

Technologizing Democracy or Democratizing Technology?  
A Layered-Architecture Perspective on Potentials and Challenges

*unpublished draft – final version to appear in  
forthcoming book from the University of Chicago Press*

Bryan Ford

January 30, 2020

**Contents**

<b>1 Introduction</b>	<b>2</b>
<b>2 Democratic Deliberation and Choice</b>	<b>4</b>
<b>3 Information Curation, Reputation, Bias, and Polarization</b>	<b>6</b>
<b>4 Access, Inclusion, and Economic Empowerment</b>	<b>9</b>
<b>5 Identity, Personhood, and Digital Citizenship</b>	<b>11</b>
<b>6 An Architecture for Digital Democracy</b>	<b>16</b>
<b>7 Conclusion</b>	<b>18</b>

# 1 Introduction

Democracy is in the midst of a credibility crisis. Some of the most well-established Western democracies have become increasingly polarized [Prior, 2013, Iyengar and Westwood, 2015] to the point of tribalism [Hawkins et al., 2018, Packer, 2018] and authoritarianism. [Browning, 2018] The information sources voters use to understand the world and make their decisions is increasingly suspect. [Woolley, 2016, Ferrara et al., 2016, Woolley and Guilbeault, 2017, Broniatowski et al., 2018, Shao et al., 2018] While democracy preaches a gospel of treating all citizens as equal, established democracies fail to protect the equality of citizens' influence at the ballot box. [Smith, 2014, Gilens and Page, 2014, Cost, 2015, Flavin, 2015, Kalla and Broockman, 2016, Samuel, 2018, Tisdall, 2018]

Outside the ballot booth, people in real democracies depend on government to protect not only their physical safety, but also their economic and social equality and human rights. Here too, established democracies fail to protect their citizens from private coercion or feudal rent-seeking structures. [Shlapentokh and Woods, 2011] They fail to ensure equal access to equal economic opportunity by accelerating transfers of public wealth to the already-rich in the face of skyrocketing economic inequality. [Keller and Kelly, 2015, Piketty, 2017] They fail to offer an adequate social safety net to protect the ability of the unlucky or disadvantaged to participate in society as equals with dignity, and they even fail even to protect many people from effective slavery. [Weitzer, 2015, Kara, 2017] As Robert Dahl asked: "In a political system where nearly every adult may vote but where knowledge, wealth, social position, access to officials, and other resources are unequally distributed, who actually governs?" [Dahl, 1961]

Many perceive tremendous potential for technology to improve democracy: for example, by making it more convenient (vote from home with your laptop or smartphone), more participatory (express your opinion more than once every few years), or more inclusive (even in the developing world smartphones have become ubiquitous). But this somewhat "techno-utopian" view, common among the denizens of the early Internet, has gradually been overshadowed by realization of the many ways in which technology can undermine democracy, either by accident or design.

Technologists have often talked about technology as somehow inherently "democratizing" – using this term simplistically to refer to technological capabilities becoming inexpensive and widely available. The unstated and evidence-free implication embedded in this use of the term "democratizing," however, is that any inexpensive and widely-available technological gadget somehow makes society automatically more democratic. Actual experience in practice seems to suggest the opposite. The evolution of "democratized" social networking capabilities into advertising-driven instruments of mass surveillance, the weaponization of "democratized" free expression capabilities into instruments of fear, chaos, and polarization, and the transformation of "democratized" financial technologies like Bitcoin into shiny objects mainly attracting money launderers and financial scammers, all offer abundant experiential evidence of how

anti-democratic a “democratizing” technology can be.

But we have also seen how technology is almost infinitely flexible and adaptable. Technology is what we design it to be. Can we design technology to be *genuinely* democratic – to support and facilitate democracy reliably rather than undermining it? This paper explores several ways in which democracy in today’s digital world increasingly depends on technology for better or worse, ways that technology is currently failing democracy, and potential ways in which technology could be fixed to support democracy more effectively and securely.

