

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Democracy, Epistocracy, and the Epistemic Problems of Politics: How Centralized Decision Making Undermines Three Fundamental Democratic Values

Permalink

<https://escholarship.org/uc/item/49q2m9xq>

Author

Richardson, Brandon Samuel

Publication Date

2020

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Democracy, Epistocracy, and the Epistemic Problems of Politics: How Centralized Decision
Making Undermines Three Fundamental Democratic Values

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

In Philosophy

by

Brandon Samuel Richardson

Dissertation Committee:
Professor Annalisa Coliva, Chair
Professor Aaron James, co-Chair
Professor Karl Schafer

2020

DEDICATION

To Shanan and Peyton,
whose constant support and encouragement made this work possible.

TABLE OF CONTENTS

	Page
AKNOWLEDGEMENTS	iv
VITA	v
ABSTRACT OF THE DISSERTATION	vii
INTRODUCTION	1
CHAPTER 1: A Trilemma of Democratic Values	8
CHAPTER 2: Competence and Competent Government	49
CHAPTER 3: Epistocracy, Competence, and the Trouble with Trial and Error	73
CHAPTER 4: Complexity, Encapsulation, and Decentralization	118
REFERENCES	149

ACKNOWLEDGEMENTS

I would like to thank my advisor, Annalisa Coliva, and my committee co-chair, Aaron James, for their guidance and feedback throughout this project.

I would also like to thank G. Michael Pace, Daniel Pilchman, and Gordon Babst for encouraging me to become a philosopher.

VITA

Brandon S. Richardson

Areas of Specialization:

Epistemology and Political Philosophy

Areas of Competence:

History of Philosophy, Philosophy of Mind, and Ethics.

Education **PhD University of California, Irvine**, Philosophy, *June 2020*
Advisor: Annalisa Coliva

MA University of California, Irvine, Philosophy, *June 2016*
Advisor: Jeff Helmreich

BA Chapman University, Philosophy and Political Science, *May 2014*

Dissertation: *Democracy, Epistocracy, and the Epistemic Problems of Politics: How Centralized Decision Making Undermines Three Democratic Values*

Committee: Annalisa Coliva (Chair), Aaron James (co-Chair), Karl Schafer

Teaching **Lecturer**

Experience *Belief, Truth, and Knowledge*
Introduction to epistemology
23 Students
Chapman University Spring 2019

Teaching Assistant

Intro to Knowledge
Upper division epistemology and philosophical writing, 2 discussion sections of 23 students each. Fall 2016

Puzzles and Paradoxes

Basic topics in philosophy, deductive reasoning, logic, and history of philosophy. Two sections of 35 students each. Fall 2017, Winter 2018, Spring 2018, and Summer 2019

Introduction to Philosophy

Basic introduction to philosophical topics. Two sections with 30 students each. Winter 2017 and Fall 2018

Introduction to Ethics

Basic topics in ethics and political philosophy. Two sections with 30 students each. Spring 2017, Winter 2019, and Summer II 2019

Think Program

Teach basic concepts of philosophy to 5th graders.

Selected

Graduate

Coursework

At UCI:

Medical Epistemology, Coliva and Bernecker, Winter 2018

Rights and Demands, Gilbert, Fall 2017

Plato, Perin, Fall 2017

History of Analytic Philosophy II, Heis, Spring 2017

Epistemology of Conspiracy Theories, Bernecker, Spring 2017

Kant's Ethics, Bencivenga, Winter 2017

Hume's Treatise, Schafer, Winter 2017

Hinge Epistemology, Coliva, Fall 2016

Set Theory, Walsh, Fall 2016

Self Knowledge, Coliva, Spring 2016

Kant's Critique of Pure Reason I & II, Heis and Greenberg, Winter/Spring 2016

Rawls and Utilitarianism, Skyrms and James, Winter 2016

Plato's Republic, Perin, Winter 2016

ABSTRACT OF THE DISSERTATION

Democracy, Epistocracy, and the Epistemic Problems of Politics: How Centralized Decision

Making Undermines Three Fundamental Democratic Values

by

Brandon Samuel Richardson

Doctor of Philosophy in Philosophy

University of California, Irvine, 2020

Professor Annalisa Coliva, Chair

The empirical evidence of widespread ignorance of voters in modern democracies has led some philosophers to question traditional justifications for democracy. These philosophers attempt to resolve this problem by arguing for limits on who can vote with the goal of increasing the competence of the average voter. Epistemic democrats attempt to assuage these worries by looking to *a priori* theorems that purport to show that democracy can be competent despite the widespread ignorance of voters. In this dissertation I argue for a fuller account of what it means to be competent in making political decisions, without which we face a problem I call the graveyard spiral. I then use this account of competence to show that epistemic democrats and epistocrats fail to give us a reason to expect competence in government, which leaves us open to the graveyard spiral problem. Because of this, I argue that these views face a trilemma of democratic values. I argue that in order for competence to be satisfied, public problems should be addressed at the lowest levels of complexity to reduce the problems with tracking the effects of decisions, but also to increase the variety of solutions to public problems. Through a process

of encapsulating public problems and decentralizing the decision making process, we can expect meta-competence to be satisfied and to avoid the trilemma.

Introduction

Democracy finds wide-spread support because it is considered a fair means to make political decisions. It gives citizens an equal say in the decision-making process and allows citizens to collectively determine the laws and policies they will live under. Recently, however, some philosophers have questioned the ability of democracy to make competent decisions. Research into voter knowledge of basic political facts—including facts about the political beliefs of active politicians and major historical events—has revealed that voters tend to be misinformed. These philosophers argue that we have a right to competent government, but that right cannot obtain if large segments of the electorate are misinformed. This brings into tension three democratic values:

1. The non-instrumental value of universal inclusion
2. The collective right to self-determination
3. The right to competent government

The purpose of this dissertation is to argue that epistemic defenses of democracy and recent arguments for epistocracy fail to give us a reason to believe that the problem of voter ignorance can be overcome, which leaves these kinds of accounts unable to maintain the three values above. I argue that through decentralization of the decision making process, we can hold all three values together.

Epistemic democrats fail to account for how democracy can be competent with a largely misinformed electorate. Epistocrats like Jason Brennan (2016) attempt to secure the right to

competent government by abandoning universal suffrage, and therefore rejecting the non-instrumental value of universal inclusion. I argue that neither can account for the right to competent government given the epistemic problems of politics (chapter 3). Indeed, it appears that any epistemic defense of democracy or epistocracy must give us a reason to believe that mistakes will be identified and corrected. I argue they cannot, and that the only available means to resolving these problems is through a decentralized system of problem solving.

Decentralization allows attempted solutions to be tracked with greater reliability and multiple solutions can be attempted simultaneously such that evidence can be gathered regarding what works and what does not.

Epistemic democrats attempt to resolve the problem of voter ignorance by appealing to three mathematical theorems. These theorems attempt to show that individual knowledge is much less important to collective decisions than we might intuitively think. Epistocrats reject these arguments as telling us much of anything about real-world democracy, and so advocate restricting suffrage in one way or another to limit the influence of voters who fail to have basic political knowledge. If epistocrats are correct, then these theorems do nothing to establish the right to competent government. Hence, epistemic democrats fail to be able to account for all three values above.

Epistocrats see the right to competent government as overriding considerations of fairness on the grounds that better decisions will be better for everyone. If restricting suffrage leads to better decisions, then we will all be better off. In chapter 1, I show that Brennan (2016) fails to establish that epistocracy will be capable of identifying errors and correcting mistakes. If we have no reason to think errors will be identified and corrected, then there is no reason to think that Brennan's account of epistocracy can maintain the right to competent government. It

appears, then, that epistocracy rejects the non-instrumental value of universal inclusion *and* fails to secure the right to competent government.

In chapter 2, I show that the intuitive conception of competence presented in Brennan (2016) only accounts for part of what it means to be competent. I utilize the conception of competence found in Sosa (2010) to establish that competence means more than performing well. Indeed, competence requires both the ability to perform well, but also the ability to aptly judge when an impending performance will be competent and to be situated such that it will appear to us when we go wrong. Hence, we can split competence into a first and second order. First-order competence will be the ability to perform well and second-order competence will be the capacity to judge when an impending performance will be first-order competent. With both first and second-order competence satisfied, meta-competence will be satisfied.

Without being competent in this way, we set ourselves up to mistake errors for corrections and corrections for errors. I call this situation the graveyard spiral after the situation in which pilots find themselves when flying in areas of poor visibility. Without the ability to orient the plane competently, pilots can quickly lose control over the plane and enter a lethal spiral that cannot be escaped unless visibility is regained. In order to escape this possibility in politics, meta-competence must be satisfied so that both competent decisions can be made and errors can be corrected when they occur.

In chapter 3 I use the conception of competence from the previous chapter to assess the likelihood that democracy and epistocracy can, or will, be competent. There I argue that both fail for similar reasons. Democracies with a largely misinformed electorate cannot be expected to make competent decisions because they fail to have first-order competence. That is, the

absence of basic political knowledge should lead us to conclude that democracy will fail to satisfy first-order competence. Epistocracy, as characterized by Brennan, fails to satisfy first-order competence as well. As Brennan admits, epistocracy on his account will be dominated by what he calls political hooligans (Brennan 2016, 209). According to Brennan, most people who participate in politics are hooligans who are inclined toward bad faith inquiry and tribalism at the expense of what is true. Because of this, they cannot be expected to admit their mistakes and correct them in the way that meta-competence requires.

I further consider whether an epistocracy with an electorate of informed, good faith inquirers can satisfy meta-competence. I argue that they cannot be expected to satisfy meta-competence because they fail to satisfy second-order competence. The basic problem they face is the inability to compare the option chosen during an election with the option that was rejected. Without the ability to gather evidence about options not chosen, there is no means to compare the actual effects of the chosen option and the one that was not. Without this ability to make direct comparisons, we are left with *a priori* intuitions about how the option not chosen would have fared compared to the actual effects of the option that was chosen. Indeed, in any case where the electorate decided that they had made an error, the possibility is always open that the option not chosen would have been worse. If we cannot generate evidence about the other option, we have no grounds to claim we made a mistake, even when things do not go the way we hoped they would.

Further, given the complex nature of political decisions, the means we have for tracking the effects of the chosen option rely on metrics and statistical analyses that do not satisfy second-order competence. Because of this, confidence in our ability to track the effects of our decisions should be weakened. When taken together, our inability to gather evidence about how other

options would have fared and the difficulties in determining the effects of the options we do choose, we should expect that making political decisions at the highest levels of complexity will not be competent, even with a well-motivated and informed electorate.

In chapter 4, I argue that in order to resolve these problems we need a means to address public problems in a way that allows for a multiplicity of solutions to be tested so that we can compare the results. Moreover, we need a means to reduce the complexity of the problems we are addressing and the solutions to those problems. I look to other domains where complexity poses problems for decision makers to find a plausible means for satisfying meta-competence in political decision making. I argue that by looking at decentralization in military tactical command and encapsulation in computer science, we can develop a basic strategy for eliminating the epistemic problems identified in the previous chapters.

Encapsulation is the idea that problems can be broken up into smaller, more manageable bits to make them easier to solve. One of the virtues of encapsulation is that it allows problems that arise in large, complex computer programs to be resolved more easily. Similarly, encapsulation allows for bugs in one part of a program to be contained and therefore easier to manage. Decentralization in military tactical command is the idea that tactical decisions are better made by those closest to the problems that confront soldiers on the battlefield. With sufficient training, battlefield commanders, including small-unit leaders, can make more effective decisions in real time than waiting for decisions to be made higher up the chain of command. In addition to giving license to lower-level leaders to make decisions, decentralized tactical command allows smaller units to take advantage of their greater ability for variety of operations than large, centrally-commanded units. The disparity of variety that exists between

large-scale, conventional forces and small guerilla units underlies the advantages of guerilla warfare against much larger and more powerful forces.

I argue that we can take these insights and attempt to apply them to the political domain. By breaking problems up into more manageable bits, we can better manage problems as they arise and better control problems when solutions fail. That is, when a graveyard spiral happens, and errors are taken to be corrections and corrections are taken to be errors, the problems that result will be contained in the encapsulated unit instead of effecting the unit whole. What's more, when we decentralize the process of selecting solutions to problems, we will be able to implement a higher variety of solutions and compare the results directly. Hence, encapsulation and decentralization allow problems to be reduced in complexity, to contain problems when we make errors, increase the variety of solutions tested, and thereby increase the evidence we have for the effects of competing solutions. Over time, we will be able to generate a great deal of evidence about what works and apply those solutions in a way that satisfies complete competence.

If we can secure complete competence in this way, and therefore the right to competent government, then we can maintain the non-instrumental value of universal inclusion and the collective right to self-determination. Once problems are addressed through decentralization, there will be a greater role for the sort of knowledge that is dispersed throughout society. That is, when problems are confronted by those closest to them, knowledge held by average citizens about their communities can be put to use in theorizing about solutions, implementing them, and assessing the results. If decentralization makes effective use of universal inclusion, the value that comes from political participation will be enhanced. As citizens are given a greater role in solving problems, they will be better situated to take part in the democratic way of life.

Furthermore, the collective right to self-determination will be similarly enhanced by giving citizens a greater role in determining the policies and programs that govern their communities. Hence, in a system of decentralized problem solving, the non-instrumental value of universal inclusion, the collective right to self-determination, and the right to competent government become mutually supporting.

Chapter 1

A Trilemma of Democratic Values

One of the more recent issues in democratic theory is the tension between democratic theorist's commitment to universal inclusion and the problem of voter ignorance. The former underlies much of democratic theory and for some theorists serves as the moral justification for democracy itself. On the other hand, there are serious questions to be raised regarding the possibility that voters are collectively up to the challenge of making the sorts of decisions modern democracies face. The goal of this chapter is to present some of the more prominent views of democracy in light of the tension mentioned above. In doing so I show that prominent views of democracy are insufficient to ground democracy in its instrumental function. On the other hand, if democracy is to be altered such that it is able to fulfill its instrumental role, then there is reason to deny the first value mentioned above. That is, if universal inclusion undermines democracy's instrumental value, then there may be reason to reject it as a justification for democracy.

This is the position of epistocrats like Jason Brennan who argue that the ignorance of voters regarding basic political facts should lead us to reject democracy with universal inclusion and replace it with some form of epistocracy, i.e. a system that takes individual's knowledge into account when apportioning political power. In such a system, voters would be required to demonstrate their competence before being granted the right to vote. The goal, then, is to raise the average competence of voters in such a way as to increase the instrumental value of government policy.

There are several points where both epistemic democrats and epistocrats seem to agree. First, they agree that one of the primary functions of government is to *do* something. That is,

government is an entity that exists to establish justice and resolve problems faced by the public. Hence, government is instrumentally valuable. Second, they agree that there isn't a unified account of what makes political decisions correct. Because of this, they seem to agree that the best way to move forward is to structure the mechanism for making collective decisions such that we can maximize the benefits of having a government while respecting the dignity and rights of each citizen. Democrats hold that universal inclusion in the decision-making process is important to respecting the dignity and autonomy of each citizen. Epistocrats, on the other hand, argue that universal inclusion is in direct tension with the right to competent government. That is, epistocrats argue that the ignorance of voters about basic political facts undermines their ability to make competent decisions.

In what follows I sketch the positions of some prominent epistemic democrats and epistocrats with the goal of showing how each proposes a way to understand the epistemic aspects of political decision making. I argue that the views presented fail as epistemic accounts of political decision making because they fail to provide an account of how mistakes can be corrected. In the absence of such an account, I argue that each view faces a trilemma of fundamental democratic values.

In the first section I present some democrat arguments for universal inclusion, self-determination, and democracy's epistemic function. This will set the stage for the trilemma presented in the final section as well as offer a look at some proceduralist views in democratic theory. In the second section I present the case for epistocracy from Brennan (2016). In the third section I consider the motivations for moving away from correctness accounts of political decisions and how they are replaced by what I characterize as a sort of reliabilism about mechanisms of collective decision making. There I raise problems for any view that attempts to

show *a priori* that any individual or group are more likely to be correct. In the final section I raise the trilemma of democratic values that any epistemic account of a political decision-mechanism must face without an account of how mistakes can be identified and corrected.

I. Democracy

The purpose of this section is to give an account of two fundamental democratic values and an account of democracy's epistemic function.¹ The purpose here is twofold. First, we need an account of universal inclusion in order to understand the epistocrat's rejection of it. Second, we also need an account of the democratic value of self-determination. So in the first two parts below I sketch some of the reasons why democrats think these values are important to our understanding of democracy. In the third part of this section, I present arguments from Estlund (2007) that show the insufficiency of pure proceduralist justifications for democracy and an attempt to model this epistemic function of democracy found in Anderson (2006).

Although it may seem obvious that universal inclusion is important to any account of democracy, philosophers have given a number of answers for why this is the case. For example, Cristiano (2011) argues that the end of political society is the establishment of justice, which is guided by a principle of equality in the advancement of interests. However, because people in political societies disagree about justice and the common good, there must be a way to resolve these disagreements in a way that preserves public equality. In other words, there must be a way that these decisions can be made while respecting others' beliefs about justice and the common good (1). Democratic authority, on this view, must be grounded in public equality, the violation of which undermines that authority.

¹ I use the term 'democratic' to denote any system that utilizes elections. Since at least some proposed forms of epistocracy utilize elections, they will be included as being democratic, broadly construed.

Here universal inclusion has instrumental value because it is fundamental to the establishment of justice.² That is, without universal inclusion public equality and democratic authority would fail to obtain. The possibility of widespread disagreement about political matters leads us to find ways of making decisions for everyone while respecting others with whom we disagree. Hence, there is a sense of ‘instrumental’ at play here that isn’t grounded in any notion of correctness with regard to the decisions that democracies ultimately make. Rather, views like Cristiano’s seem to be grounded on the idea that we cannot know the right decision in any particular case, so we need a way to make decisions that accepts a plurality of views held by the public. Thus, on these kinds of views, universal inclusion has instrumental value because it is necessary for establishing a just order, but not necessarily because it tends to be an effective method for making correct decisions.

Other views of universal inclusion have it that democracy is both instrumentally valuable for its tendency to produce just decisions but also that it has non-instrumental value. Elizabeth Anderson (2009) sees herself as following John Stuart Mill and John Dewey in holding that democracy is a way of life that establishes “equality of membership, reciprocal cooperation, mutual respect and sympathy” (214). On this view, similar to Cristiano, democracy is instrumentally valuable for its effect with regard to establishing justice. Additionally, it has non-instrumental value because Anderson sees democracy as a way of life and a culture. In other words, democracy has intrinsic value in part because of its instrumental role in creating the kind of culture that values equality of membership.

² I use ‘instrumental’ to denote anything that serves as a means to something else. For example, if voting is valuable to me because it gives me a sense of purpose, then voting has instrumental value for me. On the other hand, if voting is valuable to me merely because it is my right, then it does not have instrumental value. In that case, voting would be non-instrumentally valuable.

An example that might help to illustrate this point is the value that baseball players see in batting practice. Batting practice is certainly instrumentally valuable in sharpening the players' batting skills; however, it is easy to imagine how they may begin to enjoy practice for its own sake. That is, the players may value batting practice because it is something they have always done, regardless of its instrumental value. Similarly, Anderson thinks that voting and other forms of political participation become valuable because they become part of life for citizens in democracies. Because of this, universal inclusion has non-instrumental value in the sense that it is not valued because it gets things right, or because it is crucial to the establishment of justice, it is valuable because it is part of the democratic way of life.

The above examples of the value of universal inclusion are certainly not exhaustive. For instance, we could argue that universal inclusion is valuable because it gives individual citizens a sense of autonomy in directing the course of public affairs. However, the point of this part is to present some accounts of universal inclusion that argue for the non-instrumental value of universal inclusion. The notion of the non-instrumental value of universal inclusion here, then, is 'non-instrumental' in the sense that it is not merely valuable because the decisions of citizens are likely to be correct.

With a basic idea of why universal inclusion has non-instrumental value, we can move on to the second fundamental value of democracy, which I will call the democratic right of self-determination. What I mean to denote here is the right of the collective in democratic systems to have political authority extend over all issues so long as doing so does not violate anyone's basic rights. In other words, if universal inclusion grounds democratic authority generally, the collective right to self-determination grounds the scope of that authority. For example, we could imagine a democracy with severe restrictions on the state that only allows democratic

participation to determine speed limits. In that case, the collective right to self-determination would be limited to a subset of traffic law. In the account I give below, this right includes all issues of public interest. Similar to accounts of universal inclusion, there are several accounts of the democratic right of self-determination. In what follows I give a few examples of such formulations with the goal of establishing that self-determination is a fundamental value in democratic theory.

One conception of self-determination has it “that a people of a country have the right to work out their own constitutional and political arrangements without interference from the outside” (Waldron 2010, 397). In this sense, self-determination is a right against outside influence over the political affairs of a sovereign democracy. While it is explicitly about denying outside influence, it is clear that this conception of self-determination is about establishing the collective autonomy of democracies.

Another conception is that “within certain limits and without interference from others, [citizens in democracies] have the right and liberty to decide which political arrangements and policies should govern political life and public matters” (Bülow 2016, 761). Here, democracies have a broad right to determine which problems to pursue and how to pursue them. Most theorists seem to agree that there are limits to this right, e.g. democracies can pursue a wide range of policies so long as they do not violate anyone’s basic rights.³ Although theorists disagree on the precise limits of this right, most seem to think that it applies in all cases so long as no one’s basic rights are violated.

³ Philosophers disagree about what the precise limits are on this right. See Waldron (2010) on immigration restrictions and democratic self-determination, Altman (2005) and Bülow (2016) on felon disenfranchisement and self-determination. What matters for our purposes is that the basic conception of self-determination is established to be a core democratic value.

Anderson (2006) thinks that self-determination extends to all problems of ‘public interest’ (9). Although Anderson does not give a precise definition of what makes a problem one of public interest, she does offer an example that helps to illustrate what this is supposed to mean. Anderson gives the example of the desire amongst some evangelical Christians to have Christianity be the state religion. Since the problem is only one for evangelicals, and since the broader political community has restricted religious belief to the private sphere, it does not count as an issue of public interest.

If the public is to have autonomy over the direction of public affairs, then they will also be responsible for determining which problems are of public interest and setting the agenda for the apportionment of resources dedicated to solving these problems. If self-determination is about the political community having control over its own destiny, then deciding the issues and the order of their importance is tightly connected to the right of self-determination. Here the emphasis is on the idea that democracies ought to be able to exert control over their own destiny, and this includes solving the various problems that society faces.

What is left in this section is to give an account of democracy’s epistemic function. Other views of democracy have it that democratic authority is justified by the fairness of the procedure for making decisions. In what follows I sketch two views: the first holds that pure proceduralism is insufficient to justify democratic authority and another that attempts to give an account of the epistemic dimension of democracy. The purpose is to establish that we have good reason to hold that democracy does have an essential epistemic function and to give a preliminary account of what a model of democracy’s epistemic function might look like.

I begin this part by reconstructing the argument found in Estlund (2007) that democracy has an essential epistemic function. The argument is simple but powerful. It begins by examining the reasons we might hold that the fairness of a procedure is sufficient to ground democratic authority. Similar to views like Cristiano's above, some democratic theorists rest the case for democracy on the idea that because it is a fair method for making controversial decisions and because we do not have a unified account of what makes these decisions correct, we must look to procedural fairness to ground the authority of these decisions. In other words, what makes democratic decisions authoritative is the fairness of the procedure that produced them.

However, Estlund wonders whether procedural fairness really is enough to ground the authority of democratic decisions. Merely considering the fairness of the decision mechanism does not leave democracy as the only contender. For example, we could exchange elections for coin tosses. We can imagine a system that functioned exactly like democracy but anytime we would have an election we instead toss a coin to determine a winner. This seems unacceptable. But the reasons why are not so clear from the perspective of pure proceduralism. Coin tosses are fair, possibly fairer than any real-world election. However, it would be hard to imagine a convincing argument that we should abandon elections in exchange for coin tosses because it would increase the fairness of the procedure.

Instead, Estlund contends that the reason we would not exchange elections for coin tosses is that we believe the intelligence of voters is an important part of the procedure in addition to its fairness. Hence, for Estlund this means that democratic authority must be grounded in the, however limited, epistemic power that voters bring to bear on collective decisions. That is, even if voters are highly fallible, voting provides the guarantee that some intelligence is bearing on the question under consideration, and we prefer decisions to be made intelligently rather than

randomly (Estlund 2007, 6). Given these considerations, we can see that even if procedural fairness is constitutive of democratic authority, there is still the concern that our decisions are going to get things right. And this should lead us to hold that this epistemic dimension is constitutive of democratic authority, too.

Estlund's reasoning for bringing in epistemic concerns in addition to procedural concerns seems fair. Partly because we know that voters are engaging in inquiry regarding how they should vote; but also because governments are supposed to be, at least in part, engaged in problem solving. Thus it makes sense to hold the position that knowledge or true beliefs are required for decisions to be reliably correct. However, if the point of Estlund's argument is to establish that democracy's epistemic dimension is constitutive of its authority, then it seems any model of democracy should account for this dimension in order to fully account for its authority. Estlund introduces a test that could be used to justify the claim that democracy is better than chance at getting the answer right by looking to see how democracies fair with regard to what he calls the primary bads: war, famine, economic collapse, political collapse, epidemic, and genocide (Estlund 2007, 63). If democracy tends to avoid the primary bads we can assume that democracy is better than chance at getting the answer right on other issues.

Since Estlund is doing ideal theory he does not consider the sorts of real-world problems that voters face when making complex political decisions. All he is attempting to show is that democracy is epistemically better than chance and that all other forms of government fail the test of public reasons. So, for example, epistocracy is rejected because it would require invidious comparisons between citizens that are open to reasonable rejection. Because of this, the epistemic aspect of Estlund's theory is minimal. This is in part because it is unclear how we would determine the probabilities that a policy chosen by chance would be correct or incorrect

(Anderson 2008, 134). Thus the ‘better than chance’ standard does not seem capable of establishing the epistemic chops of democracy. This is because the thesis that intelligence is preferable because it avoids the primary bads tells us nothing about how democracy fares in more mundane cases of problem solving. If we want to establish that democracy is epistemically viable, then we need more than avoidance of the primary bads as evidence that democracy is better than chance at getting things right.

The problem that a view like Estlund’s motivates is to give an account that can justify democratic authority based on how real-world democracies fare epistemically. This would require one to look to how real democratic institutions function and evaluate the plausibility that they are epistemically reliable. Otherwise there is no way to judge in any particular case whether real democratic institutions are pulling their epistemic weight. An epistemic defense of democracy should give us a model for understanding the epistemic advantages and limitations of democratic institutions.

Several theorists have attempted to give such an account. For example, Anderson (2006) argues that John Dewey’s model of democracy provides the best model of this epistemic dimension. Dewey’s model attempts to account for three constitutive epistemic elements of democracy: (1) the epistemic value of diverse perspectives, (2) the relationship between voting and discussion, and (3) epistemic feedback mechanisms (9). The first is important here because it sees diverse experience as being an important source for novel solutions to problems faced by the public. In addition to the diversity of experience, citizens hold situated knowledge about the world that cannot be effectively collected by policy makers. Anderson remarks that this is similar to arguments made by Friedrich Hayek regarding the importance of price signals for decision making in markets. That is, central planners cannot collect all the relevant information

necessary to make decisions about production and resource allocation. The only way to effectively do this is to rely on market prices as a proxy for that information. Similarly, democracies rely on talk and voting to resolve issues of public concern (8-9).

Since problems of public interest cannot be solved using price signals, citizens in democracies must use talk and voting to communicate information to each other and to policy makers. This requires other institutions like a free press and free speech rights. When these institutions fail or function improperly, citizens would be unable to make the sorts of decisions they need to make. Hence, the ability to identify and evaluate these institutions is important to understand the epistemic function of democracy.

