

## DISTRIBUTED LEDGER TECHNOLOGY AND THE FUTURE OF DEMOCRACY

VOTER COMPETENCY, DISTRIBUTED LEDGER TECHNOLOGY, AND THE FUTURE OF DEMOCRACY

By TRAVIS RAMSAY, B.A.

A Thesis Submitted to the School of Graduate Studies in Partial Fulfilment of the Requirements  
for the Degree Master of Arts

McMaster University MASTER OF ARTS (2019) Hamilton, Ontario (Philosophy)

TITLE: Voter Competency, Distributed Ledger Technology, and the Future of Democracy

AUTHOR: Travis Ramsay, B.A. (McMaster University)

SUPERVISOR: Professor Stefan Sciaraffa

NUMBER OF PAGES: vii, 87.

## ABSTRACT

A great challenge for democracy is to account for the conflict between the ideal of self-governance and the capacity of the average person to participate in democratic decision-making. This challenge has led some observers to question the defensibility of democracy and consider other systems of social organization. I argue instead that the problem can be solved with a technologically enhanced version of Thomas Christiano's choice of aims model of democracy. I begin by setting up the voter competency problem: I describe the ideals of democracy and the role that is ascribed to citizens under traditional accounts of democracy, then proceed to a discussion of the empirical evidence that shows how unlikely it is that voters could ever adequately perform such a role. While I consider a number of alternative democratic models which attempt to reconstruct the role of citizens in a way that is consistent with their capacities and with the democratic ideal of self-governance, I find that the choice of aims model strikes this balance in a way that is most tenable. Despite this, I argue that changes to the way information is distributed in modern democracies, to do with the rise of the internet, pose a serious threat to the viability of even this model, as it is becoming increasingly difficult for voters to ascertain reliable information. The second half of the thesis offers support to Christiano's model in the form of technologically enhanced institutions. Chapter 3 provides a basic understanding of an emerging technology called distributed ledger technology, which offers a new paradigm for how information is stored, controlled, and distributed around society. The final chapter demonstrates how this technology can be used to strengthen democratic institutions so that citizens are able to truly be said to self-govern in a way that is consistent with their capacities.

## **ACKNOWLEDGEMENTS**

I would like to thank my supervisor, Dr. Stefan Sciaraffa, for his support throughout this thesis project and in my professional development. He always engaged thoughtfully with my work and never failed to adeptly point me in the right direction when I needed guidance.

I would also like to thank the people in the Philosophy Department at McMaster. Thank you to Dr. Nick Griffin for encouraging me to pursue a graduate degree in Philosophy and to Dr. Wil Waluchow for being receptive to and supportive of this somewhat atypical project. Thank you to the department staff and students for providing a warm environment in which to work and learn.

Lastly, I would like to thank my friends and family for their support and patience, which afforded me the opportunity to pursue something I was passionate about. Special thanks to Morgan for bearing with me throughout the course of this adventure and never hesitating to read over my work, no matter what else was going on.

## TABLE OF CONTENTS

Acknowledgements	iv
List of Figures and Tables	vii
<b>INTRODUCTION</b>	<b>1</b>
<b>CHAPTER 1</b>	
<b>The Problem of Voter Competency</b>	<b>4</b>
1.1 Why Organize Democratically	4
1.2 The Ideals of Democracy	7
1.3 The Role of Citizens in Democracy	9
1.4 Heuristics	12
1.5 The Miracle of Aggregation	14
1.6 Retrospective Voting	15
1.7 Why Organize Epistocratically?	19
1.8 An Alternative Model of Democracy	20
<b>CHAPTER 2</b>	
<b>Information Distribution in the Modern State</b>	<b>23</b>
2.1 What Citizens Need to Know	24
2.2 The Division of Labour	26
2.3 Delegating Tasks in the Decision-making Process	28
2.4 Truth Sensitivity, Principles of Selection, and Public Realization	30
2.5 Information Distribution in Modern Democracies	35
2.6 The Missing Pieces for Modern Institutions	39
<b>CHAPTER 3</b>	
<b>The Advent of Distributed Ledger Technology</b>	<b>41</b>
3.1 Early Ledgers and the Organization of Society	41
3.2 Out of the Frying Pan and Into the Fire	42
3.3 Satoshi Nakamoto and the Advent of Distributed Ledgers	45
3.4 Guaranteeing the Immutability of Information on the Blockchain	49
3.5 Achieving Distributed Consensus	51
3.6 Administering a Distributed Ledger	54
3.7 Ethereum and Blockchain 2.0	58
3.8 Public and Private Ledgers	60

<b>CHAPTER 4</b>	
<b>Distributed Ledgers for Defensible Democracy</b>	<b>62</b>
4.1 Modern Identity for the Modern State	63
4.2 Truth Sensitivity and Transparent Selection Principles for Online Discourse	68
4.3 Electoral Systems in a DLT-Enhanced Choice of Aims Model of Democracy	72
4.4 Incentivizing Background Knowledge and Engagement	75
4.5 A DLT Enhanced Choice of Aims Model of Democracy	77
<b>CONCLUSION</b>	<b>79</b>
<b>REFERENCES</b>	<b>82</b>

## LIST OF FIGURES AND TABLES

**Figure 1:** Peter Steiner's Cartoon, p. 36

**Figure 2:** Difference between centralized and distributed networks, p. 48

**Figure 3:** A small change to the input results in a drastic change to the hash, p. 50

**Figure 4:** hashes, labelled here as "consensus", link blocks of data in a blockchain, p. 51

**Figure 5:** A fork occurs when some nodes update to new software and others do not, p. 57



## INTRODUCTION

This project started with a deep interest in an emerging technology called blockchain technology and a sense that it could have a profound effect on the way democratic institutions were organized. Originally, the thought was that blockchain technology would allow communities to organize using a direct form of democracy. As my research got underway it became clear that there were some major problems that made direct democracy an unrealistic possibility; chief among these was the problem of voter competency. Without a competent electorate, direct democracy was an unattainable ideal. On the other hand, if voters are not capable of making political decisions, it is difficult to conceive of a model of democracy that is consistent with the idea that voters should play a significant role in determining the terms of association for their societies. As the project proceeded, it turned out that this was really the most pressing challenge for democracy. Because of this, my attention turned to finding a defensible model of democracy, i.e., one that appropriately balanced a realistic notion of the capacity of voters, with the need for them to play a fundamental role in the decision-making process. What I found was that no one had described such an account in a way that could be placed in the context of the modern state and adequately defended. Despite diverging from the original goal of describing a blockchain-based form of direct democracy, I still believed the unique features of the technology had important implications for how we should design democratic institutions. Thus, the new plan for the project became to investigate and elaborate on how blockchain technology (something I later found is better described as distributed ledger technology) could be used to design a set of institutions that could realize the goals of democracy.

Towards this end I have written four chapters. In chapter 1, I provide an account of the ideals of democracy, establish what role must be performed by the electorate in order to satisfy those ideals, and investigate the ability of voters to perform that role. In light of evidence against the tenability of democratic ideals, I consider arguments from proponents of epistocracy, i.e., the thesis that society should be ruled by the knowledgeable. I conclude that the problem of voter competency is a serious threat to the defensibility of democracy and that, if proponents of democracy are going to find a defensible model of democracy it is going to resemble Thomas Christiano's choice of aims model.

In chapter 2, I move on to examine how information is distributed around modern societies. This is important because, if we are dealing with a problem of voter competency we need a clear understanding of how people get information and what barriers might prevent them from doing so. What I reveal in this chapter is that existing institutions in most contemporary democracies do not have adequate institutions in place to distribute information to voters, even under the most tenable accounts of democracy. Instead, significant changes to how our institutions are organized would be necessary for a defensible account of democracy. Importantly, I outline four criteria that democratic institutions must meet to achieve a minimal level of tenability.

In chapter 3, I introduce distributed ledger technology. This chapter describes the history of the technology, its unique features, and the technological innovations that make them possible. This chapter prepares the reader to start thinking about the application of distributed ledger technology and to engage in a discussion about how it can be used to address the problems currently facing democracy.

In the final chapter, I return to the four criteria laid out in chapter 2 and examine how the unique features of distributed ledger technology could be used to design institutions capable of meeting them. Referring to some innovations currently being made with these technologies, I provide a sketch of a set of institutions built on distributed ledgers. Ultimately, I argue that distributed ledger technology can be used as the framework for a set of institutions that would complement a model of democracy like the choice of aims model, together forming a defensible model of democracy that is adequate in the context of the modern state.

## CHAPTER 1

### The Problem of Voter Competency

It is often taken for granted that democracy is a justified form of social organization. Typically, this belief is grounded in what Christopher Achen and Larry Bartels call the *folk theory of democracy*, i.e., that ordinary people decide the policies that will be enacted in their societies by choosing them through referenda or by electing officials that carry out their will.<sup>1</sup> Accordingly, democratic governance is supposed to be justified because it involves self-governance. In this chapter we will examine what is required for citizens to perform the kind of role that is expected of them under something like the folk-theory of democracy and if such expectations are reasonable. Ultimately, I argue that the folk-theory of democracy is not tenable because ordinary people are incapable of making the policy decisions required of them under it. I close by considering an alternative conception of democracy that may avoid the problems of traditional models.

#### 1.1 Why Organize Democratically?

There are many different explanations for why we might want democratically organized societies. For example, John Stuart Mill argued that democracy is valuable because of a special kind of engagement that occurs between a person and the political considerations relevant to their society when they have a role in defining the terms of association under which they will be ruled. He asked his reader to imagine a beneficent despot who was guaranteed to produce the best laws and create the best set of social conditions for their nation.<sup>2</sup> Even if such a thing were possible, Mill warned that without widespread political participation, all but the few with an

---

<sup>1</sup> Achen and Bartels, *Democracy for Realists*, 1–3.

<sup>2</sup> Mill, *On Liberty; Representative Government; The Subjection of Women*, 179–80.

unusual taste for intellectual exercise would see their intellect and morality increasingly eroded.<sup>3</sup> For Mill, widespread civic engagement and political participation were desirable because these activities require citizens to consider the interests of others and to try to promote the common good, this in turn means engaging with moral, social scientific, and philosophical ideas.<sup>4</sup> Thus, according to Mill, democracy is valuable because it elevates common people; helping them to become better educated, to care about the interests of their fellow citizens, and to cultivate virtues.

Alternatively, proceduralists argue that some methods of distributing power or making decisions are intrinsically good, just, or legitimate (or that others are bad, unjust, or illegitimate).<sup>5</sup> David Estlund, for example, compares the process of democratic decision-making to the process of a jury trial: In the same way that a jury trial, when done right, produces a verdict that carries moral force (in addition to legal force) by virtue of being arrived at through a legitimate process, a political decision could carry moral force or fail to carry moral force depending on the legitimacy of the process through which it was determined.<sup>6</sup> According to explanations like this, democratically organized societies are desirable because the decisions made in them are made following a process that can be morally justified.

A third explanation for why we should organize societies democratically is the perspective of democratic instrumentalism, i.e., that democratic institutions and widespread political participation are good because they lead to just, stable, and efficient outcomes for citizens as compared to alternatives.<sup>7</sup> One version of this view argues democracies maintain a strong track

---

<sup>3</sup> Mill, 181.

<sup>4</sup> Brennan, *Against Democracy*, 54.

<sup>5</sup> Brennan, 11.

<sup>6</sup> Estlund, *Democratic Authority*, 7–9.

<sup>7</sup> Brennan, *Against Democracy*, 7.

record of producing better outcomes for their citizens as compared to other systems of government. Some even think the ability of democracy to reliably protect a number of particularly urgent and widely accepted human rights, combined with the inability of other systems of social organization to do so, produces a human right to live in a minimally egalitarian democratic society.<sup>8</sup> Proponents of views like this are able to defer to robust sets of empirical evidence showing the importance of democracy to protecting human rights. Christiano’s domestic peace argument, for example, is based on several studies that provide evidence of the reliability of minimally egalitarian democracies in protecting personal integrity rights (i.e., rights not to be murdered, disappeared, tortured, or imprisoned for political reasons by the state), while societies that are not minimally egalitarian democracies do not do so.<sup>9</sup> Moreover, the empirical evidence shows:

(1) *Correlation*—a very strong correlation between minimally egalitarian democracy and the protection of personal integrity rights, as well as between nondemocracies and the failure to protect these rights; (2) *democracy as an independent variable*—the identification of democracy as a key independent variable in explaining the protection of personal integrity rights; (3) *sequencing*—a sequencing argument showing that newly minted minimally egalitarian democracies achieve their protection of personal integrity rights fully only after about five years from when they come into existence; and (4) *an explanatory model*—a model that explains how minimally egalitarian democracies protect human rights and nondemocracies do not. These four elements give us good reason for thinking that minimally egalitarian democracy has the effect of protecting human rights and other societies do not.<sup>10</sup>

So, there are a number of seemingly good reasons for organizing societies democratically. But counter arguments have been offered for every one of them to varying degrees of success. Regardless, this thesis is not concerned with the question of *why* we should organize democratically, but instead we are interested in whether we can realize the underlying values of democracy, whatever they might be. Whether or not we should organize democratically is an

---

<sup>8</sup> Christiano, “An Instrumental Argument for a Human Right to Democracy,” 143.

<sup>9</sup> Christiano, 148–49.

<sup>10</sup> Christiano, 148.

interesting question, but it does not really matter if we can not first show that it is even possible to do so. Thus, for the remainder of this thesis we turn our attention to what is required to make democratic social organization possible and how such things can be achieved.

### *1.2 The Ideals of Democracy*

The fact that it is possible to organize societies in more than one way (e.g., a democracy, a monarchy, or an oligarchy) allows us to compare one way of organizing a society to another and to determine if one is better than another. In liberal western democracies it is often taken for granted that democracy is the best form of social organization and that the only appropriate way to rule a society is through a collaborative committee made up of all that society's adults.

The idea of one-person-one-vote appeals to many of our intuitions about equality and fairness. In an organized society, someone must be responsible for making decisions and for defining the terms by which people associate. In a monarchy or oligarchy, it seems somewhat arbitrary that a single person or group should have the authority to determine the rules of association for the whole of society. Whereas, the idea that only the people – or representatives they have chosen – have a right to rule, corresponds with our deeply held belief that we ought not to be forced to follow what someone else says without ourselves having some say in the matter. It seems that the fact that all people get to participate in choosing the terms of association for a society means that those terms have a special less-arbitrary force than if they were chosen by a single monarch or other subset of society.

Imagine a group of five strangers (Alice, Bill, Carol, David, and Erin) are stranded on an island and need to come together in order to build a basic survival shelter before a devastating storm rolls in, none of the individuals has any special skills or expertise that might justify them taking charge of the situation, nor do they have any pre-existing relationships that might confer

authority of one over another. If Alice were to just start ordering the others around and assigning duties to people, we would expect Bill, Carol, David and Erin might take exception to this and we would probably think they would be justified in doing so. It does not appear that Alice has the right to boss everyone else around. Now, what if, upon realizing they must build a shelter, the group decided to elect a leader who would organize the shelter construction and assign the relevant duties? It would seem that any directive given by the elected leader related to the domain of shelter construction would be imbued with a special type of justificatory authority.

For the sake of the example, we can imagine Bill was chosen as the leader. What sorts of things would undermine Bill's authority? What if it were the case that some of the castaways were not allowed to vote? If, for example, the female castaways were not given a vote, or if only a certain ethnic group got to vote, or some other kind of limited suffrage were in place. This would undermine precisely what justified Bill's directives in the first place. Imagine that only men got to vote in our example; the votes of Bill and David would determine the leader for the whole group of five, this would also run counter to our intuition about what gave legitimate authority to the elected leader. Similarly, imagine that only some subset of the group was eligible to run as leader or to nominate candidates, or that in some other way some party exercised more control over the election than another, this too would seemingly undermine the leader's authority. Finally, it would seem that each of the five people should get an equal opportunity to voice their opinions about the leadership decision, e.g., what they are looking for in a leader, who has the qualities they want in a leader, or what would make them the best leader. Again, if one of the castaways or a subset of them got an unfair opportunity to voice their opinions, it would seem to undermine the results of the vote.