Since effective democracy depends on far more than the occasional act of voting, we explore technology’s interaction with democracy “top-to-bottom,” across multiple levels at which the ability of people to self-govern depends on behavioral practices that are heavily impacted by technology. Yes, effective democracy requires people to have both the right and ability to vote. When they do vote, they need *effective* choice, not just a choice “between tweedledum and tweedledee.” [Zinn, 2005] Technologies such as E-voting, [Moynihan, 2004, Alvarez et al., 2009, Germann and Serdült, 2017] online deliberation, [Iyengar et al., 2003, Grönlund et al., 2009, Esau et al., 2017] and liquid democracy, [Ford, 2002, Sayke, 2003, Litvinenko, 2012, Green-Armytage, 2014, Blum and Zuber, 2016, Ford, 2019b] show promise in expanding the convenience and effectiveness of democratic choice, but each bring associated risks and major unsolved challenges that we outline.

Effective democracy also requires that people live in a social and economic environment satisfying the conditions for intelligent, informed, and effective democratic choice. People need reliable information sources protected from both subversion through “fake news” and polarization through automated over-personalization. People need free expression and free association to discuss ideas and organize effectively – but they also need protection from trolls and sock puppets seeking to amplify their voices via anonymous bot armies. People need an economic environment offering them the empowerment and leisure time needed to become informed and participate deeply in the deliberative phases of democracy, and not just in the final vote. Finally, people need the digital ecosystem to be able to recognize and identify them *as people* – *i.e.*, as formal “digital citizens” – and to be able to distinguish these real people from the millions of fake accounts of bot-farmers inhabiting the Internet, [Berger, 2018, Read, 2018] without undermining effective participation through exclusionary and abuse-ridden digital identity systems.

Having examined some of the promises, failures, and unsolved challenges at each of these levels, we will attempt to sketch briefly a long-term vision of a potential architecture for effective digital democracy, layered in the classic fashion followed in network protocol architecture. [Day and Zimmermann, 1983] The following sections outline, from top to bottom, such a layered architecture for digital democracy. The top layer, which we address first, represents the highest-level functionality we seek as a primary end-goal: namely effective technology-supported self-governance through democratic deliberation and social choice. Subsequent sections address critical “building block” layers for

effective technology-supported democracy: an information layer ensuring that participants have manageable feeds of high-quality, accurate, and unbiased information as an adequate basis for deliberation and decisions; an economic foundation layer to help ensure that citizens have the baseline means and freedoms to invest the time and attention required for genuine democracy; and finally, a digital citizenship layer ensuring that technology can securely but inclusively protect the rights and resources of real people from being abused, undermined, and diluted by online fakery. Finally, in the last two sections we will briefly recap this architecture and summarize how appropriate technologies for each layer could eventually fit together into a fundamentally more solid foundation for digital democracy than exists today.

## 2 Democratic Deliberation and Choice

As networked computing technology was just emerging, visionaries immediately recognized its potential use to involve people more richly in the democratic process. [Heinlein, 1966, Tullock, 1967, Miller III, 1969] Instead of trekking to a physical polling place every few years to make a nearly-binary choice between candidates that voters have at best heard about on TV, technology promised the possibility of “virtual town halls” in which millions could observe and participate continuously in democratic deliberation processes. Bringing that online democracy vision into reality has been far more slow and fitful, however.

As a starting point, E-voting systems promise the convenience of voting from the comforts of one’s own home, or remotely from outside the country of one’s citizenship, without fundamentally changing the nature or frequency of democratic choice. [Moynihan, 2004, Alvarez et al., 2009, Germann and Serdült, 2017] Switzerland’s long-held and extensive practice of direct democracy results in citizens being asked to vote typically four or more times per year. This participatory approach in part motivated Switzerland’s pervasive adoption of voting by mail, [Luechinger et al., 2007] followed by its early adoption of E-voting. [Gerlach and Gasser, 2009, Serdült et al., 2015, Mendez and Serdült, 2017, Germann and Serdült, 2017]