Finally, Dewey's model, as sketched by Anderson, views democratic decision making as analogous to scientific experimentation. That is, citizens confront problems by theorizing about potential solutions, choosing from the available options, then testing them out. Since this process is one of experimentation, a sort of fallibilism about democratic decisions must be institutionalized in order to ensure mistakes are corrected. In order to correct mistakes, then, there must be a process for citizens to communicate back to policy makers whether a policy is failing to solve the problem it was intended to solve or that the policy solved the problem but at the cost of creating worse problems. Thus, one of the virtues of this view is that it attempts to give an account of how the public can attempt to track the effects of policy and make corrections where they are needed. This, it seems to me, is one of the most important aspects of any epistemic defense of democracy. That is, the judgement that a system of decision making is epistemically reliable should be, at least in part, predicated on its ability to identify and correct mistakes.

Although I think there are limitations to Deweyian model put forth by Anderson, it nevertheless serves as an important example of what a satisfying account of democracy's epistemic function should look like. One of these limitations is that it lacks a specific account of how citizens can track the effects of their decisions. Without such an account, we do not have reason to suspect that error correction is likely or even possible. Thus a satisfying model of the epistemic dimension of democracy should include such an account.

The purpose of this section was to argue that universal inclusion and self-determination are fundamental democratic values. Furthermore, I argued that democracy is at least partially justified by its epistemic dimension and that Anderson's attempt at an epistemic defense of democracy provides a useful example of what an account of the epistemic dimension of democracy might look like.

What I think should be taken away from these considerations is that democracy's epistemic function is foundational to holding universal inclusion and self-determination as positive values. That is, if these values obtain in the real world, part of their positive value is going to be grounded in the idea that the epistemic function of democracy is workable. If it is not, then there might be practical reasons to reject them. When considering individuals, we tend to respect their right to make personal choices autonomously, but this isn't absolute. In cases where a person is incapacitated in certain ways we tend to find it reasonable to restrict their autonomy. For example, if a person is intoxicated and attempting to drive, it would be irresponsible to respect their autonomy by allowing them to do so. We also tend to disallow individuals to enter into contracts unless certain conditions obtain, like they have reached a certain age or they are of sound mind when entering the contract. If personal autonomy isn't absolute, then we may find that self-determination and universal inclusion are contingent on

similar considerations, e.g. finding that democracies are competent to make the sorts of decisions that they do make. In the next section, I sketch the epistocratic position that attempts to undermine our confidence in democracy's epistemic reliability by looking to the empirical evidence that a large portion of the electorate are systematically in error.

II. Epistocracy

The debate between epistemic democrats and epistocrats centers on the issue of which political decision mechanism will produce the best decisions while also respecting all citizens' basic rights. Epistemic democrats like Estlund believe that democracy has an epistemic function and that this function is sufficient to justify democratic authority. Similarly, epistocrats seem to think that some form of democratic decision making is the best way to establish justice and seek the common good. However, they believe that the empirical reality of modern democracies is that voters are not epistemically equipped to make the decisions that confront them. This seems to leave them unable to understand their mistakes and leads them to repeat them, e.g. by favoring things like trade protectionism even though it is widely discredited by experts (Brennan 2016, 29). One of the points I raised in the last section was meant to draw this issue to the forefront. That is, the ability of voters (and, ultimately, policy makers) to track the effects of their decisions is crucial to establishing the authority of those decisions.

If epistocrats are correct, then there is a big problem for views like Estlund's. This is because, as Estlund notes, we seem to prefer a system that is fair but which also has epistemic value over one that is fairer but with no epistemic value.⁴ If this is true, then fairness may not be

⁴ Estlund (2007) admits that removing some issues from democratic control would likely result in better outcomes, but he thinks deciding who the experts are would require invidious comparisons between citizens that can be reasonably rejected (262).

enough to restrict the move from democracy with universal suffrage to an epistocratic system that is less procedurally fair but which has greater epistemic reliability. This is not to say that the only considerations that count in favor of democracy with universal suffrage concern fairness in the decision-making procedure. But it is also not the case that the only value that epistocracy can provide is epistemic. After all, the point of increasing the epistemic value of the decision-making process is to increase the practical value of the decisions themselves. The question will emerge, then, of how we are to balance these values against one another. In other words, does procedural fairness always trump practical utility? Epistocrats think not, and in this section I present their case for this position.

This section will unfold as follows. In the first part I present Jason Brennan's (2016) account of the right to competent government. In the second part I present Brennan's argument that voters are systematically ignorant of basic political facts and I present arguments—and some prominent criticism of them—aimed at undermining epistocratic worries about voter ignorance by looking to *a priori* mathematical theorems. In the third part I sketch Brennan's instrumental argument for epistocracy.

As we saw in the last section, epistemic democrats like Estlund see the reason for choosing democracy over a system of random decision making is a preference for the use of intelligence in the decision-making process. One reason we might prefer to use intelligence is because it increases the likelihood that our competence to make those decisions will come into play. In the everyday context, most people seem to think that their competence to make personal decisions entails the right to a certain level of autonomy. In cases where we are reliant on others to make decisions that affect us, we tend to require some proof or demonstration that those who are making these decisions are competent to do so. For example, doctors, lawyers, teachers,

pilots, and all drivers on the road are required to demonstrate a certain level of competence before they perform tasks that could potentially harm others if done incompetently. Thus in many areas we tend to require some kind of demonstration of competence to avoid unnecessary harm.

Brennan argues that this notion of competence ought to be generalized to electorates as well. Using a jury analogy, Brennan argues that there are many ways a jury can be incompetent. This would be the case when juries make decisions based on ignorance or irrational, immoral, or corrupt reasons rather than the evidence of the case. In these instances, we would, or should, conclude that the defendant is not morally required to submit to the decision of the jury. This is because the jury made the decision incompetently with disregard for the facts of the case. Similarly to juries, then, Brennan thinks that the authority of political decisions is determined in part by the process that produced the decision. That is, when a decision is produced in an unjust or incompetent way, that decision lacks moral authority and legitimacy.

Here is the generalized version of Brennan's competence principle:

Competence Principle (CP): It is presumed to be unjust, and to violate a citizen's rights, to forcibly deprive a citizen of life, liberty, or property, or to significantly harm their life prospects, as a result of decisions made by an incompetent deliberative body, or decisions made in an incompetent way or in bad faith. Political decisions are presumed legitimate and authoritative only when produced by competent political bodies in a competent way and in good faith. (Brennan 2016, 156-7)

Since the government is tasked with the establishment of justice, claims sole authority over certain decisions, it can deprive citizens of their life, liberty, and property, and because its decisions are enforced with violence or the threat thereof, we have a presumptive right to competence in government in the same way we have a presumptive right to competent juries (156). It should be noted here that the competence principle applies to individual decisions.

This means that if decision P precipitates decision Q, and P is made incompetently, that does not mean that decision Q is similarly incompetent. For instance, if a politician is elected to office in an incompetent manner, this does not entail that the decisions made by that politician are necessarily incompetent.

One of the limitations of Brennan's account of the right to competent government is that it does not tell us what competence means in the political domain. In other domains it seems clearer what competence means. One way we could account for competence is to look to obvious cases of incompetence. For instance, a voter who chooses candidates based purely on what the candidate looks like would seem to be deciding incompetently. This seems to be the same form that Brennan takes when pumping our intuitions about competence generally. For instance, he gives the example of a doctor who uses a witch to diagnose your symptoms. It is ultimately irrelevant if the diagnosis is correct, since the procedure used was not epistemically relevant to the decision made. So the doctor's diagnosis was incompetent in this case (Brennan 2016, 154). Similar considerations could come into play when trying to understand if democracy itself is generally incompetent.

There are other clues that should lead us to think that competence requires a notion of getting things right in a way that is relevant to the domain where the decision is being made. For instance, Brennan remarks that one option may be better than the other alternatives but still fail to be competent. The example he gives is of two doctors. One wishes to do right by his patients but is prone to make many mistakes and frequently prescribes the wrong medication. The other is corrupt and intends to take advantage of his patients, including killing them when he pleases. The former is certainly better than the latter but that does not entail that the former is competent (Brennan 2016, 164). This is important because this example seems to establish that competence

is something objective in any given domain, i.e. that there is some criteria other than the relative value of alternatives that makes decisions or decision makers competent. This implies a notion of correctness either with regard to decisions themselves or the procedures that produce them.

This last point will become important in chapter 3 where I examine the case for epistocracy in more detail. For now, it will suffice to assume that Brennan is correct to generalize the CP to include electorates. Moreover, that generalizing the CP provides us with *prima facie* justification for assuming that the right to competent government exists, even if we do not, at this point, know exactly how to apply it in any given case. What we are left with are plausible reasons to think that someone, or some procedure, is incompetent by recognizing that the decision-making process has no relevant connection to the decision under consideration. If we accept this method for pointing out incompetence, then we can at least attempt to point out cases of incompetence in the political domain. This is what Brennan does to argue against universal inclusion, which we turn to now.

There is a vast literature on the topic of voter ignorance. The results are interesting and a bit depressing. Brennan (2016) sketches this literature to give examples of the breadth and depth of political ignorance (23-33). While I will only give a few examples here, it should suffice to show that the problem of political ignorance appears pervasive. For instance, 40 percent of Americans are wrong about who the Allies fought during World War II,⁵ in the midst of the Cold War 62 percent believed that the Soviet Union was part of NATO, and as many as 70 percent cannot name their state's senators (Somin 2004, 4). This pattern of ignorance seems to have existed since research began in the 1930s. As Somin puts it, “[a] relatively stable level of

⁵ *Newsweek* Staff, “Take the Quiz: What We Don’t Know,” <http://www.newsweek.com/take-quiz-what-we-dont-know-66047> (accessed February 24, 2018).

extreme ignorance has persisted even in the face of massive increases in educational attainment and an unprecedented expansion in the quantity and quality of information available to the general public at little cost” (4).

In the year 2000, the American National Election Survey conducted a study that asked participants to answer 18 multiple choice questions to test their knowledge of the most recent elections. Scott Althaus, a political scientist and author of *Collective Preferences in Democratic Politics*, summarized the results as follows:

Just how high [the variance is] is made clear when we add up the number of correct answers to these questions and divide respondents into knowledge quartiles. While people in the highest knowledge quartile averaged 15.6 correct answers out of 18 possible, people in the lowest averaged only 2.5 correct answers. (Althaus 2003, 11-12: Cited in Brennan 2016, 32)

The study showed that only 25 percent of eligible voters are well-informed. The bottom 75 percent ranged from badly informed to systematically misinformed at the bottom quartile (Brennan 2016, 32). Less than half of participants knew that Al Gore was more supportive of abortion rights than George W. Bush. Brennan (2016) notes that if you were unaware of this fact and you had the choice of flipping a coin or asking a randomly chosen voter to decide which was more supportive of abortion rights, the coin toss would be more reliable (32).

If the studies on voter ignorance are reflective of the underlying reality, then there is good reason to suspect that the information used in the political decision-making process is similar to the examples of incompetence from above. For instance, if doctors were surveyed and only 25 percent of them were able to answer basic questions of medical fact, like what the heart does or what the function of the immune system is, it would likely reduce our confidence that when we visit the doctor we will get competent treatment.

Another issue that I would like to mention that is discussed by Brennan is political tribalism. Political tribalism is the idea that some voters tend to dogmatically hold views based on what other members of their political party or ideology believe. In a study by Drew Westen (Westen et al 2006, Westen 2008) showed that Republicans and Democrats demonstrate bias against members of the opposite party. In the study, participants looked at quotes from celebrities who were identifiably Republicans or Democrats, and then they were shown information that appeared to show that they were being hypocritical. Finally, they were shown further quotes meant to mitigate that apparent hypocrisy. Both Republicans and Democrats strongly agreed that the members of the opposite party contradicted themselves, but only weakly agreed that the members of their own party did. Furthermore, functional magnetic resonance imaging showed that participants' pleasure centers were activated when criticizing members of the opposite party. Their pleasure centers were activated once again when participants rejected evidence that members of their own party were being hypocritical (Cited in Brennan 2016, 38).

These results appear to show that voters are both generally ignorant of basic political facts and biased in ways not conducive to competent decision making. This leaves the epistemic democrat in a tough position. That is, since democrats like Estlund are engaged in ideal theory, they don't address the empirical issue of voter ignorance directly. All they need to do is show that even in light of possible ignorance that voters are better than chance at getting the answer right. While this may provide reason to accept democracy in theory, these considerations are much less persuasive in justifying real democracies. In fact, some of the studies mentioned above seem to show that voters are, collectively, worse than chance at getting answers right. This is at least the case with answering easily verifiable questions about politics like who the American Revolution was fought against or what countries are members of NATO. This does

not bode well for epistemic democrats who want to resist epistocratic arguments against universal inclusion.

One way theorists have attempted to undermine worries about voter ignorance is to look to three mathematical theorems that aim to show that crowds have the ability to make sound decisions even when they are largely ignorant.⁶ These are (1) the miracle of aggregation thesis, (2) the Condorcet Jury theorem, and (3) the Hong-Page theorem. Brennan considers these arguments and attempts to undermine them by showing that the conditions that need to obtain to get the desired results do not obtain in real-world democracy. I chose to summarize these views (and prominent criticism of them) here because they might amount to an answer to the epistocrat's worries about voter ignorance more directly than other epistemic defenses of democracy.

The miracle of aggregation thesis attempts to show that it is possible for a group, comprised of a majority of decision makers who are likely to make a mistake, to produce correct decisions. The way this works is that if errors in the decision-making process are randomly distributed, then these errors will cancel each other out. Once the errors of the incompetent members have cancelled each other out, the remaining decision makers will then cast the deciding votes. Thus, the miracle of aggregation thesis demonstrates that so long as errors are randomly distributed in democratic decision making then those who are more likely to get the answer correct will ultimately cast the deciding votes (Converse 2003, 381-2).

The miracle of aggregation thesis relies on the condition that all errors made in voting are randomly distributed. In order for the theorem to work, then, it would have to be shown that this

⁶ From here on these theorems will be referred to as the 'wisdom of the multitude' arguments.

is actually the case in practice. Although the question is certainly not settled, there is empirical evidence that voters do not make random errors when voting. What is actually found is that voters tend to make systematic errors (Althaus 2003, 129). This result should not be too surprising. Given the way that citizens in modern democratic systems organize into political parties or factions, and how partisan many political issues tend to become, it should be a matter of course that voters (when they do) make errors, these errors fall along partisan or ideological lines.

Even more problematic is that if a democratic system essentially reduces many—if not most—political decisions to binary, party-line votes there is a greater chance that neither side of the partisan line is correct.⁷ In fact, if this is the case—i.e. that most political decisions in modern democratic regimes reduce to merely two or three options—then it would seem that the benefits gained from having diverse perspectives within the voting base will be largely thwarted. Thus, systematic errors may well be built into the design of modern democratic systems.

The second argument in support of the ‘wisdom of the multitude’ thesis is the Condorcet jury theorem. The Condorcet jury theorem states that if the average voter in a democratic decision-making process has a probability of at least better than chance at getting the answer correct, then as the number of voters increases the probability that they will get the answer correct approaches 1 (Condorcet 1976, 48-9). This theorem shows that decision-makers need not be experts, or even highly competent. Rather, it shows that democratic decision-making only needs the members to be slightly better than chance at getting the answer right. This makes the Condorcet jury theorem highly favorable as a solution to the problem of voter ignorance. This is

⁷ This effect is similar to that which Hong and Page identify wherein groups of problem solvers who have similar models of the world converge on a limited number of proposed solutions to a given problem. The effect of this limitation is that it leaves the problem-solving group blind to alternative solutions.

because even if the majority of the population is ignorant of nearly all relevant facts, they merely have to be many and better than a coin toss to greatly increase the chances of ultimately selecting the correct answer.

In order for Condorcet's jury theorem to show how collective decision-making can be accurate, certain conditions must be met. One such condition is that voters must make their decisions in a sufficiently independent manner. That is, their votes must not be mere reflections of each other's. Because of this, some philosophers deny that Condorcet's jury theorem tells us anything about real-world democracy (Estlund 2007, 136-158). Given the issues related above about partisan voting habits, it seems all the less likely that this condition obtains in the real world.

Furthermore, in order for Condorcet's jury theorem to hold, it would have to be the case that in each voting cycle voters are better than chance at getting the answer right—i.e. the overall average between election cycles must be greater than chance as well. If this condition does not hold, and voters are in fact less likely than chance to get the answer right in either a single election cycle or on average generally, then as the number of voters increases the probability that correct decisions will be made decreases over time. Thus, in order for Condorcet's jury theorem to be a defense of democracy it would have to be shown that voters are in fact better than chance at getting the right answer. Unfortunately for proponents of the Condorcet jury theorem, there does not seem to be any way to confirm that the conditions required for the theorem to hold ever actually obtain, at least not without an account of what makes these decisions correct. In fact, some philosophers think that the Condorcet jury theorem may actually be a tool to criticize democracy. This is because if the average voter is less likely than chance to get the answer correct, then the Condorcet jury theorem would amount to a critique of democracy, not a defense

(Brennan 2016, 180). Given the evidence of voter ignorance presented above, this seems like it may, in fact, be the case.

The final argument that I address here is the Hong-Page theorem. The Hong-Page theorem states that, under certain conditions, cognitive diversity contributes more to getting correct answers in collective decision making than does increasing any individuals' reliability or expertise. The way this theorem works is that when a group of problem solvers is large enough, the best-performing individual problem solvers become similar to each other in their proposed solutions. This result restricts the set of possible solutions and therefore limits the problem solving capacity of the group. The Hong-Page theorem, then, demonstrates that cognitive diversity in a large pool of problem solvers provides greater diversity in proposed solutions than large groups with more homogenous perspectives. This is due to the diverse heuristics for identifying solutions that is correlated with identity diversity (Hong and Page 2004).

The Hong-Page theorem, too, requires that some conditions obtain in order for the theorem to apply. These conditions are as follows:

- The participants must have genuinely diverse models of the world.
- The participants must have sufficiently complex models of the world.
- They must agree on what the problem is and what would count as a solution.
- The participants must all be trying to solve the problem together.
- And they must be willing to learn from others and take advantage of other participants' knowledge. (Hong and Page 2004, 163-86: Cited in Brennan 2016, 181)

The most obvious problem for the Hong-Page theorem is similar to that of the Condorcet jury theorem. That is, one must show that the above conditions actually obtain in the real world. If they do not obtain, then the conclusion of the theorem fails to apply. One thing to note is the

third condition seems to be ruled out *prima facie* in the political decision-making process. This is because the purpose of the democratic decision-making process is to resolve precisely the question that must be assumed for this condition to obtain. If the democratic process is partially responsible for determining what the problem is and what counts as a solution, then it does not seem plausible that the theorem actually accounts for what democracies do.

A further problem for this theorem is the criticism it faces as a mathematical theorem. That is, there are questions as to whether Hong and Page actually demonstrate the truth of their conclusion. According to mathematician Abigail Thompson (2014), the Hong-Page theorem suffers from seven fatal flaws, each of which purport to show that the theorem is either mistaken or unproven. What's more, Thompson states, "Once the unnecessary technicalities are removed and basic errors corrected, the theorem is revealed to be little more than a straightforward restatement of its hypothesis" (1024). On Thompson's view, the Hong-Page theorem does not provide support "for the social applications proposed by the authors" (1024).

In addition to these issues, Brennan thinks that the Hong-Page theorem does not rule out epistocracy in the way the some theorists suggest. That is, the Hong-Page theorem only tells us that larger, more diverse groups are better than small, less diverse groups at solving problems. Because of this, the theorem does not rule out the possibility that epistocracy is preferable to democracy. As Brennan puts it, "Many heads are sometimes better than fewer heads, yet that doesn't mean that many heads are always better than fewer heads" (Brennan 2016, 184).

If the theorems summarized above do not account for what real democracies are doing, or do not mitigate the problem of voter ignorance, then the epistocrats are in a much stronger position. That is, their arguments in favor of limiting suffrage may turn out to be more plausible

in light of these sorts of arguments. This is because increasing average voter competence would have the effect of increasing the epistemic reliability of the decision-making process. For example, the Condorcet Jury theorem says that increasing the average competence of voters will increase the likelihood that their decisions are correct. If epistocrats are right that voters in real democracies are on average incompetent decision makers, then restricting suffrage could raise the average level of competence.

The question remains whether restricting suffrage would increase average competence *enough* to make the system generally competent. As was noted above, demonstrating that epistocracy is better than democracy would not entail that epistocracy is competent. Brennan attempts to show that it will be, or likely will be, competent by showing that citizens with greater knowledge of basic political facts have different preferences than the less informed. These are the so-called enlightened preferences. However, Brennan doesn't clearly establish that this alone will lead to competence *per se*.

Because those with greater political knowledge systematically support different policies than the less informed, Brennan takes this to show that there would be systematic differences in the sorts of policies chosen under epistocracy. For instance, Brennan notes that high information citizens tend to favor things like free trade, they are more accepting of gay rights and abortion, and they favor less military intervention abroad and less government intervention in the economy (Brennan 2016, 34).

Since those with enlightened preferences tend to be more in agreement with experts, there is reason to suspect that they are more likely to be right. This follows from the proposition that experts are more likely to be right in cases of expert-novice disagreement. Brennan argues for

this proposition by noting that there are several issues that economists agree on—like free trade and avoiding price controls—regardless of their ideological background, i.e. regardless of their political ideology. Brennan takes this to mean that there are cases where we can have a high degree of confidence in the expert consensus in part because that consensus cuts across the ideological spectrum. Thus, we are to take from this that expert consensus on these issues is likely based on sound empirical work. The novices who disagree on these issues, then, are *wrong* to do so and this failure is rooted in their ignorance of economics. This isn't an unreasonable position to take, however in chapter 3 I take this issue up in greater detail.

For now, let us assume that novice agreement with experts is a sign that they are more likely to be correct than those who disagree with experts, at least on issues where there is consensus amongst experts. If those with enlightened preferences are more likely to be correct than the electorate generally, then this is good evidence that epistocracy will outperform democracy. Furthermore, if the electorate's beliefs are generally found to be in conflict with experts, especially on issues where the experts are in broad agreement, then it would appear that the electorate's beliefs are likely unjustified. If the electorate tends to make decisions using unjustified beliefs, then those decisions will be made incompetently.

The instrumental argument for epistocracy, then, is that given the empirical evidence that democracies are largely misinformed about basic facts relevant to making political decisions and they are frequently wrong on issues where there is expert consensus, then democracies are largely incompetent to make political decisions. Furthermore, since the more enlightened citizens are more likely to be right on account of their tendency toward agreement with experts, allowing them greater power over political decisions will lead to better policy, more prosperity,

and more just outcomes. In other words, democracy is likely in systematic error and this can likely be overcome by eliminating universal suffrage.

One worry here is that even if the instrumental argument for epistocracy is sound, it does not tell us that epistocracy will be competent. All Brennan's argument implies is that epistocracy will be better than democracy, but being better than democracy does not entail competence. The CP tells us that the authority and legitimacy of collective decisions is constituted, at least in part, by the competence of the decision makers and the decision-making process. Brennan seems to have failed to give us a reason to think that an epistocratic electorate will be competent. This should lead us to question whether epistocracy will ultimately have the same problem of incompetence as Brennan identifies in democracy. In the next section I explore this issue as it applies to both epistocracy and epistemic defenses of democracy.

III. Reliabilism and *a priori* Bootstrapping

Both the epistemic defenses of democracy and Brennan's arguments for epistocracy seem to rely on a sort of reliabilism about collective decision-making processes that is analogous to reliabilist theories in epistemology. I begin this section by sketching a problem that reliabilist theories face. I will then attempt to show how this problem confronts epistemic defenses of democracy discussed above and Brennan's argument for epistocracy. In doing so, I hope to show that these attempts to establish the reliability of either democracy or epistocracy will fail to be able to account for how mistakes can be corrected. Further, that if error correction is understood as a constitutive element of the reliability of any epistemic process, then without an account of error correction we will fail to establish the reliability of that process.

There is a well-known problem in epistemology that confronts process reliabilist accounts of justification. Process reliabilism is a set of views in epistemology that attempt to avoid the problem of the criterion by arguing that justification for a belief is in some way conditioned on the process that produced the belief. For instance, if a true belief was produced by a process of wishful thinking, then it would fail to be justified because wishful thinking is not a reliable belief-forming process. One problem with these sorts of views is that they have been accused of requiring what Vogel (2000) calls bootstrapping.

Bootstrapping allows one to gain justification that a belief-forming process is reliable via the use of that process. For example, I can gain justification that my speedometer is accurate by looking at it while driving, seeing that it reads '45 mph', then forming the belief that I am travelling at a speed of 45 mph. Furthermore, I can then infer that since I believe that I am travelling at 45 mph at this moment, the gauge is reading accurately. Repeating this procedure over time will give me further evidence that my speedometer is accurate. The problem with this sort of reasoning is that my speedometer could have been inaccurate and I would still come to the conclusion that it is a reliable indicator of my speed.

Theories that engage in bootstrapping endorse what Cohen (2002) calls 'basic knowledge'. Theories that allow basic knowledge allow that we can gain knowledge from a belief-forming process before we know that that process is reliable. This leads to what Cohen calls "the problem of easy knowledge." There are two ways the problem of easy knowledge can arise. First it can arise by the closure problem. The closure principle states that if S knows that P entails Q, and P, then S knows that Q. To use Cohen's example, imagine there is a table before you and it appears red. From the table's appearing red you can come to know that it is red, provided your color vision is reliable. Further, by the closure principle you can come to know

that the table is not white with red lights shining on it. But, it does not seem plausible that you can come to know that the table is not white with red lights shining on it merely by the table's looking red (312).

The other way the problem of easy knowledge can arise is through bootstrapping. As noted above, the problem is that once we form a belief that P, we can reflect on that fact and infer that the belief-forming process that produced the belief that P is correct in this case. This is shown in the example of bootstrapping the reliability of my speedometer above. That is, each time I glance at the speedometer while driving and it reads 'X mph' and I form the belief that I am travelling at a speed of X mph, I can infer from this belief that the speedometer is reading accurately on this occasion. Over time I can accrue a great deal of evidence that my speedometer is reliable even if it is not.

Cohen considers whether a strategy that holds that we can know a belief-forming process is reliable *a priori* can avoid the problems associated with basic knowledge. It seems that if I could come to know *a priori* that my perceptual faculties are reliable, then I could gain knowledge via perception without thereby having basic perceptual knowledge. This would get us out of the problem of easy perceptual knowledge. Of course, this would entail that we have basic *a priori* knowledge, so it would still require a sort of bootstrapping on the *a priori* knowledge that a particular belief-forming process is reliable. That is, we would have to know *a priori* that *a priori* belief formation is reliable (Cohen 2002, 320).

Leaving the issue of basic *a priori* knowledge aside, there is still the issue that the reliability of any belief-forming process is contingent, and therefore any *a priori* knowledge of the reliability of a belief-forming process will be contingent *a priori* knowledge. Given this,

there does not seem to be any way, starting with the contingent *a priori* premise, to infer that any process is, in fact, reliable. For that we would need some independent check, otherwise we would never gain justification for the conclusion that a process is reliable beyond what is provided by the justification for the contingent *a priori* premise. This is the original bootstrapping problem discussed above. Suppose I come to know, *a priori*, that if my speedometer reads ‘X mph’ then I am traveling at X mph. Each time I check my speedometer I will gain knowledge that it was accurate this time. However, that knowledge will follow from the *a priori* knowledge that the speedometer is accurate. So each instance where I come to know my speed by checking the speedometer will fail to amount to further evidence of the speedometer’s reliability, since that knowledge is already secured *a priori*. The same will also be true for justification. Hence, the use of the process will not provide further evidence of its reliability.

