2006. *Public Understanding of Science*, 19(1), 80–97.
- Kjølborg, K. L. (2009). Representations of Nanotechnology in Norwegian Newspapers: Implications for Public Participation. *Nanoethics*, 3(1), 61–72.
- Kleinman, D. L., Delborne, J. A., & Anderson, A. A. (2011). Engaging Citizens: The High Cost of Citizen Participation in High Technology. *Public Understanding of Science*, 20(2), 221–204.
- Knight, J., & Johnson, J. (1997). What Sort of Political Equality Does Deliberative Democracy Require? In J. Bohman, & W. Rehg, *Deliberative Democracy: Essays on Reason and Politics* (pp. 279–319). Cambridge: MIT Press.
- Kuhn, T. S. (1970). *The Structure of Scientific Revolutions (2nd ed.)*. Chicago: University of Chicago Press.
- Kurath, M., & Gisler, P. (2009). Informing, Involving, or Engaging? Science Communication in the Age of Atom-, Bio-, and Nanotechnology. *Public Understanding of Science*, 18(5), 559–573.
- Kurtzberg, T. R. (2005). Feeling Creative, Being Creative: An Empirical Study of Diversity and Creativity in Teams. *Creativity Research Journal*, 17(1), 51–65.
- Kyle, R., & Dodds, S. (2009). Avoiding Empty Rhetoric: Engaging Publics in Debates About Nanotechnologies. *Science and Engineering Ethics*, 15(1), 81–96.
- Ladha, K. (1992). The Condorcet Jury Theorem, Free Speech, and Correlated Votes. *American Journal of Political Science*, 36(3), 617–614.
- Ladha, K. K. (1995). Information Pooling through Majority-Rule Voting: Condorcet's Jury Theorem with Correlated Votes. *Journal of Economic Behaviour and Organization*, 26(3), 353–372.

- Landemore, H. (2012a). *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*. Princeton: Princeton University Press.
- Landemore, H. (2014a). Beyond the Fact of Disagreement? The Epistemic Turn in Deliberative Democracy. *Epistemic Dimensions of Democracy Revisited*, (pp. 1–39).
- Landemore, H. (2014b). Yes, We Can (Make It Up on Volume): Answers to Critics. *Critical Review: A Journal of Politics and Society*, 26(1–2), 184–237.
- Landemore, H. (2017). Beyond the Fact of Disagreement? The Epistemic Turn in Deliberative Democracy. *Social Epistemology*, 31(3), 277–295.
- Laurent, B. (2008). *Engaging the Publi in Nanotechnology? Three Visions of Public Engagement*. Ecole des Mines de Paris. Paris: Centre de Sociologie de L'Innovation. Retrieved from <https://pdfs.semanticscholar.org/2680/7c75e45a080fd1e685c3da81e44915eb394b.pdf>
- Lecce, S. (2009). Should Democracy Grow up? Children and Voting Rights. *Intergenerational Justice Review*, 9(4), 133–139.
- Lee, C., & Farh, J.-L. (2004). Joint Effects of Group Efficacy and Gender Diversity on Group Cohesion and Performance. *Applied Psychology: An International Review*, 53(1), 136–154.
- Liang, X., Ho, S. S., Brossard, D., Xenos, M. A., Scheufele, D. A., Anderson, A. A., . . . He, X. (2015). Value Predispositions as Perceptual Filters: Comparing of Public Attitudes toward Nanotechnology in the United States and Singapore. *Public Understanding of Science*, 24(5), 582–600.
- Lippert-Rasmussen, K. (2012). Estlund on Epistocracy: A Critique. *Res Publica*, 18(3), 241–258.
- Lippmann, W. (1922). *Public Opinion*. New York: MacMillan Co.
- Lippmann, W. (1993[1925]). *The Phantom Public*. New Brunswick: Transaction Publishers.
- List, C., & Goodin, R. E. (2001). Epistemic Democracy: Generalizing the Condorcet Jury Theorem. *The Journal of Political Philosophy*, 9(3), 277–306.