Any technology that permits voting outside the controlled environment of the ballot booth, however, may increase risks of undetected voting fraud such as coercion or vote-buying attacks, as the incident in North Carolina recently highlighted. [Blinder, 2019, Ford, 2019a] E-voting systems present particularly critical security concerns, however, due to the risks they may present of scalable electronic attacks, *e.g.*, by an attacker anywhere in the world successfully compromising the vote-counting servers, or exploiting a security bug common to many end-user devices. [Schryen and Rich, 2009, Zetter, 2019] Further, the same efficiency and scalability that makes E-voting attractive could potentially enable attackers to coordinate large-scale voter fraud, through “dark DAOs” (decentralized autonomous organizations)

for example. [Daian et al., 2018] Some of these challenges are likely to be solvable only in coordination with “lower layers” of the technology stack, such as communications and identity layers to be discussed later.

There have been many attempts – with varying success – to get citizens involved not just in one-off votes or polls, but in true deliberative processes where participants learn about and discuss issues in depth, often with the help of domain experts. [Iyengar et al., 2003, Grönlund et al., 2009, Esau et al., 2017] Selecting participants by sortition, or randomly sampling a target community, can keep costs manageable while ensuring that the deliberative body is diverse and representative. [Fishkin, 1993, Iyengar et al., 2003, Landemore, 2018] Because the size and cost of each such representative group is small, governments and other organizations can in principle launch and run many such deliberative groups in parallel on different topics, making the process efficient and scalable. Recent computer science efforts to scale *automated* decentralized systems, such as blockchain and smart contract systems, have relied on essentially the same principle of running many small representatively sampled groups in parallel. [Luu et al., 2016, Kokoris-Kogias et al., 2018]

Some of the key benefits of democratic deliberation, however, are embodied not so much in the *outcome* of deliberation (*i.e.*, the lessons learned or the report written at the conclusion), but in the impact of deliberation on *the participants themselves*: *e.g.*, giving the participants deeper understanding of issues that affect them, and a sense of participating actively in their community and (hopefully) a feeling of having their voices heard. Deliberation in small sampled groups, however representative and scalable, have the key limitation of awarding the latter class of benefits only to the few lucky winners of the lottery. The larger population benefits at best indirectly, from the participants’ reports about their experiences and/or from the effects of better decisions hopefully being made.

The idea of *delegative* or *liquid democracy* pursues the goal of enabling *everyone* to participate in regular or even continuous online deliberative processes, while recognizing the fundamental constraint that everyone has limited time and attention. [Ford, 2002, Sayke, 2003, Litvinenko, 2012, Green-Armytage, 2014, Blum and Zuber, 2016, Landemore, 2018, Ford, 2019b] The essential idea is to give citizens the freedom to choose when and how much to participate, based on their limited attention, while *delegating* their voice on matters beyond their capacity or interests to others they trust to represent them. In essence, all participants receive an *individual* choice between direct and representative democracy, on an issue-by-issue or vote-by-vote basis.

There have been many experiments in implementing and deploying liquid democracy throughout the world over the past two decades, with promising but mixed results. [Litvinenko, 2012, Ford, 2014, Hardt and Lopes, 2015, Ford, 2019b] The most prominent and large-scale experiment in liquid democracy so far was the German Pirate Party’s adoption of the idea for online intra-party discussion via its LiquidFeedback platform. [Swierczek, 2014, Litvinenko, 2012, Behrens, 2014] Liquid democracy presents many concerns and potential risks, however.