Assuming we can derive knowledge of or justification for the reliability of an epistemic process *a priori*, there still remains the closure problem. For instance, suppose I have *a priori* justification for the belief that if X looks red then X is red. From this I can infer that X is not white with red lights shining on it. This seems implausible because I will be justified in believing that X is not white with red lights shining on it *a priori* (Cohen 2002, 321). So the closure problem does not seem to be avoided by the *a priori* strategy.

As I will show below, these problems arise for both the epistemic defenses of democracy discussed above and Brennan’s case for epistocracy. In what follows I lay out a version of the problem of the criterion that faces correctness accounts of political decisions. In the second part, I argue that the wisdom of the multitude arguments are self-defeating and ultimately fail to be able to account for error correction. In the third part, I argue that Brennan’s argument for

epistocracy suffers from a similar problem, which makes it impotent to establish that epistocracy will correct its mistakes.

Any attempt to give an account of what makes a political decision correct faces a version of the problem of the criterion. Chisholm's (1973, 12) version of this problem concerns knowledge and is formulated as follows:

(A) *What do we know? What is the extent of our knowledge?*

(B) *How are we to decide whether we know? What are the criteria if knowledge?*

According to Chisholm, there are three responses to this set of questions. We can assume an answer to (A) first, and then attempt to formulate an answer to (B). We could also assume an answer to (B), and then attempt an answer to (A). Chisholm calls the first response *particularism* and the second *methodism*. The third response is to deny that either of the other responses can answer one question in a way that does not depend on an answer to the other. This response is called *skepticism*. If we assume this is correct, then there is a problem that arises for any attempt to answer this set of questions. That is, each of the three responses requires us to beg the question against the other two possible responses.

The primary issue that will confront us in the political version of this problem is dialectical. In democratic political systems, we need a means for resolving controversial questions, and which political decisions are correct is a central concern. In order to justify our response to the political version of this problem, we will need to have grounds to offer our interlocutors. However, because each possible response to the set of questions will require an

assumption, we do not have the means to do this. We can formulate the political version of this problem as follows:

- (1) What past political decisions were correct?
- (2) What makes a political decision correct?

If we are particularists here, then an answer to (1) will not only be controversial, it will require us to assume that the methodist response is wrong at the outset. The same result applies to the methodist and the skeptic. In order to avoid this dialectical problem, the views sketched in previous sections of this chapter try to argue that one system for making political decisions will be epistemically reliable or more reliable than another. This allows them to safely avoid a controversial account of correctness while still being able to make the claim that one system is epistemically better than another.

This seems to lead us to a sort of reliabilism about political decision making. For example, Estlund supposes that the use of intelligence in decision making will make the system better than chance at getting the answer right. This, then, grounds our choice of democracy over a system of random decision making. Brennan's arguments against democracy rely on similar reasoning. Citizens in real-world democracies tend to be misinformed. This undermines the system's reliability. However, some citizens are not misinformed. Therefore, the system will be more reliable if we only allow informed citizens to vote.

If Brennan is correct that we have a right to competent government, then we should look to any proposed system's reliability to ground the competence of the system. So if we want to increase the competence of the system we should look to increase its reliability. This requires us to look at some of the features of the system that will give us clues as to whether we should

expect it will be reliable. This is how Brennan attempts to undermine our confidence that democracy is reliable and competent. That is, he shows that many citizens fail to know basic facts that are relevant to making political decisions. Another way we could estimate a proposed system's reliability is to look at the likelihood that that system will correct its mistakes. Since we are not assuming any system will be infallible, we should expect mistakes to be made. Because of this, a system's reliability should be predicated, at least in part, on its capacity to track the effects of its decisions and correct errors.

Recall that the wisdom of the multitude arguments attempt to show that so long as certain conditions obtain, the probability that a correct answer will be chosen increases either with the diversity of participants (Hong-Page theorem), increasing the average participants reliability above chance (Condorcet Jury theorem), or so long as errors are randomly distributed (miracle of aggregation).⁸ I sketched some problems for each of these views as accounts of democracy in the last section. But, if we assume that those objections fail, there is still the problem that none of these models attempts to account for how mistakes will be corrected. Instead, they assume that they will lead the electorate to the correct answer, even in cases where they do not. This is because these theories attempt to justify the belief that democracy is (or could be) a reliable process *a priori*, then extend that justification to establish the reliability of real-world democracy.

As noted above, there are problems with this sort of reasoning. Take the bootstrapping problem first. Assuming the relevant conditions obtain for one or more of the wisdom of the multitude arguments to apply, we will gain justification for the belief that democracy got it right

⁸ I think the problem raised here also applies to views like Estlund's, where *a priori* intuitions about the use of intelligence give us reason to think democracy will be better than chance.

this time. In order for the wisdom of the multitude arguments to provide justification for the reliability of democracy, we would have to be able to use this track record to amass evidence that it is reliable. However, by repeating this process over time we will not be amassing evidence that democracy is reliable. As noted above, each time a decision is made we will gain no new information about the reliability of democracy. This is because our justification for the belief that any particular decision is correct will be derived from the original *a priori* argument. Thus we cannot gain justification for democracy's actual reliability without bootstrapping. The only way to avoid this kind of circularity would be to have some independent check. But in order to do that we would face the problem of the criterion. Therefore, the wisdom of the multitude arguments cannot serve as the basis for establishing the reliability of any real-world democracy.

Now take the closure problem. Because the wisdom of the multitude arguments provide justification for the belief that a particular decision is correct, we are, via closure, justified in believing that the dissenting minority was incorrect. But the justification for the belief that the dissenters were incorrect will not be based on anything relevant to the content of the decision. That justification will be provided solely from the *a priori* wisdom of the multitude arguments and the closure principle. Hence, the closure problem here leads to what Cohen (2010) calls "basic justification", which is a corollary to basic knowledge. Again, this is an implausible method for determining whether the dissenting minority was incorrect. Given this situation, it appears that the wisdom of the multitude arguments cannot serve as a basis for establishing the reliability of democracy.

Epistemic democrats may counter that the justification supplied by the wisdom of the multitude arguments and the closure principle is defeasible. Because of this, empirical considerations can come into play as defeaters. This would make room for rational dissent.

However, there does not seem to be a way to determine when a defeater is justified or rational. Indeed, it seems that any theory of justified or rational defeaters would face another version of the problem of the criterion. Maybe this is asking too much of epistemic democrats. It is possible that there are legitimate grounds for dissent merely based on the idea that the theorem is probabilistic. Hence, it is rational to expect that some of the time the majority will be wrong. We will just be blind to when supposed defeaters are legitimate. Without some way to discriminate between justified or legitimate defeaters and unjustified or illegitimate defeaters, these arguments appear even weaker as accounts of democratic reliability, since without such an account we don't have any reason to suppose that any particular democracy is making more correct decisions than incorrect ones or whether there are defeaters for any particular decision.

If we cannot give an account of what makes a defeater legitimate because we face the problem of the criterion, we could try to resolve this problem in the same way that epistemic democrats attempt to resolve the other version of the problem of the criterion that faces correctness accounts. That is, by applying the wisdom of the multitude arguments to the dissenting minority. For instance, it is perfectly compatible with the Condorcet Jury theorem that the dissenting citizens are, on average, better than chance at getting the answer right. This is because being more likely to be right does not entail that one will be right. So if the relevant conditions obtain for the group of dissenters, then we can say that they are more likely to be right. This would provide justification for believing the defeaters are legitimate in the same way it provides justification for the belief that the majority was right.

Indeed, it seems that all the wisdom of the multitude arguments could be applied to the dissenting minority. For instance, so long as they were sufficiently diverse the Hong-Page theorem would apply. And if the minority's errors were randomly distributed we would get the

miracle of aggregation. Although, it may be the case that the wisdom of the multitude arguments do not apply to subsets of collectives, I am unaware of any reason why this would be the case. It seems that so long as the relevant conditions obtain, then the theorems should apply. Hence, any time there is dissent, and the relevant conditions obtain for the dissenting group, then that dissent represents a defeater for the belief that the majority was correct.

There will be cases where the relevant conditions obtain for the majority but not the dissenting minority. In those cases we would have no defeaters present. But there will also be cases where the conditions obtain only for the minority, too. So the only cases where the wisdom of the multitude theses give us reason to think the majority is correct without a defeater is when the relevant conditions do not obtain for the minority. This seems to be the case only for the Condorcet Jury theorem and the Hong-Page theorem. Because the condition for the miracle of aggregation to occur is that errors are distributed randomly throughout the collective, this would necessarily include the minority group, so the miracle of aggregation would always produce a defeater for the majority decision.

The question remains as to whether this can provide us with a reason to think that democracy can identify and correct mistakes. At this point it seems that it cannot. Primarily because the views attempt to show that the majority will get the right answer, but as I showed above, this more often than not is self-defeating. Anytime democracy identified an error and attempted a correction, in two of three possible outcomes there will be a dissenting minority that amounts to a defeater. Thus, we would have a defeater for the belief that there was any error at all. If the wisdom of the multitude theses are self-defeating some (or most) of the time because they justify defeaters for the majority decision, they cannot provide us with the tools necessary for establishing that democracy is likely to identify and correct its errors.

This problem seems to apply to Brennan's argument for epistocracy. In particular, the thesis that increasing average voter competence—by restricting suffrage—will increase the reliability of the electorate's decisions. This thesis is similar to the Condorcet Jury theorem in that it supposes with greater individual reliability in making decisions comes greater reliability of the collective. Thus if Brennan is correct, then by restricting suffrage to those with enlightened preferences we will get more reliable decisions. However, the issue we are trying to settle is whether Brennan can account for how an epistocratic electorate can identify and correct its mistakes. I argue that he cannot.

Recall that Brennan's argument that democracy is most likely in systematic error is based on the correlation between citizen's scores on political knowledge surveys and their policy preferences. Those who do well on these surveys tend to have systematically different policy preferences than those who score poorly. Furthermore, the preferences of those who score higher tend to be more in agreement with certain instances of expert consensus. For example, economists tend to agree that restricting labor mobility, through mostly closed borders, produces a high cost on global productivity (Brennan 2016, 192). However, voters tend to reject open-borders policy. The suggestion is that if voters knew better they would be more accepting of reduced restrictions on immigration. Hence, we would be better off if only the voters were more informed.

Brennan (2016) splits citizens into three rough categories with regard to politics: Hobbits, Hooligans, and Vulcans (4-5). Hobbits don't really have strong opinions about politics, they don't tend to have much social scientific knowledge, and they prefer to go on about their daily business unperturbed by politics. Hooligans are the "rabid sports fans of politics" who have strong views about politics. They can recite arguments for their favored position but tend to be

incapable of explaining other points of view. They tend to write off evidence that contradicts their beliefs, they are overconfident, and they despise people with whom they disagree. Further, they hold their political views to be part of their identity. Brennan says “[m]ost regular voters, active political participants, activists, registered party members, and politicians are hooligans” (4-5). Finally, vulcans are those whose beliefs are grounded in social science and philosophy. Their views are rational and they are capable of accurately explaining the points of view of others with whom they disagree. They care about politics but are “dispassionate” because they actively try to “avoid being biased and irrational” (4-5).

We might expect that epistocracy would deliver us an electorate of vulcans. But Brennan admits that epistocracy “will still feature the rule of the hooligans” because the vast majority of political participants and voters are hooligans (Brennan 2016, 207). Even so, Brennan still thinks that we will get more reliable government because even these hooligans will be amongst the enlightened, and we may still be better off under epistocracy than democracy. Fair enough. But it is unclear that this sort of reasoning can lead us to believe that epistocracy will most likely correct its errors or be generally reliable. Here’s why. Brennan’s thesis, much like the wisdom of the multitude theses, is *a priori*. That is, Brennan’s thesis that voters with enlightened preferences are more likely to be correct on matters of candidate and policy choice is not an empirical one. This is because in order to gain empirical evidence that any one group is more likely to be correct would require us to look to see who was correct the most. In other words, we would need frequency data that records the total and proportionate number of correct decisions made by different groups over a particular period of time. But this would require us to answer the problem of the criterion, which would require question begging against those who disagree.

So, all we can do is use the *a priori* claim that voters with enlightened preferences are more likely to be correct.

The same problems the wisdom of the multitude theses faced arise here, too. For instance, if we assume the *a priori* thesis is correct and apply it to any given decision, we would gain justification for the belief that the majority decision was most likely correct. As noted above, however, we cannot derive justification for the belief that epistocracy will be reliable in this way. Just like before, Brennan could argue that defeaters will arise when epistocracy goes wrong, which will provide defeaters for our justification that the majority was right. However, like in the previous cases, it looks like the *a priori* thesis applies to the dissenting minority. This is because the entire electorate falls into the category of voters whose decisions are most likely to be right. Hence, if we assume the *a priori* thesis is true, then every time epistocracy makes a decision, the very existence of a dissenting minority will provide us with a defeater for the belief that the majority was correct. So every decision will bring on its own defeater. Because of this, there remains no method for determining when epistocracy has actually made a mistake. It seems the *a priori* thesis rules out (*a priori*) any reason to think epistocracy will be capable of identifying errors and correcting them.

These issues will be covered in more detail in chapter 3, where I consider the argument for epistocracy along with some of the major epistemic problems faced by voters and policy makers. Suffice it for here, however, to show that none of the views sketched give us reason to think that democracy or epistocracy are likely to correct their mistakes, and that this is a fundamental aspect of any model of political decision making. In the next section, I raise the trilemma of democratic values that any model of political decision making ought to be able to hold.

IV. The Trilemma

So far I have tried to show that arguments for democracy and epistocracy fail to give us reason to suspect that they will be able to identify and correct their mistakes. This is because they engage in a sort of *a priori* bootstrapping that leaves them unable to account for error correction. That is, the arguments they give in support of their reliability will also apply to a dissenting minority. Because of this, it remains unclear how any of these views can account for three fundamental democratic values:

1. The non-instrumental value of universal inclusion
2. The collective right to self-determination
3. The right to competent government

Proceduralist democrats attempt to show that we can hold 1 and 2 based on considerations like procedural fairness and public equality. As Estlund shows, however, these considerations are undermined when we consider the fact that proceduralists do not have the tools necessary to show why democracy is preferable to random decision making. Epistemic democrats attempt to argue that we can hold 1, 2, and 3 together using ‘wisdom of the multitude’ arguments designed to show that ignorant collectives can still be competent. However, epistocrats like Brennan think these arguments do not actually apply, and it seems they may be right. Given the reasons why, and the facts about voter ignorance, epistocrats think that we must reject 1 in order to preserve 3. It is unclear how epistocrats see 2. That is, if we reject 1, then 2 will only apply to the electorate. So the electorate would have the right to self-determination, meaning their decisions would extend over all issues of public interest.

If we assume Brennan is correct about 3, and his arguments for epistocracy cannot account for error correction, then we should question whether epistocracy can satisfy 3. If it

cannot, as I argued above, then epistocrats cannot justifiedly hold 2 either. This is because if epistocracy is incompetent (by the error correction standard) then epistocratic decisions will fail to be authoritative and legitimate. Hence, the electorate's right to have authority over all issues of public interest would fail to be legitimate.

V. Conclusion

If what I have argued above is correct, then in order to hold all three values above we need a model of political decision making that can account for how errors will be corrected. This will require an account of what it is that decision makers are doing epistemically. That is, a model of the epistemic situation of decision makers that takes into account all the epistemic limitations faced by decision makers. Where there are problems, the model of decision making ought to be able to account for how those problems will be mitigated. Furthermore, the model will need to be able to account for how defeaters can be found to be legitimate without relying on the same reasoning about who will get the answers right. This will likely entail a quite different approach to political philosophy than the views discussed here, although Anderson (2006) gives us a sort of basic blueprint for how to construct such an account.

In chapter 2, I offer an expanded notion of competence that explicitly accounts for error correction. Further, I argue that without this kind of competence, we face a problem I call the graveyard spiral. In chapter 3, I examine the epistemic problems faced by voters and show that epistocracy cannot be expected to be competent on the expanded notion of competence developed in chapter 2. In chapter 4, I expand further on the epistemic issues faced by voters and offer a solution that can hold all three democratic values simultaneously.

Chapter 2

Competence and Competent Government

In the first chapter I argued that epistemic arguments for epistocracy and democracy face a bootstrapping objection. That problem calls into question the possibility that these sorts of arguments could ever serve as the basis for establishing the competence of either system. If Brennan (2016) is correct that we have a right to competent government, which I find intuitively plausible, then the bootstrapping problem is a reason to reject his argument for epistocracy. That is, if these epistemic arguments cut no ice, then there is no reason to think that epistocracy will outperform democracy with regard to competence. Further, since Brennan does not provide a clear account of competence it is not clear that, even without considering the bootstrapping problem, epistocracy will be competent. The purpose of this chapter is to argue for a conception of political competence using the general framework of competence in Sosa (2010). In doing so, I provide further argument for the claim made in the previous chapter: that competence requires error correction and identification.

By establishing a fuller account of competence, we will be better able to judge when a system can be expected to be competent. This will offer us a better test of a system's competence and give us greater confidence that the right to competent government can be established. That is, we will be able to understand the conditions that must be present to maintain the three democratic values.

In the first section I partially reconstruct Brennan's argument against democracy on the grounds that it is incompetent. In the second section I reconstruct Sosa's account of competence and apply it to the cases that Brennan uses to pump our intuitions about competence in political

decision making. In the third section, I argue that if epistocracy (or democracy) is to be competent, then error identification and correction should be reasonably expected to occur. Finally, I argue that without error correction and identification we face the graveyard spiral problem.

I. Brennan (2016) on Competence

Early in the book, Brennan argues that most citizens in modern democracies are incompetent to make political decisions. It seems to follow quite naturally that we should suspect these people to collectively make incompetent decisions. As such, Brennan thinks that our rights are violated when these people are allowed to make decisions that have a real (and potentially harmful) effect on the rest of us. This is the basis for the claim that democracy violates our right to competent government. What I want to focus on here is on what makes democracy incompetent and epistocracy competent—or at least less incompetent.

Simply put, Brennan's argument is that democracy, by allowing large numbers of incompetent people to vote, fails to be competent on that account. His solution, apparently, is to remove the clearly incompetent in exchange for the indeterminately competent. That is, Brennan offers no positive account of competence, but, rather, some platitudes about competence in order to pump our intuitions and lead us to the conclusion that democracy is clearly incompetent and epistocracy will be a step in the right direction. It is not clear, however, if Brennan thinks that epistocracy will result in competent government, partly because he does not tell us what the conditions are for competence but also because there is evidence that so-called "hooligans"—i.e. politically active citizens who systematically form opinions and make decisions in epistemically arbitrary ways—will be in the majority even in epistocracy (Brennan 2016, 207). Without such

an account, we remain in the dark about whether epistocracy amounts to a solution to the problem that Brennan identifies.

Brennan relies on Jilian Craigie's list of criteria for patient competence in medical ethics as a basis for judging competence. The list is meant to denote the relevant conditions that must obtain for a patient to be competent to consent to medical treatment. Here's the list:

- i. Patients must be aware of the relevant facts
- ii. They must understand the relevant facts
- iii. Patients must appreciate the relevance of those facts for their own particular case
- iv. Patients must be able to reason about those facts in an appropriate way (Craigie 2011, 326-33: Cited in Brennan 2016, 162)

At first glance, this list seems to be quite inappropriate for determining competence in political decision making. This is because the 'relevant facts' in politics are controversial and part of the political decision making process itself. This is not to say that there aren't clear cases of facts that do not seem to have any relevance to political decision making—e.g. the fact that the Battle of Hastings happened in 1066 doesn't seem particularly relevant for any contemporary political issue. Nonetheless, it is not clear how being aware of the relevant facts, reasoning about them appropriately, etc., entails that one will be competent to make political decisions.

The examples that Brennan gives are meant to demonstrate that the above four conditions on patient competence can be used as a razor to identify cases of incompetence in other domains. For instance, he uses the example of a plumber who decides that the best treatment for clogged pipes is to mow the lawn (Brennan 2016, 163). In that case, it seems clear that the plumber has either missed some facts about how plumbing works or he has failed to reason appropriately

given the evidence. But it does not seem possible to demonstrate that a plumber is *competent* given the four criteria above. Take the plumber example again. If the plumber had all the relevant facts, and reasoned appropriately about them, say, by recognizing that the pipe was full of hair and that the solution was to clear the obstruction, this alone does not establish that the plumber is competent. If that were all there was to plumber competence, then anyone who could recognize the problem and deduce the solution (i.e. remove the obstruction) would not need a plumber to begin with. The reason we call a plumber is because the plumber has the skill, and the tools, to carry out such a treatment, which is (or should be) central to any conception of plumber competence. So it does not seem to be the case that competence can be identified using the four criteria above, even though they may provide a means to point out clear cases of incompetence in domains outside medical ethics.

Since Brennan's comments on competence cannot do anything more than show obvious cases of incompetence, it cannot serve as a basis to demonstrate that anyone is, or will be, competent, including epistocracy. Further, if the problem that Brennan is attempting to resolve is the systematic violation of our rights through incompetent government, then we do not yet have the resources to claim that he has solved the problem. That is, we need an account of competence to base any judgment that epistocracy, or any form of government, will ever reach the level of competence such that it does not violate our rights.

II. Toward a Positive Account of Political Competence

In "How Competence Matters in Epistemology", Ernest Sosa lays out a theory of competence that I believe will be useful to the present discussion. Sosa uses this account of competence to establish a distinction between skeptical scenarios and normal cases of perception. That is, he is

attempting to offer an account of what goes wrong in skeptical scenarios. What goes wrong, according to Sosa, is that certain conditions on epistemic competence are absent in skeptical cases. This will be useful for our purposes because the difference between ordinary cases of perception and skeptical scenarios may provide a basis from which to judge the reliability of a political process. That is, if competence can be established in the political case in the same way it is established in the case of perception, then an analogy between skeptical scenarios and a failure of competence in the political domain may be established. This should be a good test, then, to decide when and how a political decision mechanism can be, or will be, competent.

Sosa (2010) defines competence as a “disposition of an agent to perform well” (465). Similarly, then, competences will be structured in the same way as other dispositions. That is, there will be three components that must all be satisfied for the disposition to obtain. These components are constitution, condition, and situation. Since these three components are the same for skill-based competences, I will follow Sosa and use the example of an archer to fill in the details. The constitution condition accounts for the “inner” skill of the archer. The condition component accounts for the condition the archer must be in for any given performance to be competent, e.g. the archer must be sober and awake. Finally, the situation component accounts for the external situation of the archer, e.g. good light and sufficiently low wind speed (465). When all three—condition, constitution, and situation—are satisfied, then the agent will have a complete competence. So an example where an archer can be expected to make a competent shot is one where the archer has the appropriate skill, she is sober and alert, and her target is clearly visible and within the range of her bow.

The satisfaction of the competence triad above for archers is associated with the truth values of what Sosa calls trigger-manifestation conditionals, e.g. *if she takes this shot, she will*

likely hit the target. In other words, if our archer satisfies the above competence triad, then we should suspect that she will hit the target. Since we are interested in competences generally, it makes no difference if our archer would fail to make a competent shot in cases where the competence triad is not satisfied, e.g. she is taking a shot at a target that is far outside the range of her bow. Hence, we are interested in how certain competences are manifested under certain conditions, primarily because competence, like dispositions, require certain conditions to obtain in order for the relevant competence or disposition to manifest.

Since we are interested in establishing when a competence is likely to manifest, we want to know under what conditions we can truthfully say that a particular performance was a reflection or failure of the relevant competence. Even with the competence triad satisfied, there are still things that can go wrong such that the manifestation of a competence can be thwarted. For instance, imagine our archer takes a shot and a random gust of wind blows the arrow off target. This would not demonstrate a failure of competence on the part of our archer. This is because the failure of the arrow to hit the target is the result of the gust of wind, not a failure of competence to manifest. A good test of competence is one where the result is appropriately connected to the performance. Hence, we need to distinguish between the accuracy of the performance, the skill involved in that performance, and the connection between the skill and the accuracy.

Sosa introduces three conditions for identifying competence in a given performance. The first is adroitness, which is the skill manifested in the performance. Second, there is the accuracy of the performance. And third there is the aptness of the performance, i.e. an archer's shot is apt just in case the shot is accurate and the accuracy manifests the adroitness of the shot (Sosa 2010, 466). When our archer makes an apt shot, it is accurate because of her skill as an archer. In a

case where the shot is blown off course by a gust of wind, then blown back by another such that the arrow hits the target, then that shot will fail to be apt because the resulting accuracy was caused by the wind, not the adroitness of the shot (466). Any performance that is apt in this way will be a manifestation of first-order competence.

What is not yet clear is the connection between the competence triad and the three A's. It seems at this point that the competence triad and the three A's come apart. For instance, a poorly situated archer can still make an apt shot. A case like this would be one where the shot was long, but just within range, and our archer hit the bullseye. It seems in that case she has manifested her competence as an archer against the odds, so to speak. She made the shot when we shouldn't suspect that she would, even being aware of her skill. So there will be cases where an agent must judge the situation to be one where their competence can be manifested.

In competitive archery, for example, each archer must take a shot when it is their turn. An archer's judgment about the likely aptness of an impending shot is irrelevant to whether the archer should take the shot, since the only risk is in failing to take the shot. However, this does not generalize to all situations. Following Sosa, we can imagine a hunter who must decide when a potential shot will be apt. Failure to do so could result in missing the shot or merely wounding the animal. Thus the hunter must make a judgement whether the competence triad is satisfied as part of the process of estimating the risk of any particular shot. For example, imagine our huntress is stalking a deer in the early evening just after the sun has set. She knows that there is limited daylight left to make the kill and must decide if taking a shot now is worth the risk. In other words, she must judge whether a potential shot will be apt before making the shot. If she decides to take the shot and it isn't apt, then she has failed to judge accurately. Sosa calls this kind of competence, that of the judgment about the aptness of an impending performance,

second-order competence. Hence, to satisfy second-order competence agents must aptly judge that their impending performance will be apt. Thus, we now have a basic conception of first and second-order competence. When a performer satisfies both first and second-order competence, we will call it meta-competence.

The importance of meta-competence for epistemology is the same as it is for archery. That is, an agent may form an accurate and adroit belief—i.e. a true belief that is the result of an otherwise reliable process—but that belief may fail to satisfy meta-competence. An example of this from traditional epistemology would be Gettier cases. For example, when an agent in fake barn country sees a real barn, he forms the true belief ‘this is a barn’. However, because the vast majority of apparent barns are actually only barn façades, the truth of the belief is the result of luck rather than the belief forming process. In other words, if the agent was presented with a barn façade he would have formed the same belief about the object before him. Hence his true belief in the first case, while accurate, is not an example of complete competence because he failed to recognize that he was not situated in an appropriate way. Sosa calls this animal knowledge, and contrasts it with reflective knowledge. Animal knowledge, then, is accurate and adroit belief without meta-competence, whereas reflective knowledge is apt belief that also satisfies meta-competence (Sosa 2010, 471).

In order to have reflective knowledge, on this view, one must have an apt belief on the second-order judgment that the first-order belief will be apt. For instance, if our agent in fake barn country is aware he is in fake barn country, and is familiar enough with the façades to distinguish them from real barns, then his belief that ‘this is a barn’ will satisfy meta-competence so long as it is apt. Sosa calls this second-order aptness ‘apt grasp’ (472). So in the barn façade example, our agent’s belief would satisfy meta-competent grasp if and only if the first-order

belief is apt and the second-order judgment is apt (or presupposed) that the first-order belief is apt (472). Hence, reflective knowledge entails the first-order belief will be meta-competently grasped by the agent.