- Luskin, R. C., & Fishkin, J. S. (2002). *Deliberation and "Better Citizens"*. Center for Deliberative Democracy. Retrieved from Center for Deliberative Democracy: <http://cdd.stanford.edu/2002/deliberation-and-better-citizens/>
- Luskin, R. C., Fishkin, J. S., & Hahn, K. S. (2007). Consensus and Polarization in Small Group Deliberations. *American Political Science Association* (pp. 1–19). Chicago: Center for Deliberative Democracy. Retrieved from <https://cdd.stanford.edu/2007/consensus-and-polarization-in-small-group-deliberations/>
- Luskin, R. C., O'Flynn, I., Fishkin, J. S., & Russell, D. (2012). Deliberating across Deep Divides. *Political Studies*, 1–20. doi:<http://doi.org/10.1111/j.1467-9248.2012.01005.x>
- Lynn, L. E. (1999). A Place at the Table: Policy Analysis, Its Postpositive Critics, and the Future of Practice. *Journal of Policy Analysis and Management*, 18(3), 411–424.

- Lyons, K., & Whelan, J. (2010). Community Engagement to Facilitate, Legitimize and Accelerate the Advancement of Nanotechnologies in Australia. *Nanoethics*, 4(1), 53–66.
- Macnaghten, P. (2010). Researching Technoscientific Concerns in the Making: Narrative Structures, Public Responses, and Emerging Nanotechnologies. *Environment and Planning A*, 42(1), 23–37.
- Macnaghten, P., & Guivant, J. S. (2011). Converging Citizens? Nanotechnology and the Political Imaginary of Public Engagement in Brazil and the United Kingdom. *Public Understanding of Science*, 20(2), 207–220.
- Macnaghten, P., Kearnes, M. B., & Wynne, B. (2005). Nanotechnology, Governance, and Public Deliberation: What Role for the Social Sciences? *Science Communication*, 27(2), 268–291.
- Macoubrie, J. (2006). Nanotechnology: Public Concerns, Reasoning and Trust in Government. *Public Understanding of Science*, 15(2), 221–241.
- Majone, G. (1989). *Evidence, Argument, and Persuasion in the Policy Sciences*. New Haven: Yale University Press.
- Martí, J. L. (2006). *The Epistemic Conception of Deliberative Democracy Defended: Reasons, Rightness and Equal Political Autonomy*. Retrieved from <http://www.trinitinture.com/documents/martimarmol.pdf>
- MBIE. (2014). *A Nation of Curious Minds (He Whenua Hihiri I Te Mahara): A National Strategic Plan for Science in Society*. Wellington: Ministry of Business, Innovation & Employment.
- Medina, J. (2012). Hermeneutical Injustice and Polyphonic Contextualism: Social Silences and Shared Hermeneutical Responsibilities. *Social Epistemology*, 26(2), 201–220.
- Mill, J. S. (1861). *Considerations on Representative Government*. London: Parker, Son and Bourn.
- Miller, C. C., Burke, L. M., & Glick, W. H. (1998). Cognitive Diversity Among Upper-Echelon Executives. *Strategic Management Journal*, 19(1), 39–58.
- Milliken, F. J., & Martins, L. L. (1996). Searching for Common Threads: Understanding the Multiple Effects of Diversity in Organisational Groups. *Academy of Management Review*, 21(2), 402–433.
- Mitchell, R., & Nicholas, S. (2006). Knowledge Creation in Groups: The Value of Cognitive Diversity, Transactive Memory, and Open-Mindedness Norms. *The Electronic Journal of Knowledge Management*, 4(1), 67–74.
- Mitchell, R., Nicholas, S., & Boyle, B. (2009). The Role of Openness to Cognitive Diversity and Group Processes in Knowledge Creation. *Small Group Research*, 40(5), 535–554.
- Mohammed, S., & Dumville, B. C. (2001). Team Mental Models in a Team Knowledge Framework: Expanding Theory and Measurement across Disciplinary Boundaries. *Journal of Organizational Behaviour*, 22(2), 89–106.
- Moore, A. (2014). Democratic Reason, Democratic Faith, and the Problem of Expertise. *Critical Review: A Journal of Politics and Society*, 26(1–2), 101–114.

- Moschella, D. (2015, September 2). *The Pace of Technology Change is Not Accelerating*. Retrieved from Leading Edge Forum: <https://leadingedgeforum.com/publication/the-pace-of-technology-change-is-not-accelerating-2502/>
- Munshi, D., Kurian, P. A., Morrison, T., & Morrison, S. L. (2016). Redesigning the Architecture of Policymaking: Engaging with Māori on Nanotechnology in New Zealand. *Public Understanding of Science*, 25(3), 287–302.
- Nelson, R. R. (1977). *The Moon and the Ghetto: An Essay on Public Policy Analysis*. New York: W.W. Norton.
- Nichols, T. (2017). *The Death of Expertise: The Campaign Against Established Knowledge and Why It Matters*. New York: Oxford University Press.