One important concern with liquid democracy is that different delegates, freely chosen by proxy voters, will necessarily exercise different amounts of voting power in the deliberative process. [Blum and Zuber, 2016] Concern for such effects seems to be supported by the German Pirate Party’s experience of one delegate accumulating (apparently by accident) an outsize share of voting power. [Becker, 2012] Many of these risks may be mere artifacts of immature implementations of liquid democracy, however, with weaknesses that are important but easily fixed. The concentration of power the Pirate Party experienced, for example, may be attributable to the LiquidFeedback software’s allowing voters to choose *only one* person to delegate their *entire* vote to, artificially creating a “winner-take-all” scenario in which almost-but-not-quite-as-popular delegates lose out completely. Other formulations of liquid democracy allow voters to split their voting power among multiple delegates, [Ford, 2002, Boldi et al., 2011, Ford, 2019b] enabling delegated power to spread among *all* the delegates each voter trusts instead of concentrating on a few global winners. Other recent work defines multiple-delegation mechanisms with provisions specifically designed to limit the inequality of delegated voting power. [Gölz et al., 2018] It is not yet clear whether such provisions are strictly necessary, however, or what the attendant costs and tradeoffs might be.

Other concerns that are less fundamental but equally critical in practice center on the immature technology implementations of current online deliberation and liquid democracy platforms, almost all of which rely on a single centralized server, whose compromise could undetectably corrupt the entire democratic process. The experience of Italy’s “Five Star” movement, widely suspected to embody more of a techno-autocracy than a democracy facilitated by the software platform designed and run by a father-son duo, illustrates the risks inherent in centralized platforms. [Horowitz, 2018, Loucaides, 2019] There is growing interest in building liquid democracy systems on decentralized blockchain and smart contract platforms, [Agarwal, 2018, Crichton, 2018, Zhang and Zhou, 2019] but these experiments and the platforms they build on are still immature, and subject to the same critical security, privacy, and voting fraud risks that apply to E-voting systems as discussed above.

### **3 Information Curation, Reputation, Bias, and Polarization**

Voters cannot make informed decisions without access to good information, together with the time and motivation to digest it – the key prerequisite to effective democracy that Robert Dahl terms “enlightened understanding.” [Dahl, 1989] Throughout most of human history, information was scarce and precious. The digital world stands this problem on its head, creating the equally serious but opposite problem of too *much* information, with only inadequate, insecure, and essentially undemocratic mechanisms for users to filter, curate, and mentally process that information.

The USENET was the first global, decentralized public forum online allowing anyone to read and post messages

and discuss practically any topic. [Hauben and Hauben, 1997, Templeton, 2001a] In its time, USENET was intensely exciting and empowering to many, and was an early entrant in the long line of digital technologies frequently referred to, rightfully or not, as “democratizing.” [Hill and Hughes, 2007, Blumler and Gurevitch, 2001] The USENET is now largely forgotten, not because it stopped working, but because it worked *too well*, reliably broadcasting signal and noise together and rendering both nearly uncensorable. Spam – both the term and the online practice – were invented on USENET, [Templeton, 2001b, Templeton, 2003] and precipitated its effective downfall as uncontrolled spam and trolling finally sent most “netizens” scurrying away to more protected forums on centralized platforms. [Templeton, 2001a]

But trading the uncontrolled chaos of the decentralized USENET, for the protection of professional moderators and opaque filtering algorithms owned by profit-motivated technology companies, may in hindsight have been a Faustian social bargain. Centralized technology platforms like Facebook and Twitter did give people freedom to communicate among their friends with greater protection from spam and trolls. These platforms had their own heyday of being called “democratizing” – especially around the time of the Arab Spring. [Comminos, 2011, Howard and Hussain, 2013] But spammers and trolls learned to adapt and abuse these platforms, leading to the online forum governance and exclusion problems detailed elsewhere in this volume. [Caplan, 2018, Gangadharan, 2018, Farrell and Schwartzberg, 2018, Cohen and Fung, 2018] Further, the effective concentration of information-filtering power into opaque and unaccountable algorithms, designed and run by a few profit-motivated technology giants, represents a crucial threat to democracy in its own right.