An issue arises here about when we can justifiably say that we have second-order competence. In normal cases of perceptual belief, we seem to tend toward assuming we are properly situated unless there is some reason to suppose otherwise. In the first fake barn case above, the agent was not aware of the existence of barn façades in the area and was therefore unaware of the risk in forming a false belief merely based on the appearance of the barn alone. Without some evidence to alert him to the existence of the barn façades, there would be no reason for him to suspect a failure to satisfy second-order competence. Thus reflective knowledge on this account requires that one be situated such that if the first-order belief failed to be apt, there would be some sign that this was the case. Sosa writes, “In order to be competent on the second-order, one must be so situated that sooner or later there *would* be tell-tale signs if the relevant first-order competence were absent” (473). Sosa thinks that this is what goes wrong in standard skeptical scenarios. That is, potential knowers are unaware, and unable to avoid being unaware, of the absence of first-order competence. In the fake barn case, the agent would continue believing that what he was seeing were barns, even though they were mostly barn façades, because there would be no tell-tale signs that his first-order competence was absent.

On this competence-based epistemology, error identification is paramount to establishing that any belief satisfies meta-competence. What this means is that without the possibility of the ‘tell-tale’ signs that our first-order competence is absent, we face the possibility that we are effectively in a skeptical scenario. In other words, we will be under the impression that our first-order beliefs are apt, when in fact we are in error. It is also possible that even in cases of a

failure of a belief to be meta-competent we can still have apt first-order beliefs, as was the case in the first barn-façade case above. The issue, then, is when we should care about meta-competence; since, arguably, many of our everyday beliefs about the world fail to be meta-competent without any serious negative outcomes.

Suppose I am attempting to lose some weight so I begin tracking my caloric intake using a food scale. If the food scale is inaccurate, yet I am still losing bodyweight over time, then there will fail to be tell-tale signs that my belief about my scale's accuracy fails to be apt. Nevertheless, since I am getting the results that I seek, there is not going to be any reason, in this case, to care about the meta-competence of my food-scale beliefs. There are obvious cases, however, where we will probably want more assurance before we act on a similar assumption. If I am a competitive athlete, I may want to know for sure that my food scale is accurate because my caloric intake is highly important to my performance in the sport. Even though there is the chance that an inaccurate scale will not greatly affect my performance, I may not want to leave it to chance that it will not. In that case, I may test the scale regularly to ensure that it remains accurate over time to a specified degree. Hence, one thing that might determine how much we care about meta-competence is the stakes involved in holding a particular belief.

There may be other issues, too. For example, it may not be possible to be meta-competent about certain beliefs that we are nevertheless required to hold in order to act. For example, I may be late for an event and form the belief that it will be faster to avoid freeway traffic by taking surface streets. Since I do not take the freeway, I will never know whether the route I took was, in fact, faster. I just needed to form a belief about which would likely be faster and act on it. Since I will never get to compare the time it would have taken me to get to the event via the freeway, I will never know if my belief was mistaken. If I end up making it on time, it will not be

of any benefit to further inquire about the belief. I can just accept that my belief was sufficient enough to get me where I needed to be on time, regardless for whether it was *actually* true. In the next section I map the above account of competence onto the political domain and consider the value of meta-competence therein.

III. Competence and the Electorate

If Sosa's conception of competence accounts for our intuitions regarding epistemic *and* skill-based competence, then we have a basic framework for judging when individual decisions and decision makers are, or will be, competent. Since political decisions of the sort that voters make appear similar in many ways to everyday epistemic performances, we can use this framework to judge the political decision making of both individual voters and the system of decision making as a whole. One difference is that the analysis of competence detailed in the previous section and the analysis of political decisions is that political decision making is a diachronic process wherein previous performances influence later performances. In other words, when we form beliefs about what happened with regard to the outcomes of past decisions, these beliefs will influence how we assess situations and evidence in the future. For example, the effect of price controls in the past has led economists to generally oppose price controls going forward. If our beliefs about the past effects of price controls are mistaken, then the possibility opens up that going forward we will be making an error.

In this section I map the account of competence laid out in the previous section onto political decision making. Further, I argue that this account of competence can be used to assess the performances of individual voters as well as the political decision making system as a whole. I also argue that political decisions are high stakes such that political systems should seek to

satisfy meta-competence. Finally, I explain the graveyard spiral problem and show that systems for political decision making face this problem if they do not satisfy meta-competence.

Recall that in the general framework for first-order competence of belief, the belief in question should be accurate (true) and that accuracy must be a manifestation of the adroitness of the process for forming that belief. In politics, we want true beliefs because they lead to good decisions. That is, holding true beliefs about the situation and the options available seems to be paramount to making good political decisions. So first-order competence in political decisions is aimed at achieving some political, social, or economic goal via our beliefs about the situation, which themselves manifest the adroitness of the process that led to the decision we ultimately make. First-order competence for voters, then, will be performances that aptly achieve some sought after end.

One issue here is how we understand what makes a decision accurate. That is, it is possible that a decision achieves some goal but does so at the cost of creating a new problem or exacerbating an existing problem. So aptness on the first-order seems to require that decisions made toward some end do not cause unintended consequences that make the situation worse all things considered or undermine the original goal itself over time. Imagine our huntress from before. She is stalking a deer shortly after sunset. She takes a shot that hits the target, but which does not kill the animal, merely wounding it and allowing it to escape. In that case, the shot was accurate to a degree, but only accurate enough to make the situation worse than if she had missed the shot completely (or not taken it at all). That is, if she had missed or not taken the shot, she would have been in the same position as she is in now that she has taken the shot, i.e. she is going home empty handed. The deer is all the worse off for it, and she no better off for it. Hence, with some degree of accuracy, we can make a situation worse off than it would have otherwise

been had we not acted. In the political case, then, aptness on the first-order should require that the situation is not made worse off for having acted, even if the action achieves some partial accuracy in achieving the intended goal.

Second-order competence in political decision making will be similar to second-order competence in the case of the huntress from the previous section, too. That is, voters will have to judge when a decision will be apt, which will require voters to make a judgment about the situation and the options available to address that situation. In the case of the huntress, she must consider things like the timing of the shot (if the animal is standing still or moving), the range of the bow and the distance to the target, and her own condition such that there is a high probability that the shot will be apt. If strong winds are blowing or there is not enough daylight left to see the target, she must withhold her shot until the situation gets better or abandon the shot altogether. In the political domain, decision makers must similarly make judgments about whether any proposed policy is likely to have the intended effect and whether it will cause new or exacerbate existing problems. Imagine a proposal to raise the minimum wage is selected in order to alleviate some of the economic hardships of low-wage workers. If the policy raises the minimum wage too high—with the result being large scale increases in unemployment for low-wage workers—and the voters did not realize this was the likely result, then that decision would fail to satisfy second-order competence. Similarly, if voters did not consider the connection between minimum wage and unemployment at all because they were ignorant of this connection, then that decision, too, would fail to satisfy second-order competence. In both cases the voters fail to realize that, in this situation, the decision they made would, or are likely would, fail to be apt.

In order for a political decision to satisfy meta-competence, the decision in question should satisfy both first-order and second-order competence. However, it is not clear that meta-competence is always necessary for every political decision to be considered appropriate. Before we considered the possibility that we are often forced by circumstance to hold beliefs that do not satisfy meta-competence. In those cases, there may be practical issues that prevent us from ensuring that our first-order competences are present. For instance, I may form the belief that the strange sound outside my window last night was just the wind. It may be the case that it was not the wind at all, but rather a neighbor stumbling home after a night out at the bar. Because I have no way to uncover this fact, the belief that I do hold is sufficient for me to move on because I don't have any reason to engage in further inquiry. In the political domain, there may some cases like this, but they seem to be rare. That is, since political decisions can have such broad effects on the population, there is reason to suspect that all or most political decisions are high stakes for someone.

The situation for political competence, then, seems to be that in any case where the stakes are high for someone, then that decision is high stakes. As such, analysis of the stakes should take into consideration any person that will be affected by policies or decisions made by elected and unelected government officials. This will include non-citizens residing in foreign countries, since foreign policy decisions can have large-scale impacts on the citizens of foreign countries. Hence, the responsibility for making apt decisions falls on each voter, even though the risk for any single voter making an error is quite small—i.e. if one voter makes an error it is unlikely to matter. On the other hand, when voters make systematic errors, there may be dire consequences. Hence, each voter has a responsibility to take the competence of their decisions seriously.

Brennan makes a similar point against democracy with an analogy with air pollution. When voters don't take their responsibilities seriously, they do something analogous to polluting small amounts into the air. As individuals it may not matter if we neglect our responsibilities to making apt decisions at the ballot box or preserving the environment, but collectively this lack of responsibility can have serious negative effects. He writes, "We have every incentive to free ride on others' efforts, externalize the cost of our biases onto others, and pollute democracy with our uninformed, misinformed, or irrational votes" (Brennan 2016, 144). Here Brennan is primarily talking about voters' failure to satisfy first-order competence, but the point stands for failures to satisfy second-order competence as well. That is, without second-order competence, we would never know that we were failing to satisfy first-order competence. Hence, meta-competence is essential to establishing the presence of first-order competence.

It is possible that individual decisions may be accurate, but fail to satisfy meta-competence. In these kinds of cases, decision makers may make a decision that ultimately achieves the relevant goal, but not because of the adroitness of the process that produced the decision. That is, the success of the decision happened in spite of, say, the epistemic process used by voters to come to the decision they made. Brennan thinks that the incompetence of any particular decision made by voters does not entail the incompetence of all decisions made downstream from that decision (Brennan 2016, 160-1). So voters can make incompetent decisions during an election, but still get competent government—e.g. when government officials override the will of the electorate. These kinds of cases will be analogous to the archer who makes a poor shot, but the arrow is blown off its original course by a gust of wind such that it hits the bullseye. Although the shot turned out to be accurate, that accuracy did not manifest adroitness. These are like Gettier cases in traditional epistemology, wherein the accuracy of the

result is not caused by the relevant performance. Any system for making political decisions that functions in this way fails to satisfy first and second-order competence, and should be labelled generally incompetent.

Additionally, we can separate two levels of analysis on competence regarding political decision mechanisms. First we can attempt to check whether individual voters tend to be competent. Brennan attacks democracy on these grounds, i.e. by showing that voters tend to be ignorant of basic political facts, we can conclude that these voters fail to satisfy first-order competence. This is because without the relevant facts, the accuracy of any decision will not be a manifestation of the adroitness of the voters' performance. Further, we can attempt to check whether a system as a whole tends to be competent. If the majority of voters lack first-order competence, then there is strong reason to suspect that the system itself will not satisfy first or second-order competence. This is because without the relevant facts of the situation, which we are assuming the majority of voters do not have, then it is not possible to establish second-order competence. There is also the further issue of whether or not voters who do satisfy first-order competence are also able to satisfy second-order competence. I will not attempt to give a full account for checking whether voters (or systems) satisfy second-order competence here, but it should be noted that determining second-order competence will likely be done by examining the epistemic processes that first-order competent voters use when making decisions. That is, if they are making decisions in epistemically responsible ways, e.g. by assessing the relevant evidence of the situation at hand and making their decisions based on this evidence rather than epistemically irrelevant reasons, then those voters may satisfy second-order competence. Hence, just in case voters tend to satisfy both first and second-order competence, the system may be rightly called competent. I address these issues in detail in chapter 3.

One of the virtues of satisfying meta-competence is that it ensures that errors can be identified and (hopefully) corrected. Without meta-competence, voters put themselves in a situation where they will be unaware when they are about to make an error, and likely unable to identify errors once they make them. This is because meta-competence requires there to be the possibility of tell-tale signs when first-order competence is absent. Without these tell-tale signs, voters will assume that their decisions are apt when they are not. Further, if voters continue to make decisions based on the falsely perceived accuracy of previous decisions, then the hope that their mistakes will be corrected quickly fades. This problem is similar in many ways to the situation that some pilots find themselves in when flying in areas of poor visibility, known as the graveyard spiral.

The graveyard spiral is a situation that primarily affects pilots who are licensed to fly by Visual Flight Rules (VFR). VFR-rated pilots typically fly by visual references like the terrain below and the visible horizon. VFR-rated pilots use these references as their primary source of information regarding the orientation of the plane. When a VFR-rated pilot enters an area of poor visibility—known as Instrument Meteorological Conditions (IMC)—the visual references can become obscured, and the pilot will quickly lose the ability to determine the orientation of the plane. The worst case scenario is one where the plane enters a turn that the pilot is unable to get out of, often resulting in the plane crashing.

There are two common ways this can happen. First, a VFR-rated pilot who accidentally enters IMC can attempt to turn the plane around and return to an area of clear visibility—known as Visual Meteorological Conditions (VMC). One problem that can occur here is that as the pilot turns the plane, the fluid in the inner ear shifts, creating friction that causes the sensation of turning in the pilot. As the turn progresses the fluid in the inner ear stops moving. Once the fluid

in the inner ear stops moving, it may feel to the pilot as though the plane is no longer in the turn since the movement of the fluid is what produces the sensation of turning. This can lead the pilot to conclude that the plane has exited the turn and returned to flying straight again, even though the plane is still in a banked turn. Recall that because we are assuming flight in IMC, there are no visual references for the pilot to use in orienting the plane. This can, if the pilot is not careful, lead him to attempt to reenter the turn, thus turning the plane even harder. The result of this is the plane entering a spiral, of which the pilot is unaware. If this mistake is not corrected, the plane will begin to quickly accelerate and lose altitude simultaneously. Because the pilot is now in a spiral and cannot determine this using usual visual reference, correcting the problem becomes very difficult unless the plane reenters VMC. If not corrected in time, the plane will crash into the terrain.

Another way the graveyard spiral can happen is when a VFR-rated pilot attempts to fly straight through IMC. This might happen if the pilot falsely believes that it is only a small patch of IMC and that they can make it safely to the other side by flying straight and level. Here the problem is slightly different. Because the pilot is not attempting to turn around, they will only be tasked with keeping the plane straight and level until they reach the other side of the IMC. If the pilot is not careful, the plane can begin to drift out of flying straight and level and into a slight banked turn. This is sort of like driving your car down a straight stretch of highway at highway speed then trying to continue driving straight with your eyes closed. Even though it feels as though you are continuing to drive straight based on the feel of the car's movement and the angle of the steering wheel, eventually your car will begin to drift out of your lane and possibly off the road. Similarly, the plane will slowly begin to bank, which will lead the plane into a slight turn.

Because the plane enters the turn slowly, the fluid in the pilot's inner ear will not register this change, and the pilot will be led to believe that the plane is still flying straight and level.

Once the plane has entered the turn, it will begin to lose airspeed and altitude because of the change in direction. When the pilot notices the altimeter showing the plane is losing altitude, and given the pilot's assumption that he is flying straight and level, he may attempt to pull back on the yoke to increase altitude. However, because the plane is in a banked turn, this action will only cause the turn to tighten. When this happens it will feel to the pilot as though the plane is gaining altitude, but the altimeter will show the opposite. That is, the sensation of the plane gaining altitude and the plane tightening a turn feel similar enough that the pilot cannot distinguish between them without visual references. Because of this, the altimeter will show an increase in the loss of altitude in spite of the pilot having the sensation of the plane climbing. The pilot now has a problem, either the plane is in a turn verging on a spiral or the altimeter is faulty. Due to the powerful sensations coming from the vestibular system telling the pilot that the plane is climbing straight and level, often pilot's will disregard the readings of their instruments and attempt to fly on feel alone. This, much like the previous example, will lead the pilot to crash rather quickly as the plane's turn continues to tighten until either the plane impacts the terrain or reenters VMC and the pilot is able to correctly orient the plane again.

One further issue to note here is that even in cases where the pilot suspects they are in a banked turn, it can be very difficult to get out of the turn in IMC. This is because, as already noted, when the plane slowly enters the turn, the fluid in the inner ear does not register the change in direction. So the pilot does not feel any sensation of turning. If the pilot nevertheless realizes that he is in a turn, when he does correctly exit the turn the fluid in the inner ear will shift quickly. If the plane exits, say, a left turn and begins flying straight again, to the pilot it will

feel as though the plane has now entered into a right turn. Because these sensations are so powerful, they can convince the pilot that he has overcorrected, leading him to reenter the left turn as he tries to correct the falsely perceived overcorrection. Once back in the left turn, it will feel as though the plane is flying straight. Without correction, this, too, will result in the plane crashing into terrain unless the plane exits the IMC such that the plane can be brought back under control.

Taken together, there are some interesting parallels between political decision making and VFR flying in IMC. First, by flying into IMC the pilot is not in a position to determine whether any attempt to alter the orientation of the plane will be apt. He just isn't situated to make this sort of judgment. This is because he has two sources of information that are jointly incoherent and no way to determine which information is correct. First, he has the powerful sensations produced by the vestibular system providing him with information about the orientation of the plane and its likely direction. Second, he has the instruments telling him something that contradicts these sensations. Hence, the pilot must decide which set of information to base his actions on. If he is wrong he puts himself in grave danger. Unfortunately, in either case, it will not be clear if this decision was correct until the plane either exits the IMC or crashes. Because of this, the pilot is in a situation where his second-order competence is paramount. If he does not satisfy second-order competence, attempting to exercise his first-order competence could lead him further into peril. Hence, for the pilot, first-order competence is grounded in his ability to orient the plane using visual references, without which he is left with no tell-tale signs that he has made an error until it is too late. The absence of his first-order competence is something that he cannot recognize until the effects of his errors are catastrophic.

In order to avoid this sort of situation in politics, where errors are taken to be corrections and corrections taken to be errors, there must be some reason to think that political decision makers are capable of satisfying second-order competence. In order to satisfy second-order competence, these decision makers must be able to recognize their errors through the tell-tale signs that their first-order competence is absent when it is. That is, when their first-order decisions fail to be apt there will be no evidence this is the case without being situated such that second-order competence can be satisfied. Without this, political systems will be in the same position as the VFR-rated pilot flying in IMC, they will make decisions without any indication for when they have made an error until it is too late to make corrections.

Recall that Brennan argues for the right to competent government, for which he formulates the following principle:

Competence Principle: It is presumably unjust, and to violate a citizen's rights, to forcibly deprive him of life, liberty, or property, or to significantly harm his life prospects, as a result of decisions made by an incompetent deliberative body, or decisions made in an incompetent way or in bad faith. Political decisions are presumed legitimate and authoritative only when produced by competent political bodies in a competent way and in good faith. (Brennan 2016, 156-7)

As I argued above, decisions made in situations where second-order competence cannot be satisfied are incompetent. This does not preclude the presence of first-order competence, since it is possible that competence is present on the first-order but not the second. But, if our understanding of complete competence is apt performance with meta-competence, then any decision made without the possibility of satisfying second-order competence will be an incompetent decision. According to the competence principle, it appears that in order for an incompetent, or bad faith, decision to be unjust it must deprive someone of life, liberty, property, or significantly harm their life prospects. But this does not seem correct. That is, decisions seem

to be unjust if they are made incompetently whether they *actually* harm anyone or not. This is because, without meta-competence, there is always the danger that a decision *will* harm someone because decision makers aren't in a position to judge whether their decision will or will not harm someone. Hence, meta-*in*competent decisions are always at risk of putting someone in harm's way. This seems, at least to me, to be an injustice. In the case of a VFR-rated pilot who intentionally flies a plane with passengers into IMC, *without* knowing the dangers of doing so, acts incompetently. It also seems reasonable to say that this is an injustice even if no one is *actually* harmed. This is because the pilot's ignorance put passenger's lives at serious risk.

There is also the issue of how we are to interpret 'bad faith' here. One way to understand a bad faith political decision is one that intentionally harms some individual, or group, in the way described in the principle. Another way might be when a decision is knowingly made in an incompetent way. That is, when a decision is made by decision makers who know they are not situated so as to satisfy meta-competence might be a bad faith decision. For instance, a VFR-rated pilot who intentionally flies a plane with passengers into IMC, with full knowledge of the dangers of doing so, acts in bad faith because he knows he is incompetent to fly in those conditions. Therefore, he knowingly puts himself and his passengers in danger, which is unjust. It does not seem necessary that these actions end in physical harm to the passengers for them to be unjust. That is, passengers may rightly call foul even if the plane does not end up in a graveyard spiral or crash. The mere possibility of the graveyard spiral, and the likely inability of the pilot to get out of it if it did happen, seems to be enough to call the pilot's actions unjust.

Whether we conclude that these kinds of decisions are unjust or not, Brennan's argument against democracy on the basis of incompetence seems to hinge on the thesis that they are.

Brennan's argument depends crucially on the idea that democracy likely systematically violates

our rights because voters in modern democracies tend to fail to satisfy first-order competence, which is due to their ignorance of basic political facts. Note that if Brennan thought systematic harm was a requirement for there to be a systematic violation of our rights, he would only need to list the harms that resulted from incompetent decisions to demonstrate that democracy systematically violates our rights. Because he does not do this, it seems reasonable to assume that Brennan's account of the right to competent government does not require identifiable harms to occur, only a likelihood that they will. This seems to be further supported by the last part of the principle, where it states that government decisions are only legitimate and authoritative if they are made competently and in good faith.

If Brennan thinks this is enough to judge democracy unjust, then it seems reasonable to assume that any system that is systematically meta-*in*competent will likely violate our rights. If all this holds up, then a system that is unlikely to violate our rights will be one that can identify and correct errors because it satisfies meta-competence. Hence, meta-competence is required for any system of political decision making to be competent and just.

IV. Conclusion

I argued that Sosa's account of competence can be used to establish a framework for judging the competence of both voters and political decision mechanisms. That account distinguishes between first-order competence, which is the combination of accuracy and adroitness so as to satisfy aptness. Competence on the second-order is satisfied when we aptly judge when our first-order competence will likely be manifested and when it has manifested in past performances. This requires an agent to be situated such that there will be tell-tale signs that the first-order competence is absent. If both first and second-order competence is satisfied, then meta-

competence is satisfied. Meta-competence, then, is required to avoid skeptical scenarios and Gettier cases in traditional epistemology. In the political domain, meta-competence is required to avoid the graveyard spiral problem, where errors are mistaken for corrections and corrections are mistaken for errors. Finally, if we accept the basic premise of Brennan's argument that justice requires competence in government, then a just government will be one that satisfies meta-competence.

Assuming democracy is not able to satisfy complete competence because it fails to satisfy first-order competence, then so far my account is in agreement with Brennan's view that democracy is likely incompetent. What remains to be seen is whether epistocracy can be expected to satisfy complete competence. This is the subject matter for the next chapter. There I argue that we should expect epistocracy, like democracy, will likely fail to satisfy first-order competence.

Chapter 3

Epistocracy, Competence, and the Trouble with Trial and Error

In the last chapter I proposed a general account of competence that can be used to analyze competence in political systems. In this chapter I utilize this account of competence to try and get clear about the differences between epistocracy—as construed in Brennan (2016)—and democracy with regard to competence. I argue that there is no significant difference between the competence of epistocracy and democracy so long as the primary problem solving mechanism of epistocracy remains the same as in democracy, i.e. so long as voters are forced to choose between two or more options in a system of centralized decision making. In the first part of this chapter I try to make clear what it is that voters are expected to be competent at when voting, i.e. what they are actually epistemically responsible for doing. In the second section, I turn to epistocracy as it is described by Brennan where I argue that epistocracy does not satisfy complete competence.¹ In the third section, I consider whether further limitation on suffrage will resolve the issue. I argue that it cannot, and, further, that so long as government is committed to centralized decision making the threat of the graveyard spiral will be present.

I. Competence and Competent Voters

Recall that in the last chapter I argued that the account of competence in Sosa (2010) can be used to assess political competence. Like dispositions, understanding when a competence is present requires us to get clear about what we should expect to occur under particular conditions. For example, the mug on my desk will likely break if dropped from ten feet above the floor, but it

¹ For simplicity, from here on I will use ‘complete competence’ to denote satisfaction of both meta-competence and of the competence triad—constitution, condition, and situation. In cases where the failure of complete competence is due to a failure of second-order competence it will be noted as such.

likely will remain intact if dropped from only an inch above the floor. Similarly, with epistemic performances there are cases where a putative knower will fail to have knowledge due to facts about their situation or their constitution. Competence with regard to making political decisions will have a similar component since political decision making is primarily an epistemic performance. That is, voters are required to gather information, reason about that information, and make political decisions given the available options for policy and candidate choice.

What is not clear is what this actually entails, i.e. we do not yet know what voters are responsible for doing. For instance, there are different ways to make decisions and different sets of information that can be assessed by voters during the process of forming beliefs about the decisions they face. If a voter is an expert in some politically relevant field, then they may be more inclined to look at technical studies to inform their decision making. On the other hand, some voters who are novices with regard to all politically relevant fields may make decisions using totally different sets of information. It is not clear how we should assess this plurality of differing epistemic backgrounds and methods of information gathering in making general claims about the competence of political decision-making systems.

It could be the case that there are several methods of inquiring about politics that satisfy complete competence. An example of this might be a voter who stays up on economic news and is attentive to the prescriptions of expert economists. In that case, there might be reason to think that that voter is competent because of *who* they listen to when gathering information to inform their political beliefs. That is, a voter might be well attuned to the best experts in relevant fields and be properly situated to make competent decisions because of this. Because of the possibility that there is actually a range of ways that voters can satisfy competence, I will assume that there is no one single method for a voter to be competent. What matters is whether they actually are

competent using the method they use to make political decisions. Hence, there will be a range of methods, with some satisfying complete competence and some that do not.

If we look to other domains, we can see that this situation is quite common. That is, different people may use different methods to get more-or-less the same results. For instance, different coaches use different training protocols for athletes competing in the same sport. If we can identify the range of methods that meet the basic standard for competence in politics then we can assess which voters will likely meet the minimum standard for competence. Some of this work was done in the last chapter, where I argued that satisfying competence requires the ability to track the effects of policy in order to identify and correct errors. But this level of analysis does not tell us what to expect from the various ways that voters form beliefs about politics.

Below I try to establish a range of plausible epistemic responsibilities that could be required of voters, which will include the highest and lowest plausible levels of epistemic responsibility.² This will allow the broadest range of plausibly competent methods to be considered and help to triangulate where it is on the spectrum epistocracy falls, which will allow a better informed comparison between the expected competence of democracy and epistocracy. First, I want to start with methods that will meet the highest plausible standard of epistemic responsibility. On what I will from here on call maximalism about epistemic responsibility, voters will be maximally epistemically responsible for their political decisions. This will require voters to have a considerable understanding of the various complexities of the policies they support or reject. For instance, voters will need to have a significant understanding of the

² Epistemic responsibility here is meant to denote the beliefs that an agent is personally responsible for, i.e. the beliefs that arise in the agent rather than beliefs that are merely reflections of the beliefs of peers or experts. Being epistemically responsible for a belief requires an agent to expend some epistemic labor in the formation of that belief.

intricacies of how a policy will work to achieve its intended ends, how a policy will interact with the broader system of policies and laws that are already in place, the likely social costs and benefits of these policies, and to be personally epistemically responsible for tracking and understanding the effects of their decisions over time. This will require them to track the effects of both policies they support and those they reject.

Maximalism about epistemic responsibility will require voters to undertake a great deal of work to understand subjects like economics, domestic as well as international politics, the operations of domestic government, policy making, social and political theory, and philosophy. On its face, maximalism as sketched here seems to be unreasonably demanding. Most voters don't have the time or resources to fully investigate all these areas to the best of their ability. At best, voters may have a general understanding of one or more of the above subjects, but few (if any) will have a considerable understanding of them all. However, maximalism is meant to represent what an ideal electorate might look like, i.e. good faith political inquirers whose beliefs are guided by evidence, which they can understand. It may not be likely to be achieved in the real world, but it nevertheless stands in stark contrast to what we have come to expect from voters in modern democracy.³ On maximalism, then, voters must be able to actually understand expert analysis themselves in such a way that they could raise valid objections and accurately reconstruct expert testimony regarding a policy position. This standard would also require voters to be able to generally adjudicate between experts such that, at least in the majority of cases, they would be able to decide which expert had the better argument in cases of expert disagreement. This is not to say that they will be able to decide who is *right*, but, rather, be able to construct their own argument for why one expert's position is stronger because they actually understand the

³ I should note here that this standard is similar to what Brennan calls *vulcans*, which I discuss in greater detail below (Brennan 2016, 5).

reasons being given by each expert. So they will not need to be experts *per se*, but something like a level below expertise. They might, then, be understood as journeyman with regard to political decision making.