- Nisbet, M. C., & Scheufele, D. A. (2009). What's Next for Science Communication? Promising Directions and Lingering Distractions. *American Journal of Botany*, 96(10), 1767–1778.
- Olsen, B. J., Parayitam, S., & Bao, Y. (2007). Strategic Decision Making: The Effects of Cognitive Diversity, Conflict, and Trust on Decision Outcomes. *Journal of Management*, 33(2), 196–222.
- Origg, G. (2012). Epistemic Injustice and Epistemic Trust. *Social Epistemology*, 26(2), 221–235.
- Paddock, L. (2010). An Integrated Approach to Nanotechnology Governance. *UCLA Journal of Environmental Law*, 28(1), 251–290.
- Page, B. (2015). That Same Old Song: Somin on Political Ignorance. *Critical Review*, 27(3–4), 375–379.
- Page, B., & Shapiro, R. (1993). The Rational Public and Democracy. In G. Marcus, & R. Hanson, *Reconsidering the Democratic Public* (pp. 35–64). University Park: Pennsylvania State University Press.
- Page, S. E. (2007). *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*. Princeton: Princeton University Press.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- Parkinson, J. (2003). Legitimacy Problems in Deliberative Democracy. *Political Studies*, 51(1), 180–196.
- Pateman, C. (2012). Participatory Democracy Revisited. *Perspectives on Politics*, 10(1), 7–19.
- Pelled, L. H. (1996). Demographic Diversity, Conflict, and Work Group Outcomes: An Intervening Process Theory. *Organization Science*, 7(6), 615–631.
- Petersen, A., Anderson, A., Allan, S., & Wilkinson, C. (2009). Opening the Black Box: Scientists' Views on the Role of the News Media in the Nanotechnology Debate. *Public Understanding of Science*, 18(5), 512–530.
- Petterson, B. (2004). The Second Mistake in Moral Mathematics Is Not about the Worth of Mere Participation. *Utilitas*, 16(3), 288–315.
- Pidgeon, N., & Rogers-Hayden, T. (2007). Opening Up Nanotechnology Dialogues with the Publics: Risk Communication or "Upstream Engagement". *Health, Risk, and Society*, 9(2), 191–210.

- Pielke, R. A. (2007). *The Honest Broker: Making Sense of Science in Policy and Politics*. Cambridge: Cambridge University Press.
- Plato. (375BC [2007]). *The Republic*. (M. Lane, R. Kamtekar, Eds., & D. Lee, Trans.) London: Penguin Books.
- Powell, M. C., & Colin, M. (2008). Meaningful Citizen Engagement in Science and Technology: What Would It Really Take? *Science Communication*, 30(1), 126–136.
- Powell, M., & Kleinman, D. L. (2008). Building Citizen Capacities for Participation in Nanotechnology Decision Making: The democratic virtues of the consensus conference model. *Public Understanding of Science*, 17(3), 329–348.
- Priest, S. H., & Greehalgh, T. (2011). Nanotechnology as an Experiment in Democracy: How Do Citizens Form Opinions about Technology and Policy? *Journal of Nanoparticle Research*, 13(4), 1521–1531.
- Quirk, P. J. (2014). Making it up on Volume: Are Larger Groups Really Smarter? *Critical Review: A Journal of Politics and Society*, 26(1–2), 129–150.
- Quong, J. (2010). The Distribution of Authority. *Representation*, 46(1), 35–52.
- Rahimpour, M., Rahimpour, M., Gomari, H., Shirvani, E., Niroumanesh, A., Saremi, K., & Sardari, S. (2012). Public Perceptions of Nanotechnology: A survey in the magacities of Iran. *Nanoethics*, 6(2), 119–126.
- RAND Corporation. (2015). *Standards for High Quality Research and Analysis*. Retrieved from RAND Corporation: http://www.rand.org/standards/standards_high.html
- Rawls, J. (1997). The Idea of Public Reason. In J. Bohman, & W. Rehg, *Deliberative Democracy: Essays on Reason and Politics* (pp. 93–130). Cambridge: MIT Press.
- Retzback, A., Marschall, J., Rahnke, M., Otto, L., & Meier, M. (2011). Public Understanding of Science and the Perception of Nanotechnology: the Roles of Interest in Science, Methodological Knowledge, Epistemological Beliefs, and Beliefs in Science. *Journal of Nanoparticle Research*, 13(12), 6231–6244.