Despite the public’s retreat to proprietary platforms, technology researchers never lost interest in finding decentralized solutions to the noise and abuse problems that defeated USENET. The proof-of-work algorithm underlying Bitcoin, [Nakamoto, 2008] for example, was originally proposed as a mechanism to combat spam, by requiring an E-mail’s sender to prove to have spent considerable computational effort preparing it. [Dwork and Naor, 1992] Other clever decentralized algorithms could in principle efficiently pick a set of *guides*, who find and recommend content compatible with the tastes of a given user, out of an ocean of bad content and fake accounts. [Yu et al., 2009]

The encryption tool PGP (“Pretty Good Privacy”) popularized the idea of decentralized social trust networks in its “Web of Trust” model. [Stallings, 1995, Abdul-Rahman, 1997] Many decentralized content governance and filtering algorithms subsequently built on the idea of trust networks. [Kamvar et al., 2003, Mislove et al., 2008, Yu et al., 2008, Tran et al., 2009, Viswanath et al., 2012] However, actually building trust networks with PGP or other decentralized tools never caught on among the public or even the tech-savvy. Even if decentralized social networks had caught on widely, it is doubtful whether the social networks constructed by the popular centralized platforms actually have the critical properties of trust networks required by decentralized content filtering algorithms. [Mislove

et al., 2007, Viswanath and Post, 2010, Ghosh et al., 2012, Messias et al., 2013]

Another fundamental problem with the social or trust network approach lies in the basic premise that it is desirable for voters to perceive the world through a lens filtered by their immediate social relationships, a practice widely suspected (though not conclusively proven) to create an “echo chamber” effect [Barberá et al., 2015, Dubois and Blank, 2018] and contribute to social polarization [Iyengar and Westwood, 2015] and tribalism. [Hawkins et al., 2018, Packer, 2018] People who mostly rely on – and most trust – information filtered through their social network may also be more inclined to perceive bias in information sources *not* filtered by their social tribe. [Saez-Trumper et al., 2013, Eberl et al., 2015, Kaye and Johnson, 2016, Budak et al., 2016, Ribeiro et al., 2018, Farrell and Schwartzberg, 2018] Without discounting the popularity and appeals of social communication, it seems clear that the digital ecosystem is missing an objective, unbiased, and usable source of “big picture” information and perspective.

One promising idea is to employ the sampling methods discussed above [Iyengar et al., 2003, Landemore, 2018] to the problem of information curation and filtering. For example, we might try to design newsfeeds whose topic and viewpoint choices are chosen through some deliberative curation process, by members of a representative sample population, to ensure diversity and objectively avoid bias. While this approach seems worth exploring, it presents further challenges.

A small representative group might conceivably be effective at choosing among and selecting information on topics already of widespread interest. A sample population is much less likely to be effective, however, at identifying *rare* topics of not-yet-widely-recognized importance, or at finding valuable but obscure information about such topics. This is an instance of the perennial “needle in a haystack” problem, or of the “rare event” problem in statistics.

Here again, liquid democracy ideas may be useful in synergy with sortition methods. In an online forum dedicated to information gathering and curation, suppose we initially give voting power only to the members of a small sample population. However, we allow these sampled voters to *delegate* their voice selectively – in whole or in fractions – to others *outside* the representative group who they deem trustworthy and knowledgeable on particular topics. This delegation could enable the small original sample population to spread and multiply their information-gathering power, effectively recruiting a much larger crowd of assistants and advisers to their aid, while preserving the sampling-based diversity and democratic representativeness of the group’s composition and perspectives.

Any approach to information filtering and curation runs into the fundamental problem of accounting (or not) for expertise. We generally expect information from domain experts to be more reliable and trustworthy precisely because experts are supposed to know more about the domain. Being able to identify and utilize domain expertise increases in importance as topics and policy questions become more complex and deeply technical. Experts may also bring domain-related biases, however. One obvious example of such bias is the tendency of technology developers to see



the positive uses of their systems and algorithms far more readily than the negative risks their designs carry.