This example brings to light the ideals that are embedded in our intuitions about what makes democratically organized societies preferable to other forms of social organization. Accordingly, a society is democratic and produces legitimate outcomes if it meets a set of conditions. First, it is ruled by the people, i.e., all minimally competent adults have the opportunity to participate in the deliberations and decision-making procedures that determine the terms of association which their society will follow.<sup>11</sup> Second, all citizens are political equals, they have the same participatory rights in the political system and no individual or group has a greater degree of control over the political decision-making procedure than any other.<sup>12</sup> Finally, each citizen has the opportunity to voice their opinions, to have their opinions heard by others, and to deliberate on subjects of public concern in open discussion.<sup>13</sup>

### *1.3 The Role of Citizens in Democracy*

It is these same sorts of intuitions that form the basis of what Achen and Bartels call *the folk theory of democracy* and what Robert Dahl called *populistic democracy*.<sup>14</sup> These traditional democratic doctrines emphasize the importance of popular sovereignty. For example, as Dahl describes his notion of populistic democracy, “whenever policy choices are perceived to exist [in a community], the alternative selected and enforced as governmental policy is the alternative most preferred by the members.”<sup>15</sup> Alternatively, Schumpeter says democracy “realizes the common good by making the people itself decide issues through the election of individuals who

---

<sup>11</sup> Christiano, *The Rule of the Many*, 3.

<sup>12</sup> Christiano, 3.

<sup>13</sup> Christiano, 3.

<sup>14</sup> Achen and Bartels, *Democracy for Realists*, 21.

<sup>15</sup> Achen and Bartels, 21.

are to assemble and carry out its will.”<sup>16</sup> In both cases we can see that central importance is placed on the role of citizens as bringing about specific policy decisions.

The problem with views like these is that they place an extraordinary epistemic burden on voters, as all policy decisions rest entirely on the decisions made by the electorate. Compounding this, the complexities of the modern state mean that informed decision-making requires in-depth knowledge of an overwhelming array of facts and social issues. Because of these difficulties, there has been a long history of skepticism about the ability of electorates to make informed policy decisions. But the middle of the twentieth century saw two important developments in the study of democracy that would give greater theoretical and scientific grounding to this earlier scepticism. First, economists like Anthony Downs conducted theoretical studies on collective choice which revealed much about how voters make decisions.<sup>17</sup> Second, political scientists and sociologists became armed with the technology of survey research for the study of public opinion and electoral politics.<sup>18</sup> These advances led researchers to find “time and time again... that the opinions and behaviour of ordinary citizens comported poorly with the expectations derived from democratic theory as they understood it.”<sup>19</sup>

If voters are to perform the role that is demanded of them in the classic doctrine of democracy, not only do they need to understand what the issues are (a complex enough problem on its own); they also need to know what outcome is in their own best interests, what candidates stated positions on the issues are, which candidates are likely to follow through with their stated position if they are elected, and what likelihood the candidates have of being elected.

---

<sup>16</sup> Schumpeter, *Capitalism, Socialism and Democracy*, 250.

<sup>17</sup> Achen and Bartels, *Democracy for Realists*, 23.

<sup>18</sup> Achen and Bartels, 23.

<sup>19</sup> Achen and Bartels, 23.

Every election year American National Election Studies (ANES) conducts a survey of general knowledge questions related to that year's election issues (subject matter that would have been widely debated during the election time). It is generally considered to be the most thorough scientific study of the U.S. electorate.<sup>20</sup> The survey for the 2000 Gore-Bush election consisted of 30 multiple choice questions, 19 of which had only 2 or 3 possible options; a score of 8.5 out of 30 would be expected as a result of random guessing.<sup>21</sup> The following is a list of some of the survey results:<sup>22</sup>

- The average score on the test was 14.3 correct answers out of 30 (47.6%)
- 12 of the 30 questions were answered correctly more than 50% of the time
- 47% knew Bush was a conservative
- 38% knew Gore was a liberal
- 41% knew Bush was more likely to favor jobs over the environment than Gore
- 37% knew crime rates decreased from 1992-2000
- 37% knew spending on the poor increased from 1992-2000
- 51% knew Gore was more supportive of gun control than Bush
- 46% knew Gore was more supportive of abortion rights than Bush
- Only 15% could name at least one candidate from the house of representatives in their own district and only 4% could name a second

Similarly disappointing results can be found regarding subsequent elections as well as concerning international political issues. Ilya Somin goes into extensive details on this in the first chapter of *Democracy and Political Ignorance*. While many people do not vote at all, Ilya Somin notes that 35 percent of American voters are “know-nothings” who have little to no political knowledge.<sup>23</sup> The

---

<sup>20</sup> Somin, *Democracy and Political Ignorance*, 30.

<sup>21</sup> Somin, 33.

<sup>22</sup> Results taken from table 1.4 in, Somin, 31–32.

<sup>23</sup> Somin, 33.

conclusion that ought to be drawn from these results is that clearly voters cannot play any role in a set of institutions requiring them to have robust knowledge of political facts and issues.

As results like these have become more and more common theorists have used a number of different strategies to reconstruct the role of voters or otherwise deny that voters are as incompetent as empirical research suggests. Even if it is true that voters are not competent enough to play the role accorded to them under idealist accounts of democracy, there may be other ways to understand how voters are responsible for the positive track record of democracies. In the following sections we will explore some of these criticisms and reconstructions.

#### *1.4 Heuristics*

One suggestion is that even though voters tend to be unable to demonstrate their knowledge about political candidates and policies on opinion surveys, they are nonetheless able to make rational electoral decisions by using information shortcuts or heuristics. These could take the form of cues from trusted individuals or groups, inferences based on political or social stereotypes, or generalization based on personal experience.<sup>24</sup> Some political scientists have treated heuristics and information shortcuts as a boon to idealistic views of democracy, thinking they are a way of reconciling the role demanded of voters with the empirical evidence showing the lack of knowledge about politics.<sup>25</sup>

Psychologists, however are less optimistic: these types of heuristics are related to cognitive biases. Human cognitive biases are a very well-documented concept in psychology, they involve judgement errors that humans make due to irrational thought processes which allow them

---

<sup>24</sup> Achen and Bartels, *Democracy for Realists*, 38.

<sup>25</sup> See, Popkin, *The Reasoning Voter*; and Sniderman et al., *Reasoning and Choice* for examples of these kinds of arguments.

to make sense of the complex world around them.<sup>26</sup> Some examples include, *confirmation bias* which occurs when people take in only those facts and opinions that support their established viewpoints and ignore all evidence to the contrary; *bandwagon effect*, a cognitive bias that makes people more willing to take up a particular position the more commonly held that position is in their society; and *intergroup bias*, which involves the human propensity towards forming groups, our favoritism of people in our own groups, and our dislike of people in groups other than our own.<sup>27</sup> Cognitive biases make it easier to manage the complexities of the world. Like cognitive biases, heuristics do not necessarily lend themselves to ordinary people successfully making rational decisions.

Achen and Bartels cite the example of Gerald Ford eating a tamale during a primary campaign appearance in Texas where he made the embarrassing gaffe of failing to remove its cornhusk wrapper before consuming it. Mexican-American voters could have seen this as a failure to show familiarity with their culture which could be related to a lack of ability to relate to their ethnic group, something that could in turn be later used as a heuristic or information shortcut when it came time to vote.<sup>28</sup> While it is possible the Ford's gaffe was related to his indifference towards the Mexican-American population and his policy preferences when it came to issues that were important to them, it could also have no connection whatsoever. If a mistake like this became a sound bite or was picked up by Ford's opposition, it could be to the detriment of the opinion of him among Mexican-Americans. It could easily result in a heuristic or information

---

<sup>26</sup> Ruth, "Cognitive Bias."

<sup>27</sup> Ruth.

<sup>28</sup> Achen and Bartels, *Democracy for Realists*, 39.

shortcut for some Mexican-American voters, regardless of how whether or not it provides any insight into where Ford stood on issues that would have been important to this group of voters.

Heuristic and information shortcuts are not necessarily reliable when it comes to electoral decision-making. Empirical evidence still shows that ordinary voters (using whatever heuristics are available to them) are not very successful at mimicking the policy preferences of the well-informed.<sup>29</sup> Yet when “ordinary people are exposed to intensive political education and conversation on specific policy issues, they often change their minds.”<sup>30</sup> While heuristics and information shortcuts help people come to decisions without being able to point to the specific corresponding political facts, it does not appear that this conduces to making electoral decisions that match policy preferences.

### *1.5 The Miracle of Aggregation*

The idea behind theorems of collective rationality is that, even with a group of individuals that is almost entirely ignorant, as long as there is a small percentage of the group that is competent, then the group as a whole should perform as if it was entirely competent. The larger the group gets, the more likely this is to work. Although this sounds strange, it is mathematically sound.

Imagine we have 100 people who have to pick the winner of an upcoming cricket game between the red team and the blue team and they have to do so via the outcome of a secret ballot without any collaboration. Unfortunately, 90 of them have no knowledge of cricket whatsoever; they’ve never seen it played, they know nothing about the rules, and have no knowledge of the history of competitive cricket (for the sake of example we will stipulate that they have no biases

---

<sup>29</sup> Achen and Bartels, 40.

<sup>30</sup> Achen and Bartels, 40.

about favourite colour or jersey styles). However, 10 members of the group are experts who know that the red team is awful and has never won a game and the blue team is a star-studded super team (it is essentially a foregone conclusion that the blue team will win). Based only on mathematics, if you ran an experiment like this, the 90 know-nothings should cancel out each other's votes, 45 would choose red and 45 would choose blue, the 10 experts would vote correctly, and the group as a whole would choose the right outcome (45 votes for red, 55 for blue). Even with some room for variance, more often than not, the almost entirely uninformed group should be able to consistently choose the right answer.

The problem with applying this concept to real-world elections is that it relies on a major assumption that simply is not true in the real-world, i.e., that people will make errors independently from each other.<sup>31</sup> In the example above we had 90 know-nothings who were not allowed to collaborate prior to casting their ballot, on average they would guess right and wrong in equal proportions. In real-world situations this would be a bad assumption to make, as people have opportunities to influence each other and to be disproportionately influenced by others to make errors in ways that are not independent. As Achen and Bartels put it, "When thousands or millions of voters misconstrue the same relevant fact or are swayed by the same vivid campaign ad, no amount of aggregation will produce the requisite miracle; individual voters' 'errors' will not cancel out in the overall election outcome."<sup>32</sup>

### *1.6 Retrospective Voting*

The idealistic democratic theories like the folk theory or populist democracy require voters to have a significant degree of political knowledge. We have shown above that there are

---

<sup>31</sup> Achen and Bartels, 41.

<sup>32</sup> Ibid.

no shortcuts that can be used which will allow ill-informed voters to perform as if they were properly informed, not on the individual scale through heuristics, nor on the group scale through the miracle of aggregation. But what if the problem is not how much voters know, but instead that our expectations of what they need to know are simply too high?

This is the thought behind retrospective voting accounts of democracy. Rather than needing to have a robust knowledge of the array of complex social and political issues in contemporary states, voters just need to be able to tell if the incumbent government has done a good job and are deserving to stay in power, or if they have done a bad job and should be removed from office. According to a theory like this, voters can exert a substantial amount of control over their governments, despite having little knowledge of social or political issues.<sup>33</sup>

In a retrospective theory of democracy voters decide whether the incumbents should stay in power based on whether their own welfare – and the welfare of their fellow citizens – improved or worsened as a result of the governing party's policy decisions. *Economic Voting* is a type of retrospective voting that narrows the burden of knowledge placed on voters even further, suggesting voters need only assess whether they (and their fellow citizens) are better or worse off solely in terms of economic wellbeing.<sup>34</sup> At first glance, this seems like a viable strategy for reorganizing the task of voters into something plausible. How hard could it possibly be to know if your current state of wellbeing is better or worse than it was previously?

As it turns out, this is not as easy as it sounds. A person's economic wellbeing at any given point in time is a product of more than just the policy decisions of the incumbent government. Global economic trends have significant influence over national economies and the decisions of

---

<sup>33</sup> Achen and Bartels, 91.

<sup>34</sup> Achen and Bartels, 146–47.



previous governments can have economic impacts for years after their tenure has come to a close.

As early as 1957, Anthony Downs discussed how challenging these sorts of calculations would be:

So far we have glibly spoken of voters computing their party differentials and performance ratings without pointing out how difficult such computation is. In order to find [their] current party differential, a voter in a two-party system must do the following: (1) examine all phases of government action to find out where the two parties would behave differently, (2) discover how each difference would affect [their own] utility income, and (3) aggregate the differences in utility and arrive at a net figure which shows by how much one party would be better than the other. This is how a rational voter would behave in a world of costless information... In the real world, uncertainty and lack of information prevent the most intelligent and well-informed voter from behaving in the fashion we have described.<sup>35</sup>

To be able to make competent decisions, even solely on retrospective economic terms, voters have to be capable of parsing out the difference between economic changes that incumbents are responsible for and those that they could not have controlled and they need to have an acute sense of how other candidates might have performed differently.<sup>36</sup> On the surface, retrospective voting appears like a viable way that a less informed electorate could make a meaningful contribution to how political decisions are made, but when we think about what is actually involved in retrospective voting, it is hardly any more simple than what is demanded in idealist versions of democratic theory.

Once again, when we review the empirical data, voters do not have a good track record of making rational electoral decisions using retrospective voting.<sup>37</sup> Voters often punish incumbents for misfortunes that are well beyond their control. Things like droughts, floods and even shark attacks have been seen as good reasons to remove incumbents.<sup>38</sup> Achen and Bartels argue that retrospective voting is often blind retrospection, “voters, ignorant about evidence and

---

<sup>35</sup> Downs, *An Economic Theory of Democracy*, 45–46.

<sup>36</sup> Achen and Bartels, 102-15.

<sup>37</sup> This empirical data is detailed in chapters 5, 6, 7 of Achen and Bartels, *Democracy for Realists*.

<sup>38</sup> Achen and Bartels, *Democracy for Realists*, 143.

causation, but supplied with a tale of incumbent responsibility, will punish incumbents whenever their subjective well-being falls below some fixed standard, regardless of whether or not their pain is in fact traceable to incumbent's policies."<sup>39</sup>

In the last few sections we have established that ordinary voters do not have the knowledge needed to perform the function required of them under a version of democratic theory like the folk theory of democracy. We also showed that shortcuts like heuristics or aggregative miracles are unlikely to succeed in making things any easier for them. Finally, we showed that retrospective voting, which attempts to limit what voters need to know to make policy decisions, is not substantially easier for voters to perform than the other theories we have considered so far.

Despite the difficulty we have had reconciling the role of the electorate with the ideals of theories like the folk theory of democracy, we know that democracies produce better outcomes for citizens than theocracies, monarchies or tyrannical dictatorships. Considering the capacity of the average voter this may lead us to wonder if these differences are due to the participation of the electorate or despite it. It would seem that in order to produce the kind of results that are desirable, we have had to limit the role of the voter to the extent that it is unclear that we are doing any more than paying lip service to democratic ideals. Recall the three ideals of democracy that we identified at the beginning of the chapter:

1. All minimally competent adults have the opportunity to participate in the deliberations and decision-making procedures that determine the terms of association which their society will follow.
2. All citizens are political equals, they have the same participatory rights in the political system and no individual or group has more control over the political decision-making procedure than any other.

---

<sup>39</sup> Achen and Bartels, 144.

3. Each citizen has the opportunity to voice their opinions, to have their opinions heard by others, and to deliberate on subjects of public concern in open discussion.

If the function that voters have been reduced to is to rein in the real political decision-makers when they get out of control (which at best seems like what their competency would allow), then it seems that none of these ideals are really being met in anything more than the most superficial degree. It is due to these kinds of incompatibilities that Thomas Christiano says that “the ideals of democracy to which many adhere are impossible or even incoherent under modern conditions.”<sup>40</sup> If this is the case, the other thing that ought to be considered is whether the results democracies produce, which are good relative to theocracies and dictatorships, could not be better under some alternative form of social organization.