On the other end of the spectrum will be minimalism about epistemic responsibility, which I see as providing minimal, yet plausible, epistemic demands on voters. On this standard, voters will merely be required to align their support for policies that have been given a reasoned defense, along with a critique of rival positions, by an expert in the relevant field. This condition will give a sort of basic legitimacy to any policy position that satisfies it. What this means, roughly, is that so long as there is some expert who has done the epistemic work regarding policy analysis to legitimize it, then there is legitimacy to that policy position regardless of the epistemic status of the non-expert supporter of that policy position. Essentially, on minimalism, there is no technical epistemic work to be done directly on policy or candidate choice by non-experts. Voters only need to listen to experts and make decisions based on the testimony of experts they trust.

This is how many expert-lay relationships work in other domains. For instance, when I go to the doctor I am reliant on the doctor's understanding of, say, bacterial infections and antibiotics to justify to myself (or others) why I am taking a particular antibiotic. In other words, I do not need to really understand antibiotics or bacterial infections to be fully justified in accepting the prescribed treatment. I could undertake the task of learning about how different infections respond to different antibiotics to better understand why I was prescribed a particular drug, but it is not generally accepted that it is my responsibility to do so. It is the doctor's

responsibility to properly diagnose my condition and prescribe the correct treatment.⁴ This is not to say that we are totally off the hook, epistemically speaking, when we go see the doctor. But there does not seem to be, generally, a norm of epistemic responsibility to undertake a course of study any time an expert recommends a treatment for some problem of ours, whether it is medical advice or otherwise. We might be responsible for choosing a good doctor or for determining whether the doctor we do choose is properly credentialed. If I were maimed by a surgeon who operates out of his garage, I might be at fault for choosing to undergo that treatment. On the other hand, if I am maimed by a doctor who is properly credentialed and operating under proper conditions, then I would seem to avoid any fault or accusations of irresponsibility for the outcome. In that case I would have acted as any other responsible consumer of medical treatment would.

On minimalism, then, voters will have some epistemic responsibility. Namely, they will be responsible for the epistemic labor required in selecting the experts whose policy analysis they accept. First-order competence on minimalism will be the ability to effectively pick out who is and who is not an expert on any given issue. Second-order competence, then, will be the situational component that allows these voters to judge when their first-order competence is present and to determine when they have made a mistake on their first-order judgments about who is and who is not an expert. It should not be controversial to note that very few policies or

⁴ One might object that we are responsible for our bodies and therefore we have a responsibility for making sure that the treatments we are prescribed by doctors are correct. Depending on the stakes involved, we may look for a second opinion or embark on a bit of research to better understand our condition. But these kinds of cases seem to only arise when the stakes are high for us, as might be the case with a terminal illness. Generally speaking, we do not feel responsible for understanding how a treatment works unless there is some possibility that an error by the doctor will cause us considerable harm. It may be the case that this same reasoning extends to politics—i.e. political decisions are high stakes and therefore we have good reason to look into policies beyond taking the word of experts—but I am here assuming the lowest reasonable standard for epistemic responsibility. It may well be the case that the nature of political decisions entails all political decisions are high stakes (I made a brief argument for this conclusion in the last chapter), but for our purposes here I will assume that the low bar for political responsibility as sketched here is at least plausible.

policy platforms are universally accepted by the relevant experts. Hence, there is room here for voters to exercise their epistemic powers in the selection of experts. This seems to parallel the case of choosing a properly credentialed surgeon above. That is, voters only need to be able to judge which experts are best situated to give correct advice in the same way consumers select doctors who are best situated to correctly diagnose medical conditions. Further, the experts who produce the analysis that is relevant to making political decisions tend to be known public figures. This means that their credentials, and often their reputation in their field, can be ascertained by an internet search.

On this standard, then, voters will only be required to support positions that are also supported by a properly credentialed expert in the relevant field. Since we are assuming that to have a legitimate policy position here only requires voters to support policies that are also supported by an expert, the primary epistemic labor voters have here is determining who the experts are and what positions they support. Similar to how coaching in various sports works, there are many ways to train for a sport and each athlete must submit their will to the advice of a properly credentialed coach. Thus, voters will be required to find experts who are properly credentialed and who are capable of offering reasoned defenses of the positions they support.

Since the competence manifested by voters on minimalism only regards their choice of experts, there is a disconnect between the results of policy and the epistemic labor undertaken by voters. That is, voters are not themselves responsible for determining when a policy has been a failure; they are responsible for responding to expert testimony that this is the case. This is similar to how a patient responds to their doctor's assessment of a prescribed treatment. If the treatment does not seem to be working, the doctor will make the judgment about what to do. It is not the patient's responsibility to track the results of the treatment and determine if the doctor has

made an error. It is the patient's responsibility to follow the orders of the doctor to ensure the treatment is carried out as prescribed. In the same way, then, voters will be charged with submitting to the judgments of experts, which they voluntarily selected.

There are two things I should note here. First, minimalism is not intended to license voters in selecting experts in biased ways such that they only listen to experts who offer justifications for their preconceived ideas about policy or candidate choice. The reason for this is that voters here do not have the epistemic background to make the sorts of judgments that would justify them in this kind of cherry-picking. They are reliant on experts to form their opinions in the first place, and are therefore incompetent to form the relevant judgments about the technical aspects of policy analysis. Voters still have a responsibility to choose experts based on who is likely to be correct in their analysis of policies and candidate platforms. The point of the low standard is to allow the possibility that those who have very little social scientific knowledge themselves are still able to make competent political decisions because of the advice they are given by experts. In the same way that I can make competent decisions about my health while having almost no medical knowledge, voters may be perfectly able to make competent political decisions because they make careful decisions about the experts whose advice they accept.

Second, I do not think minimalism is likely to result in competent political decisions generally. This is primarily because the list of experts who publish opinions on political issues are many and who is and who is not an expert on political issues is controversial. So there is a sort of disanalogy here with the medical examples given above. This is because doctors go through similar training and licensing in order to practice medicine. The same is not true for politics because it is inherently an interdisciplinary field. As noted in the discussion of maximalism, taking on personal epistemic responsibility for understanding all the relevant details

of policy and candidate choice requires having some basic level of understanding of many different fields of study. There are certainly doctors who specialize in treating certain diseases or who specialize in a particular area of the body—e.g. oncologists and endocrinologists, respectively—they nevertheless are required to undergo basic training in all the systems of the body. The case is much different in politics, where economists are not required to train in sociology and sociologists are not required to train in economics, unless a specialization requires it. This is not to say that there are not economists who have a solid understanding of sociology, but just that being an economist does not require one to also be a trained sociologist.

Although it may not be completely clear at this point what voters should be doing epistemically, we now at least have a basic idea of what might be required of them if complete competence is what we are after. That is, there is plausibly a plurality of methods of inquiry that can be used to make competent political decisions, which should fall somewhere between the two standards sketched above. Further, we can now attempt an analysis of the methods of political inquiry that are likely to dominate epistocracy in order to assess the likelihood that they will satisfy complete competence as described in the previous chapter. Below I argue that epistocracy faces serious problems with regard to competence primarily because of the epistemic character of the voters.

II. Epistocracy and Political Hooligans

One of the central claims that Brennan makes is that in modern democracy most voters are epistemically ill-equipped to make competent political decisions. This is due to the fact that many voters are ill-informed about basic facts regarding politics. As noted in the last chapter, this leaves many voters incompetent on the first order. Brennan's response is to suggest different

strategies for mitigating the influence of these first-order incompetent voters. The primary suggestion is restricting the suffrage of the incompetent. On Brennan's view, voters who pass the test⁵—whatever it turns out to be—will then be those who are situated to make (more) competent political decisions. The problem here is that Brennan seems to admit that the epistocratic electorate he imagines will be dominated by political hooligans (Brennan 2016, 207). Although Brennan speculates that hooligans in epistocracy may be better, epistemically speaking, than their democratic counterparts—since after all they did make it through the transition to epistocracy—they will still have the basic character of political hooligans that Brennan describes in the opening pages of *Against Democracy*.

The purpose of this section will be to gain a better understanding of hooligans and attempt to infer from Brennan's characterization of them if we should think they are likely to satisfy complete competence. I begin by looking at Brennan's assessment of hooligans in order to first understand how it is that Brennan understands them. This will give us a general idea of what it means to be a hooligan and what their epistemic character is like. Brennan's characterization of hooligans is this:

Hooligans are the rabid sports fans of politics. They have strong and largely fixed worldviews. They can present arguments for their beliefs, but they cannot explain alternative points of view in a way that people with other views would find satisfactory. Hooligans consume political information, although in a biased way. They tend to seek out information that confirms their preexisting political opinions, but ignore, evade, and reject out of hand evidence that contradicts or disconfirms their preexisting opinions. They may have some trust in the social sciences, but cherry-pick data and tend to only learn about research that supports their own views. They are overconfident in themselves and what they know. Their political opinions form part of their identity, and they are proud to be a member of their political team. For them, belonging to the Democrats or Republicans, Labor or Tories, or Social Democrats or Christian Democrats matters to

⁵ Brennan considers several options for how we might go about eliminating, or reducing, the effect of incompetent voters on the system, but stops short of fully endorsing any one of them (Brennan 2016, 208-222).

their self-image in the same way being a Christian or Muslim matters to religious people's self-image. They tend to despise people who disagree with them, holding that people with alternative views are stupid, evil, selfish, or at best, deeply misguided. Most regular voters, active political participants, activists, registered party members, and politicians are hooligans. (Brennan 2016, 5)

There are several key issues that should stand out given the discussion in the previous section.

Hooligans seem to be inclined toward biased assessments of evidence, they seem to fail to understand contrary points of view, they avoid evidence that they think undermines their previously-held beliefs, and they seem to overestimate the strength of their positions on policy and candidate choice.

Indeed, Brennan's portrait of hooligans seems to imply that they utilize a method akin to what C.S. Peirce called the method of tenacity (Peirce 1877, 7). Peirce held that there were four methods for "fixing" belief—i.e. to form a stable belief on some issue—the method of science, the method of tenacity, the method of authority, and the *a priori* method. According to Peirce, when we encounter evidence that one of our beliefs is mistaken it causes a sort of irritation. This irritation leads us to further inquire to relieve ourselves of this doubt and return to an equilibrium between our beliefs and the evidence. The method of tenacity requires one to cling to beliefs and avoid contradicting evidence in order to avoid the irritation of doubt. Peirce says,

If settlement of opinion is the sole object of inquiry, and if belief is of the nature of a habit, why should we not attain the desired end, by taking any answer to a question which we may fancy, and constantly reiterate it to ourselves, dwelling on all which may conduce to that belief, and learning to turn with contempt and hatred from anything that might disturb it? (6)

This brief characterization of the method of tenacity seems strikingly similar to the characterization of the practices and tendencies of hooligans. According to Brennan, political hooligans cherry-pick evidence that confirms their beliefs, they avoid contrary evidence, and

they do not really understand contrary positions held by their rivals. It seems that they are prone to cling to their views in spite of what the evidence might say. The question that concerns us here is whether this method is likely to satisfy complete competence and whether it falls within the range of plausibly competent methods discussed in the previous section.

Peirce clearly thinks the method of tenacity is problematic, although he praises it for its strength in achieving its end, i.e. it seems to him to be a powerful means to fix one's beliefs. But, as Peirce notes, the method of tenacity is not concerned with a strict adherence to the evidence. Rather, this method is primarily concerned with the alleviation of doubt and the accompanying peace of mind that comes with it. Since our idea of complete competence requires mistakes to be identified and corrected, it would seem that the end of using this kind of method is directly contrary to that goal. We can ask which is more likely, that a hooligan will accept that their beliefs were false and their decisions in error or that they will ignore evidence that they were in error and cling to their beliefs in spite of the evidence? It seems clear from the description above that hooligans are not good-faith inquirers and we should not expect them to assess the results of their decisions in a fair manner.

An interesting thing to note here is that Peirce thinks the method of tenacity will likely fail to be useful over time (Peirce 1877, 7). He thinks that we are unlikely to make it through daily life without realizing other people hold different views than we do and that their views are often as well founded as our own. This will undermine our ability to maintain the beliefs we try to cling to. However, in the case of politics, and in the case of hooligans in particular, we already expect that others will hold different views than we do. Indeed, in democratic politics we vote primarily *because* there are others with whom we disagree. We need a method to settle this disagreement. In politics, we accept that others disagree and often sort ourselves into factions of

likeminded people. The result is the very ‘team sports’ aspect of democratic politics that hooligans seem to have perfected. That is, hooligans treat politics not as a serious epistemic inquiry into what we should do, but rather as a sort of sport where winning for the group is prioritized over actually being right.

Based on what we’ve said so far, it does not seem like the method utilized by hooligans is likely to meet even the minimalist standard. This is because even there we require voters to make a good-faith attempt at discovering which experts are actually right. If I am correct in thinking that the methods used by hooligans do not have this as an end, then we have reason to reject the possibility that hooligans will satisfy complete competence. On this point there is an objection that I would now like to consider. The objection is that even if hooligans do not, themselves, engage in good-faith inquiry, it is still possible that they take advice from the best experts in spite of their flawed methods. Essentially, hooligans may be misguided in their methods, but their intentions are to do what’s best, even if they do not know what is best. So it is possible that complete competence is satisfied, but for reasons other than the epistemic character of the voters. One way this might be true is that hooligans take the advice of experts, who are presumably good-faith inquirers, and so their beliefs are conditioned by expertise nonetheless.

This objection is akin to an argument that minimalism *could* produce decisions that satisfy complete competence. That is, on minimalism voters are not themselves responsible for doing any epistemic work regarding particular policy questions and if minimalism is at least plausible, then the methods of hooligans are at least plausible, too. This is because the case of hooligans parallels the case of mostly ignorant voters whose beliefs are totally determined by the beliefs of experts, except hooligans merely have more beliefs of their own on policy and candidate choice. They should not be faulted for going beyond what is required on minimalism,

even if their doing so requires them to utilize the method of tenacity. All they need is to show that their positions are defended by experts and then their positions are made legitimate.

I think this view is wrong primarily because good-faith inquiry is a condition on minimalism, but also because in cases where hooligans are correct in their decisions, they are not correct for the right reasons. Take the example of a person who only takes advice from doctors who agree with their self-diagnosis, which we will imagine here is derived not from medically relevant evidence but from some other source, say, a dream. In that case, if the patient turns out to be right about their condition they are still not in any sense satisfying complete competence with regard to that belief. Although they were right about their diagnosis, it was not because of their epistemic performance. An agent may be correct to believe some proposition, but wrong to believe it for the reasons that they do. Since they believe for the wrong reasons, their belief fails to be apt since the accuracy of the belief does not manifest the adroitness of the performance; it is more-or-less accidental. Further, selecting experts based on whether or not they espouse the same beliefs as you destroys the purpose of seeking expert advice.

If the beliefs held by hooligans are determined by their biases, then there is never any room for their first-order competence to manifest. What's more, cherry-picking evidence is an example of an intentional short-circuiting of the epistemic process. Indeed, cherry-picking evidence seems to be a corrupted use of genuine inquiry, i.e. one must know what evidence to avoid and which to seek out in order to maintain beliefs in this manner. Thus, they must use something akin to what Pierce calls the method of science to sort evidence in a way that allows them to maintain their beliefs. If they are wrong about what evidence to avoid, the method of tenacity will fail. If there is complete competence displayed by hooligans, it would only seem to be manifested in the successful maintenance of their beliefs in light of counterevidence. While

hooligans likely do meet the condition on minimalism of holding positions defended by experts, they undermine that condition in the manner in which they select experts.

It should be clear at this point that hooligans do not satisfy maximalism about epistemic responsibility. They do not engage in good-faith inquiry, they do not really understand alternative positions, and, because of this, they likely do not even really understand their own positions. Their methods of inquiry appear to be primarily intended to maintain the strength of their previously-held beliefs and not to discover what is actually true. Further, they are not responsible for the epistemic work done to justify the beliefs about policy that they hold. Finally, their only real epistemic labor is done for the purpose of maintaining their beliefs, even in spite of counterevidence, and the certainty they display is not a reflection of the evidence.

We now have two conclusions regarding hooligans. First, the epistemic methods they seem to use do not even satisfy the low standard of epistemic responsibility, which we have here merely assumed could satisfy complete competence. Second, because of the flawed methods they use, they cannot reasonably be expected to satisfy complete competence. When they are correct on any issue they are not responsible for that outcome and when they are wrong they will likely fail to accept this fact. This leaves hooligans in the position of failing to satisfy first and second-order competence. They fail on the first order because, even when they are right, it will not be the result of their first-order competence manifesting. They fail on the second order because second-order competence cannot be satisfied without a relevant first-order competence. Recall that second-order competence is competence with regard to the judgment that one's first-order competence is present. If there is no first-order competence, then there cannot be a second-order competence. Finally, without the satisfaction of complete competence, the threat of the graveyard spiral will always be present. Hence, epistocracy, as presented by Brennan, will face

the threat of the graveyard spiral. Indeed, with what we have said so far there does not seem to be any reason to think that this threat is reduced at all on the switch from democracy to epistocracy. On both systems, we should expect the majority of voters to lack first-order competence with regard to political decisions.

Given that we should not expect an epistocracy dominated by political hooligans to be competent, I now want to attempt to get clear about what a competent electorate might look like. Brennan has some ideas about what will increase the competence of government, and his arguments give us some idea for where we might begin. In the first section I assumed it is at least plausible for minimalism to result in competent decisions by voters. In the section that follows I present an argument against this assumption. The reason for this is twofold: First, competence cannot fully be inferred by looking only at the epistemic character (the constitutional component) of voters, but must also include the epistemic problems they face (the situational component). We already ruled out hooligans from being competent decision makers because they fail to have the proper constitutional component required to satisfy competence. Once we look to the situational component we can with greater precision understand what is required by complete competence with regard to political decisions. Second, Brennan would likely reject minimalism as sketched here. Although he does not consider it as I have described it above, he argues against a corollary notion, which states that voters need only align themselves with a political party in order to make competent decisions. This view is sometimes referred to as the party identification shortcut. Roughly, the view is that political parties offer a sort of epistemic shortcut that allows voters to make decisions with minimal political knowledge. On this view, instead of voters being required to make decisions about experts, they are required to make decisions about political parties or other “thought leaders” (Somin 2013, 106-35). The differences will turn out to be

minimal with regard to the likelihood that either minimalism sketched here or the party identification shortcut will result in competent decisions.

III. Competence and Competent Methods

Recall that on minimalism, there are going to be very minimal epistemic demands on voters. They are responsible for engaging in inquiry to determine who is and is not an expert on the relevant issue and who amongst the identified experts are most likely to be right in cases of expert disagreement. Getting these questions right is the object of their first-order performance. With regard to second-order competence, voters must be situated such that they can effectively judge when their first-order competence is present prior to a performance and that their first-order competence has been present in past performances. Success here would offer relatively low-information voters the ability to make competent decisions while having little-to-no political or social scientific knowledge. In what follows I argue that low-information voters will not be properly situated to satisfy second-order competence, and therefore will not be capable of satisfying complete competence. First I begin with an assessment of what exactly is required for these low-information voters to satisfy second-order competence and show that we should expect them to fail. I also note some comments Brennan makes on this issue with regard to the party identification shortcut and show that it faces the same problem. With these two methods ruled out as being reasonably expected to satisfy complete competence, we can infer that a greater understanding of politics and social science is required to make competent political decisions.

The question that this section is meant to answer is the following: What does it take to satisfy second-order competence on minimalism? It would seem that in order for second-order competence to be present, voters must be situated in such a way that when they do go wrong in

their first-order performances, there would be some sign that would tell them this is the case. We can imagine simple cases where a voter makes a decision based on an expert's advice and the outcome more-or-less matches what the expert predicted would happen. We can imagine other cases, with the same set-up, but where the expert's prediction seems to have clearly failed. In those cases it would seem that the voter really only has to pay attention to what has come about following a decision and be capable of understanding the basic difference between successful and failed predictions. That is, if an expert predicts, say, that unemployment will go down if policy P is enacted, and unemployment goes up after P is enacted, then it would seem as though the expert's prediction has failed. In that case, the voter can take this information on in order to make better decisions in the future.

A problem arises when we realize that a failed prediction does not necessarily entail that the expert has made an error. In the case where the expert predicts that unemployment will go down following the enactment of P, it could be the case that, given the evidence, P really should have been expected to lower unemployment. The failed prediction could have been caused not by a failure of the expert to properly analyze the evidence and reason correctly from that evidence to form a prediction about what will happen, but rather it would have been a failure of the evidence to pick out what would happen. In other words, a failed prediction can be caused by insufficient or misleading evidence. Since in these kinds of cases it is not clear that the expert has actually made an error, there does not seem to be any grounds for the voter who acted on that expert's advice to claim she has made an error.

Consider the following example. Suppose you begin to present with a cluster of troubling symptoms and decide to go see the doctor. The doctor takes some samples, runs a series of tests, and decides that, based on the test results, you have disease D. As it turns out, disease D is

somewhat common and easily treated with medication. Several weeks later, after completing your prescribed treatment, the symptoms of disease D have not abated. After a reevaluation the doctor discovers that the testing equipment used in the previous evaluation had delivered a false positive for disease D and you actually have another condition altogether. It is not clear that in this case the doctor has made any errors. He carried out the correct tests given your symptoms, read the results accurately, and prescribed the correct treatment given the diagnosis. The error in the diagnosis was the result of a false positive, which provided the doctor with misleading evidence about your condition. Hence, in this case a failed performance by the expert demonstrated little, if anything, about the expertise of the doctor or his competence to diagnose patients with your condition. Both you and the doctor were, in some sense, victims of bad luck.

In this case you really only have grounds to judge that an error was made at all because you experience the symptoms yourself, but this will not always be the situation. Imagine a similar case where you present to the doctor with a different cluster of symptoms. Given the particular symptoms you have no tests are needed and the doctor diagnoses you with condition C. The doctor prescribes medication for C and sends you on your way. Unfortunately, you do not actually have C, but instead you have a slightly rarer condition that sometimes presents with the same symptoms as C. As luck would have it, the treatment for the condition you actually have is the same as the treatment for C. After the treatment is completed you no longer have any symptoms. You conclude that the doctor was correct.

In this case, the doctor was wrong about his diagnosis, but nevertheless correct in his prescribed treatment. Neither the doctor nor you are aware of the error because there are no signs that an error has been made. In this case, and the last, the evidence that an error has been made is misleading enough to lead you to false conclusions about the doctor's performance. That is, in

the first case you have evidence that the doctor made an error because your symptoms did not respond to treatment. It was only after the reevaluation that the testing error was discovered. If not for that discovery, the evidence would have led you to believe that the doctor did make an error. In the second case, there is no evidence that the doctor made an error, even though he did, and so whatever beliefs you form about the doctor's performance will be based on misleading evidence. If these beliefs are utilized in future decisions about which doctor to see the next time you are ill, you may be making an error. In the political case, this seems to be an obvious problem that would have to be overcome. In other words, low-information voters who might be able to make competent decisions via minimalism will need a means to discriminate these kinds of cases.

The conclusion we should draw from these examples is that in order for low-information voters to satisfy second-order competence they must be situated such that they could effectively determine when experts have in fact made errors and cases where there is only the appearance of an error. What's more, they must be capable of effectively determining when an expert has been correct rather than merely appearing to have been correct.

One might object that these voters could discriminate between appearance and reality by looking to the testimony of other experts. It seems likely that in many cases disagreeing experts would happily point out that others were wrong on some issue, especially in cases where the secondary expert appears to have been correct. There are two problems with this objection. First, even if we can safely say that the secondary expert is right—that the first expert has made an error—this does not tell us whether the voter made an error in accepting the testimony of the first expert. It could be the case that the voter made the best decision in selecting the first expert given the evidence, but that expert's advice turned out to be false nevertheless. In that case the voter's

first-order performance would have been a success despite its leading to a false belief. This would be like a skilled archer taking an easy shot (for them) and missing because of a defect in the arrow's construction. We would not count that against the archer even though she missed an easy shot. Similarly, then, voters need some means to determine if they made an error in their first-order performance or if the best expert just happened to be wrong this time. Second, we cannot just assume that the secondary expert is correct. Indeed, this whole process begins anew, with a first-order performance regarding the testimony of the secondary expert. If the voter's first-order competence is not present, they will be set up to make another mistake. If their first-order competence is present, it cannot be demonstrated that it is by another first-order performance. This would be bootstrapping on the presence of first-order competence. Hence, voters on the low bar face challenges that cannot be overcome by repeating first-order performances in order to bootstrap justifications about past performances.

It would seem that in order to make these sorts of judgments voters must have some knowledge beyond that which is required for them to determine who the experts are and who is most likely to be right when they disagree. Brennan seems to agree when he argues against the party identification shortcut. He argues that even if the party identification shortcut lessens the epistemic demands on voters, they will still need to know whether their party's policies will result in good outcomes. To do this, voters will need to have some idea of how likely the party's chosen means will achieve their stated ends (Brennan 2016, 196). Similar to the case I have made here against minimalism, these sorts of shortcuts only lessen the demands on first-order judgments, e.g. which party to vote for or which expert advice to accept. They do not offer anything to ameliorate situational demands that second-order competence requires. Hence,

attempts to lessen the epistemic load on voters' first-order performances do not increase the likelihood that complete competence will be satisfied.

IV. Vulcanocracy

So far I have attempted to show that epistocracy as presented by Brennan faces serious difficulties regarding competence. Like democracy, epistocracy likely would result in a first-order incompetent electorate, primarily because of the epistemic character of the voters who would likely dominate epistocracy. Following Brennan, we should expect that a first-order incompetent electorate will likely violate our right to competent government. However, a committed epistocrat could respond to these problems by doubling down on epistocracy. That is, one could argue for further restricting suffrage to only those who could, in some way or other, demonstrate that they are knowledgeable and that their belief forming processes are oriented toward truth. One of the three political archetypes Brennan identifies might come closest to satisfying these criteria. In the first part of this section I present Brennan's characterization of vulcans. I then attempt to offer an account of a central epistemic problem faced by these voters and attempt to draw some conclusions about the likelihood that they will satisfy complete competence. I conclude that they likely will not satisfy complete competence, and offer an explanation for why this is the case.

Brennan's characterization of Vulcans is this:

Vulcans think scientifically and rationally about politics. Their opinions are strongly grounded in social science and philosophy. They are self-aware, and only as confident as the evidence allows. Vulcans can explain contrary points of view in a way that people holding those views would find satisfactory. They are interested in politics, but at the same time, dispassionate, in part because they actively try to avoid being biased and irrational. They do not think everyone who disagrees with them is stupid, evil, or selfish. (Brennan 2016, 5)

Based on this characterization it would seem that so-called vulcans are likely to avoid the problems Brennan identifies for voters in modern democracy and the problems I have raised here for hooligans in epistocracy. This is because vulcans are not ignorant of basic political facts and they make a genuine effort to form beliefs based on the evidence they have available to them.

Even though vulcans seem to be in a better epistemic position to make political decisions than hooligans and low-information voters, it still is not clear what exactly we should expect from them. On the one hand we could expect them to fall closer to minimalism. In that case it would seem as though they might avoid the bootstrapping objection leveled against the minimalist standard in the last section. On the other hand, they could have a greater responsibility for their decisions and fall somewhere closer to the middle or higher end of the epistemic responsibility spectrum. Below I argue that vulcans must meet a higher standard of epistemic responsibility than minimalism.