- Richardson, H. S. (1997). Democratic Intentions. In J. Bohman, & W. Rehg, *Deliberative Democracy: Essays on Reason and Politics* (pp. 349–382). Cambridge: MIT Press.
- Richey, S. (2013). Random and Systematic Error in Voting in Presidential Elections. *Political Research Quarterly*, 66(3), 645–657.
- Rittel, H. W., & Webber, M. M. (1973). Dilemmas in a General Theory of Planning. *Policy Sciences*, 155–169.
- Roberge, M.-É., & van Dick, R. (2010). Recognizing the Benefits of Diversity: When and How Does Diversity Increase Group Performance? *20(4)*, 295–308.
- Roco, M. C., Harthorn, B., Guston, D., & Shapira, P. (2011). Innovative and Responsible Governance of Nanotechnology for Societal Development. *Journal of Nanoparticle Research*, 13(9), 3557–3590.
- Rogers-Hayden, T., & Pidgeon, N. (2006). Reflecting Upon the UK's Citizens' Jury on Nanotechnologies: NanoJury UK. *Nanotechnology Law and Business*, 3(2), 167–180.
- Ross, K. G., Shafer, J. L., & Klein, G. (2006). Professional Judgments and "Naturalistic Decision Making". In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R.

- Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 403–420). Cambridge: Cambridge University Press.
- Salas, E., Rosen, M. A., Burke, C. S., Goodwin, G. F., & Fiore, S. M. (2006). The Making of a Dream Team: When Expert Teams Do Best. In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 439–453). Cambridge: Cambridge University Press.
- Sandberg, J. (2011). “My Emissions Make No Difference”: Climate Change and the Argument from Inconsequentialism. *Environmental Ethics*, 33(3), 229–248.
- Sardar, Z. (2010). Welcome to Postnormal Times. *Futures*, 42(5), 435–444.
- Schiavone, G., Mamei, M., & Boniolo, G. (2015). Epistocracy for Online Deliberative Bioethics. *Cambridge Quarterly of Healthcare Ethics*, 24(3), 272–280.
- Schwarz-Plaschg, C. (2018). Nanotechnology is Like... The Rhetorical Roles of Analogies in Public Engagements. *Public Understanding of Science*, 27(2), 153–167.
- Serota, M., & Leib, E. J. (2013). The Political Morality of Voting in Direct Democracy. *Minnesota Law Review*, 97(5), 1596–1620.
- Shapiro, I. (1999). Enough of Deliberation: Politics is about Interests and Power. In S. Macedo, *Deliberative Politics: Essays on Democracy and Disagreement* (pp. 28–38). New York: Oxford University Press.
- Sherman, B. R. (2016). There's No (Testimonial) Justice: Why Pursuit of a Virtue is Not the Solution to Epistemic Injustice. *Social Epistemology*, 30(2), 229–250.
- Sinnott-Armstrong, W. (2005). It's Not My Fault: Global Warming and Individual Moral Obligations. In W. Sinnott-Armstrong, & R. B. Howarth (Eds.), *Perspectives on Climate Change: Science, Economics, Politics, Ethics (Advances in the Economics of Environmental Resources, Volume 5)* (pp. 293–315). Bingley: Emerald Group Publishing Limited.
- Smith, H. M. (2014). The Subjective Moral Duty to Inform Oneself Before Acting. *Ethics*, 125(1), 11–38.
- Somin, I. (2013). *Democracy and Political Ignorance: Why Smaller Government is Smarter*. Stanford: Stanford University Press.
- Somin, I. (2014). Why Political Ignorance Undermines the Wisdom of the Many. *Critical Review: A Journal of Politics and Society*, 26(1–2), 151–169.
- Stanford History Education Group. (2016). *Evaluating Information: The Cornerstone of Civic Online Reasoning*. Stanford: Stanford University. Retrieved from <https://sheg.stanford.edu/upload/V3LessonPlans/Executive%20Summary%2011.21.16.pdf>
- Stebbing, M. (2009). Avoiding the Trust Deficit: Public Engagement, Values, the Precautionary Principle and the Future of Nanotechnology. *Bioethical Inquiry*, 6(1), 37–48.
- Stehr, N., & Grundmann, R. (2011). *Experts: The Knowledge and Power of Expertise*. New York: Routledge.
- Stilgoe, J. (2007). *Nanodialogues: Experiments in Public Engagement in Science*. London: Demos.

- Sturgis, P. (2014). On the Limits of Public Engagement for the Governance of Emerging Technologies. *Public Understanding of Science*, 23(1), 38–42.