### *1.7 Why Organize Epistocratically?*

Jason Brennan thinks the answer to the problem of voter competency is to limit suffrage to only those who are competent, or at the very least to give more votes to the competent.<sup>41</sup> He does not believe there are any good non-instrumental reasons to organize societies democratically: he denies that democracies have enlightening and ennobling effects on their citizenry as Mill claimed; that there is anything special about democratic decision-making procedures that make them uniquely just, good, or legitimate; and he argues against a series of more minor non-instrumental arguments for why we should organize democratically.<sup>42</sup> From this series of arguments, Brennan concludes that if the only reasons to organize democratically are instrumental reasons and if epistocracies produce better outcomes than democracies (which they

---

<sup>40</sup> Christiano, *The Rule of the Many*, 105.

<sup>41</sup> Brennan, *Against Democracy*, 211.

<sup>42</sup> Brennan, *Against Democracy*.

are likely to, due to the proven incompetence of the electorate), then we have a moral obligation to institute epistocratic political institutions.<sup>43</sup>

I do not pursue this kind of strategy in this thesis. I think, even if Brennan is right that there are only instrumental reasons for organizing democratically, placing limitations on suffrage is a radical answer to the problem of voter competency which would probably result in violence and social upheaval. Advocating for such a strategy would be irresponsible if there are still things that can be done to improve existing forms democratic social organization. The reason for bringing Brennan's arguments to bear on this topic was because they help illustrate the stark reality of how serious a problem the competency of the electorate is for our conceptions and justifications of democratic institutions. Instead, in this thesis I pursue a strategy for dealing with the problem of voter competency that will involve the identification of a more tenable model of democracy and the advocacy of a complementary set of institutions that will enable citizen self-governance that is consistent with the ideals of democracy and the realities of the modern state.

### *1.8 An Alternative Model of Democracy*

Thomas Christiano has put forward an alternative model of democratic decision-making called *The Choice of Aims Model* which restructures democratic institutions in such a way that is designed to lower the knowledge burden that is placed on voters, while ensuring that they play a vitally important role in the political decision-making process. The choice of aims model decreases the amount voters need to know, because it has them choose the overall aims for society that governments pursue, as opposed to the specific policies that should be put into place.<sup>44</sup> Accordingly, under the choice of aims model:

---

<sup>43</sup> Brennan, 141–42.

<sup>44</sup> Christiano, *The Rule of the Many*, 169–70.

Citizens choose the ultimate ends that the society is to pursue. For example, citizens choose what is fair and what is unfair, they choose how much risk they wish to submit to foreign relations, and they may choose what the aims of the education are as well as the basic features of a system of health care. However, citizens do not choose the means by which to bring about these ends. Though they are to choose what is fair and what is unfair, they are not the choosers of the policies that bring about fairness in economic relations. Though they may choose that there be a certain level of control of environmental damage, they do not choose the means for achieving this control.<sup>45</sup>

Thus, the choice of aims model sees the government act as a kind of trustee to the citizens insofar as they exercise their expertise to carry out the goals of the society.<sup>46</sup>

Christiano's model has two major strengths, both of which are necessary in my opinion, for an account of democracy to be justified: (1) he succeeds in decreasing the knowledge burden on voters by requiring that they only need to understand what the aims of society should be, not the details about how this will be achieved; and (2), he does this without diminishing the role of the electorate to the extent that they are entirely unable to play a role significant enough to meet the ideals of democracy that made it justifiable in the first place. Unlike the other attempts to reconcile voter competency with the ideals of democracy, Christiano's account leaves room for all minimally competent citizens to participate in choosing the terms of association of the society they live in; for them to have the same participatory rights in the political system and the same degree of control over that system; as well as to voice their opinions, to have their opinions heard by others, and to deliberate on subjects of public concern in open discussion.

While Christiano's model achieves things that none of the other models we considered are able to, the one area where he may receive criticism is that it is still the case that voters are not competent enough to be adequate choosers of aims or that even if they are, powers exist that are capable of undermining and manipulating the ways they get the information needed to

---

<sup>45</sup> Christiano, 170.

<sup>46</sup> Christiano, 171.

perform even the simplified role that Christiano sets for them. The remainder of this thesis is ultimately about mustering some institutional support for Christiano’s model that will make it invulnerable to these kinds of criticisms and bring it up to date with the realities of the modern state.

It is my position that the problem of voter competency is a problem of information distribution. While Christiano’s model strikes a reasonable balance between voter participation and the ideals of democracy, more can be done to make sure voters have the information they need to choose the overall aims for their societies. In the next chapter I will explain and defend my claim that the voter competency problem is a problem of information distribution, I will also explain how this presents a particular problem for the contemporary democratic states. In the second half of this thesis, I go on to explain a revolutionary technology and how it can be used to correct the problems of information distribution that plague modern societies and limit the ability of ordinary citizens to contribute to choosing the terms of association of the societies they live in.

## CHAPTER 2

### Information Distribution in the Modern State

In the previous chapter we investigated the disparity between what the ideals of democracy demand of voters and what typical citizens are actually capable of. What we found was that most models of democracy are unable to account for the role of citizens in a way that is consistent with their capacities, while still ensuring they contribute in a meaningful way to choosing the terms of association of the society they live in. One model that navigates this most plausibly is Thomas Christiano's choice of aims model of democracy which describes the role of citizens as to choose the overall aims of society, but not the means by which they will be achieved. By limiting the citizens role to choosing only the overall aims of society Christiano lowers the epistemic burden placed on citizens but ensures they still play a very important role in deciding the terms of associations for their societies.

On the face of it, Christiano's model strikes a plausible balance between the need for citizens to be involved in the decision-making process and being realistic about the degree to which they can ascertain and apply political information. But we ought not take it for granted that just because Christiano appears to strike a balance that is ostensibly consistent with democratic ideals, that our existing institutions will accommodate his model. It is possible to have a *prima facie* plausible account of democracy but lack the kinds of institutions necessary to facilitate it.

To illustrate the point that I am making here, let's consider an extreme example of a state that claims to be a democracy that is consistent with the choice of aims model: they hold elections that involve citizens choosing a set of aims and maintain administrative, judicial, executive, and legislative branches that will carry out those choices with sincerity. However, in this imaginary state, the government controls all the information that is distributed throughout the state, has

the ability to censor whatever information it chooses, and has made it illegal to publicly contradict information that has been distributed by the state. Accordingly, a plausible account of democracy needs to be complemented by an appropriate set of institutions if we are to realize a defensible form of democracy. If we have a model but no appropriate set of institutions, then there may be a strong practical case to be made by the proponents of alternative forms of social organization, like epistocracy.

The goal of this chapter will be to show that existing institutions do not enable citizens to choose the aims of society and thereby realize the ideals of democratic governance. Towards this end, I will delineate what sorts of things a voter would need to know to competently perform their role under a model like Christiano's, I will then proceed to discuss the challenges associated with the dissemination of information throughout modern democratic states, and close by showing how our existing institutions fail to adequately distribute truth sensitive information to citizens in a way that would allow them to properly perform their role, even under the most plausible models of democracy. Ultimately, I argue that, due to challenges for distribution of information in the modern state, significant institutional reform would be necessary to make even our most plausible models of democracy viable.

### *2.1 What Citizens Need to Know*

A model of democracy must explain how citizens can be understood as self-governing. As was discussed in the previous chapter, most models of democracy have difficulty reconciling the capacities of the average citizen with the demand for self-governance. Part of the reason these other models fall short is because they fail to acknowledge the separable components of law and policy. On this point, Thomas Christiano says, "a law is in some ways much like an intention to act. Intentions to act include aims that the intended action is meant to promote as well as a plan by



which the aims are realized.”<sup>47</sup> Likewise, the laws and policies which regulate societies involve basic aims which they are meant to promote as well as a plan by which they will be achieved.<sup>48</sup> In order to perform a role that could be reasonably understood as self-governance, voters need not determine both the aims of society and the means for achieving them. Instead, by determining just the overall aims of their society, citizens are able to realize self-governance.

Christiano defines the role of citizens under the choice of aims model as follows: “Citizens must choose packages of ends. They must choose many different goods, combined with their costs, that they think society ought to pursue, and they must rank the various different combinations of values and costs so that their choices can be represented by a preference ordering.”<sup>49</sup> Christiano envisions voters making choices between robust political party platforms that offer a range of packages of aims which they will pursue if elected. Political parties attempt to accommodate the range of interests and concerns in society, they devise reasoned views on aims of society, and persuade citizens to their views in competition with other parties throughout an election campaign.<sup>50</sup> The electoral process is a massive deliberation: parties put out position papers, their leaders give speeches, debates are held, interest groups put pressure on parties to integrate their views, and all this promotes discussions among regular citizens who ultimately decide who wins seats in the legislative assembly.<sup>51</sup> Once this is determined members are committed to pursuing – and bargaining with other representatives – for the conception of aims

---

<sup>47</sup> Christiano, 167.

<sup>48</sup> Christiano, 167.

<sup>49</sup> Christiano and University of Arkansas Press, “Democracy and Social Epistemology,” 69–70.

<sup>50</sup> Christiano, *The Rule of the Many*, 245.

<sup>51</sup> Christiano, 246.

for which they have been elected.<sup>52</sup> Voters are part of the agenda setting for their society and have an equal voice on choosing its aims.

To perform their role in this process, voters need to have knowledge of what social arrangement is in their best interest and what is in the best interest of the rest of their society, the societal aims which would realize those interests, and the corresponding candidate or party preference which would be most likely to bring about that set of societal aims. Moreover, competent participation involves making choices that are based on *reasoned preferences*, i.e., ones for which the underlying evaluative beliefs are both connected by means of good reasoning to a substantial number of other beliefs and are not systematically or deeply false.<sup>53</sup> Accordingly, under the choice of aims model (and the same would be true of any model of democracy that gives a meaningful enough role to citizens to be defensible), institutions must exist that will ensure the widespread ability of citizens to cultivate reasoned preferences about the aims that should be pursued in their society. To be sure, this will require a baseline of relevant background knowledge and incentive to engage consistently in political discourse.

## *2.2 The Division of Labour*

One of the most difficult trade-offs that must be managed by democratic institutions is the division of labour. In complex societies, while all voters should be able to form reasoned preferences about the aims of their society, not everyone can be an expert on the vast array of knowledge that is necessary for making decisions about how the aims of society will be achieved. Moreover, there are other things that must be accomplished in a society that require different sets of expertise and most people do not have the time, energy, or cognitive capacities to be adept

---

<sup>52</sup> Christiano, 246.

<sup>53</sup> Christiano and University of Arkansas Press, "Democracy and Social Epistemology," 81.

at more than a couple of them. Thus, democracies have labour markets that encourage different people to specialize in each of the various areas that are required for society to function. This means that while some people will be experts in each of the various facets required for a society to operate, a limited number will have the expert knowledge needed to inform political decisions. Accordingly, most modern democracies have a large bureaucratic body that includes experts in all the various areas that are relevant to the policy decisions that must be made by the state. On top of this, experts from outside the government have opportunities to weigh-in on decisions when the public is consulted or in publications that will publicly critique policy proposals and decisions.

The difficulty with this is largely what we explored in the previous chapter, the demands of the ideals of democracy seem to require that each citizen have knowledge on many historical, political, economic and social scientific issues, as well as some understanding of the sciences, the state of play in law and policy, and the way compromises are made for the formation of legislation; this simply is not possible given the division of labour required in the modern state. Even under accounts of democracy that significantly diminish the role of citizens, it is unlikely that the average person will acquire the knowledge needed to participate competently on their own, especially if their role in society does not reward expertise in the domains related to political decision-making. The difficulty created by the division of labour led Anthony Downs to conclude that, “in general, it is irrational to be politically well-informed because the low returns from data simply do not justify their cost in time and other scarce resources.”<sup>54</sup>

---

<sup>54</sup> Downs, *An Economic Theory of Democracy*, 259.

### *2.3 Delegating Tasks in the Democratic Decision-Making Process*

Of course, voters are not entirely on their own when it comes to acquiring and assimilating political information. Institutions that distribute information are like shortcuts to the relevant information that is necessary to make reasoned decisions. A person left strictly to their own devices to procure all the relevant information to be properly informed would be lost in a sea of data with an incredible cost of acquisition. Imagine if people had to personally research topics and query various government bodies to get information about changes in the economy, demographics, policy implementations and so on. Finding out all the things one would need to know would likely be impossible. In reality, newspapers, media outlets, government publications, university research publications, political parties, think tanks, interest group associations, web blogs, social media, and general public discourse, all contribute to the dissemination of relevant information throughout the population, effectively lowering the cost of information acquisition by magnitudes.

Downs put forward the following list of the main steps involved in rationally deciding how to vote and then voting:

1. Gathering information relevant to each issue upon which important political decisions have been (or will be) made.
2. For each issue, selecting from all the relevant information gathered that which will be used in the voting decision.
3. For each issue, analyzing the facts selected to arrive at specific factual conclusions about possible alternative policies and their consequences.
4. For each issue, appraising the consequences of every likely policy in light of relevant goals.
5. Coordinating the appraisals of each issue into a net evaluation of each party running in the election.

6. Making the voting decision by comparing the net evaluations of each party and weighting them for future contingencies.
7. Actually voting or abstaining.<sup>55</sup>

In the Downsian view of democracy, each step of this process – except for the final act of actually voting or abstaining – can be delegated to someone else.<sup>56</sup> In other words, at each step other than the final one, the voter can take advantage of various institutions in their society and shift the cost of the step onto them. Thus, if a voter wants to gather information about the costs and benefits of nuclear energy, they can either do the research themselves to discover how nuclear fission works and to develop a first hand understanding of all the risks involved and their likelihood of occurring, or they could rely on an expert on nuclear energy to appraise the relevant information and to assess the related policies. Likewise, a union worker could go through each step of the process for each issue and come to a nuanced understanding of all the relevant information or they could skip all those steps and rely on their union to make all the appropriate evaluations, assuming that by virtue of being a member of the union, the union will advocate for the decision that has the worker's own best interest in mind. Of course citizens can utilize a complex network of delegations of these steps to numerous different information distribution institutions (e.g., newspapers, social media platforms, research papers etc.) to arrive at their purportedly reasoned preference about which party to vote for.

To be clear, Christiano's model does not map perfectly on to Downs' list of the steps in the voting process and it is probably not true that all the first 6 steps could be delegated to some other person or institution under the choice of aims model. For example, it is unlikely that

---

<sup>55</sup> Downs, 209.

<sup>56</sup> Downs, 209.

delegating step 6 – “making the voting decision by comparing the net evaluations of each party and weighting them for future contingencies” – would be consistent with a competently made decision under the choice of aims model, as it is essential under the model that voters are responsible for choosing the aims of society. Simply deciding to cast a vote that has been decided by another party will not likely be enough here. Nonetheless, somewhere in the list a line can be drawn between steps that are involved in the curation and analysis of data and the steps that must be done by the voter so that they can truly be understood as making the choice of their society’s aims. For our purposes we are mostly concerned with the delegation of steps related to the curation and analysis of data. For these steps, Downs’ statements about delegation hold across both the traditional model that Downs was working under as well as Christiano’s model.

#### *2.4 Principles of Selection, Truth Sensitivity, and Public Realization*

It is important to understand that there are costs associated with delegating steps in the process to other parties. When a part of the decision-making process is delegated to another party, just like an individual voter would have to, that party must make decisions about how to gather, select, analyze, evaluate, and transmit data. There is a vast amount of data available and these other parties are also limited by the amount they can process and how they can effectively distill it in a way that makes the delegation process fruitful for the individual delegating the part of the decision-making process in question. This is going to be governed by the delegated party’s *principles of selection*, i.e., the rules they employ to determine what data to make use of and what not to make use of.<sup>57</sup>

---

<sup>57</sup> Downs, 210-12.