If we assumed that vulcans were only responsible for satisfying the minimalist standard, they may be in a position to avoid the bootstrapping objection. Because low-information voters are totally reliant on experts to inform their beliefs, they do not have any other means to assess their first-order performances than to carry out further first-order performances. Vulcans are situated such that this does not seem to be the case for them. Since they have more background understanding of things like social science, history, and philosophy, they have more tools available to them when making judgments about their first-order performances on the minimalist standard. What this means is that when a vulcan accepts some expert testimony, they are not wholly reliant on either that expert or other experts to tell them when they have made a mistake. They would be able to draw on the perceived results of their decision and their background knowledge to make these assessments. Unlike the low-information voter, the vulcan is situated

such that their background knowledge can be used to assess expert testimony without having to repeat the same first-order performance—which was the selection of experts.

Assuming the minimalist standard, vulcans would be required to decide who is and who is not an expert on the relevant issue and decide which experts are most likely correct when there is disagreement. When checking a past decision to see if there has been an error, they could draw on their own background knowledge and other sources of evidence regarding the situation in order to make that judgment. This is so even if expert testimony is used as evidence to form their second-order judgments that their first-order performance was apt. Although they are not experts themselves, vulcans should be able to assess expert testimony in ways that the novice cannot. An example of this in another domain is the difference between a patient's reliance on the testimony of doctors and a nurse's reliance on the testimony of doctors. In the former, patients often have no basis to make any judgments about a doctor's testimony because they have little (if any) background knowledge regarding medicine. In the latter, nurses do have background knowledge that is relevant to making judgments about the testimony of doctors. This might, for instance, allow them to detect errors in doctor testimony that would go unnoticed by the novice. Hence, nurses are in a significantly different epistemic position than the novice when it comes to evaluating doctor testimony. Similarly, then, we should expect vulcans to be better situated both to make judgments about expert testimony, but also past first-order performances. This would seem to make vulcans much more likely to satisfy second-order competence.

If the above holds, then it appears vulcans can avoid the bootstrapping objection that was leveled against the low-information voter on the minimalist standard. However, since it is assumed that in order to avoid that objection vulcans would need to draw on their own background knowledge, they are not really on the minimalist standard at all. Since they utilize

their background knowledge they will be more epistemically responsible for their decisions than the low-information voter. Although they accept expert testimony as evidence, it does not necessarily determine their beliefs in the same way it would in the case of the novice. Hence, the epistemic independence from experts required to overcome the bootstrapping objection leaves vulcans with a higher degree of epistemic responsibility for their decisions.

Given that vulcans are not merely epistemically responsible for the selection of experts like in the case of minimalism, they will face somewhat different problems than those sketched already. They will not merely need to identify when they have made an error selecting experts, but also identify false background beliefs that have led them to make errors. Unlike hooligans, vulcans appear best situated to do this, though they do face some serious epistemic challenges. Below I explicate the central problem that vulcans—or any voter, for that matter—must overcome to satisfy complete competence and avoid the threat of the graveyard spiral.

Trial and error is a common and useful strategy for solving problems, both everyday and technical. Suppose I can't seem to find my car keys. I make a list of all the places I imagine they could be. Because I am pressed for time, I order the list by the places that I think are most likely to be where I left them. If I think they are most likely on my desk, I check there first. If they aren't there I continue down the list. The mark of success of one of these trials is obvious, I find my keys. Similarly, automotive mechanics sometimes use trial and error when they can't seem to identify the problem from merely considering the symptoms your car is exhibiting. If the check engine light shows that cylinder two is detecting a misfire, then the mechanic must determine which of several possible causes is responsible. Bad compression in the cylinder, faulty spark plugs or wires, or issues with fuel injection could be the culprit. Tests are run in order to determine the cause, which sometimes requires parts to be replaced, and the car tested,

before it is known that that part is the faulty part. When the correct solution is carried out, it will be obvious because the cylinder in question will begin firing properly. It is obvious when I find my keys and it is obvious to the mechanic when the misfire has been resolved.

Trial and error can also be used to test strategies. For instance, there are many strategies one can use to beat the dealer in blackjack. If I suppose that I am competent to use two such strategies, A and B, I can test these out by playing. If strategy A produces a loss over many trials and strategy B produces a profit after a similar number of trials, then I can conclude that B is a preferable strategy for playing blackjack. I can even go back and check these results again by conducting further trials. The ability to experiment, and the certainty that results from being able to count the dollars in my wallet at the end of the night, offer a reasonable expectation that I can decide which strategy is better than the other.

Another important point to note here is that in order for trial and error to be a useful problem solving strategy, there must be some means to confirm when a trial has been a success. If I was unable to count my money at the end of a long session of blackjack, I would not know whether that strategy proved successful. I require some evidence—beyond merely believing that I correctly applied the strategy—to know whether it has been a success. If we are blind to a success (or failure) we will be lost as to what we ought to do going forward. Worse, if we are misled into thinking that a success has been a failure or a failure a success, the purpose for using trial and error will be thwarted. This is what happens when complete competence is not satisfied. That is, when we fail to be situated such that there will be signs that our first-order competence is not present when it isn't, we will be in the same position as the user of trial and error who is unable to correctly identify successes and failures.

In the last chapter, I argued that without complete competence, decision makers will face the graveyard spiral. Recall the graveyard spiral is the situation pilots who are rated to fly by Visual Flight Rules (VFR) can find themselves in when they enter an area of poor visibility. Since VFR-rated pilots utilize visual cues like the visible horizon to orient the plane, clear visibility is essential to successfully controlling the plane. When a VFR-rated pilot enters Instrument Meteorological Conditions (IMC)—i.e. conditions of poor visibility—and these visual cues become obscured, the pilot loses the ability to effectively orient the plane. A dangerous result of this situation is the tendency of the pilot to unknowingly enter a turn. If the pilot realizes that the plane is in a turn, there are only a few options that the pilot can try to retake control of the plane's trajectory. I use the example of a VFR-rated pilot who finds themselves in trouble as they attempt to conduct a timed turn to return to Visual Meteorological Conditions (VMC)—i.e. an area of clear visibility—to show how trial and error is not a useful strategy under these circumstances.

Pilots can get themselves into trouble when conducting a timed turn to get back to VMC. A timed turn is a practiced strategy to get out of IMC, which involves turning the plane a particular number of degrees for a particular amount of time to put the plane onto a heading that will lead the plane out of IMC. When the plane begins to turn, the fluid in the pilot's inner ear begins to shift, which causes the sensation of turning in the pilot. After some time in the turn, the fluid in the inner ear will cease shifting. When this happens, it will feel to the pilot as though the plane is no longer in the turn, even though the instruments the pilot has at his disposal will show that the plane is still turning. Because these sensations are very powerful, they put the pilot into a position of having to determine which stream of evidence to trust. If he makes an error, the result can be a graveyard spiral, which will cause the plane to crash rather quickly. So there

are two strategies here the pilot can go with. He can decide his instruments are likely giving a false reading or assume his sensations are what should not be trusted. In the former, he will be led to attempt to reenter the turn, which, because he is already in a turn, will cause the turn to tighten possibly putting the plane into a spiral. In the latter, he will continue on with the strategy already in place, i.e. continue in the current turn until the time has elapsed and begin flying straight again until the plane exits IMC. In either case, the pilot will not be aware of whether he was correct in his decision until the plane either exits IMC or impacts the terrain. In the latter case, the error will not be obvious to the pilot until it is far too late.

The situation in politics is a little more complicated. When we make political decisions, especially at the federal or central government level, the high degree of complexity will leave us with different means for determining the success or failure of policy. There will likely be different metrics that can be used to analyze the results of policy and vulcans will have to be able to navigate these different metrics sufficiently to recognize success and failure. The level of credence voters will have in their beliefs about the effects of policy will be conditioned on their credence in the metrics used as evidence in those judgments. Since we are assuming vulcans are not experts, we should expect that they will require evidence provided through expert testimony to make these sorts of judgments.

Further, we need an account of what it means to make an error here. For instance, in the case of testing blackjack strategies, we would need some reasonable belief that we properly employed the strategy before we could make a judgment about whether it was a good strategy. Moreover, merely implementing the strategy properly does not seem to tell us anything about whether the strategy was a good one. Indeed, there does not seem to be any way to infer from the successful implementation alone whether the strategy worked. Thus, we can be right about

our successful implementation of the strategy, but still lose money employing it. So there are really (at least) two ways to go wrong when employing the trial and error method. First, we can incorrectly employ the intended strategy, which will result in the failure of the trial to tell us anything about whether that strategy is preferable to another. Second, we can properly implement the strategy while still losing money. In that case, we can draw some conclusions about the strategy so long as we are confident that we implemented it properly. Using the framework of competence, we can put implementation of the strategy as an example of a first-order competence and the situational component that grounds second-order competence will be the ability to tell when we have gone wrong in our implementation of the strategy. Satisfaction of complete competence for trial and error problem solving, then, will require the proper implementation of a strategy and being situated such that when we fail at the level of implementation it will be apparent to us.

The above represents one horn of the problem of trial and error. There is a second competence necessary for successful application of the method. This second level is the selection of one strategy (or policy) from alternative options. Similar to the blackjack example, politics requires us to choose one strategy (in the form of a policy or candidate platform) to employ at a time. Generally speaking, when we vote on issues there are two or more options, only one of which can be chosen by the electorate. Unlike the blackjack example, where we can go back as often as we like to test the available strategies, there is no such possibility in politics. Voters can certainly repeal a law they find unsatisfactory, but there is not the (near) unlimited ability to experiment with different strategies as in games like blackjack. Because elections are regular events, and not continuous processes, we must make an *ex ante* judgment about the situation as it exists now and the policy or candidate options available to resolve problems with

that situation. Then, once a policy is enacted, we must make *post hoc* judgments about the effects of that policy and attempt to decide if we have made an error in making the choice that we did.

Take the blackjack example again. If I only have enough money to carry out one trial (which is a series of games) per year, and my goal is to find the best strategy as quickly as possible, then I will need to make a judgment about which strategy will most likely be profitable for me. There are several considerations that come into play here. I must attempt to determine which strategy is best generally speaking, but I must also consider which strategy I am most likely able to implement correctly. If the best strategy is too complicated for me or if I am not familiar enough with it to expect that I can implement it competently, I may need to implement a simpler strategy that I am more comfortable using. Since we are assuming that I am only able to carry out a single trial, in the *post hoc* assessment of my performance on these *ex ante* decisions will require me to determine three things. First, I will have to make a judgment about whether I implemented the strategy correctly. Second, I will need to judge whether the decision to implement the strategy I did implement, at the cost of implementing the others, was correct. For instance, if I barely produce a profit after correctly implementing the strategy, I may regret not taking a chance testing out one of the more complex strategies. And, third, I will need to judge whether the strategy produced a profit and how that profit compares to my expectations. For instance, if after a long trial I barely come out ahead, then I may judge that the strategy largely underperformed despite its correct application. This will allow me to make better choices regarding strategy selection the next time I head to the casino.

In politics, the first of the *ex post* judgments above seems to be fairly easy to assess. If a policy is enacted to lower taxes, then success on the first judgment above will depend on whether taxes, in fact, went down. The second two, on the other hand, are not as straight forward.

Consider an example from the political domain and the options available for making the second judgment above. Imagine we want to decrease the unemployment rate by manipulating the generosity of unemployment benefits. There are some economists who believe that there is a positive correlation between the generosity of unemployment benefits and the unemployment rate (Claveau and Mireles-Flores 2014). In other words, as the generosity of benefits goes up the unemployment rate does so too. Some economists explain this fact by citing the ‘job-search effect’. According to the job-search effect, increasing the generosity of unemployment benefits leads people currently on unemployment to feel less urgency in searching for a job. This is because as the generosity of benefits increases so does their threshold for genuine job searching. On the other hand, there are economists who believe that there is another effect, the ‘entitlement effect’, which says that as the generosity of unemployment benefits increases, the threshold for those outside the labor force—i.e. those not counted in unemployment statistics—for getting a job decreases. Thus they are more likely to begin searching for a job. The reasoning here is that the opportunity cost of not having a job, and thereby ruling out the possibility of getting those generous benefits, goes up. As these people enter the workforce the unemployment rate goes down, all things equal. Thus, if we want to take advantage of the job-search effect we should decrease the generosity of unemployment benefits and if we want to take advantage of the entitlement effect we should increase the generosity of unemployment benefits.

Since these are two distinct, mutually exclusive options for lowering the unemployment rate by manipulating the generosity of unemployment benefits, we can set about deciding which option we should choose. We’ll call manipulating the job-search effect to our benefit option A and manipulating the entitlement effect to our benefit option B. Now we just need to collect some evidence on which option is most likely to give us the results that we want, namely a

decrease in the unemployment rate. Let's suppose that the evidence shows that the job-search effect dominates the entitlement effect—meaning it has more causal influence over the unemployment rate—and so we decide that A is better than B and choose A. Now that we have selected A, we sit back and wait for the results.

Let's suppose that unfortunately A does not succeed in bringing down the unemployment rate. Instead, the rate actually increased. There seems to be a couple options for determining whether we have made an error in deciding that A is better than B. The range of possible cases seems to be as follows:

Case 1: A confounding factor F is identified and determined to have caused A to fail.

Case 1a: The confounding factor F was identified in the *ex post* analysis but was not known prior to choosing A as a possible confounding factor, i.e. we did not consider the possibility of F preventing A from succeeding in this case because the evidence didn't support that conclusion.

Case 1b: The confounding factor F was known to us prior to choosing A, but the evidence led us to believe that it was not likely to have the causal influence over the success of A that it in fact did have or that it would have had a similar effect on B had we chosen it instead.

Case 1c: The confounding factor F was known to us prior to choosing A, but the evidence led us to believe that it was not likely to have the causal influence that it did have (on A or B), but in light of the results regarding A we determined that F would not have had the same causal influence over B, i.e. there is no reason to think that B would have been affected at all by F.

Case 2: No confounding factor F is identified to have caused A to fail. It just failed and there isn't a clear answer as to why it failed.

In each of the possible cases above, we followed the evidence and rationally formed our *ex ante* judgment that A is better than B. However, policy A failed to do what we expected it to do, namely decrease the unemployment rate. The question that we need to answer is: Where did we go wrong and how do we make adjustments for future cases?

Regarding Case 1a, we can reevaluate our evidence in light of the results about how policy A went. We can even attempt to control for F in the evidence that we used in our initial judgment to better account for F in future cases. But it does not appear to be the case that we made any mistakes here. The evidence was insufficient to predict the outcome, and so all we can do is take this into account when making decisions in the future.

Regarding Case 1b, we can reevaluate our evidence here too with the goal of better controlling for the causal impact of F. Similarly to Case 1a, then, there does not seem to be any error on our part. We followed the appropriate methods for making the judgment but the policy failed. Furthermore, since in this case we determined that if F was to have any causal impact on policy A, it would have had a similar causal impact on policy B, the *ex ante* judgment that A is better than B appears to remain the correct judgment even in light of the new evidence. That is, the failure of policy A does not give us any reason to believe that policy B would have done any better.

Case 1c is somewhat different than the last two cases we have considered. In this case, we determined that F would not have been a confounding factor had we chosen policy B. Since this is the case, it seems as though we can attempt to call our *ex ante* judgment into question; namely, by asserting that we should have chosen B instead. But, like the other two cases above, this would not tell us that we made a mistake. It would merely tell us that the evidence we had when making the *ex ante* judgment that A is better than B was insufficient to accurately predict the outcome. To conclude that because A failed we should have chosen B would be a mistake. This is because it is possible that B would have turned out worse. So here again it appears that we have made no mistakes when making the initial judgment. Moreover, since we stipulated that the initial judgment was based on the evidence and that we followed all relevant norms of

epistemic rationality and assuming—as seems to be the case in any political decision—that some citizens’ *ex ante* judgment would have favored policy B over policy A, we can still claim that those citizens were wrong in their initial judgment. They didn’t follow the evidence and the relevant norms of epistemic rationality that lead the electorate to the judgment that A is better than B.

Finally, regarding Case 2, there does not seem to be any way to judge what happened. All we can do is adjust our evidence in light of the results and try to incorporate that information into future decisions, i.e. we can engage in *post hoc* recalibrations of our evidence to account for the failure of A. But, again, this tells us nothing about whether we made any mistakes.

It seems that when we choose one option over another in single trials, thus excluding all other options from a test, we eliminate the possibility of evidence that we made an error in selecting the option we did select. This seems to be the case in the blackjack example, too. If, because of regret, I form the belief that I should have attempted a more complex, and therefore more risky, strategy, that belief will not be based on evidence that things would have gone better if I did. If vulcans are committed to rationality and evidence, as Brennan describes, then they may have to face this situation with a sort of skepticism. That is, they will have to accept that they cannot competently judge whether they have made an *ex ante* error without some evidence that can tell them how things would have gone otherwise.

An objection that could be raised here is that experts can use models, at least in some cases, to estimate what would have happened in counterfactuals. If these models provide justification for making *ex ante* judgments about which policies are more likely to succeed, it seems that we could use them to judge what would have happened otherwise. I will not deny

that models provide some justification, but if that justification were sufficient to judge what would have happened in counterfactuals, we would not need to do *post hoc* analysis for the policies we did choose. We would just look at what the model predicted and form our beliefs based on the model's predictions about how things would go instead of inquiring about how things, in fact, turned out. Since we do not do this, we should assume that, whatever justification models do provide, it is insufficient to tell us what happened regarding policies we did enact. If this is true, then we cannot accept that models would provide sufficient justification to secure beliefs about what would have happened in counterfactuals.

Indeed, if we look at the issue through the lens of competence, we can see that the motivation for looking to models for help in making these judgments is that without them we are not situated so as to satisfy second-order competence. The problem as I have described it so far is a problem for second-order competence, i.e. second-order competence requires that we are situated such that we can judge when our first-order competence is (or was) present. When we consult models, we are attempting to situate ourselves such that we can justifiedly claim that our past first-order performances (in this case, the selection of one policy or candidate over another) were apt. However, because of the nature of counterfactuals, and the models that purport to inform us about them, the models do not *themselves* satisfy second-order competence. In other words, there is no way to be situated such that we could say when a counterfactual claim is false or misleading. All we can do is hope to construct the model properly such that all our assumptions are correct. But even assuming our assumptions are correct, we would have to assume that the correct application of a method (or strategy) is enough to judge that the outcomes of the use of that method (or strategy) were a success. Further, the judgment that our assumptions are correct needs some means of checking otherwise we fail to be situated to satisfy

second-order competence on that judgment, too. Thus the problem reasserts itself. We are attempting to justify the belief that we satisfy second-order competence in a way that does not satisfy second-order competence. If we cannot satisfy second-order competence, then we will be blind when our first-order competences are not present (if they ever were) with regard to the results of models. Thus, two issues arise with the use of models. First there is no means to satisfy second-order competence with regard to our judgments about the assumptions models require. Second, there is no way to be situated so as to judge when the models have gone wrong.⁶ Hence, we are not in a position to claim that a model justifies our beliefs about counterfactuals—and therefore our past first-order performances—because the models suffer from the same problem they were summoned to defeat.

Curiously, if we assume the rational requirement that our beliefs must stick close to the evidence, we actually undermine our ability to make the sorts of judgments needed to assess our past decisions. If we are not situated such that our second-order competence can be satisfied because we have no access to evidence that would tell us we made an error, then complete competence cannot be satisfied here. In the blackjack example, the mitigating strategy for this problem would be more experiments with different blackjack strategies. Instead of repeating the same trial next time, I could try a more complex strategy and see how things turned out. If after a sufficiently long trial I decided that the strategy is too complex, I can return to my previous, less-complex strategy. If it was not too complex, and it provided more winnings than the less-complex strategy, I will ditch the old strategy and play the new one, and so on. The structure of political decision making, at least as it is in modern democracies, prevent us from conducting

⁶ Indeed, it is unclear if these kinds of claims even have truth values.

this sort of experimentation because elections require voters to select one option at the cost of the others.

Another objection that might be leveled here is that voters need not be motivated by discovering the *best* strategy; they only need to find solutions that work to resolve the problem. In the original blackjack example, I stipulated two things that might alter our judgments about how vulcans might address this problem. The first stipulation was that the goal of the experimentation was to find the best strategy. The second stipulation was that I was only able to conduct a trial once a year. Both were intended to mirror the situation in which political decision makers find themselves. That is, the first was meant to represent the electorate's intention of finding the best strategy for resolving the problems they want resolved and the second was meant to represent that they get to make (or revise) these decisions only so often. If my motivation for playing blackjack was just to not lose money on my yearly visit to the casino, then my decision to experiment might change. For instance, if my original, less-complex strategy offered me winnings that I depend on, which are sufficient for my purposes, I might be inclined to stick with that less-complex option. This might be because I cannot afford to lose money on any trip to the casino. In this situation, I would not be able to make any judgments about whether I was making an error by not selecting another blackjack strategy. All I can do is attempt to refine my playing skills with the less-complex strategy and hope that this increases my winnings over time. In the same way that utilizing a less-complex blackjack strategy might suit me better because of my skill level, vulcans could stick to their guns, so to speak, so long as things appear to be getting better. In other words, vulcans could attempt to refine their beliefs about policy and candidate choice over time with the only purpose of improving *outcomes* instead of selecting the best policies to solve some problem. So instead of favoring certain types of policy solutions, they

would only look to outcomes to judge whether they had made errors. This would get them out of the problem identified above, where the choice of one policy or candidate at the exclusion of the others prevents them from gathering evidence about the options they did not choose. This way the *ex ante* decision would be deemed correct or incorrect based on the results *simpliciter*, rather than whether another option would have been better.

This could offer voters a way to satisfy second-order competence. That is, we could stipulate that first-order performances are only to be judged by their real-world effects. If voters are situated such that they could competently judge when a decision has made the situation better (or worse), then they could make judgments about the aptness of their first-order performances. However, in cases where voters identify an error in this way, there will be no evidence to support a change in strategy. This is because, as already noted, there is not a rational—or evidence based—basis for changing strategies without competently formed beliefs about counterfactuals. If I began losing money playing with my current blackjack strategy, and the only evidence I have supports the current strategy over other available options, then it is not clear how I ought to proceed. I could begin a series of experiments to test out new strategies and attempt to determine which one is the best fit for my skill level. Although, if I do not have the financial means to carry out this type of experimentation, I will have strong incentives to either quit playing altogether or double down on refining my current strategy in order to turn my losing streak around.

An example from politics would be a case where an anti-poverty policy was enacted. If that program did not produce results, we could continue to refine that policy until it produced the results we wanted by reexamining the evidence and our assumptions about the evidence. Therefore, voters have the option of continuing to refine their reasoning processes with regard to the evidence and their strategy or they could abandon the evidence that they have in order to

pursue another option. In the former, this could entail reexamining background beliefs to try to root out beliefs suspected of being false, both generally and regarding the anti-poverty policy. But, it isn't clear how many failures would be required to prompt this sort of reflection or how it should proceed.⁷ In the latter, voters would have to disregard evidence in (at least) some cases. It does not seem clear how a decision made in this way could be rational. Further, the move to disregard evidence seems to amount to a change in strategies. That is, the strategy for making decisions would have to shift from strict adherence to the evidence to one where evidence is given less weight. This would push the decision making process toward randomness rather than rationality. This seems unacceptable. The only tenable option, then, is to attempt to refine the original strategy or cease attempting these sorts of decisions altogether. For now I will assume the only tenable option is refinement.

If we settle on refinement as our strategy, then we need to be able to judge what effects resulted from our decision. Note this is the third judgment from before regarding whether our chosen strategy produced a profit. This means we need a causal story that connects the decision of a candidate or policy to the observed results. Further, there will need to be a means to elucidate progress over time. This will likely require some metric(s) that will identify trends over time, i.e. whether the problem is getting better or worse. In the blackjack case this is determined by the size of my stack of chips at the end of the night. As I refine my use of my chosen strategy I can look to how much more money I am making on average as I progress. Presumably, I will reach a point where my winnings top out. Once this point has been reached, I will have maximized the value of the strategy given my skill level. If we are making decisions

⁷ There seems to be a species of underdetermination that might pose a problem here. When we follow the evidence and make a rational decision that fails, it will likely not be clear where our error lies. It could be in the background beliefs that we use to interpret data or in the data itself. The complexity of the decision will probably determine how difficult it will be to uncover the error.

that have effects that range over all of society, then we will need to use metrics capable of offering reliable evidence that the situation is getting better (or worse). For example, if we want to increase economic activity we might look to proxies like GDP growth over time to see if our chosen economic policies are beneficial. Hence, these metrics might offer a means to make judgments about how a policy or candidate decision turned out.

Although this seems a more promising method of judging our past decisions, it is not clear that this really gets us out of the problems with competence identified above. Measurements are not as precise in social science as they seem to be in other domains. They attempt to measure the things we want measured, but they have limitations that are of epistemic interest here. Two such problems with using proxies to accurately measure the performance of policy decisions is that they are subject to *operational distortion* and *surrogation*. Operational distortion occurs when agents alter their behavior in order to increase the performance of a proxy rather than true performance of a policy or program (Bloomfield 2015). For instance, if a school is being rated by its students' performance on standardized tests, teachers could alter their teaching methods so as to maximize student test scores in a way that does not improve their actual learning. In these kinds of situations, the goal becomes manipulating students' test scores, rather than increasing actual learning or academic performance, in order to ensure the perceived success of a school or educational program.

Surrogation is where agents begin to believe that performance indicators are all that really matters (Choi, Hecht, and Tayler 2012, 1136). To take the standardized test example again, surrogation would obtain if teachers actually began to believe that increasing average test scores was more important than students' true academic performance. In the case of major policy decisions, both surrogation and operational distortion act to corrupt how we judge the

success or failure of policies because the unmeasurable effects of policy decisions are often ignored. Hence, the actual effects of our decisions can become obscured, which will, in turn, undermine our ability to competently judge our past performances.

In addition to the above issues, proxies or performance indicators are subject to Campbell's law: "The more any qualitative social indicator is used for social decision-making, the more subject it will be to corruption pressures and the more apt it will be to distort and corrupt the social processes it is intended to monitor" (Campbell 1979, 85). Some examples of Campbell's law are the highly publicized manipulation of VA hospital waitlists to show decreased wait times,⁸ the use of body counts in Vietnam instead of other indicators of military success,⁹ and the intentional overproduction or underproduction of goods in the Soviet Union to meet quantitative production goals¹⁰ (Campbell 1979, 86). All of these considerations should lead us to question the level of epistemic justification we can derive from proxies for judging the success or failure of a policy.

Another issue that confronts us here is that the policy effects on proxies run together and influence one another. This complicates the possibility of tracking how individual decisions influence our proxies. In addition to the difficulty in untangling these effects, what makes these judgments correct is that we employed the relevant methods properly and analyzed all available evidence. Thus, it appears the problems with second-order competence arise once again. The issue with metrics looks a lot like the problems noted above for modeling counterfactuals. We

⁸ See, for example, Donovan Slack, "VA bosses in 7 states," 2016.

⁹ Campbell notes that the use of body counts as a quantitative measure of military success in Vietnam likely contributed to the Mai Lai massacre.

¹⁰ For example, Campbell recounts how if the indicator for a production goal was total weight of goods produced, factories would produce the largest items they could (e.g. producing only the largest nails in a nail factory). If the indicator was total number of items produced they would produce only the easiest items (e.g. the smallest nails in the nail factory) they could without retooling the production line.

will fail to have second-order competence in our judgments about the accuracy and significance of the measurements. We would need to be situated such that we could tell that we were going wrong either in applying the method of measurement or in our belief that the measurement represented what it purports to represent. This is because we could be going wrong in either without any tell-tale signs this is the case. Note, this inability to satisfy second-order competence here will undermine our judgment that we are satisfying second-order competence on judgments that are downstream, so to speak, of this kind of evidence—e.g. if we are wrong about the effect of our decisions regarding monetary policy because our measures of inflation deliver misleading evidence, there would not be the tell-tale signs that these policies were failing. One option here would be to conduct multiple measurements using different methods or assumptions. This would offer a collection of plausible measurements on any given situation. But it isn't clear how this would help. It is an attempt to satisfy second-order competence by carrying out further first-order performances of the same variety, i.e. to check if *this* measurement of GDP is accurate by attempting further measures with the exact same epistemic status just pushes the problem out a step.