- Tetlock, P. E. (2005). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton: Princeton University Press.
- Timmerman, T. A. (2000). Racial Diversity, Age Diversity, Interdependence, and Team Performance. *Small Group Research*, 31(5), 592–606.
- Torgerson, D. (1985). Contextual Orientation in Policy Analysis: The Contribution of Harold D. Lasswell. *Policy Sciences*, 18(3), 241–261.
- Torgerson, D. (1986). Between Knowledge and Politics: Three faces of policy analysis. *Policy Sciences*, 19(1), 33–59.
- Tribe, L. H. (1972). Policy Science: Analysis or Ideology? *Philosophy and Public Affairs*, 2(1), 66–110.
- Turner, S. P. (2003). *Liberal Democracy 3.0*. London: SAGE Publication Ltd.
- van der Vegt, G. S., & Janssen, O. (2003). Joint Impact of Interdependence and Group Diversity on Innovation. *Journal of Management*, 29(5), 729–751.
- van Giesen, R. I., Fischer, A. R., & van Trijp, H. C. (2018). Changes in the Influence of Affect and Cognition over Time on Consumer Attitude Formation toward Nanotechnology: A Longitudinal Survey Study. *Public Understanding of Science*, 27(2), 168–184.
- van Knippenberg, D., & Schippers, M. C. (2007). Work Group Diversity. *Annual Review of Psychology*, 58, 515–541.
- van Knippenberg, D., Carsten, K., & Homan, A. C. (2004). Work Group Diversity and Group Performance: An Integrative Model and Research Agenda. *Journal of Applied Psychology*, 89(6), 1008–1022.
- Vanderheiden, S. (2006). Climate Change and the Challenge of Moral Responsibility. *Journal of Philosophical Research*, 32, 85–92.
- Vandermoere, F., Blanchemanche, S., Bieberstein, A., Marette, S., & Roosen, J. (2011). The Public Understanding of Nanotechnology in the Food Domain: The Hidden Role of Views on Science, Technology and Nature. *Public Understanding of Science*, 20(2), 195–206.
- Wagle, U. (2000). The Policy Science of Democracy: The Issues of Methodology and Citizen Participation. *Policy Sciences*, 33(2), 207–223.
- Waldron, A. M., Spencer, D., & Batt, C. A. (2006). The Current State of Public Understanding of Nanotechnology. *Journal of Nanoparticle Research*, 8(5), 569–575.
- Wegge, J., Roth, C., Neubach, B., Schmidt, K.-H., & Kanfar, R. (2008). Age and Gender as Determinants of Performance and Health in a Public Organization: The Role of Task Complexity and Group Size. *Journal of Applied Psychology*, 93(6), 1301–1303.
- Whitehouse Office of Science and Technology Policy. (2014). *The 2015 Budget: Science, Technology, and Innovation for Opportunity and Growth*. Washington, D.C.: Whitehouse. Retrieved from

<https://www.whitehouse.gov/sites/default/files/microsites/ostp/Fy%202015%20R&D.pdf>

- Wildavsky, A. (1979). *Speaking Truth to Power: The Art and Craft of Policy Analysis*. Boston: Little, Brown and Company.
- Wynne, B. (1996). May the Sheep Safely Graze? A Reflexive View of the Expert–Lay Knowledge Divide. In S. Lash, B. Szerszynski, & B. Wynne (Eds.), *Risk, Environment and Modernity: Towards a New Ecology* (pp. 44–83). London: SAGE Publications.
- Yadav, Y. (1986). Political 'Science': Positivist method and philosophic critiques. *The Indian Journal of Political Science*, 47(4), 502–517.
- Yates, J. F., & Tschirhart, M. D. (2006). Decision-Making Expertise. In K. A. Ericsson, N. Charness, P. J. Feltovich, & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp. 421–438). Cambridge: Cambridge University Press.
- Yi-Fan Su, L., Cacciatore, M. A., Scheufele, D. A., Brossard, D., & Xenos, M. A. (2014). Inequalities in Scientific Understanding: Differentiating Between Factual and Perceived Knowledge Gaps. *Science Communication*, 36(3), 352–378.
- Young, I. M. (2001). Activist Challenges to Deliberative Democracy. *Political Theory*, 5, 670–690.
- Young, I. M. (2004). The Deliberative Model. In C. Farrelly (Ed.), *Introduction to Contemporary Political Theory: A Reader* (pp. 227–231). London: Sage Publications.