Consider the following example, you are a professor in a university department and you want to know how many people in your philosophy department are familiar with Immanuel Kant's *Critique of Pure Reason*. But because you are so busy you do not have time to investigate it yourself, so you delegate the task to one of your graduate students. If you were to conduct the investigation yourself you would have asked all the people in the entire department (staff, graduate students, and undergraduates) whether they have read at least some part of the book. Ideally, your delegate would go around and do just that. But we can imagine that the principles of selection employed by the graduate student might be different than your own. Perhaps the student thinks being familiar with *The Critique of Pure Reason* requires reading the whole book cover to cover, so they ask the entire department whether they have read the whole book, consequently they report back a much lower number than you would have got had you done the investigation yourself. Alternatively, we could imagine that the student is a big proponent of Kant and they want people to think he is much more relevant in contemporary philosophy departments than he actually is, because of this disposition towards Kant they deliberately ask a question that will result in a high prevalence of "familiarity" with Kant, so they ask respondents if they have merely heard of Immanuel Kant, consequently they report a much higher number back to you than you would have got if you conducted the investigation yourself. Finally, let's imagine the student has a very full schedule, just like you, they have costs associated with acquiring information, and they have to make decisions concerning how to go about doing so, because of this they ask the same question you would have, but they just ask whatever people they find around the department office in a single afternoon, rather than exhaustively polling the entire department, which would have taken much longer. Because the majority of people around the office were department staff and graduate students, your delegate finds that almost everyone

they asked had read at least some part of *The Critique of Pure Reason* and they report back a different result than what you would have got yourself. This example shows how differences in the selection principles of individuals and those they delegate information curation processes to can result in those individuals acquiring different information than they would have acquired had they curated it themselves.

Likewise, there are risks associated with voters delegating steps in the decision-making process to other parties. Namely, that voters will make decisions based on information gathering and assimilating processes that are not akin to their own. However, these can be mitigated if voters have a good sense of the selection principles of those they delegate to and how they might differ from their own. This means having transparent information distribution institutions, for which voters understand the biases, which have reputations that are well-formed over time, and are thoroughly vetted by members of the public. On the other hand, if the institutions involved in the distribution of information throughout society are opaque or have no reliable reputation, then voters will be unable to know how well the selection preferences of such institutions align with their own or to adjust for them when it comes time to make decisions.

A related concern is the problem of discourse failure. Discourse failure occurs when discourse produces political positions that are traceable to truth insensitive processes, i.e., ones that disregard the best available reasons, understood as those that define the state of investigation in the relevant reliable scholarly disciplines.<sup>58</sup> The worry is that we cannot take the deliberation of citizens seriously if they do not acquire relevant information, from disciplines that they are not experts in themselves, through processes that are reliably truth sensitive.<sup>59</sup> Here we

---

<sup>58</sup> Christiano, "Rational Deliberation among Experts and Citizens," 31.

<sup>59</sup> Christiano, 31.



can think of the example of the recent phenomenon of ignoring or disregarding the opinions of experts on climate change, when making deliberations about policy related to the environment. When this happens, it represents a discourse failure, as people who deny the reality of climate change and oppose the pursuit of social aims that would curb it, are arriving at political positions that can be traced to processes that disregard the current state of investigation in the environmental sciences. Accordingly, a democracy needs institutions that ensure the proper integration of “the parts of the division of labour that are concerned with arriving at sophisticated and well-developed understandings of the social world” with the “parts that are concerned with decision-making.”<sup>60</sup> This will be true whether voters are choosing the specific policies or merely the social aims.

Finally, voters need some meta-knowledge about their institutions. Not only do institutions need to promote ideals of democracy, but voters need to be able to see that this is the case. Being told that you are being treated equally, that governments are accountable to citizens, that the outcome of an election was as it is being reported, or that your participations matters, is one thing. Being able to see that you are an equal, that your government is accountable to you, that people actually elected the government that is in power, or how one’s participation impacted their society, is something altogether different. When people have knowledge like this latter set about their institutions, they have reason to be engaged and to perform a role in society that is consistent with democratic ideals.

To summarize what we have discussed up to this point in the chapter, Christiano’s choice of aims model provides an account of democracy that is *prima facie* more tenable than other

---

<sup>60</sup> Christiano, 32.

accounts because it describes how democracy can be organized such that voters have a meaningful role in the decision-making process which is compatible with their capacities given the limitations caused by the division of labour. However, even though the choice of aims model lowers the epistemic burden placed on voters, it does not eliminate it entirely. For the choice of aims model to be successful it needs to be complemented by a set of institutions that: (1) ensure the widespread ability of citizens to cultivate reasoned preferences about the aims that should be pursued in their society (this is going to involve incentivizing the formation of necessary background knowledge and continued engagement in political discourse); (2) enable voters to delegate steps in the decision-making process such that they can be sure about the similarities and differences between the delegates selection principles and their own; (3) promote political discourse that is truth sensitive; and (4) that not only promote the ideals of democracy, but that can be seen by the public to do so.

As we have already discussed newspapers, media outlets, government publications, university research publications, political parties, think tanks, interest group associations, web blogs, social media, and general public discourse all contribute to the dissemination of relevant information throughout democratic societies. The question that we will investigate in the remainder of this chapter is whether, in contemporary democracies, these institutions are able to do this in a way that is consistent with the criteria outlined above.

### *2.5 Information Distribution in Modern Democracies*

Throughout most of the 20<sup>th</sup> century mass media outlets adopted journalistic norms which arose in part as a response to the use of propaganda in World War I.<sup>61</sup> Traditional journalistic norms include that journalists should:

- Seek and report the truth.
- Never deliberately distort facts or context.
- Act independently, maintaining that their highest and primary obligation is to serve the public.
- Deny favored treatment to advertisers, donors or any other special interests, and resist internal and external pressure to influence coverage.
- Distinguish news from advertising and shun hybrids that blur the lines between the two.
- Be accountable and transparent; ethical journalism means taking responsibility for one's work and explaining one's decisions to the public.<sup>62</sup>

Norms like these were sustained by local and national oligopolies who controlled the dominant information distribution technologies of the time, i.e., print and broadcast.<sup>63</sup> However, with the rise of the internet, the decreased cost of entry for new competitors, many of which reject or ignore these norms, has “undermined the business models of traditional news sources that had enjoyed high levels of public trust and credibility.”<sup>64</sup> This creates major problems when it comes to maintaining a set of institutions that meet the criteria set out in the previous section.

The 1993 Peter Steiner comic pictured below captures the idea beautifully. The comic depicts a dog sitting at a computer telling another dog, “On the internet, nobody knows you’re a dog.”

---

<sup>61</sup> Lazer et al., “The Science of Fake News.”

<sup>62</sup> “SPJ Code of Ethics - Society of Professional Journalists.”

<sup>63</sup> Lazer et al., “The Science of Fake News.”

<sup>64</sup> Lazer et al.



**Figure 1:** Peter Steiner's Cartoon.<sup>65</sup>

Whatever playful images this might conjure up in our minds, it deftly points to a serious problem when it comes to the distribution of important information in internet age societies. On the internet nobody knows who you are. As an individual, you can participate in discourse anonymously, you can sabotage conversations, you can make multiple accounts and create (or even contract out) false support for your ideas, and you can censor ideas that run counter to your own.

Take the popular web community Reddit for example: Reddit bills itself as the “front page of the internet,” a play on the idea of the front page of a newspaper, meaning it has the most important stories you will find across the entirety of the web. On the site, individual posts gain prominence based on an algorithm that weighs the number of upvotes it gets from individual user accounts against the number of downvotes it gets from individual accounts. More upvotes results

---

<sup>65</sup> Cavna, “‘Nobody Knows You’re a Dog.’”

in a post getting more prominence on the site and becoming more visible to people who visit it, while downvotes can decrease the prominence of a post, effectively acting as a mild form of censorship for unpopular ideas. The difficulty with this is that there are no restrictions on the number of accounts an individual can create or operate, so if a person were dedicated enough they could create hundreds of accounts and artificially increase the prominence of their own posts. Compounding this, there are subforums on Reddit that exist for the sole purpose of requesting help from others to make your posts more prominent, people agree to upvote another's post in exchange for upvotes on their own post. This is just one example of how the pseudonymity of the internet creates problems of discourse failure in institutions that distribute information throughout modern societies. As long as you are not doing anything flagrantly illegal, you can do and say what you want on the internet because there is a tenuous link between your personal identity and your actions online. Even if your username or IP address gets banned for some particularly unsavory actions this can be easily side-stepped with multiple accounts or virtual private networks (VPNs) that can be used to obfuscate your identity and location in the world.

Recently, the rise of *fake news* has created challenges for news media consumers and the institutions tasked with the dissemination of information throughout society. Fake news can be defined as fabricated information that mimics news media in form, but not in organizational process or intent,<sup>66</sup> and is intentionally generated to mislead readers.<sup>67</sup> Of course, the phenomenon of fake news is not new. Historically, propaganda and rumour mongering have been a familiar part of political discourse. What is new is the ease at which disinformation can be

---

<sup>66</sup> Lazer et al., "The Science of Fake News," 2.

<sup>67</sup> Qayyum et al., "Using Blockchain to Rein in The New Post-Truth World and Check The Spread of Fake News," 1.

generated using artificial intelligence and machine learning, then disseminated using social networks and targeted advertising.<sup>68</sup>

To exacerbate these problems private corporations have unprecedented access to people's private information through pervasive information and communications technologies (PICTs) and can sell it to whatever interested party is willing to pay for it, allowing them to manipulate public opinion more efficiently than ever before. PICT is variously termed ubiquitous computing, pervasive computing, everywhere, and ambient intelligence. These names characterize it well, as it is a technology that is embedded in the devices and software applications that we use daily which operates in ways that usually go undetected by the casual observer to collect data that can be transmitted and used for a whole assortment of future purposes.<sup>69</sup> PICT tracks the web browsing activity of people and the things they do on web-based applications and the data is used under the pretense of improving user experience by tailoring that experience specifically to a user-generated dataset. Data can then be pooled with other existing data to create a huge dataset that reveals otherwise unknown facts about an individual's life, behavior, preferences, and predilections.<sup>70</sup> When it comes to the distribution of political information PICTs are without a doubt at play. As Pimple explains, "the collection of massive amounts of private information makes possible the wide-spread manipulation of the behavior of whole populations."<sup>71</sup>

But seriously undermining a democracy does not even require the manipulation of whole populations. When campaign strategies can target specific individuals in the population with misinformation, disinformation, and fake news, and they can do so with intimate details about

---

<sup>68</sup> Qayyum et al., 1.

<sup>69</sup> Pimple, *Emerging Pervasive Information and Communication Technologies (PICT)*, 2.

<sup>70</sup> Pimple, 2.

<sup>71</sup> Pimple, 3.

their predilections and biases, then energizing or demoralizing just a few hundred people in an important swing state or riding, could mean the difference between winning and losing a close election. Do this several times in several key states or ridings, and it could mean the difference in an election that should have never even been close.

### *2.6 The Missing Pieces for Modern Institutions*

In real-world interactions you build a lasting reputation based on how you behave with others: if you have a history of lying, people learn to be skeptical of the things you say; if you try to pass rumours off as facts, people learn to wait for further confirmation before acting on something you say; if you have a history of making reckless decisions, people learn not to defer to you to make decisions on their behalf. Traditionally, analogs existed for the institutions in society responsible for the distribution of information: if you are a news source that has a track record of espousing extremist perspectives and selecting only data from questionable sources, people who are not interested in those kinds of selection principles can avoid your influence; if you are a reporter who scoffs at widely accepted journalistic norms, you get sanctioned by your professional community and risk the loss of your job or a diminished level of trust with the public. Today, things are increasingly disanalogous. The rise of the internet and the subsequent changes to business models and technologies used in the spread of information in society have brought new challenges when it comes to ensuring that voters have the information required to perform their role in even the most tenable models of democracy like the choice of aims model. Thus far, institutions in democracies have failed to respond adequately.

For voters to be able to choose the overall aims of society democracies need institutions that: (1) ensure the widespread ability of citizens to cultivate reasoned preferences about the aims that should be pursued in their society (this is going to involve incentivizing the formation of

necessary background knowledge and continued engagement in political discourse); (2) enable voters to delegate steps in the decision-making process such that they can be sure about the similarities and differences between the delegates selection principles and their own; (3) promote political discourse that is truth sensitive; and (4) that not only promote the ideals of democracy, but that can be seen by the public to do so. The evidence provided in this chapter demonstrates that there are serious challenges facing the ability of democratically organized societies and their institutions to allow voters to participate in a way that is consistent with the ideals of democracy. As things stand even the viability of choice of aims model of is threatened. Accordingly, significant institutional reform will be required to address the challenges of the modern state and guarantee voters are able to play a significant role in determining the way their societies are organized.

In the next half of this thesis I introduce an emerging information technology called distributed ledger technology and consider its potential to support the types of institutions that would be needed to complement a model of democracy like Christiano's. In chapter three I describe the unique features of the technology, while in chapter four I investigate the relevant ways the technology is currently being used and consider the potential of future applications as applied directly to the challenges that have been described above.



## CHAPTER 3

### The Advent of Distributed Ledger Technology

#### *3.1 Early Ledgers and the Organization of Society*

The use of ledgers to record public and private information is at least as old as recorded history and is one of the earliest known uses of written language.<sup>72</sup> Sometime between 4000 and 3500 BCE in Mesopotamia, as societies became increasingly urbanized and advancements were made to human productivity, a need arose for more sophisticated means of administration.<sup>73</sup> The use of ledgers enabled societies to record information for the management and administration of increasingly complex societies. With this technology, what were previously customs could be made into explicitly codified laws, knowledge could be more efficiently accumulated and passed to people in different places and times, and bookkeeping allowed for the development of complex economies. Furthermore, bureaucratic polities were made possible as governments could formulate formal identities for citizens and keep records on raw materials, manufactured goods, duties, tributes, inventories, and expenditures.<sup>74</sup>

As useful as ledgers are for organizing complex societies, the shortcoming of the technology, when used in these contexts, is that it almost always involves a significant element of trust between multiple parties. It is certainly possible to maintain a ledger without having to trust anyone at all, for example I can maintain a private record of all the furniture I own and this will not involve trusting any third parties or intermediaries. But, when ledgers are used to exchange information between two or more parties, then trust becomes a concern.

---

<sup>72</sup> Goody, *The Logic of Writing and the Organization of Society*, 49.

<sup>73</sup> Glassner, Bahrani, and Van de Mieroop, *The Invention of Cuneiform*, 179.

<sup>74</sup> Fischer, *A History of Writing*, 22.

When third parties or intermediaries are involved trust can be a factor in the record keeping processes at a number of different points: we must trust that information is reported and recorded correctly when first added to the ledger, we must trust that the custodian of the information will keep it secure, and we must trust that the custodian will not use or alter the information in any way in which they have not been given authority. Accordingly, the way we share information (i.e., in ledgers or records) not only allowed us to organize society in ways that were not previously possible, but also influenced the kinds of institutions we needed to support this more complexly organized society. Namely, the organization of society using a medium of information exchange that relies on trust, encouraged the development of institutions that mitigate the risk involved in trusting intermediaries and third parties, i.e., highly centralized institutions and third-party custodians.

## *2.2 Out of the Frying Pan and into the Fire*

Centralized institutions and trusted third parties are designed in a way that mitigates the risks associated with using trust-based ledgers to exchange information. While they can be successful at this, they open up another suite of serious problems; in the following paragraphs I point out four kinds.

First, there is the problem of data misuse; this occurs when a centralized institution or trusted third party uses – or allows another party to use – data for some purpose other than for what they have been explicitly been given consent. For example, this is what happened in the Facebook–Cambridge Analytica data scandal that occurred in 2014-2015, when Cambridge Analytica used data from millions of Facebook user profiles without consent in an attempt to

influence public opinion over elections and referendums.<sup>75</sup> Another example would be if someone who had special access to people’s personal information (such as a doctor or police officer) used that access to investigate their friends or family for non-work related reasons.