This was the situation for low-information voters on minimalism. Recall they were unable to satisfy second-order competence with regard to the selection of experts because they had to select more experts to use as evidence regarding the success of previous performances of expert selection. If this sort of bootstrapping is illegitimate there, it should be just as illegitimate here. Hence, it seems that—in addition to the problems of operational distortion, surrogation, Campbell's law, untangling the causal chains that flow from electoral decisions to changes in data, and competence regarding our judgments about the reliability of our measurements—there is also the problem that bootstrapping is required to justify the use of measurements as evidence

in our judgments about our electoral decisions. Hence, the attempt to lessen the burden on decision makers by looking to the simpler refinement strategy does not avoid the problems with second-order competence.

The conclusion that we should draw from the above considerations is threefold. First, restricting suffrage will not do enough to generate complete competence. Selecting vulcans may increase the first-order competence of the electorate and we may therefore get more good decisions. However, it will not be enough to satisfy complete competence. Second, we should expect that a democracy of low-information voters, an epistocracy of hooligans, and an epistocracy of vulcans will all face the graveyard spiral problem. This would seemingly entail that our right to competent government is being violated in the same way our rights would be violated by a VFR-rated pilot intentionally flying us into IMC. If complete competence cannot be satisfied under these conditions, then no decisions made under them can avoid this result. Third, Brennan's instrumental argument for epistocracy appears to no longer be tenable. This is because in order for epistocracy to be preferable to democracy it must satisfy competence and outperform democracy.¹¹ But, as I have argued, looking at the performance of one strategy without evidence regarding how alternative strategies would have performed under the same conditions requires bootstrapping. Hence, whatever beliefs we form about the relative performance of epistocracy and democracy will have to be based on counterfactual imaginings that do not satisfy second-order competence. Indeed, in order to even formulate a judgment we would need to utilize the same metrics and models that do not satisfy complete competence.

¹¹ Note that Brennan (2016) does not argue for an epistocracy of vulcans, he argues for an epistocracy of hooligans. If hooligans do not satisfy first-order competence, as I have argued, then they do not seem to be in any better position than low-information voters in modern democracy with regard to competence (207).

V. Conclusion

What I have attempted to show here is that Brennan's argument for epistocracy faces serious problems that do not admit of simple or obvious solutions. I argued that voters need some social scientific knowledge to avoid the bootstrapping objection regarding expert selection. I also argued that hooligans, as Brennan describes them, give us no reason to suspect that they will be capable of satisfying first-order competence because of the methods they use to secure their beliefs. I argued that an electorate of vulcans, with their greater social scientific knowledge, will not satisfy complete competence because they fail to satisfy second-order competence. Finally, given these considerations we can conclude that Brennan's instrumental argument for epistocracy fails. Indeed, if the competence principle requires that decision makers are able to identify and correct errors, then decision makers must satisfy complete competence. If decision makers cannot satisfy complete competence, then we face the threat of the graveyard spiral.

I want to close with a few remarks about the graveyard spiral problem and what is next. The graveyard spiral problem arises when pilots cannot tell that they are making fatal errors until it is too late. When we centralize problem solving, increasing the complexity of the problems under consideration, we undermine our ability to recognize error. If the tell-tale signs that we have made errors are not present, or not recognizable as such, we should take very seriously the likelihood that the decisions we make today could have disastrous effects in the future. Moreover, the tendency of modern politics to focus on short-term effects, often driven by politicians' electoral ambitions, will lead us to increasingly ignore the long-term effects at our peril. When a VFR-rated pilot in IMC attempts to fly by feel alone, he may, after pulling back on the yoke, mistake the sensation of a tightening turn for the plane climbing. When this happens, he may be comforted by the sensations that result, although he has just put himself into

a deadly scenario. The ‘long-term’ consequence here is catastrophe, while his chosen evidence leaves him thinking the problem has been resolved in the short term. If we cannot avoid this sort of problem with modern democratic structures for problem solving, then we should look for methods that would allow solutions to be reduced in both breadth and complexity. This might allow for more experimentation and less reliance on macro-level metrics to guide us.

In the next chapter I consider some options along these lines. There I attempt to demonstrate that the three democratic values can all obtain if we pursue problem solving strategies that avoid the problems sketched here. This will require us to look to the complex systems sciences and attempts to resolve complex problems in other domains. I argue that through a process of encapsulation and decentralization, we can mitigate or eliminate the epistemic problems above and eliminate the graveyard spiral problem.

Chapter 4

Complexity, Encapsulation, and Decentralization

In the last chapter, I discussed the epistemic problems associated with political decision making in both democracy and epistocracy. This is the principle issue this chapter will attempt to resolve. In doing so, I present an alternative account of political decision making that offers a means to resolve the epistemic problems discussed in the last chapter, while also maintaining the other two values that comprise the trilemma of democratic values discussed in the first chapter. In the first section I recapitulate the trilemma of democratic values. In the second section, I discuss the issue of complexity as it arises in the study of complex systems and two central complex systems concepts. I argue that the observed behavior of complex systems undermine voters' ability improve their decision making over time. Indeed, these considerations seem to undermine the idea that voters can control the sorts of things they wish to control. In the third section, I look to other domains where complex systems are dealt with and argue that encapsulation and decentralization of the decision making process are a means to maximize the competence of political decision makers.

I. The Trilemma Revisited

In chapter one I argued that there are three fundamental democratic values that cannot be held consistently without good reason to think that errors could be identified and corrected: (1) the non-instrumental value of universal inclusion, (2) the collective right to self-determination, (3) the right to competent government. In the first chapter of this dissertation, I argued that democrats, in particular epistemic democrats, cannot hold (3) given the problems with the 'wisdom of the multitude' arguments meant to justify the competence of mostly ignorant crowds.

Epistocrats, who attempt to resolve that problem, cannot hold (1) and likely would need to amend (2) to only include the electorate—i.e. the ‘knowers’. In chapter three I argued that the epistemic problems faced by all forms of centralized decision making, in particular epistocracy, cannot maintain (3) in the way Brennan (2016) argues. The goal of this chapter is to establish that the problems with competence can be resolved in a way that preserves, or enhances, the other two values.

Recall in the first chapter I argued, following Anderson (2009), that universal inclusion has value because it is part of the democratic way of life, rather than its value merely being conditioned on the likelihood that voters will settle on the ‘correct’ answers. Epistocrats like Brennan see universal inclusion as negatively valuable on instrumental grounds, since voters are mostly ignorant of basic facts relevant to making political decisions. Thus, epistocrats reject non-instrumental arguments for universal inclusion because they think instrumental considerations dominate non-instrumental considerations. In the last chapter I identified the epistemic problems with attempting solutions to public problems through a process of centralized decision making as being the core of the epistemic problems faced by voters. Further, that these problems are not resolved by merely increasing the average voter’s political knowledge. If eliminating universal inclusion is not required in order to sufficiently increase the competence of the electorate, then there are no grounds to eliminate it. Hence, if we think universal inclusion has non-instrumental value, then we have *prima facie* reason to retain it.

In order to justify retaining it, however, we need a reason to believe that it will not fall to the epistemic problems discussed in the last chapter. Indeed, we need an account of democracy that can utilize the inclination to participate in politics such that universal inclusion does not become a liability. This will require a system of problem solving that relies less on specialized

knowledge, and which allows the sort of knowledge that average citizens hold to be utilized to produce, test, and evaluate problem solving strategies.

The collective right to self-determination was meant to denote the rights of the collective to have political authority extend over all issues that face the public so long as doing so does not infringe on anyone's basic rights. Epistocrats will likely need to amend (2) in order to account for the right of the restricted-suffrage electorate to have this authority. In order to maintain this right for the whole collective, the view I present below can guarantee that it is not restricted to only a subset of the population. Indeed, if universal inclusion can survive on the view presented here, then the collective right to self-determination should follow as a matter of course.

Finally, there is the right to competent government. This is the central problem this dissertation attempts to address. I argued in the last chapter that democracy and an epistocracy of hooligans fail to satisfy complete competence because they fail to satisfy first-order competence. An epistocracy of vulcans might be able to satisfy first-order competence, but they fail to satisfy second-order competence for two reasons. First, they do not have evidence regarding the choices they did not make, and therefore are incompetent to judge whether the option they did select was the better option. Second, because the effects of large-scale decisions require the use of metrics and models that do not themselves satisfy second-order competence, they cannot be used to supplement the absence of second-order competence regarding judgments about past first-order performances. Further, I argued that without second-order competence we have no reason to believe that mistakes will be identified and corrected, which raises the threat of the graveyard spiral. Since the arguments in Brennan (2016) cannot account for how an epistocracy of political hooligans can satisfy complete competence—and because the ignorance of the electorate undermines their competence in democracy—I argue below that the only means

to achieve complete competence will be through a decentralized approach to public problem solving.

II. Complexity and Complex Systems

Concepts like ‘complexity’ and ‘complex systems’ as we will be using them are somewhat difficult to nail down precisely. This is, at least in part, because there is not a unified science of complex systems (Hooker 2011a, 4). Instead, we have a set of behaviors that are observed in systems when they reach a sufficient level of complexity. It is these features of complex systems that will be discussed below. This will give us a basic idea of the kind of problems complexity poses to political decision making and offer us a means to address them.

Although there are not agreed upon necessary and sufficient conditions for ‘complex system’, we will assume they have some or all of the following features. First, the system, taken as a unit, has many subunits. These subunits are, at least some of the time, interdependent and their interactions are nonlinear at least some of the time (Ricklefs 2011, 534). Furthermore, the properties observed at the unit level are generated by the actions of subunits. Since the unit-level properties supervene on the behavior of the subunits, any changes at the unit level are the result of changes in the behavior or properties of the subunits; although, changes in the subunits need not result in changes at the unit level. In adaptive complex systems, the subunits react to changes in the environment and to changes at the unit level, allowing new properties to emerge at the unit level. So there is feedback from the environment and the unit system that feeds back into the subunits, sometimes altering their behavior. Lastly, subunits can change their own behavior such that new unit-level properties emerge without any preceding change at the unit

level or in the environment (535). This means that, in at least some complex systems, the subunits self-organize to create new unit-level properties.

There are many other features of complex systems discussed in the literature,¹ but I only focus on two such concepts here. These are nonlinearity and self-organization. The purpose for introducing these two properties only is that the former has clear epistemic implications and the latter forms the foundation for how we might go about getting out of these problems.

Nonlinearity, then, is the property of a system where a change to a system input variable does not necessitate a proportionate change in an output variable. In other words, if an input variable is changed by a factor f , then there need not be a change in the output variable by f . What is more, the relationship between variables that constitute the system may also be nonlinear. Thus, nonlinearity is not merely a property of unit systems, but also of the relationship between subunits. Additionally, nonlinearity underlies a system's sensitivity to initial conditions. Sensitivity to initial conditions is where small differences in the state of a system can have large effects on the trajectory of the system's evolution. This underlies the characteristic unpredictability of complex systems and poses several problems for political decision making.

Self-organization is the spontaneous emergence of order in complex systems.² As noted above, in some complex systems subunits can alter their own behavior or properties such that new properties or behaviors emerge at the unit level. Social phenomena often have this character, as in the way markets emerge endogenously, i.e. without external or top-down interventions. The economist Friedrich Hayek is famous for noting the emergence of markets as

¹ See Hooker (2011a, 20-40) for general discussion of the properties of complex systems and the difficulties in giving precise general definitions of them.

² There is not a unified definition of self-organization, but this seems to capture the idea as it is portrayed in the literature. See Hooker (2011b, 206-212) for discussion of different accounts of self-organization and its relation to other complex systems concepts like emergence and level formation.

the result of self-organized, voluntary behavior (Hayek 1976). Self-organization, then, generates either new order or complexity at the unit level. Below I argue that self-organization may form the basis for how to go about mitigating the problems posed by complexity, but first I discuss the epistemic problems posed by nonlinearity.

In the last chapter I focused on the *ex post* judgments about policy or candidate selection and how models cannot help us with regard to second-order competence there. Here it seems that nonlinearity poses a problem for our *ex ante* judgments about policy or candidate choice. This is because there are significant problems with prediction in nonlinear systems. As Bishop (2011) notes, nonlinear models cannot be confirmed or disconfirmed in the same way linear models are confirmed or disconfirmed. With linear models, there are two piecemeal strategies for confirmation or disconfirmation. First, we can hold fixed the input data and attempt to refine the model to bring it closer to the behavior of the target system. The idea here is that if the model is faithful to the target system, refining the model will produce a closer match between model behavior and target system behavior. Second, we can hold the model fixed and attempt to refine the input data to bring it closer to the behavior of the target system. Here the idea is that if the model is faithful, making the input data more accurate will bring model behavior closer to the target system's behavior. If the model does not converge with target system behavior using these two strategies, then the model is considered disconfirmed (118-19). However, with nonlinear models this is not possible because small changes in input data or to the model can cause non-proportional shifts in model behavior that may or may not converge with the target system (118-19). Since we expect this to happen, these piecemeal strategies do not function with nonlinear models as they do with linear models. That is, failing to converge following the piecemeal strategies does not tell us whether the model is actually faithful to the target system. Bishop

notes that this poses a problem for modeling the effects of public policy. By running models over new data sets as a policy's effects unfold opens up the possibility that model behavior will diverge from target system behavior and there is no guarantee that the new model forecasts will match the old forecasts.

One result of this situation is that after policy is enacted, we should expect that re-running the model over new data will produce new forecasts. If these alterations to original forecasts are expected to happen, then our confidence in model forecasts used for *ex ante* judgments about policy or candidate choice should be tempered. Moreover, if there are competing models used to make *ex ante* judgments about which policy or candidate platform to prefer, which then produce more than one option for how a policy or candidate platform will perform, then we face an underdetermination problem regarding evidence we already have reason to believe is unreliable. The issue here, again, is with the ability to make competent decisions under these conditions. If we know our evidence is insufficient to tell us what will happen, and there are problems in assessing these *ex ante* decisions for correctness in *ex post* analysis, then it does not seem that competence should be expected to obtain here. Further, since models are generally unreliable to make precise or often even remotely correct predictions, then we should attempt to make decisions in ways that do not rely so heavily on models where possible.

Indeed, Orsekes and Belitz (2001) argue that the use of models to make predictions for use in public policy can be misguided because these "predictions are always wrong, in the sense that we don't expect models to be precisely correct in all respects" (25). This is not to say that models cannot be useful tools for formulating public policy, but that their reliability is limited and this is often not reflected in the use of model predictions for public policy. Orsekes (2003) goes so far as to argue that models should not be used for long-term predictions because they

undermine public confidence in science when their predictions are inaccurate, which is often. Furthermore, the difficulty in predicting unlikely events, which are often highly significant to the public (like, for instance, economic crashes and war), should also temper our confidence in our beliefs about what will happen after a policy is enacted or a candidate is elected. The inability to account for these kinds of rare events often leads models to make overly optimistic predictions (Orsekes and Belitz 2001, 28-9). What is more, the methods used for model calibration, like using historical data, can lead to implicit, but sometimes false, assumptions about the causal framework of the system. That is, it will lead modelers to assume that the processes observed in the historical data will continue into the future (31).

Another means to make *ex ante* decisions is to look to the past effects of policy to generate generalizations to be used in future cases. For instance, generalizations like “reducing taxes increases tax revenues” or “minimum wages hurt those they are intended to help” or “high taxes slow economic growth” are ubiquitous in politics. It isn’t clear how we should take these generalizations when looking at them through the lens of competence. This is because, as noted above, complex systems are notoriously difficult to predict and the evolution of complex systems is often determined in large part by the conditions present when some change occurs, i.e. these systems exhibit sensitivity to initial conditions. Since this is the case, it is unclear how much weight we should put in these kinds of generalizations. Further, these kinds of generalizations face the closure problem discussed in the first chapter. That is, if S knows that P entails Q, and P, then S knows that Q. So, for example, if S knows that the tax cuts didn’t pay for themselves, and the tax cuts not paying for themselves entails tax revenues would have been higher otherwise, then S knows that tax revenues would have been higher otherwise. Moreover, because tax revenues are connected to economic growth, S would have to know something about what

economic growth would have been otherwise. If S does not have competent beliefs about these counterfactuals, then S does not know that the tax cuts did not pay for themselves.

If we believe that we have competent beliefs about causation in politics, then we need some reason to think that we have competent beliefs about counterfactuals. Without this, we will merely be generating abductive inferences that have competitors and no means to test them. Imagine you and a friend are discussing a recent boxing match where Jones, your favorite fighter, lost a fight you both expected him to win. You both agree that Jones looked tired after the fifth round, which is uncharacteristic of Jones, so you reason that fatigue is responsible for his loss. Further, you note that Jones had done an enormous number of media appearances during fight week, which you conclude was the cause of Jones' loss. Your friend, however, notes that Jones has been taking minor acting roles in movies lately, which she reasons interrupted his training and therefore caused him to lose the fight. Both explanations are plausible, and each may in fact be partially true. But in either case, you and your friend are committed to counterfactuals where Jones would have won the fight if not for the fatigue-inducing factors you identified. However, there is no reason to think that this is true. First, Jones' fatigue could have been caused by something besides the factors you and your friend identified. Second, Jones could have lost no matter what kind of shape he was in before the fight. Without some means to test these kinds of beliefs, we are stuck with interminable disputes about what happened. If the evidence we rely on for our causal generalizations in order to make *ex ante* decisions fares similarly, we will be stuck in the situation of relying on intuition to tell us which scenario is more likely. Indeed, it does not seem like actual learning through experience is possible when our intuitions determine which causal story we believe. This is especially so if past judgments of this sort shape our intuitions now.

In simpler systems, the ability to learn through trial and error is obviously possible. For instance, if we are attempting to figure out how to unclog a drain, we can quickly learn that removing the obstruction will cause the drain to function properly again. In complex systems, this is not always possible (or likely). In part this is because the system itself reacts to our attempts to control it. A change in policy or law can have the effect of controlling the thing we want controlled in the way we want it controlled, or the system can react to resist these attempts. Drug policy in the U.S. is an example of this, where anti-drug policy attempts to eradicate markets for the production, distribution, and sale of illicit drugs. However, illicit drugs are ubiquitous in American society despite these attempts at control. Another example is the relationship between U.S. tax rates and tax revenues (measured as a percentage of GDP). Economist Antony Davies shows that although the top marginal tax rate, the effective corporate tax rate, Medicare and Social Security tax rates, and average tax rates have changed over the years from 1954-2009, tax revenues hovered at around 18% of GDP during that same period (Davies 2013). Davies concludes from this that we need a simpler tax code, which will eliminate the tax avoidance strategies currently available to tax payers. In both of the examples here, it seems that the capacity to self-organize is at the bottom of these failed attempts at control. That is, the components of the system adapt to new conditions, which in turn thwart attempts at control. If self-organization can thwart political controls *and* generate the enormous, legitimate markets that we have for goods and services, then it seems that self-organization is a very important concept. In what follows, I focus on the ideas presented in Allen (2011), which use evolutionary theory to model how firms adapt to increase their fitness by increasing internal ‘micro-diversity’ through self-organization.

What I want to show here is that the problems faced by firms confronted with a changing environment are similar in form as those we have been discussing here and in the last chapter. That is, firms must adapt to changes in their environment in order to remain competitive. This process requires decision makers to attempt solutions to problems and develop new ways of doing things in order to increase the firm's fitness. This requires a process of trial and error, where possible solutions are generated, selected, and implemented. This process can happen from the top-down, where upper management attempts to find optimal solutions to problems. Alternatively, firms can attempt to utilize self-organization within the firm to generate heterogeneity in solutions, i.e. to experiment with possible solutions and either retain or reject them based on performance.

Allen argues that the strategy that utilizes self-organization at lower levels to solve problems outperforms the optimization strategy. This view is derived from Allen's work on 'evolutionary drive' in complex systems. This is the view that micro-diversity drives evolution through small changes in interactions between component parts of a system and the environment. Essentially, small differences at the micro-level lead to differences in how systems interact with their environment. Over time these small differences allow some systems to survive environmental changes while others die off. Further, this sort of co-evolution does not generate optimal solutions or strategies, since the environment is constantly changing. Hence, evolution here is seen as an on-going process that does not select for optimization but rather for "good enough" (Allen 2011, 787).

In biological evolution, micro-diversity is generated by genetic variation that leads to phenotypic variation. Phenotypic variation is what drives the selection process whereby fitter organisms reproduce and survive where others do not. This leads to the conclusion that the

mechanisms that generate this variation lead organisms to “explore the fitness landscape,” e.g. when a population of organisms is sufficiently diverse in phenotypic traits, changes in the environment suppress or amplify these traits allowing some to survive while others do not (Allen 2011, 788). Populations without sufficient phenotypic diversity will be unable to survive environmental changes over time. When applied to human systems, the evolutionary drive view sees the most important factors as those that generate micro-diversity in the processes that determine practices within a firm or organization. Similar to biological evolution, the environment (market) selects for ‘fitter’ firms while suppressing the less fit. However, the variation that allows firms to evolve is generated by individuals within the firm itself. Hence, “the guiding premise is that successful organizations require underlying mechanisms that continuously create internal micro-diversity of ideas, practices, schemata, and routines—...so that they may be discussed, possibly tried out and then either retained or rejected” (787). Additionally, Allen argues that because we are incapable of fully understanding how complex systems like firms operate in relation to their environment, there is no chance that the top-down strategy of optimization can result in long-term success. This is because the likelihood of finding the optimal strategy is low for any given problem at any given time and the changing environment requires firms to be dynamic in how they react to these changes. Hence, the capacity to adapt to a changing landscape is what increases a firm’s fitness in the long run and internal micro-diversity and heterogeneity in solutions to problems is how to achieve this end.

If self-organization is the key to the evolution of firms, then it stands to reason that a similar strategy would be useful in generating solutions to public problems. Indeed, it seems that the standard means for addressing public problems in modern democracies looks more like the optimization strategy than the learning strategy where heterogeneity in solutions allows firms to

settle on what works. That is, when the public looks to the central government to generate policy to address some national problem, they are attempting to select a single optimal solution to a complex problem. But, as I argued in the last chapter, attempting to solve problems in this way undermines our ability to recognize success and failure, which undermines our ability to realize when we have made an error. This raises the threat of the graveyard spiral. To avoid this situation, public problems should be solved in a way that utilizes self-organization in order to generate heterogeneity in solutions to problems. In order for self-organization to work in the context of firms and other organizations, top-level managers must give up some control in favor of allowing lower-level employees to attempt solutions. In the context of politics, it is not exactly clear yet how self-organization can be harnessed. In the next section, I argue that self-organization can be utilized only when problems are broken up into smaller, more manageable pieces where diversity in attempted solutions can be realized.

III. Encapsulation and Decentralization

So far we have seen that complex systems pose significant problems for political decision making. Nonlinearity undermines our attempts to control systems, our ability to project how decisions will fare, recognize error in the aftermath of decisions, and undermines our confidence that we can rationally make decisions to solve problems at the highest levels of complexity. The only way to avoid these problems is to reduce the complexity of the decisions voters, and other political decision makers, are expected to make. When we make big, complex problems smaller we make them easier to understand and solve. In this section I lay out the case that encapsulation and decentralization are the keys to reducing or eliminating the epistemic problems of politics and increasing the likelihood that decision makers will satisfy competence in decision making. In the first part of this section I briefly discuss the concept of encapsulation

as this term is used in computer science. Next, I discuss some of the ways modern militaries utilize decentralized decision making to increase tactical effectiveness on the battlefield. Finally, I take some of the lessons learned by military leaders when dealing with complexity and attempt to outline some ways decentralization might take form.

Encapsulation is a fairly simple concept. When computer scientists are confronted with a great deal of complexity in computing, one of the most basic strategies for reducing these problems is through encapsulation. Encapsulation is the process by which a large, complex system is broken up into smaller, more manageable pieces (Green and Leishman 2011, 139). An example of this is how programmers write large programs with many lines of code. Because no program is error free, programmers needed a means to manage problems that arise as a result of errors made by the programmer, but also from errors that arise due to unexpected interactions between bits of code. Encapsulation allows programmers to isolate units of code from other parts of the program, which decreases the amount of debugging required. Additionally, these encapsulated units of code can then be reused in other areas where similar functions are required (141).

Taken as a general problem solving strategy, encapsulation would allow for big problems to be broken up so that more than one solution could be tested. This achieves one goal, which is a reduction in complexity and opens up the possibility for diversity in problem solving. However, it does not tell us what to do about issues of scale. That is, we still need a decision making process that treats these units as discrete entities to be solved independently. It would not make much sense to break problems up in this way and then put it up for national vote to select the best strategy for them all. Instead, we want to make decisions about how to solve each unit on its own, using information and resources available to those who are attempting to

implement solutions. Hence, decisions about what to do should be made by those attempting solutions at the level of the encapsulated unit.

If we encapsulate public problems and we are not looking for a single, optimal solution, then we should decentralize the problem-solving process. This means we should decentralize the decision-making process. Thus, instead of looking to politicians with large-scale policy agendas, we should look to address public problems at the lowest scale possible.

Decentralization of the decision-making process has several virtues. First, it will allow more solutions to be tested to find what actually works. Successful solutions could then be adopted in other areas where similar problems arise. Second, it will reduce our reliance on macro-level metrics to track the effects of our decisions. This will increase our ability to track the actual effects of our decisions, which will in turn increase the likelihood that we recognize errors when we make them. Third, it could, depending on how a problem is broken up, lead to faster reaction times from decision makers since they will not be required to convince the national public, or a large political body, that a solution is preferable or not working. So decisions can be made faster, and rescinded faster when they do not work, than they would be at the level of national policy. Fourth, decentralization will allow knowledge held by the public about their own communities, and the problems within them, to be utilized in ways that it is not under centralized decision making. For instance, instead of taking all people living in poverty as a single unit, defined by a small number of variables, and attempting to solve their problems with a single, national policy, voters could with greater precision deal with the problems posed by poverty that may be unique to their community. This will increase the micro-diversity of solutions, which will allow the process of solving public problems to evolve as problems evolve over time.

Recall the two horns of the epistemic problem of trial and error. First was the ability to recognize when a strategy (or solution) has been implemented correctly and to be situated such that second-order competence can be satisfied. The second horn concerned our ability to competently judge whether a mistake had been made in selecting the option that was, in fact, selected at the cost of other options. There I argued that decision makers are unable to determine when they made an error in their *ex ante* judgments because they lacked evidence regarding how the option not taken would have fared. Under a system of decentralized decision making, the ability to experiment with a variety of solutions largely eliminates this problem. That is, so long as there is diversity in problem solving there will be an ever increasing body of evidence about the past performance of attempted solutions. Furthermore, because the effects of these decisions will not be as far reaching, it will be easier to track how these solutions fare.

Indeed, increased decentralization will give the public a greater ability to assess these results qualitatively rather than only quantitatively, using metrics that may or may not track what they purport to track. For example, if homelessness is the problem under consideration, it will be more readily observable to decision makers whether their attempted solution was having any effect. Under a system of large-scale decision making, these kinds of observations tell us little about how a policy fares since local conditions tell us little about how a policy or law is faring nationally. Hence, the ability to assess solution attempts in this way will increase the likelihood that we recognize when we have made an error. Decentralization, then, will situate us such that we can be confident that we can detect errors and correct them and, depending on the form decentralization takes, we can react quickly when solutions are failing in order to remedy them.

What is more, the graveyard spiral problem will be much more easily avoided. There are two reasons for this. First, decentralization alters our situation such that we can expect to be

competent on the second-order, which requires us to be situated such that there will be signs when our first-order competence is—or will be—absent. Second, because decentralization requires that solutions be tried out on encapsulated problems, when a solution fails to work it will not cause system-wide problems. In other words, when we decentralize decision making, we reduce the risk associated with any one decision. Hence, when political decisions do go bad, the negative effects will be contained rather than affecting the system as a whole.