Related to data misuse is data leakage (also sometimes referred to as data breach). This occurs when a trusted third party or centralized institution intentionally or unintentionally allows the information they are in custody of to be accessed by people it was not meant to be accessed by. This usually involves some kind of error on the part of the data custodian, often in the form of a failure to properly secure data. A recent example of this is the 2017 Equifax data breach. Equifax is a global data, analytics, and technology company that offers credit and demographic data and services to business as well as credit monitoring and fraud-prevention services directly to consumers.<sup>76</sup> While they do offer other more nuanced services, in essence, they are a third party that acts as an intermediary between individuals or companies and the personal credit information of individual people. In the 2017 data breach, the full names, Social Security numbers, birth dates, addresses of over 143 million individuals were stolen by hackers who exploited a security flaw in the data custodian’s website.<sup>77</sup>

A third problem associated with the reliance on centralized institutions and third party intermediaries is that they could be creating inefficiencies in markets or participating in rent-seeking behaviours if they are profiting from fees that are charged to users for a service that might not need to exist if institutions were organized around a medium of exchange that did not depend so heavily on trust. For example, PayPal is an intermediary that makes money by vouching for the

---

<sup>75</sup> “Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach | News | The Guardian.”

<sup>76</sup> “Equifax | Credit Bureau | Check Your Credit Report & Credit Score.”

<sup>77</sup> Goodin, “Why the Equifax Breach Is Very Possibly the Worst Leak of Personal Info Ever.”

integrity of payments that are made for the purchasing of online goods and services; something that is made necessary by the inherent difficulty of transacting online with trust-based monetary institutions. With our current set of institutions, a consumer's options include divulging credit card information directly to a merchant who they may have little reason to trust, sending a payment without any guarantee about whether they will receive the product they are purchasing, or using a third party like PayPal. The concern noted here is that the fees being paid out to an intermediary are inefficient and could be understood as rent-seeking, i.e., involving the accumulation of wealth while not adding anything of value to the economy. This will become clearer as we progress through this chapter and begin to understand the possibility of trustless peer-to-peer financial institutions.

The final problem that can arise when we rely on third parties or centralized institutions is corruption and the abuse of power. To illustrate this problem, I will point to the 2007-2008 global financial crisis. It is widely accepted that a major contributor to this crisis was the failure of the institutions responsible for the oversight of the financial system; imprudent mortgage lending standards were allowed, regulators failed to do their jobs, institutions acted recklessly and took on too much risk in the pursuit of big payoffs, and opaque shadow banking practices were allowed to grow.<sup>78</sup> The outcomes of their irresponsible institutional practices had profound economic impacts on people around the world that will continue to be felt for years, but might not have occurred if we were not forced to rely so heavily on centralized institutions which could be corrupted or abuse their power.

---

<sup>78</sup> United States, *The Financial Crisis Inquiry Report*, xvi-xxii.

The sections of this chapter up to this point have provided context for understanding the role of ledgers in how our societies are organized and outlined the pitfalls that come about when we are forced to rely on centralized institutions and third parties. With this context in place, the next several sections of this chapter will describe distributed ledger technology as an alternative medium for the exchange of information. This will allow us to move forward with a working knowledge of the technology and understand its implications for how we will exchange information in the future and how it could be used to address the problems facing our existing democratic institutions.

### *3.3 Satoshi Nakamoto and the Advent of Distributed Ledgers*

Most of the pieces of technology and cryptography that enabled the first distributed ledgers existed in the 1980s and people had theorized about how they could be used to enable a digital payment system at that time as well.<sup>79</sup> But commentators argue that the background conditions did not exist for the technology to catch on at that time. Instead, it was not until people's confidence in existing institutions was shaken by the 2007-2008 financial crisis that the idea of decentralized institutions built on distributed ledgers started to gain recognition as something that could be of value to society.<sup>80</sup> Thus, it is argued that it is no mere coincidence that the Bitcoin whitepaper was published in late 2008 by the pseudonymous individual or group known as Satoshi Nakamoto. In it, Nakamoto describes the world's first peer-to-peer electronic cash system and the technology that makes it possible, a technology that would come to be known as blockchain technology due to its system of linking blocks of timestamped data together in a

---

<sup>79</sup> Narayanan, *Bitcoin and Cryptocurrency Technologies*, 2016.

<sup>80</sup> Narayanan.

continuous chain.<sup>81</sup> Four months later, in January 2009, the bitcoin network went live, and the world had its first distributed ledger.

The founding innovation of blockchain technology is that it enables the creation of a database (or ledger) of information that can be shared across a network of computers in different locations without the need for centralized network administration or centralized data storage. If we think of this in comparison to a traditional paper ledger or even a modern centralized digital database, each of these older information technologies requires some kind of central authority or administrator in charge of what is added to the ledger, as well as how the information is stored, who can access or make changes to the database, and how the network is secured.


Above we discuss some of the pitfalls of relying on centralized institutions or third parties to enable the transfer of information between transacting parties in a society organized with mediums of information exchange that are trust based. In the following section we make similar consideration specifically regarding information databases. Consider how the average university or large workplace runs their information databases; typically, users are given credentials that allow them to access the network, once logged-in they can utilize the network according to whatever permissions they have been given by a network administrator, which might include the ability to upload files to a centralized database for storage purposes. Those files are stored on a central server that is maintained and administered by an Information Technology (IT) department.

This way of organizing an information database is great if you have a centralized authority like a university or bank or an employer who wants to maintain ultimate control over what can be done on the network and restrict who can do it. Because the associated IT department (who is

---

<sup>81</sup> Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System."

employed and given its mandate by the central authority) gets full control over the access permissions of individual users and over the files that are stored in the database, the network is used only in the way that the people who own it want it to be used.



**Figure 2:** Difference between centralized and distributed networks.<sup>82</sup>

But there are also some downsides to the centralized model of information databases. First, they tend to have only one or very few points of failure. This means that the network is vulnerable to certain types of cyber-attacks. This is why institutions like government departments, universities, corporations and hospitals have increasingly been the targets of cyber-attacks that look to disable the central servers and hold them ransom. But perhaps more important to this thesis, a major weakness of these types of information databases is that they involve major discrepancies in power and they require a relationship of trust to exist between users and the central administrators. While someone might be fine trusting that their employer will not intentionally delete files that are integral to the execution of their job, other arrangements that involve centralized information databases may be more precarious. For example, if I am a bank

---

Adapted from an image found in, Baran, "On Distributed Communications," 2.

customer using an automated teller machine to send a monetary transaction from one account to another in a country experiencing a financial crisis, I may not relish being in the position of having to take it on faith that the money will arrive where it supposed to.

Of course, the alternative to the centralized model discussed above is one without any central administration. Rather than having a client-server model where we have client nodes (computers) on the network being served by a categorically different server node with administrative permissions, every node in the network that supports a distributed information database would have the same network permissions. That is to say, on a distributed network, all the computers can participate in all the same aspects of the network as any other computer on the network in all the same ways. Accordingly, there is no central authority that controls the network, there is no central point of failure that is vulnerable to attacks (in the distributed system, if one node goes down all the others keep on running without issue), and there is no need for any prior relationship of trust between parties that are utilizing the network. This final feature of a distributed information database will become more apparent once the inner workings of the technology are better understood.

Prior to the invention of distributed ledger technologies in the form of the Bitcoin blockchain in 2009, there was no way to achieve a shared information database that did not rely on a centralized administrator. In fact, without understanding how blockchain technology works, it may be hard to even conceive of how an information database without a central administrator would work. How would you guarantee that the information had not been tampered with after being added to the database? How would you ensure that things being added to the database are actually what users are trying to add to the database? Who would be in charge of securing the information and administering the network and how would we know that they are doing what is



in the best interest of the network? There are three key technical hurdles that had to be overcome to make these things possible: guaranteeing the immutability of information on the ledger, achieving distributed consensus, and administering a distributed database. I will deal with each of these in turn in the following three sections.

### *3.4 Guaranteeing the Immutability of Information on the Blockchain*

One of the major hurdles when conceiving of a database with no central authority is the concern that once information is added to the database, one of the many nodes on the network (all of which have the same permissions or capabilities) could access data that had been added to the ledger and alter or remove it. In a traditional centralized network this would not be possible, as under this kind of arrangement there would be a central administrator who can override the actions of other network participants or unilaterally lock certain users out of certain network capabilities. Thus, one of the primary problems that must be solved before a distributed ledger can be achieved is how to guarantee the immutability of content once it is added to the ledger.

In the Bitcoin whitepaper, Satoshi Nakamoto describes a system that uses cryptographic hash functions to link blocks of data in an ever-growing chain so as to prevent the altering of data stored on the ledger.<sup>83</sup> A hash function is a mathematical algorithm designed to convert a piece of data of any size (an input) into a string of letters and numbers of a standardized size (called an output, digest, or hash).<sup>84</sup> Importantly, these are one way functions that are infeasible to invert.<sup>85</sup> If input  $X$  is put into hash function  $Y$ , then  $Y$  will always produce the exact output  $Z$ . However, if

---

<sup>83</sup> Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 1–3.

<sup>84</sup> Balasubramanian, Rajakani, and IGI Global, *Algorithmic Strategies for Solving Complex Problems in Cryptography*, 67.

<sup>85</sup> Balasubramanian, Rajakani, and IGI Global, 67.

any change whatsoever is made to  $X$ , then the hash function  $Y$  will produce a radically different output  $Z^1$ .

Input		Output/Hash
McMaster Philosophy		5266 668B BA23 FFF1 F9D9 BF31 EC39 B6DC
MacMaster Philosophy		F149 70AE FA33 BDE3 D566 11A2 23A0 0E0E
mcmaster philosophy		5787 7B69 1DB9 E38F FF72 A867 00F9 2984
Mac phil		FE1B B3D3 9324 F971 B439 9B19 0BEF E584

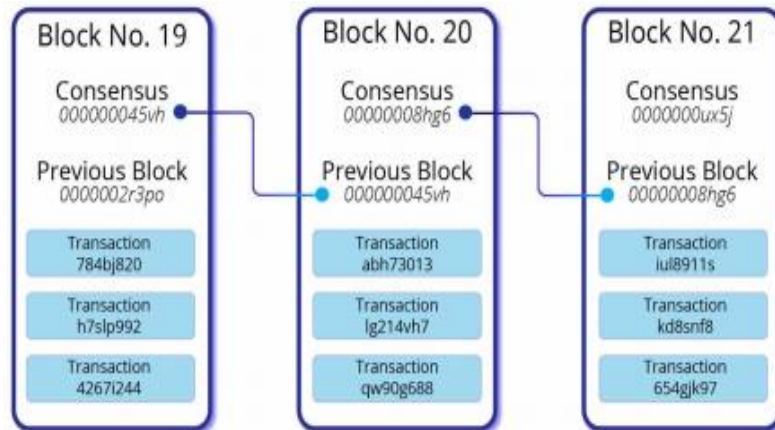
**Figure 3:** A small change to the input results in a drastic change to the hash.<sup>86</sup>

When a user wants to add something to a blockchain ledger they broadcast it to the other nodes on the network. At regular intervals data that has been broadcast to the network is grouped together in a block and timestamped in accordance with when it was added. All the data in a block is used as the input and ran through a hash function thereby producing a string of letters and numbers of a standard size as an output.<sup>87</sup> The output (or hash) of all the data is published with the block and will be referenced by the next block in the chain. Thus, if anything in a published block were to be altered at a later time it would result in a radical change in the output of that block’s hashed data and would no longer match up with the referent in the subsequent block, causing a chain reaction throughout the blockchain that would be infeasible for an attacker to try to cover up and would be rejected by the network.<sup>88</sup>

<sup>86</sup> These hashes were generated at: “Online Hash Calculator — String & File Hash Generator.” For the sake of style, the second half of each hash has been omitted.

<sup>87</sup> Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” 2–3.

<sup>88</sup> Nakamoto, 3.



**Figure 4:** hashes, labelled here as “consensus”, link blocks of data in a blockchain.<sup>89</sup>

### 3.5 Achieving Distributed Consensus

The linked hashes prevent the data in a blockchain ledger from being tampered with after the fact, but how do we know that the data being added in the first place is actually representative of what users are broadcasting to the network? That is to say, what is stopping people from tampering with data before it gets immutably published? Somehow the network must be able to come to consensus about what belongs on the ledger.

To accomplish this, blockchains incorporate a consensus mechanism, i.e., a computing process that allows the network to come to agreement on what is supposed to be added to the ledger. In normal circumstances, there could always be individual incentives that compel people to try and get a different set of data added to the ledger than what was actually broadcast by network users. The most obvious example being a ledger recording monetary transactions between people, like the Bitcoin ledger. Even a traditional bank account balance is just an entry on a ledger controlled by a financial institution, the Bitcoin ledger is just the same, but without the centralized institution. If an attacker could get the network to publish a data set that reflected

---

<sup>89</sup> “ACT-IAC White Paper,” 3.

that they were sent a large amount of money, when they were not actually, then (for all intents and purposes) they would have that money and there would be no administrative authority who could go back and correct the error. To avoid this, distributed ledgers use consensus mechanism to realign the incentives of network participants so that the only worthwhile actions they have on the network are honest ones, i.e., ones that work towards publishing data to the ledger as it was actually broadcast by users. It is this realignment of incentives that removes the need for a pre-existing relationship of trust between parties. Blockchains are trustless because they align incentives such that an individual using the network knows that acting in an untrustworthy way using this medium of information exchange is infeasible. We explain the technical aspects of how this works next.

The most common and first consensus mechanism to be used in blockchain is proof-of-work (PoW). This consensus mechanism requires that at regular intervals nodes on the network that elect to participate in the block confirmation process (also called miners) gather together all the information that was broadcast to the network and compete against one another to solve a complex mathematical puzzle which is difficult to solve but easy to verify, all for the right to be the publisher of the next block of data added to the chain.<sup>90</sup> These puzzles are so difficult to compute that it actually takes a non-trivial amount of energy (in the form of electricity) for a computer to solve them. Whatever node is first to solve the puzzle gets the opportunity to propose the next block of data that will be added to the blockchain and, if successful, is rewarded with the blockchain's native cryptocurrency, e.g., a pre-set number of bitcoins in the case of the Bitcoin

---

<sup>90</sup> Tapscott and Tapscott, *Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business and the World*, 31.

network. Thus, we get the name “proof-of-work”, because a node must prove they did the work to solve the puzzle in order to receive the reward for publishing a block on the blockchain.

However, before the puzzle solver’s proposed block is added to the blockchain (and before they get the corresponding reward), other nodes on the network perform a process akin to checking the work of the new block proposer. To do this they compare the hash of the data that they found being broadcast to the network to the hash of the data in the block being proposed by the puzzle winner. If there is a discrepancy between the data that the rest of the network found being broadcast and what is in the proposed block, then it will be obvious that something in the proposed block is wrong and it will be rejected. At that point, the miners will return to another round of puzzle solving until they come to consensus on the new block.

In summary, here are the simplified steps for adding data to a proof-of work blockchain:

1. Data is broadcast to the network by users
2. Data is grouped in blocks based on when it is broadcast to the network
3. Miners compete to solve a highly-complex mathematical puzzle, this takes a lot of energy (in the form of electricity) and is costly to miners
4. The first miner to complete the puzzle gets the opportunity to propose the next block that will be added to the blockchain and receive a reward for doing so
5. The other miners in the network confirm the block by inputting their data set into the hash function and confirming that they get the same hash as the proposed block. Here there are two possible outcomes:
  - i. If their data set returns the same hash, this means the data in the proposed block is identical to the data in their own block, so they confirm the validity of the block and it is added to the blockchain
  - ii. If their data set returns a different hash, this means the data in the proposed block is different than what has actually been broadcast to the network, so the proposed block is rejected, and miners return to step 3

While this sounds fairly complicated, it is simply a mechanism designed to rearrange the incentives of network participants so that they only act in accordance with what is best for the network. Because miners have to expend resources to solve the puzzle and because the network will be able to easily catch someone altering the data that is to be added to the ledger, there is

actually a cost that will be incurred to try to behave dishonestly and it will have no chance of working. Rather than trying to lie, which will never be successful, one would be better off not participating at all, or better yet, participating as an honest node and having the opportunity to earn the reward for proposing the correct block.

In summary, blockchains are made immutable by linking chains of blocks of data with hashes and they come to consensus on what should be added to the blockchain by using a consensus mechanism that realigns incentive such that the incentive of individual users is aligned with what is good for the network as a whole. While it should be noted that there are other consensus mechanisms that achieve this in different ways, they have the same goal and the same basic strategy of incentive realignment. The final step to understand the basics of how a blockchain achieves a distributed ledger is understanding how the blockchain is administered.

### *3.6 Administering a Distributed Ledger*

In a traditional centralized network, things like security upgrades and changes to software are conducted by the associated centralized administration, e.g., a company's IT department can install new software on computers on the network, implement new security measures, and perform maintenance as needed. This is a necessary operation in a computer system as the environment the computers are operating in is always changing and critical errors in programs can be found that may need to be fixed. In a distributed ledger there is no central authority that can make unilateral changes to the nodes on the network, nonetheless things like security updates and removal of bugs still need to be implemented.

To understand how this works on a distributed ledger it is important to understand that what makes a computer a node on a given distributed network is simply that it runs the latest version of a piece of software that executes the processes necessary for the network to operate.

Thus, if you are running a node on the current bitcoin network as of May 5<sup>th</sup>, 2019, you are running a piece of software called *Bitcoin Core Version 0.18.0*.<sup>91</sup> For a change to be made to the network (e.g., a bug-fix, a security upgrade, or a change in one of the operational processes that dictates how the Bitcoin network operates), people in the Bitcoin community can propose changes and independent software developers can implement these changes by making changes to the next version of bitcoin software. This whole process is done in an open-sourced, community driven way (as there is no one who owns bitcoin); so other community members review the changes, check for bugs or other problems, and discuss the impact the changes would have on the network.

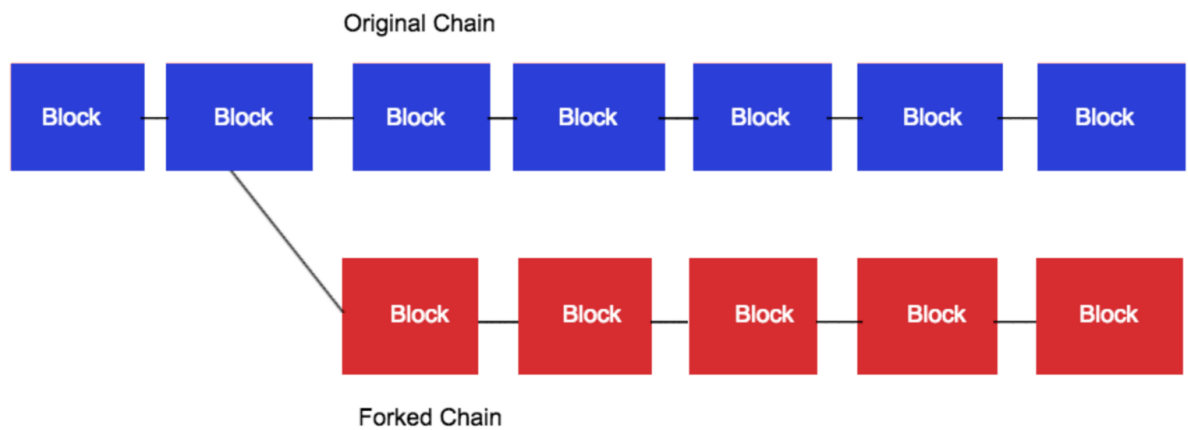
When a change is simple, like a small bug fix, people typically review it and then the new version of the software is released. In order for the changes to take effect, the nodes on the network simply need to install the new version of the software and then they continue operating as they used to (but with the new changes in place). However, on occasion, there can be a divide in the community if a controversial change is proposed to the network. Let's say for example that it is proposed that blocks of data on the blockchain need to be made smaller so that transactions can be verified more quickly. There could be one part of the community that wants this change to be made because it will increase the utility of the network (faster transactions are more useful than slower transactions), however there could be a second part of the community that wants to conserve the network in the way it was originally designed by Satoshi Nakamoto. Such a divide will be resolved by two processes.

First, miner nodes on the network will accept or reject the changes by deciding whether or not to install the version of the software that contains them. Technically speaking, if any portion

---

<sup>91</sup> "Version History - Bitcoin Core."

of the network fails to install the new software and thereby does not implement the changes, then they will affectively create a separate blockchain that continues forming blocks according to the processes set out in the old software. Meanwhile, the portion that accepted the changes and ran the new software will be forming blocks according to the new processes. If this occurs, the blockchain forks; where we used to have one continuous chain of blocks with a completely shared history, we now have a chain of blocks that share their history up to the point where the network was divided, then a fork at the point where some people implemented the new software and others did not, and two distinct chains with two now distinct native currencies, continuing on from that point.



**Figure 5:** A fork occurs when some nodes update to new software and others do not.<sup>92</sup>

A fork will occur even if a small subset of the nodes inadvertently fails to update their software. But usually when a fork occurs (whether inadvertently or due to disagreement), this is resolved based on community sentiment. While there could be a divide among the network’s miners about which is the true blockchain, this is often resolved according to the recognition of the greater community about which side of the fork represents the direction the network’s users

---

<sup>92</sup> Ray, “Blockchain Forks.”



want the network to go. If people do not use one of the forks to transact on it will quickly become defunct. Consider our earlier example: *side A* of the fork made changes because they wanted the blockchain to process transactions more quickly, something that gives the blockchain more utility, while *side B* wanted to conserve the original vision of the creator of the technology. If, when users go to transact on the network, they mostly care about how the network performs, and not about the vision of an anonymous creator, then they are likely to transact on *side A*. If not enough people are using *side B*, then the miners on that network will be expending energy to receive a reward (*side B's* native currency), that may not be valuable enough to warrant doing the work. If this happens *side B* will quickly become defunct, as miners will give up mining a cryptocurrency that is not valuable. In the rare case that both sides of the fork continue to be used then we will simply have two distinct blockchains that, while sharing a history (or set of blocks) up to a certain point, are now completely unique; it will have its own unique set of mining nodes and its own unique set of data processes.

So, the process of administration in blockchain is an open, community-based process. Anyone can propose changes and anyone who can write the code can participate in coding the changes into a new version of the blockchain software. Miners then consent to the changes by running the new version of the software. In the rare case that there is widespread disagreement between miners, this is resolved by public opinion about what set of processes people want to use for whatever use-case they have in mind. In even rarer circumstances, if there is enough support for both chains, this can ultimately result in two separate blockchains from that point on.

We have now covered the basic elements of blockchain technology that allow it to achieve a successful distributed ledger: data is added to the blockchain according to a consensus that is achieved by realigning incentives of network participants so that they always have reason to be

good actors on the network and never bad ones, once data is added to the ledger it is protected from tampering through the use of hash functions that link data and easily detect alterations, and the blockchain is administered in an open community based way that allows the community to consent to changes at multiple levels. Before we continue to the final chapter of this project and begin applying this new medium of information exchange to democratic institutions, I want to add a couple more recent developments in distributed ledger technology to our suite of knowledge. These are the innovations made possible by blockchain 2.0 and the distinction between private and public blockchains.

### *3.7 Ethereum and Blockchain 2.0*

The creation of the Ethereum blockchain represents a watershed moment in the history of distributed ledger technology. In 2013 Russian-Canadian Vitalik Buterin proposed the idea of a blockchain that could “specify any functionality that is possible to program into a turing-machine, an abstract model of a computer that is believed to be capable of computing any function that can be computed at all.”<sup>93</sup> This means that (given enough time) it could execute any computer process that could be written in any modern computing language whatsoever. This represented a massive increase in the sorts of processes that could be done on a blockchain, which, prior to Ethereum, were limited to a fairly small instruction set.<sup>94</sup>

Ethereum ushered in a second wave of blockchain projects that allowed for much more complex data to be stored on the blockchain and opened up the possibility of two important types of blockchain use-cases: distributed applications (d-apps) and smart contracts. Distributed Applications are computer applications that run across a distributed network of computers rather

---

<sup>93</sup> Narayanan, *Bitcoin and Cryptocurrency Technologies*, 2016, 263.

<sup>94</sup> Narayanan, 263.

than on a single computer or server. This allows for the creation of applications just like any other computer or smart phone application (e.g., google chrome, twitter, a banking app etc.), except for the difference that they can take advantage of the features of a distributed ledger (e.g., decentralization, trustlessness, immutability etc.).<sup>95</sup> Smart contracts, on the other hand, are computer programs that secure, enforce and execute the settlement of recorded agreements between people or organizations, while minimizing the need for trusted intermediaries.<sup>96</sup>

Think of a smart contract as a computer coded contract between two or more parties that will automatically execute when certain conditions are met. There are a wide range of ways smart contracts can be used, but a simple example would work like this: say I finance the sale of a car to someone, we could have a smart contract in place that integrates an internet of things (IoT) device that is built into the vehicle's key fob and/or ignition, the smart contract could be coded so that if a previously agreed upon car payment is not received on a predetermined day of each month then the smart contract automatically disables the integrated IoT device that allows the vehicle to start. If you wanted to allow for a 7-day grace period or an allowance of one late payment per year or any other specification that could be written in a computing language, this would all be possible. The complexity of the smart contract is really only limited by what one can imagine and what one can code. Among other things, smart contracts have already been used to create online betting (or prediction) communities, complex decentralized organizations, and to facilitate loans between strangers; all without the need for a central authority or third party.

---

<sup>95</sup> Tapscott and Tapscott, *Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business and the World*, 117–25.

<sup>96</sup> Tapscott and Tapscott, *Blockchain Revolution*, 101.

### *3.8 Public and Private Ledgers*

Up to this point we have only discussed public blockchains. Most of the earliest blockchains were public blockchains, anyone was allowed to participate in all aspects of the blockchain without any special permissions or authority given to a single party or subset of the user base. More recently various institutions, organizations, and industries have started to recognize the value proposition of distributed ledgers. However, many of the use-cases these organizations have in mind do not lend themselves to completely permissionless networks, because of this there has been a marked increase of interest in private and semi-private distributed ledgers.<sup>97</sup> Private (or permissioned) blockchains maintain many of the same benefits of traditional distributed ledgers, but they limit access to some or all of the network permissions to people according to what they want them to be able to do or see on the network.<sup>98</sup>

Some of the reasons an organization might want to utilize blockchain as opposed to a traditional centralized database would include if they wanted the information on the ledger to be visible and transparent to the public, while guaranteeing that records would not be tampered with, but they only wanted their own staff to be able to add information to the ledger. In another scenario an entire industry might have reason to have a shared database where no single party has more authority than another. Think of a group of stakeholders in the pharmaceutical supply chain, you might have manufacturers, shippers, merchants, consumers, and regulators all who have a stake in information about the supply chain and who also do not necessarily have a strong relationship of trust; a distributed ledger could be the ideal solution, but a completely public one may not. These are some of the areas in which private or semi-private blockchains could be the

---

<sup>97</sup> Tapscott and Tapscott, 67–68.

<sup>98</sup> Tapscott and Tapscott, 67.

answer. When it comes time to discuss the use of distributed ledgers in democratic institutions in the following chapter, it will often be the case that a private or semi-private ledger is the best suited solution.

One final clarification that is necessary for a proper understanding of this technology is a terminological one. The first type distributed ledgers were made using blockchain technology, it is still by far the most common and well-tested type. However, recently developers have begun experimenting with alternative strategies for achieving ledgers without central authorities, ones that should not be categorized as blockchain, but share many of the same features and deal with the same kinds of trade-offs in achieving distributed ledgers. The grouping of all these different types of technologies (including blockchain) is called Distributed Ledger Technology (DLT). We will use this more general term throughout the remainder of this thesis, except for when what is being said refers specifically to blockchain technology.

This concludes our introduction to the technical aspects of distributed ledger technology. Distributed ledgers are a new medium for the exchange of information, they offer a possibility that could not ever be realized previously; the ability to exchange trustworthy information without having to rely on any intermediaries. In the next chapter we investigate how this new medium can be brought to bear on the problems faced by democratic institutions.

## CHAPTER 4

### **Distributed Ledgers for Defensible Democracy**

The choice of aims model is more tenable than other existing models of democracy. It strikes a fine balance between decreasing the epistemic burden placed on voters, while ensuring they play an integral role in the formulation of the terms of association for their societies. However, a good model of a system of social organization needs to be complemented by a set of institutions that can support it. In particular, even with a decreased epistemic burden, voters need institutions that do the following:

1. Incentivize the formation of the background knowledge and continued engagement with political discourse necessary to make reasoned preferences about the aims of their society;
2. enable voters to delegate steps in the decision-making process such that they can be sure about the similarities and differences between the delegates selection principles and their own;
3. promote political discourse that is truth sensitive; and
4. not only promote the ideals of democracy but can be seen by the public to do so.

Changes to the way information is distributed in contemporary democracies, mostly due to the rise of the new information technologies and the ways they can be used to manipulate public opinion, have eroded the ability of existing institutions to meet these criteria. However, by working with the technical understanding of distributed ledger technology (DLT) that was established in the previous chapter and referencing projects currently being developed in this space, we will provide a sketch of a set of institutions that could satisfy the conditions outlined above. Ultimately, I argue that DLT powered institutions can be used to complement Christiano's choice of aims model of democracy, forming a viable account of democracy.

#### *4.1 Modern Identity for the Modern State*

It seems sensible that the first step in designing a set of DLT-based institutions that could support a model like Christiano's would be to establish a system of universal digital identity for citizens that would enable them to engage with their other institutions. In most modern democracies a person's identity (as far as the state is concerned) is made up of a highly fragmented set of government issued physical documents with corresponding entries in a number of government databases. Even though identity systems like this are usually poorly organized and have a low degree of interoperability, they allow people to interact with their state and its various institutions (and vice versa). Our identities are precursors for our ability to interact with the state; among other things, they allow us to receive nationalized healthcare, enroll in state sponsored education, open a bank account, participate in elections and referendums, pay taxes, and receive various social benefits. The information contained in our identities dictates how we will be received by the various institutions we interact with, e.g., a person's age affects their eligibility to receive a retirement pension or to buy certain controlled substances, their address could affect the way they are taxed by the government, and their credit history could impact the way they interact with financial institutions. A person's identity, as far as state recognized identity is concerned, is a set of entries in one or more databases that dictates the ways in which they will be recognized by a state (and its various institutions) and tracks the ways the two bodies are eligible to interact with each other. Identity plays a foundational role in the mediation of the relationship between an individual and their state. Likewise, digital identity based on DLT will form the foundation of DLT powered institutions.

Over the last decade the country of Estonia has become widely regarded as a world leader in digital governance, an important part of this is their unique approach to identity.<sup>99</sup> Estonian citizens are issued a mandatory<sup>100</sup> ID-card that has an embedded microchip; it “holds basic information about the card holder as well as two certificates – one to authenticate identity and one to provide a digital signature – and a personal identification number (PIN) of their choice.”<sup>101</sup> Citizens then use these cards to vote, submit their taxes, access banking services, travel throughout the European Union, apply for social security benefits, review medical records, access national health insurance, use public transit, and even for students (and parents) to access school assignments and report cards.<sup>102</sup>

Why is Estonia comfortable and willing to host so much of their institutional infrastructure through a single digital apparatus? The system is built on blockchain technology; all electronic activity is verified mathematically on a blockchain without system administrators or government staff.<sup>103</sup> Recall from the previous chapter, blockchains link blocks of data together in a way that prevents alteration of data once it is added to the ledger and makes it easy to audit specified actions that are made on the network. This gives the Estonian system total transparency and accountability; stakeholders can see who accessed what information, individuals can verify the integrity of their own records without the need for a third party, and once something is recorded it cannot be manipulated by anyone including government staff.<sup>104</sup> The same technology that makes it possible for the

---

<sup>99</sup> Tapscott and Tapscott, *Blockchain Revolution*, 197.

<sup>100</sup> “ID-Card.”

<sup>101</sup> Tapscott and Tapscott, *Blockchain Revolution*, 198.

<sup>102</sup> “ID-Card”; Tapscott and Tapscott, *Blockchain Revolution*, 198.

<sup>103</sup> “KSI Blockchain.”