I now want to turn to examples of decentralized decision making in modern militaries. Although militaries are usually thought of as very hierarchical, top-down organizations, modern militaries have learned a lot about complexity through experience on the battlefield and this has led them to look to decentralize certain decision making processes. Since war is chaotic and unpredictable, it stands to reason that the successes observed in military tactics could offer some insight into how political decision making could be decentralized. In what follows I briefly discuss the history of decentralized tactical command on the battlefield. This will show the similarities between our discussion and warfare theories that favor decentralized command. I then discuss three theories of warfare that utilize concepts that are relevant to complex systems and show that they avoid the problems associated with complexity by increasing micro-diversity at the small-unit level, increasing reaction times to problems, and making better use of latent resources.

Decentralization of tactical command has its roots in the concept of *Auftragstaktik*, which was developed by the Prussian army after military losses to Napoleon (Nelson 1987, 22). The problem the Prussian army dealt with was the increased effectiveness of cannon warfare, which forced units to spread out on the battlefield to avoid concentrated fire. Because of this, the small-unit leaders were unable to communicate effectively with higher-level commanders to pass

information and receive orders. Since the Prussian army at the time utilized a strict centralized decision-making system—where small units only took action after receiving orders from above—they could not respond quickly to changes on the battlefield. This led to many disasters, where small-unit leaders were forced to make decisions that they were not competent to make, often after waiting for orders until it was too late to act effectively. The solution was to train small-unit leaders to make decisions without orders from above (22).

This brings us to the distinction between strategy and tactics. Strategy is the overall aims of an operation and tactics denote the means to achieve those strategic goals (Kometer 2003, 43). Higher-level leaders, then, would give orders to achieve some strategic goal and lower-level commanders would plan and carry out the mission. With sufficient training to carry out these tasks, small-unit leaders can then use their knowledge of the situation as it unfolds to make decisions effectively. Decentralization of tactical command has found widespread success and is used by the U.S. military today. Indeed, there are three theories of warfare discussed by Ryan (2011) that deal with problems relating to complexity that may be useful here. These are maneuver warfare, network centric warfare, and irregular warfare.

Maneuver warfare is a theory that emphasizes agility and movement on the battlefield to generate uncertainties that can be exploited (Ryan 2011, 738-9). Further, it rejects centralized decision making and utilizes the ability of warfighters to change tactics quickly to open up new opportunities to exploit as the enemy reacts to changes on the battlefield. Similarly, decentralized decision making in politics could make use of this sort of real-time decision making such that opportunities can be exploited as they arise. This would be a major change from the way problems are solved through politics currently, where decisions are made slowly and often suffer from compromises made out of political expediency.

Network centric warfare focuses on the application of communications technology to increase the movement of information relevant to making decisions on the battlefield. This allows decision makers to better make use of latent resources, to better connect adjacent units, and to reduce bottlenecks in information flows. Ryan describes network centric warfare as “an integrating concept, concerned with how to coevolve a network of capabilities to deliver the best operational effect” (Ryan 2011, 738). In a decentralized system of political decision making, the use of modern communications technology could be used with greater efficacy, allowing latent resources in communities to be utilized. This could take the form of individual citizens offering their skills, knowledge, or other resources in a system of sharing not unlike the ‘sharing economy’ that has emerged over the last several years. This could take the form of individuals linking up through applications designed to connect individuals in need with those who can help. The public is an enormous reservoir of latent resources and knowledge, which could be put to use in solving problems under a system that rejects centralized decision making.

Finally, irregular warfare is a theory of asymmetric warfare, which aims to describe the ways smaller forces exploit advantages of scale against a much larger conventional force. The way this works is irregular forces attempt to increase the “fine-scale complexity” (or micro-diversity) by operating in small, autonomous units (Ryan 2011, 738). In other words, they use tactics that take advantage of their smaller and more maneuverable units to undermine the advantages of the larger force. Irregular warfare takes advantage of the multiscale law of requisite variety, which states that there is a tradeoff between the complexity of actions an organization can take and the scale of those actions (Bar-Yam 2004). Highly relevant to this discussion is that Bar-Yam argues that multiscale variety poses a problem for traditional attempts to solve public problems. He thinks that attempting to solve problems of high variety—in which

he includes poverty, Third World development, healthcare, and education in the US—through the “dominant ‘holistic’ but still centrally planned approach is insufficient” (43). Hence, increasing the variety at small scales through encapsulation and decentralization should increase our ability to address complex problems in ways that the traditional means for addressing these problems cannot.

In the next section, I discuss a few possibilities for how decentralized problem solving could take form by looking at some examples that are already being employed. Further, I consider some of the issues that might arise when theorizing about how decentralization should take form. What should be kept in mind here is that decentralization is geared toward generating evidence about what works in a way that can be tracked more effectively so that solutions can be continuously refined over time. Since the goal is competence in decision making, and because self-organization is the principle means for achieving the sort of variety that is required here, I assume that any public problem will need to be assessed individually to determine how it ought to be addressed. For instance, federal systems seem to already be set up in many ways for this sort of problem solving, but, depending on the problem under consideration, there may be differences in how a system that relies on decentralized decision making will be put into effect. In other words, it may be the case that some problems are better dealt with through state governments, but others may be better dealt with at the community level. Still further, there will be cases that cannot be decentralized.

One final thing to note before moving on is that the encapsulation and decentralization approach offers us a means to test democratic institutions. In the first chapter of this dissertation I made the claim that an epistemic defense of democracy, or any form of government for that matter, should offer us a means to assess the epistemic advantages and limitations of democratic

institutions. We can now look to institutions and ask the following questions. First, does the institution suffer from the epistemic problems discussed in the last chapter? Second, if so, can that institution utilize encapsulation and decentralization? If, after examining an institution, we determine that the answer to both questions is ‘yes’, then that institution should make adjustments toward encapsulation and decentralization—constantly refining this process until problems are sufficiently addressed—to bring it in line with complete competence. If the institution cannot utilize encapsulation and decentralization, then that institution should continue operating as it does not so long as there is not an alternative that can avoid the epistemic problems from the last chapter. An example of an institution that most likely cannot be reformed to utilize decentralized decision making would be the military. This is the case both because much of what the military does, including what it is capable of, is classified and because much of what militaries do concerns overseas operations. Another example is central banks like the Federal Reserve. This is the case because there is no means to decentralize it while maintaining its primary function, which is to conduct national monetary policy and oversee banking operations among other things.

IV. Semper Gumby

Although there is much we can learn from the way modern militaries deal with complexity, it is not yet clear how exactly encapsulation and decentralization will work in politics. One thing that should be mentioned here is that I am not arguing for a change in the form of government.

Epistocrats like Brennan (2016) believe that the problems surrounding voter competence in modern democracy should lead us to a change in the form of government, i.e. a change in who holds political power and how that power is distributed (167). Since self-organization already rejects top-down planning, there is no need to consider changes in how governments are

structured because self-organization occurs under all systems of government, e.g. even under totalitarian regimes black markets arise spontaneously. Here I am proposing a change in the scale and scope of government, i.e. if the arguments made in the last chapter and the present one hold water, then we have good reason to suspect that governments are attempting to solve problems in ways that cannot satisfy competence, and therefore violate the competence principle. ‘Scope’ here refers to the sorts of issues where the government has license to intervene and ‘scale’ will concern the level of government where public problems are addressed. As mentioned above, considerations of competence should lead us to question where exactly the state should intervene and how far reaching those interventions should be.

If we accept the premise that centralized decision making leads to problems with competence, then we have *prima facie* reason to look for a means to mitigate or eliminate those problems wherever we can. Since we have no empirical evidence of how a system of decentralized decision making will perform, we will need to be willing to explore many possibilities and test these possibilities to gain evidence about what works. It stands to reason, then, that however a move to decentralization might unfold will be a practical matter where theory will only be marginally useful. We can make general claims about our ability to track the effects of decisions and problems that arise when mistakes go unrecognized, but without actually testing out different methods for decentralized problem solving we will not really know how it works. Since this is the case, the same sort of reasoning will apply here as in the argument for decentralization, i.e. we need to run tests to generate evidence that we can use to make refinements over time. One of the assumptions of this chapter is that decentralization will produce such evidence, and so a similar recommendation should apply here too: different approaches should be tested and results compared.

The details of how decentralization should proceed will likely require consideration of details specific to the problem and the region where the problem arises. The purpose of this chapter is to lay out a basic strategy for establishing a more suitable method of solving problems given the epistemic problems of politics. Moreover, the goal is to secure the right to competent government while keeping universal inclusion, rather than a system-based approach that attempts to identify the proper forms of government *a priori* and then implement them. The goal, then, is to create a situation where the apparatus for solving public problems has the capacity for perpetual refinement and adjustment to meet the needs of the citizenry. “Semper Gumby” is a concept adopted by military personnel that means “always flexible” as a play on the mottos of the different military branches like *semper fidelis* and *semper fortis*. Semper Gumby, then, is used as a reminder to military personnel to expect changes to the mission at any moment and to be ready to react appropriately when confronted with those changes. Hence, semper Gumby may serve as a motto for decentralized public problem solving because we are attempting to generate the kind of adaptiveness that problem solving in complex systems seems to require. If the goal is effective problem solving, then flexibility and adaptiveness is going to be essential to actually making progress towards identifying solutions that work.

In what follows I primarily discuss some examples of non-state based decentralized problem solving, namely social entrepreneurship, but also some potential options for how the state might be involved in promoting self-organization to solve problems. The distinction between strategy and tactics can be useful here. Instead of eliminating the central government from the problem solving business, we could put it in the position of defining the ends while delegating the “tactical” decisions to those closest to the problems. There are several ways this could work. First, the central government could apportion resources and delegate the actual

problem solving to representatives at the community, local, or regional level, where decision makers would have broad license to make decisions about how to confront problems the public wants addressed.

Another way this could take form is through a process of guided self-organization, with the central government being responsible for incentivizing and supporting citizen-based problem solving. The goal here would be to generate a system where the public draws on its own resources and knowledge to solve problems. This could take the form of a sort of sharing economy, where citizens lend their skills, knowledge, or other resources to others in need. The central government could create incentives for this kind of activity, leading citizens to take a greater role in addressing their community's problems. What is more, the communications technology widely available to the public could be of great value to this sort of system, allowing those with immediate needs to be in contact with other citizens who can help in real time. Ultimately, if we are looking for ways to increase the competence and efficacy of public problem solving, we need to look at all available means.

Another possible form decentralized problem solving could take is a dedicated civil service. There are many ways that a dedicated civil service could be organized. It could work similarly to the way in which the National Guard or military reserve functions, i.e. an interconnected network of units staffed by both full time and part-time workers. However, instead of military training members would either be trained, or selected because they are already trained, in a needed skill, then members could be put to work in the service of the community. A combination of volunteer workers as well as paid staff could be utilized to draw on the skill and knowledge of the community to solve problems. The types of problems that a dedicated civil service could solve would be limited in some ways. They would not be able to change laws or

use coercion. But their goals could be set by the local community through a process of voting and town-hall-style meetings. This would require that the civil service be given broad license over local problem solving in the way discussed above, where the central government apportions resources and the local representatives select tactics.

If the process of agenda setting occurred through local voting procedures and town halls, the community served could be engaged directly with those in charge of the service. In addition to participating through the agenda setting process, the service could be open to members of the community to join as either volunteers or as paid members. The civil service, then, will allow for genuine feedback mechanisms to form through online communication or through in-person interaction. Thus, the civil service could allow the public to take part in the generation, implementation, and evaluation of the results of solutions.

The service could clean up polluted areas, it could provide training for out of work members of the community to help them find jobs, it could offer health care services, or it could organize construction crews to repair damaged houses, roads, and schools that had been neglected or underfunded before. In each case, the service would look to the community for guidance in determining how resources should be used and to judge whether the problem was solved sufficiently. Further, if the community had control of the service's agenda through a real-time voting process—e.g. each member of the public has five votes that can be used to rank options, vote on whether a project should be taken up, abandoned, or has been sufficiently completed—then they would be able to directly control how the service is used and how resources are apportioned.

Another avenue similar to the citizen-based problem solving strategy is market-based problem solving via social entrepreneurship. Social entrepreneurship can be roughly defined as the use of start-up business ventures to solve social, cultural, and environmental problems. Social entrepreneurship can take the form of a for-profit business or non-profit business, but in either case the primary goal is to produce some social benefit. Probably the most famous example of social entrepreneurship is the Grameen Bank founded by Muhammad Yunus in Bangladesh. The Grameen Bank has inspired other institutions to experiment with micro-lending all over the globe, including in the United States. It utilizes partnerships with major corporations and government to deliver billions of dollars in no-collateral loans to impoverished villages in Bangladesh. It has served over 78,000 villages in Bangladesh and 64 percent of borrowers who had been customers of the bank for five years or more no longer live below the poverty line (Yunus 2007, 51-2).

Although the Grameen Bank is an example of social entrepreneurship, it is more specifically an example of a social business. Social business, a subset of social entrepreneurship, is a business venture where all profits generated are put back into the company. Investors recoup their investments but do not gain any dividends or interest from their investment. In every other way, a social business is like a regular profit-maximizing business. They compete in the market and must cover their costs in the same way a traditional business does. However, their primary goal is to produce a social benefit. One way this is done is providing goods or services to a disadvantaged community. Another is to form a for-profit business that is owned by a disadvantaged community such that the profits are used for the benefit of the community. In both types of venture, the community plays a role in communicating to the business, either as customers or as shareholders, what the needs of the community are and how they can be best

served. Therefore, social business is another form of decentralized problem solving that is compatible with the encapsulation and decentralization approach.

Another possible tool for decentralization is the use of consumer behavior to direct profit-making corporations to engage in public problem solving. Since consumers have a great deal of collective power over the market, they could use their spending habits in order to incentivize large companies to aid social businesses and a dedicated civil service through joint projects, joint fundraising, or access to corporate distribution networks. Consumers could organize into consumer unions to strategically spend their money. One way this could work is for these unions to create lists of businesses that meet a particular standard for corporate behavior. Companies would have to get on these lists before union members would patronize them. Businesses would have to demonstrate their willingness to give back to the communities where they operate in order to get the business of union members. Thus, consumer unions could act as a way to influence the behavior of businesses so that more resources are used for socially beneficial purposes.

What this means is that the state will be primarily concerned with developing strategies geared toward cultivating an environment where people will voluntarily engage in socially beneficial behavior. This will likely require creating incentives for individuals to start, for example, a social business rather than a profit-maximizing business. These incentives could come in the form of tax relief, funding opportunities like grants, or even beneficial regulatory arrangements designed to encourage investment in socially beneficial ventures. The central government would also likely play a crucial role in the initial funding and implementation of a civil service. Since the resources necessary for creating a dedicated civil service able to serve every community would likely be quite large, the central government would provide the best

means to accumulating those resources. Although the process of funding and implementation could also utilize public-private partnerships to expand the number of options available for the initial creation of a civil service system. However, as noted before, the civil services should ultimately be governed locally by the communities served.

The central government may also play a role in incentivizing individuals to join or volunteer with the civil service. Since the civil service will likely require general labor, the state could offer various forms of tax incentives to those who volunteer a certain number of hours each year. Another means to encourage citizens to take part in the civil service would be to offer individuals the opportunity to learn a skill or trade that can be used in the private sector. For instance, the civil service could offer on-the-job apprenticeships to people looking to learn a skill while also helping disadvantaged communities. If a civil service provided basic construction services to the community—e.g. renovating an underfunded school—people looking to learn a useful skill like carpentry could work as apprentices to gain employment experience and learn a trade. This sort of program might appeal to high school graduates who are looking at other options besides college or who want to take a break before going off to college. If the funding is available, college tuition could even be offered to those who sign on to work with the civil service for a predetermined period, similar to how the GI Bill works for US military members.

Whatever form encapsulation and decentralization might take, it seems to be the best bet for increasing the likelihood that complete competence can be satisfied by decision makers. This seems to be true for three reasons. First, encapsulation and decentralization allow for more solutions to be attempted for problems with the same character. This will increase the evidence decision makers have when evaluating their own choices, which will allow them to contrast the effectiveness of their chosen solutions with solutions implemented elsewhere. This will allow

improvements to be made both in judging past decisions but also with regard to impending decisions. There will still be cases where decision makers are faced with the same epistemic problems they face under a system of centralized decision making, but those cases should become increasingly rare over time as more problems are addressed and evidence gathered.

Second, by reducing the size and scope of the problems being addressed, which will decrease the size and scope of solutions, decision makers will be able to better understand the problems and how the solutions are supposed to work. There will still be use for social science and expertise in problem solving, but because problems are smaller and more manageable decision makers will need to rely less on the testimony of experts to form their opinions on what they think will work. Moreover, since they will be making more decisions concerning their own community's problems, they will likely have more relevant knowledge concerning local or regional particularities that are relevant to solution generation. Further, if decentralization takes the form of a system where individuals actively participate in helping each other with their personal knowledge, skill, or resources the latent knowledge held by the public can be better put to use. Indeed, many people have skills and knowledge that can be useful to others but they often lack the platform to connect with those who need it.

Third, encapsulation and decentralization reduce the risk of any one decision failing and making the situation worse. Because of this, the threat of the graveyard spiral will be reduced or eliminated. Even when decision makers mistake errors for corrections and corrections for errors, the problems that result will be contained and will be less likely to affect the system as a whole.

V. Concluding Remarks

Epistemic democrats cannot secure the right to competent government because of the problems with the wisdom of the multitude arguments and because they fail to account for error correction. Epistocrats, in their attempt to resolve these issues, fare similarly. Without an account for how mistakes can be corrected there is no reason to suspect that they will be. This problem is further compounded by the epistemic character of so-called hooligans, who do not engage in good faith inquiry and who will likely dominate epistocracy. What's more, I argued that a closer examination of the epistemic problems of politics undermines our confidence that even an electorate of well educated, good-faith inquirers would be able to surmount these problems. This should lead us to suspect that we are unable to identify and correct mistakes when we make them under a regime of centralized decision making. Encapsulation and decentralization offer us a means to mitigate or eliminate these problems. Because encapsulation and decentralization seeks to put solution generation, selection, and implementation in the hands of those closest to the problems, universal inclusion as a valuable aspect of the democratic way of life is enhanced. Additionally, the collective right to self-determination is similarly enhanced because citizens will have greater control over the political destiny of their communities.

Finally, if anything positive is accomplished with this dissertation I hope that it is to push philosophers and political theorists to consider in more detail the epistemic problems of politics in order to gain a greater understanding of what it is that democracies are capable of achieving in their current form. In *The Fatal Conceit*, Friedrich Hayek wrote, “[t]he curious task of economics is to demonstrate to men how little they really know about what they imagine they can design” (1988, 76). Similarly, it may be the case that closer examination of the epistemic challenges posed by politics will lead us to reconsider traditional models for solving public

problems. Whether encapsulation and decentralization is the answer can only be determined by test, but even if it is not, then a new paradigm of theorizing about politics that genuinely takes the epistemic concerns presented in this dissertation into account cannot be reasonably avoided.

References

- Allen, Peter M. "Complexity and Management." In *Philosophy of Complex Systems*, ed. Cliff Hooker, 783-808. Amsterdam: Elsevier, 2011.
- Althaus, Scott. *Collective Preferences in Democratic Politics*. New York: Cambridge University Press, 2003.
- Altman, Andrew. "Democratic Self-Determination and the Disenfranchisement of Felons." *Journal of Applied Philosophy* 22, no. 3 (2005): 263-73. doi:10.1111/j.1468-5930.2005.00309.x.
- Anderson, Elizabeth, "The Epistemology of Democracy," *Episteme* 3 (1-2):8-22, 2006.
- . "Democracy: Instrumental vs. Non-Instrumental Value." In *Contemporary Debates in Political Philosophy*, ed. Christiano, Thomas, and John Christman. Malden, MA: Wiley-Blackwell, 2009.
- . "An Epistemic Defense of Democracy: David Estlund's *Democratic Authority*." *Episteme: A Journal of Social Epistemology* 5, no. 1 (2008): 129-39. doi:10.1353/epi.0.0025.
- Bar-Yam, Yaneer. "Multiscale Variety in Complex Systems." *Complexity* 9, no. 4 (April 22, 2004): 37–45. <https://doi.org/10.1002/cplx.20014>.
- Bishop, Robert C. "Metaphysical and Epistemological Issues in Complex Systems." In *Philosophy of Complex Systems*, ed. Cliff Hooker, 105-136. Amsterdam: Elsevier, 2011.
- Bloomfield, Robert J. "What Counts and What Gets Counted." SSRN Electronic Journal, 2015.
- Brennan, Jason. *Against Democracy*. Princeton, NJ: Princeton University Press, 2016.
- Bülow, William. "Felon Disenfranchisement and the Argument from Democratic Self-Determination." *Philosophia* 44, no. 3 (2016): 759-74. doi:10.1007/s11406-016-9722-y.
- Campbell, Donald T. "Assessing the Impact of Planned Social Change." *Evaluation and Program Planning* 2, no. 1: 67-90, 1979.
- Chisholm, Roderick M. *The Problem of the Criterion*. Milwaukee: Marquette University Press, 1973.
- Christiano, Thomas. "Reply to Critics of The Constitution of Equality." *Journal of Ethics and Social Philosophy* 5, no. 3 (2017): 1-15. doi:10.26556/jesp.v5i3.117.
- Choi, Willie, Gary Hecht, and William B. Tayler. "Lost in Translation: The Effects of Incentive Compensation on Strategy Surrogation." SSRN Electronic Journal, 2009.
- Claveau, François & Mireles-Flores, Luis, "On the Meaning of Causal Generalisations in Policy-oriented Economic Research." *International Studies in the Philosophy of Science* 28 (4):397-416, 2014.
- Condorcet, Marquis de. "Essay on the Application of Mathematics to the Theory of Decision-Making." In *Condorcet: Selected Writings*, edited by Keith M. Baker. New York: Macmillan Press.
- Converse, Phillip E. "Popular Representation and the Distribution of Information." In *Information and Democratic Processes*, edited by John A. Ferejohn and James H. Kulinski. Urbana: University of Illinois Press.
- Cohen, Stewart. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research* 65, no. 2 (2002): 309-29. doi:10.1111/j.1933-1592.2002.tb00204.x.
- . "Bootstrapping, Defeasible Reasoning, And *A Priori* Justification." *Philosophical*

- Perspectives* 24, no. 1 (2010): 141-59. doi:10.1111/j.1520-8583.2010.00188.x.
- Craigie, Jillian. "Competence, Practical Rationality and What A Patient Values." *Bioethics* 26, no. 6 (2011): 326–33. <https://doi.org/10.1111/j.1467-8519.2009.01793.x>.
- Davies, Anthony. "Tax Reform: Simplification is a Good Policy, with Anthony Davies." YouTube video, 39:25. October 7, 2013. <https://www.youtube.com/watch?v=Mzy20p9UNu8&t=944s>.
- Estlund, David. *Democratic Authority: A Philosophical Framework*. Princeton, NJ: Princeton University Press, 2007.
- Green, David, and Tania Leishmann. "Computing and Complexity—Networks, Nature, and Virtual Worlds." In *Philosophy of Complex Systems*, ed. Cliff Hooker, 783-808. Amsterdam: Elsevier, 2011.
- Hayek, Friedrich A. *Law, Legislation, and Liberty*. Vol. 2. The University of Chicago Press, 1976.
- . *The Fatal Conceit: The Errors of Socialism*. Edited by William Warren Bartley. London: Routledge, 1988.
- Hong, Lu, and Scott E. Page. "Groups of diverse problem solvers can outperform groups of high-ability problem solvers." Proceedings of the National Academy of Sciences of the United States of America. September 17, 2004. Accessed October 23, 2017. <http://www.pnas.org/content/101/46/16385.full>.
- Hooker, Cliff. "Introduction to Philosophy of Complex Systems: A." In *Philosophy of Complex Systems*, ed. Cliff Hooker, 3-90. Amsterdam: Elsevier, 2011.
- . "Conceptualising Reduction, Emergence and Self-Organisation in Complex Dynamical Systems." *Philosophy of Complex Systems*, 2011, 195–222. <https://doi.org/10.1016/b978-0-444-52076-0.50007-9>.
- Kometer, Michael. "The Strategy of Control: Centralized vs. Decentralized Control of US Airpower." *Defence Studies* 3, no. 2 (January 2003): 36–63. <https://doi.org/10.1080/1470243032000214269>.
- Nelsen II, John T. "Auftragstaktik: A Case for Decentralized Battle", *Parameters* 17, no. 3 (September 1987), 21-34.
- Oreskes, Naomi. "The Role of Quantitative Models in Science." In *Models in Ecosystem Science*, eds. C.D. Canham, J.J. Cole, W.K. Lauenroth. 13-31. Princeton: Princeton University Press, 2003.
- Oreskes, Naomi, and Kenneth Belitz. "Philosophical Issues in Model Assessment." In *Model Validation: Perspectives in Hydrological Science*, eds. M.G. Anderson and P.D. Bates, 23-41. London: John Wiley and Sons, 2001.
- Pierce, Charles S. "The Fixation of Belief." *Popular Science Monthly* 12, (November 1877), 1-15.
- Rickles, Dean. "Econophysics and the Complexity of Financial Markets." In *Philosophy of Complex Systems*, ed. Cliff Hooker, 531-569. Amsterdam: Elsevier, 2011.
- Ryan, Alex J. "Military Applications of Complex Systems." In *Philosophy of Complex Systems*, ed. Cliff Hooker, 783-808. Amsterdam: Elsevier, 2011.
- Slack, Donovan. "VA Bosses in 7 States Falsified Vets' Wait times for Care." USA Today. April 08, 2016. Accessed May 18, 2018. <https://www.usatoday.com/story/news/politics/2016/04/07/va-wait-time-manipulation-veterans/82726634/>.
- Somin, Ilya. "When Ignorance Isn't Bliss: How Political Ignorance Threatens Democracy."

- Policy Analysis 525: 1-27, 2004.
- . *Democracy and Political Ignorance: Why Smaller Government Is Smarter*. Stanford, CA: Stanford University Press, 2016.
- Sosa, Ernest. "How Competence Matters In Epistemology." *Philosophical Perspectives* 24, no. 1 (2010): 465–75. <https://doi.org/10.1111/j.1520-8583.2010.00200.x>.
- Staff, Newsweek. "Take the Quiz: What We Don't Know." Newsweek. March 29, 2011. Accessed May 18, 2018. <http://www.newsweek.com/take-quiz-what-we-dont-know-66047>.
- Thompson, Abigail, "Does Diversity Trump Ability? An Example of the Misuse of Mathematics in the Social Sciences." *Notices of the American Mathematical Society* 61:1024-30, 2014.
- Vogel, Jonathan. "Reliabilism Leveled." *The Journal of Philosophy* 97, no. 11 (2000): 602. doi:10.2307/2678454.
- Waldron, Jeremy, "Two Conceptions of Self-Determination." In *The Philosophy of International Law*, ed. Besson, Samantha, and John Tasioulas. Oxford: Oxford University Press, 2013.
- Westen, Drew. *The Political Brain: The Role of Emotion in Deciding the Fate of the Nation*. New York: Perseus Books, 2008.
- Westen, Drew, Pavel S. Blagov, Keith Harenski, Clint Kilts, and Stephan Hammann. "The Neural Basis of Motivated Reasonings: An fMRI Study of Emotional Constraints on Political Judgment during the U.S. Presidential Election of 2004." *Journal of Cognitive Neuroscience* 18:1947-58, 2006.
- Yunus, Muhammad. *Creating a World without Poverty: Social Business and the Future of Capitalism*. New York: PublicAffairs, 2007.