<sup>104</sup> “KSI Blockchain”; Tapscott and Tapscott, *Blockchain Revolution*, 199.



bitcoin network to transfer hundreds of millions of dollars of value daily, also makes it possible for Estonia to provide transparent and efficient services to their citizens centered on digital identity. Estonia regained its independence from the Soviet Union in 1991. It could be that the experience of being annexed by a state that was infamously opaque and manipulative of information is part of the reason why the country has built a system of institutions that prioritizes transparency, accountability, and reduced bureaucracy. It could also be that the small democracy gives us incredible insight on how to address the problems undermining democracy around the world.

As was discussed previously, a major problem for modern democracies is a lack of transparency and accountability in information distribution and political discourse that takes place on the internet. Internet use has become an increasingly important part of the day-to-day lives of citizens in democracies. People are banking, becoming educated, getting news, and socializing online at an increasing rate. But the internet was built without an identity layer, i.e., there is no plainly available connection between an individual and their activity online. Without this essential layer, websites offering services have had to offer workarounds that act in place of a pervasive system of identity, ultimately resulting in a patchwork of identity one-offs.<sup>105</sup> The workarounds for the internet's missing identity layer have become commonplace in our internet use and are treated as a normal part of the user experience, but in addition to the issues they create for democratic institutions, they also create caches of private data primed for theft and abuse from hackers. In response to the problem of identity on the internet, Kim Cameron, former Chief Architect of Access for Microsoft, wrote:

---

<sup>105</sup> Cameron, "The Laws of Identity," 1.

“Hundreds of millions of people have been trained to accept anything any site wants to throw at them as being the ‘normal way’ to conduct business online. They have been taught to type their names, secret passwords and personal identifying information into almost any input form that appears on their screen. There is no consistent and comprehensible framework allowing them to evaluate the authenticity of the sites they visit, and they don’t have a reliable way of knowing when they are disclosing private information to illegitimate parties. At the same time, they lack a framework for controlling or even remembering the many different aspects of their digital existence.”<sup>106</sup>

These concerns led Cameron to believe that a “unifying identity metasystem” was needed for the internet.<sup>107</sup>

By extrapolating on the advances made to identity in Estonia we can imagine a system that solves the issue of the missing identity layer of the internet, while at the same time creating the necessary foundation for a system of democratic institutions built on distributed ledger technology. People would use some combination of existing state issued identification (e.g., passports and driver’s licences), unique biomarkers (e.g., fingerprints and iris scans), and randomly generated numbers that only they know, to produce a unique cryptographic key. The relationship between the individual – by way of their state issued identity and biomarkers – and this digital identity would be established on a distributed ledger. The cryptographic key could then be stored on a digital storage device, an ID card like the Estonian one, and/or a mobile device like a smart phone and protected with PIN in each instance. This digital signature would act as a universal identifier, people would no longer need a separate set of login information for every web service, they would have a single digital identity that was immutably linked to their physical identity and integrated into their interactions with the other institutions in their society, similarly to as in the Estonian system.

---

<sup>106</sup> Cameron, 1.,

<sup>107</sup> Cameron, 1.

An identity metasystem could be built without using DLT. In fact, most people who use computers regularly have some experience with a couple early attempts at this. Namely, Google and Facebook are acting as identity protocols when external online services or web applications prompt you to login using your Google or Facebook account. The problem with these identity metasystems is that they are centralized, meaning they have the shortcomings characteristic of centralized systems like those discussed in chapter 3, e.g., the need for trust between the centralized body and its users, security risks associated with centralization, and power imbalances between the central authority and users. An internet identity layer built on DLT would have no central authority thereby avoiding the pitfalls of centralization. Moreover, a DLT-based identity meta layer could be easily integrated with a broader set of institutions built on the technology.

Currently there are several start-ups and large corporations developing projects like what I describe above. A partnership between Factom Inc. and IPRD Solutions Inc. carried out a pilot project that was funded by the Bill and Melinda Gates Foundation,<sup>108</sup> which saw the use of biomarkers to create personal health information records built on blockchain technology for patients in South Africa so that they could have a single health record that could be used without impediment across a network of previously unintegrated healthcare facilities.<sup>109</sup> A company called WISEKey, based out of Switzerland describes their WISEID system as, “an open source software, audited by network participants and enabled by several types of biometrics [with] a decentralized digital Identity... Users control their ID and have the equivalent of a digital passport, able to build reputations across Web3.”<sup>110</sup> In addition to these, big players like IBM, Microsoft, The Rockefeller

---

<sup>108</sup> “Awarded Grants.”

<sup>109</sup> Texas Bitcoin Conference, *Hardy Helburn*.

<sup>110</sup> SA, “WISEID.”

Foundation, and GAVI the Vaccine Alliance, are all working on projects related to the development of DLT-based identity.<sup>111</sup>

#### *4.2 Truth Sensitivity and Transparent Selection Principles for Online Discourse*

The decentralized nature of DLT makes it well suited to balance the interests of multiple stakeholders across multiple jurisdictions, as there is no central authority who has the ability to abuse a position of power or give their interests more consideration than others. While DLTs are typically very open, facilitating transparency and accountability, they also offer incredible customization in how they are designed. Because of this, they can be open and accountable in some areas, but made to be incredibly secure and private in others. Thus, it would be a mistake to think this means surveillance or the erosion of personal privacy in online activities. Instead, a DLT-based digital identity would extend the features of our real-world identities to our online interactions.

In the real-world, people have private lives and public lives. In their private lives they should be able to expect a certain degree of privacy, many people view this as a fundamental human right.<sup>112</sup> On the other hand, certain public actions are not deemed to fall under the protection of a right to privacy. Accordingly, if you were to send a letter to your friend or family member, you would expect that (under normal circumstances) it would not be read by anyone along the way. Likewise, if you had a private conversation in your home with your partner, it would be reasonable to expect that no one would be listening in. On the other hand, many actions that are carried out in the public sphere carry no such expectation. In fact, for many it would be counter

---

<sup>111</sup> “ID2020 | Home”; “Blockchain for Digital Identity - IBM Blockchain”; “Decentralized Identity – Own Your Own Identity.”

<sup>112</sup> For arguments grounding privacy as a fundamental right see: Çinar, “The Right to Privacy in International Human Rights Law” or Najeeb Alamuor, “Right to Privacy in Law (Sanctity of Public Law).”

intuitive to think they would; for example, if a person published an opinion piece in a newspaper or made a political campaign speech.

Blockchain based identity has been described as “a black box of identity.”<sup>113</sup> The idea is that once your digital identity is created, you control the parameters of who can access the information contained in it and what exactly they access. This is made possible by a cryptographic protocol known as zero-knowledge proofs. Imagine you had a safe with a wheel combination lock, and your friend asked you to prove you knew the combination. One way would be to tell your friend the combination and have them open the lock themselves. But maybe you do not want this friend to know your safe combination, yet you still want to prove *you* know it. You could accomplish this simply by entering the combination without them seeing the numbers and then showing them the open safe. In effect, the lock acts like a mechanical third party that verifies your knowledge of the combination without requiring you to divulge it to your friend. Zero-knowledge proofs are a cryptographic version of the lock, i.e., a protocol that acts as third party which can verify facts. However, instead of trusting a person or group with this role, zero-knowledge proofs require trust only in mathematics. In terms of a DLT-based system of identity this cryptographic protocol allows people to prove certain details related to their identity without revealing any additional information about themselves.<sup>114</sup>

Accordingly, under a DLT-based system of identity, if you want to share some piece of information with another party, you can do so. But revealing one piece of information does not mean divulging everything about yourself. If you need to verify you are the age of majority you

---

<sup>113</sup> Tapscott and Tapscott, *Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business and the World*, 14–16.

<sup>114</sup> For more information about the mathematics of zero-knowledge proofs see: Fagin, Naor, and Winkler, “Comparing Information without Leaking It”; and Huberman, Franklin, and Hogg, “Enhancing Privacy and Trust in Electronic Communities.”

can do so, but you could do so without even having to reveal your actual birthdate. You could verify that you have an accreditation that indicates expertise on a particular subject without revealing your name. If you were going to vote in an election, you could prove you were yet to cast a vote and that you were eligible to vote in a particular jurisdiction without revealing any more details about your identity (and you could later verify that your vote was recorded and counted properly). In many ways, distributed ledger-based identity affords people better privacy and more nuanced control over their personal identities. We will return to a discussion of how DLTs impact the electoral process, but for now I would like to turn our attention to how a system of identity like this would affect information distribution and political discourse on the internet.

DLT-based identity allows us to create an identity layer for the internet. By doing so we can extend features of real-world identity to our online interactions. Chief among these is reputation building and accountability for the things one says and does online. Traditionally, news media outlets operated on business models that encouraged the building of reputations as reliable sources, adherence to journalistic norms, and accountability for the things they published. As we discussed in chapter 2, diminished barriers to entry and anonymity on the internet have eroded these norms resulting in the easy spread of information designed for the express purpose of misleading people. Moreover, the lack of an identity layer allows people to create multiple identities and thereby generate the appearance of support for their own ideas or the pretense of widespread opposition to rival ideas. Establishing an identity layer for the internet and linking it immutably with our physical identities would address these problems. It would no longer be the case that, “on the internet no one knows you’re a dog.”

Importantly, this can be done while respecting the distinction between public and private actions. Just like in our real-world interactions, there are online interactions that are public and

others that are private. Sending an email to a friend is clearly a case of a private action, so is researching a subject of interest, or a person's internet browsing habits. On the other hand, publishing an online news story and advertising are clear cases of public actions. Of course, there are cases that may not be as clear as these. Just like we have disambiguated matters of what ought to be counted as public and what ought to be counted as private in difficult real-world cases, these decisions will be a matter of social decision-making processes.

DLT-based identity can make online discourse more truth sensitive and make the selection principles of various information sources more transparent by ensuring that those things that are considered public actions (like publishing news stories and advertising) can be traced to their source. As a citizen who is trying to gather information for political decision-making, one would be able to clearly identify the source, look into their reputation, and even search for the other public actions that have been made by this particular individual. This would give people a clear idea of the biases of potential information sources to which they might delegate steps in their decision-making process to and allow them to make deliberate decisions about where they want to get information from. On the other side, parties who disseminate information throughout society will be held accountable for how they do so, the spread of fake news and disinformation would be easily traced back to the individuals and groups responsible for it, and these actions could leave a permanent impression on their identity. Something that they could not undo by simply creating a new identity in an opaque system of identity one-offs like we currently have on the internet.

In terms of maintaining truth sensitivity in political discourse (i.e., that people's political positions should be arrived at through a process that considers the best available reasons, understood as those that define the state of investigation in the relevant scholarly disciplines),

built-in reputation systems will encourage truth sensitive discourse simply because people will be accountable for what they say, if they have a tendency to ignore the state of investigation in relevant disciplines, this will become apparent and can be ignored or approached with scepticism.

Throughout this thesis, and particularly in the second chapter, I have maintained the position that the institutions responsible for the dissemination of information have been negatively affected by the internet and have grown increasingly unreliable. This is made especially clear in the rise of the spread of disinformation and the lack of transparency in the selection principles of information sources. In this section I have demonstrated how these issues could be addressed with a DLT-based identity system that would extend our identities into our online activities, thereby re-establishing journalistic norms and accountability to information distribution institutions.

#### *4.3 Electoral Systems in a DLT-Enhanced Choice of Aims Model of Democracy*

One of the key criteria that we have identified for a set of institutions that would be needed to facilitate a model of democracy like Christiano's choice of aims model, is that people not only need institutions that promote democratic ideals, but they also need to be able to see that this is the case. This means that people need to be able to see that they are treated as an equal in their society, that their participation matters, that their vote is recorded as it was cast, and that their government is acting in accordance with what they were elected to do.

Christiano's model of democracy relies on citizens casting votes to determine what the aims of their society should be, but not the means of how those aims will be achieved. Instead, determining how exactly the various chosen aims will be achieved is up to the elected government. How this gets cashed out in Christiano's view is that citizens choose from robust sets



of packages of social aims which are presented in the form of political party platforms.<sup>115</sup> In my view, there are at least two areas in this process where DLT can be used to improve the way in which citizens are able to see that their institutions are upholding democratic principles. First, in the electoral process itself and second, in the accountability of the government after it has been elected.

Just like our government issued ID mediates interactions between ourselves and our state's institutions today, DLT-based ID will mediate interactions with DLT-based institutions in democracies in the future. Using their DLT-based digital identities citizens of the democracies of the future will be able to securely vote from their personal computer or mobile device using a digital signature that only they can produce (remember this is already being done in Estonia). They will do so on a system built on a distributed ledger that is open-source (i.e., its inner workings can be reviewed and audited by the public) and that is not controlled by the incumbent government or any other individual party. Citizens will be able to review that their vote was counted and recorded as it was cast. The possibility of tampering with election results and stuffing ballot boxes will be a thing of the past, as no one party will have more control over the process than any other, and just like in the bitcoin network, a consensus mechanism will be in place that guarantees that information is only recorded to the ledger exactly as broadcast by voters.

When a government is elected, the specific aims that were chosen by the electorate could be codified into a smart contract that would track the commitments governments have made in their platforms and their progress in achieving the corresponding aims that were chosen. If they achieve a particular aim, they could verify it by referencing the relevant legislation and citizens

---

<sup>115</sup> Christiano and University of Arkansas Press, "Democracy and Social Epistemology," 69–71.

could sign off on whether the referenced articles constitute achievement of the chosen aim. For example, imagine that voters chose the following as part of the overall aims for their society: reducing tobacco use and the desire to meet all the United Nations Sustainable Development Goals. These societal aims, along with all the others that were chosen by voters, could be codified in a smart contract on a public ledger that would represent a digital agreement between citizens and their elected officials and would be reviewable by the public through a digital interface. If the government passes legislation that actually brings about one of the aims, e.g., they successfully reduce tobacco use by increasing taxes on tobacco, then they could reference the legislation that reflects this in the smart contract. Citizens would be able to see that the actual aim they chose for their society was achieved and could be directed towards the piece of legislation that brought it about. However, if the government failed to bring about some of the aims chosen by voters, e.g., they met only 10 of the 17 UN sustainable development goals, they could at least indicate the progress they made in achieving the overall aim, what steps they made to try to fully complete it, why they failed, and what their plans are to achieve the remaining 7 goals in the future. In effect, this would act as a roadmap for how the government is progressing towards realizing the aims that are chosen by citizens.

Beyond this, the use of public funds could be tracked down to the very dollar. This is a natural fit for DLT, as its original purpose was to act as a ledger for monetary transactions. Several non-profits are exploring strategies like this for improving transparency in the use of donor contributions.<sup>116</sup> The same principles could be applied to hold governments accountable for how they use taxpayer's money and to reduce corruption.

---

<sup>116</sup> "Gavi and Germany Partner to Harness Blockchain Technology."

There is a great degree of flexibility in the specifics of how features of DLT-based institutions like these might be designed but the point is DLT is well suited to hold electoral institutions and governments to a high standard of transparency and accountability. Following strategies like these, a system could be designed that truly reflects the idea that citizens are choosing the destination for their societies and governments are merely charged with getting them to that destination. Importantly, citizens will not just be told that their institutions are a certain way, but they will be able to confirm it for themselves.

#### *4.4 Incentivizing Background Knowledge and Engagement*

In order to make reasoned preferences about the aims of society, citizens do need a certain amount of background knowledge and they need some engagement with political discourse. Christiano's model decreases the epistemic burden placed on voters by requiring them to only make decisions about the aims of society, as opposed to the aims and the means of achieving them. Other characteristics of a DLT-enhanced choice of aims model of democracy that are likely to result in an electorate with more background knowledge and more interest in political engagement are the transparency, accountability, and sense of being part of society that values democratic ideals. These characteristics are likely to result in populations being less apathetic about their role in the political decision-making process and being more willing to spend their time engaging in politics.

However, I think this issue can be more directly addressed in a system of DLT institutions, while also contributing to an integral process in any distributed ledger, i.e., its consensus mechanism. As was discussed in the previous chapter, a consensus mechanism is used in all DLTs to guarantee that the information added to the associated ledger is exactly what was broadcast to the network. In the bitcoin network, this prevents people from trying to have false information

added to the ledger regarding how much bitcoin is in their wallet. In a system of DLT institutions it would prevent people from being able to manipulate any kind of data as it was added to the ledger. In the system I am proposing, citizens use a distributed application (d-app) to access their DLT based institutions; whatever business they have with a particular institution they do so through a single d-app that they access using their digital identity. This d-app would allow them to vote, file their taxes, apply for employment insurance, register ownership of a vehicle, and so on. But it could also act as a platform for people to engage in political discourse and complete educational modules that would provide them with the breadth of background knowledge needed to make reasoned preferences about the aims they would like society to pursue. As they complete modules they would receive a badge that would act as a digital certificate that could be integrated in to their DLT-based identity and shared whenever and with whomever desired. Whenever a user performs some action on this d-app, it would relegate a portion of the computing power used to perform this action to power the consensus mechanism for the entirety of the network. Moreover, because this would help maintain the DLT network, that provides incredible utility to society, users would be financially rewarded in proportion to the amount of computing power they contribute to the network. Different actions could be incentivized differently, and controls could be put in place to stop people from spamming the network. Accordingly, a citizen may receive a very small reward for filing their taxes or querying a record on the ledger, a somewhat larger reward for posting or replying to a discussion topic on the political discourse forum, and an even greater reward for completing a module in the educational platform. Depending on the economic conditions in the society in question, this could be a way of distributing a universal basic income to citizens. Under a system like this, citizens could be financially incentivized to gain necessary background knowledge and to engage in political

discourse, while performing a useful function in supporting the network that maintains their society's institutions. In essence, by running the d-app on their device, citizens would be running a node of the network that supports the distributed ledger infrastructure for their state.

#### *4.5 A DLT Enhanced Choice of Aims Model of Democracy*

The choice of aims model of democracy is prima facie more tenable than other models of democracy due to the fact that it balances the need to give voters a meaningful role in determining the terms of association for the societies they live in with a realistic account of the capacities of the average citizen. But a model like this needs to be complemented by a set of institutions that can support it. Specifically, I have identified that these institutions need to promote truth sensitivity in political discourse, allow citizens to assess the similarities and differences between their own selection principles and those of the information sources they use in political decision-making, incentivize the formation of background knowledge and interest in political discourse such that citizens can make reasoned preferences about the aims of their society, and be able to be seen by citizens to promote the ideals of democracy.

Distributed ledger technology is open, auditable, trustless and transparent by design. It allows disparate stakeholders to maintain a reliable record without the need for a centralized authority. In the case of institutional design, it opens up design strategies that would not have previously been possible. In this chapter I have described one possible strategy for using these technologies to create a set of institutions that embody criteria set out for institutions that would support a choice of aims model of democracy.

Ultimately, I have envisioned a set of institutions that integrates our online identity with our real-world identity in a way that preserves the distinction between private and public actions. By doing this we can extend the same kind of norms that have existed for information that is

distributed in real-world interactions to information that is distributed online, i.e., reputation building and accountability for public actions, as well as privacy and security for private actions. The extension of these features to online discourse should help to curtail the spread of disinformation and fake news and promote truth sensitivity in political discourse. Moreover, by using a single interoperable digital identity (like the one being used in Estonia) which allows citizens to interact with their society's institutions and is used for voting and holding elected officials accountable for achieving the aims of society that were chosen by citizens, our institutions can embody the ideals of democracy in a way that can be clearly seen by citizens. The network that supports this can be maintained by a consensus that is achieved by an associated distributed application through which citizens interact with institutions and are rewarded proportionately for doing so. As citizens perform actions on this application they act as a node on the distributed network that supports their institutions and they will be incentivized to participate in political discourse and to ascertain the background knowledge needed to form reasoned preferences about the aims of society. Thus, we have conceived of a set of institutions that meets the previously laid out criteria and together with Christiano's choice of aims model, forms a defensible account of democracy.

## CONCLUSION

There are two main things I would like to accomplish with my concluding remarks; first I will provide a summary of the arguments that have been made throughout this thesis and, second, I would like to offer some thoughts regarding what I see as the importance of this project to work in political philosophy.

The primary objective of this theses was to determine if DLT could be used to enhance the institutions of democratic societies, such that they could be combined with a model of democracy to provide a realistic account of democracy for use in the context of the modern state. The first step in this process was to determine what a viable model of democracy might look like. To arrive at this picture, we provided a discussion of various democratic ideals and considered the capacity of voters to perform the role necessary for them under the ideals of democracy. What we found was that it was quite difficult to provide an account that described a role that was significant enough to voters so that could be said to be participating in determining the terms of association for the societies they lived in, while at the same time constituting a role the was realistic given the capacities of the average voter. We ultimately concluded that Thomas Christiano's choice of aims model struck this balance better than other models and represented the best chance at a defensible account of democracy.

From here we wanted to determine just what the electorate needs to know in order to adequately play their role in the choice of aims model and whether something like this was possible in the modern state. Through this discussion we determined that it should not be taken for granted that voters will be able to do what is required of them under the choice of aims model, instead we developed a set of criteria that institutions should have so as to complement

Christiano's model. In the second half of chapter 2, we showed that changes to the way information was being distributed throughout society (mostly having to do with the rise of the internet) were making it more difficult for voters to get the kind of information they needed to participate in the political decision-making process. This led us to the conclusion that the current institutional framework in most modern democracies would not adequately support a model of democracy like Christiano's, instead, even though Christiano's model lowers the epistemic burden placed on voters, significant institutional reform would be needed to ensure its viability.

The second half of the thesis shifts towards trying to solve the problems outlined in the first half. The third chapter delivers a description of DLT, providing the reader with an understanding of unique features of the technology and how they are achieved. This knowledge gives the reader the opportunity to start thinking about how the technology could be used to address issues in institutional design and understand the sketch of institutions that is provided in the final chapter. With an understanding of DLT in hand we proceed to chapter 4 where we provide an account of how the technology could be used to enhance democratic institutions such that they would be able to meet the criteria laid out in chapter two. Ultimately, we argue the DLT-enhanced institutions would complement the choice of aims model of democracy and, as such, it would be defensible in the context of the modern state.

One of the interesting things about DLT is its degree of customizability. The work I have done here shows just one way the technology might be used to design institutions with a particular goal in mind. If you do not agree that Christiano's model is one worth designing institutions to support or you are not even a proponent of democracy, you should still see the usefulness of the technology for institutional design in general. It is my belief that DLT is going to be an important part of the institutions of the societies of the future and I have tried to



demonstrate this to the reader. For political philosopher's who are interested in questions regarding the institutions of actual societies in the here and now, DLT is going to be an important tool going forward and something they should take the time to think about and understand.

References

- Achen, Christopher H., and Larry M. Bartels. *Democracy for Realists: Why Elections Do Not Produce Responsive Government: With a New Afterword by the Authors*. Princeton Studies in Political Behavior. Princeton, New Jersey: Princeton University Press, 2017.
- “ACT-IAC White Paper: Enabling Blockchain Innovation in the U.S. Federal Government.” ACT-IAC, October 17, 2017. <https://www.actiac.org/act-iac-white-paper-enabling-blockchain-innovation-us-federal-government>.
- “Awarded Grants.” Bill & Melinda Gates Foundation, March 5, 2019. <https://www.gatesfoundation.org/How-We-Work/Quick-Links/Grants-Database>.
- Balasubramanian, Kannan, M. Rajakani, and IGI Global, eds. *Algorithmic Strategies for Solving Complex Problems in Cryptography*. Hershey, Pennsylvania (701 E. Chocolate Avenue, Hershey, Pennsylvania, 17033, USA): IGI Global, 2018.
- Baran, Paul. “ON DISTRIBUTED COMMUNICATIONS: I. INTRODUCTION TO DISTRIBUTED COMMUNICATIONS NETWORKS.” RAND CORP SANTA MONICA CALIF, August 1964. <https://apps.dtic.mil/docs/citations/AD0444830>.
- “Blockchain for Digital Identity - IBM Blockchain,” June 14, 2019. <https://www.ibm.com/blockchain/solutions/identity>.
- Brennan, Jason. *Against Democracy*. Princeton: Princeton University Press, 2016.
- Cameron, Kim. “The Laws of Identity,” n.d., 12.
- Cavna, Michael. “‘Nobody Knows You’re a Dog’: At 20, Web Cartoon as True as Ever.” *Washington Post* (blog), July 31, 2013. <https://www.washingtonpost.com/blogs/comic-riffs/post/nobody-knows-youre-a-dog-as-iconic-internet-cartoon-turns-20-creator->

peter-steiner-knows-the-joke-rings-as-relevant-as-ever/2013/07/31/73372600-f98d-11e2-8e84-c56731a202fb\_blog.html.

Christiano, Thomas. “An Instrumental Argument for a Human Right to Democracy: An Instrumental Argument for a Human Right to Democracy.” *Philosophy & Public Affairs* 39, no. 2 (March 2011): 142–76. <https://doi.org/10.1111/j.1088-4963.2011.01204.x>.  
———. “Rational Deliberation among Experts and Citizens.” In *Deliberative Systems: Deliberative Democracy at the Large Scale*. Cambridge: Cambridge University Press, 2012.  
———. *The Rule of the Many: Fundamental Issues in Democratic Theory*. Focus Series. Boulder, Colo: Westview Press, 1996.

Christiano, Thomas, and University of Arkansas Press. “Democracy and Social Epistemology.” *Philosophical Topics* 29, no. 1 (2001): 67–90.  
<https://doi.org/10.5840/philtopics2001291/28>.

Çinar, Özgür H. “The Right to Privacy in International Human Rights Law.” *Journal of Information Systems*, n.d., 12.

“Decentralized Identity – Own Your Own Identity.” Accessed July 10, 2019.  
<https://www.microsoft.com/en-us/security/technology/own-your-identity>.

Downs, Anthony. *An Economic Theory of Democracy*. New York: Harper, 1957.

“Equifax | Credit Bureau | Check Your Credit Report & Credit Score.” Accessed May 4, 2019.  
<https://www.equifax.com/personal/>.

Estlund, David M. *Democratic Authority: A Philosophical Framework*. Princeton, N.J: Princeton University Press, 2008.

Fagin, Ronald, Moni Naor, and Peter Winkler. "Comparing Information without Leaking It."

*Communications of the ACM* 39, no. 5 (May 1, 1996): 77–85.

<https://doi.org/10.1145/229459.229469>.

Fischer, Steven R. *A History of Writing*. Globalities. London: Reaktion, 2001.

"Gavi and Germany Partner to Harness Blockchain Technology," February 19, 2019.

<https://www.gavi.org/library/news/press-releases/2018/gavi-and-germany-partner-to-harness-blockchain-technology/>.

Glassner, Jean-Jacques, Zainab Bahrani, and Marc Van de Mierop. *The Invention of Cuneiform: Writing in Sumer*. Baltimore: Johns Hopkins University Press, 2003.

Goodin, Dan. "Why the Equifax Breach Is Very Possibly the Worst Leak of Personal Info Ever." Ars

Technica, September 8, 2017. [https://arstechnica.com/information-](https://arstechnica.com/information-technology/2017/09/why-the-equifax-breach-is-very-possibly-the-worst-leak-of-personal-info-ever/)

[technology/2017/09/why-the-equifax-breach-is-very-possibly-the-worst-leak-of-personal-info-ever/](https://arstechnica.com/information-technology/2017/09/why-the-equifax-breach-is-very-possibly-the-worst-leak-of-personal-info-ever/).

Goody, Jack. *The Logic of Writing and the Organization of Society*. Studies in Literacy, Family, Culture, and the State. Cambridge [Cambridgeshire] ; New York: Cambridge University Press, 1986.

Huberman, Bernardo A., Matt Franklin, and Tad Hogg. "Enhancing Privacy and Trust in Electronic Communities." In *Proceedings of the 1st ACM Conference on Electronic Commerce - EC '99*, 78–86. Denver, Colorado, United States: ACM Press, 1999.

<https://doi.org/10.1145/336992.337012>.

"ID2020 | Home." ID2020. Accessed July 10, 2019. <http://id2020.org/>.

"ID-Card." e-Estonia. Accessed July 10, 2019. <https://e-estonia.com/solutions/e-identity/id-card/>.

“KSI Blockchain.” e-Estonia. Accessed July 10, 2019. <https://e-estonia.com/solutions/security-and-safety/ksi-blockchain/>.

Lazer, David M. J., Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, et al. “The Science of Fake News.” *Science* 359, no. 6380 (March 9, 2018): 1094. <https://doi.org/10.1126/science.aao2998>.

Mill, John Stuart. *On Liberty ; Representative Government ; The Subjection of Women : Three Essays*. London : Oxford University Press, 1912.

<http://archive.org/details/onlibertyreprese00milluoft>.

Najeeb Alamuor, Rabi. “Right to Privacy in Law (Sanctity of Public Law).” *Journal of Law, Policy and Globalization* 80, no. 1 (2018): 1–10.

Nakamoto, Satoshi. “Bitcoin: A Peer-to-Peer Electronic Cash System,” n.d., 9.

Narayanan, Arvind. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*.

Princeton, New Jersey: Princeton University Press, 2016.

———. *Bitcoin and Cryptocurrency Technologies: A Comprehensive Introduction*. Princeton, New Jersey: Princeton University Press, 2016.

“Online Hash Calculator — String & File Hash Generator.” Accessed May 8, 2019.

<https://www.pelock.com/products/hash-calculator>.

Pimple, Kenneth D., ed. *Emerging Pervasive Information and Communication Technologies*

(*PICT*). Dordrecht: Springer Netherlands, 2014. <https://doi.org/10.1007/978-94-007-6833-8>.

Popkin, Samuel L. *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. 2nd ed. Chicago: University of Chicago Press, 1994.

Qayyum, Adnan, Junaid Qadir, Muhammad Umar Janjua, and Falak Sher. "Using Blockchain to Rein in The New Post-Truth World and Check The Spread of Fake News."

*ArXiv:1903.11899 [Cs]*, March 28, 2019. <http://arxiv.org/abs/1903.11899>.

Ray, Shaan. "Blockchain Forks." Hacker Noon, November 29, 2017.

<https://hackernoon.com/blockchain-forks-b0dca84db0b0>.

"Revealed: 50 Million Facebook Profiles Harvested for Cambridge Analytica in Major Data Breach | News | The Guardian." Accessed May 4, 2019.

<https://www.theguardian.com/news/2018/mar/17/cambridge-analytica-facebook-influence-us-election>.

Ruth, Michael. "Cognitive Bias." In *Salem Press Encyclopedia of Health*. Salem Press, 2018.

<http://libaccess.mcmaster.ca/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=ers&AN=100259337&site=eds-live&scope=site>.

SA, WISEKey. "WISEID." WISEKey | The Human-Driven Technology Platform. Accessed July 15, 2019. <https://www.wisekey.com/wiseid/>.

Schumpeter, Joseph A. *Capitalism, Socialism and Democracy*. 1. Harper colophon ed., [Nachdr.].

New York, NY: HarperPerennial, 2006.

Sniderman, Paul M., Richard A. Brody, Philip E. Tetlock, and Henry E. Brady. *Reasoning and Choice: Explorations in Political Psychology*. Reprint. Cambridge: Cambridge Univ. Press, 1999.

Somin, Ilya. *Democracy and Political Ignorance: Why Smaller Government Is Smarter*. Stanford, California: Stanford Law Books, an imprint of Stanford University Press, 2013.

"SPJ Code of Ethics - Society of Professional Journalists." Accessed June 24, 2019.

<https://www.spj.org/ethicscode.asp>.

Tapscott, Don, and Alex Tapscott. *Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business and the World*, n.d.

———. *Blockchain Revolution: How the Technology behind Bitcoin is Changing Money, Business, and the World*. Toronto, Ontario, Canada: Portfolio, an imprint of Penguin Canada, a division of Penguin Random House Canada Limited, 2016.

Texas Bitcoin Conference. *Hardy Helburn*, 2018.

<https://www.youtube.com/watch?v=rf5lg6kEMeg>.

United States, ed. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States*. Authorized ed., 1st ed. New York, NY: Public Affairs, 2011.

“Version History - Bitcoin Core.” Accessed May 10, 2019. <https://bitcoin.org/en/version-history>.