Further, there is the fundamental question of *who decides* who is an expert, on what grounds, and whether that expert-selection process can be called “democratic” in any sense. Neither ordinary citizens, nor professional politicians without domain knowledge, are necessarily good at distinguishing experts from smooth-talking charlatans. [Edens et al., 2012, Gemberling and Cramer, 2014] But professional organizations and certification systems, in which yesterday’s experts vet and choose tomorrow’s experts, are subject to narrow groupthink, gradual inbreeding, and cultural ossification. [Collins and Evans, 2002, lib, 2003, Kotzee, 2012] Organizations that vet, fund, or reward experts for their work may become disconnected from and unaccountable to the broader public. [Brewster, 2003, Reich, 2018]

Can we find more democratic and accountable ways to recognize and vet experts and the critical role they play in both producing and evaluating information serving the broader public on deeply technical topics? One observation that may be useful is that although non-experts may have trouble distinguishing top experts from lesser experts or charlatans who merely speak the language, it might be more feasible to rely on people merely to identify others with *more expertise than themselves* in a domain. This observation suggests a variation on the delegation ideas explored above: ask members of a community or a representative group to identify a few other people each – inside or outside the original group – who they consider trustworthy and to have more expertise than themselves on some topic. Sort the resulting group by delegated voting weight, eliminate (say) the bottom half, and repeat. The hypothesis, yet to be developed out and tested, is that each iteration of this process will use progressively better (“more-expert”) information to narrow the population of candidate experts progressively, and ensure that charlatans will be discovered and eliminated at some level at which the genuine experts can reliably distinguish them. [Ford, 2019d]

Yet another important question is how we can democratically finance and reward journalism and the production of good information. [Cagé, 2018, Lee et al., 2018, Farrell and Schwartzberg, 2018, Bernholz, 2018] We defer exploration of this topic to the next section.

## **4 Access, Inclusion, and Economic Empowerment**

Real voters in functioning democracies aren’t just disembodied decision-making entities in an academic’s theoretical model, but must make decisions and participate (or not) in the context of the real environment they live in. The opportunities and constraints their environment provides – including their education, social networks, money, and free time – has significant practical impact on their effective inclusion or exclusion in political and civic society. [Landemore, 2018, Gangadharan, 2018, Ananny, 2018]

Many practical factors can present exclusionary barriers to the act of voting, such as to voters who have no identity

card, [Hicks et al., 2015, Hajnal et al., 2017, Highton, 2017] no home address, [Feldman, 2006] or past criminal convictions. [Manza and Uggen, 2008] Timing and other logistical factors may affect voter turnout, though in complex and often-unclear ways. [Quinlan, 2015]

Excluding citizens from voting, however, is merely the most blunt and crude way to compromise Dahl’s democratic criterion of inclusiveness. [Dahl, 1989] For effective democracy, citizens also need good information (see above), the time to digest and discuss it and form their preferences, and the time and opportunity to participate in controlling the agenda – whether by attending town hall meetings, joining political demonstrations, calling their representatives, etc. People also need the requisite education and political culture to have a sense of *what democracy is*. [Cho, 2015] Those struggling to survive while juggling three precarious part-time jobs [Standing, 2011] may reasonably consider voting, let alone taking the time required for *informed* voting and participation, a luxury they cannot afford.

Given the evidence that political participation is linked with economic inequality, [Armingeon and Schädel, 2015, Filetti, 2016] which has been growing uncontrollably, [Piketty, 2017] we may justifiably consider the economic equality and well-being of voters to be as essential to effective democracy as voting itself. This raises the question of whether the economic foundations of today’s democracies are adequately “democratic” – and if not, to what extent the proper use of technology could improve that situation.

In the developing world experiencing the highest global inequality, mobile phones have become surprisingly ubiquitous. [Aker and Mbiti, 2010] There is evidence this penetration has helped mitigate inequality [Asongu, 2013a] and stimulate financial development. [Asongu, 2013b] This ubiquity of mobile devices could potentially offer a technological foundation for further projects to improve financial inclusion, as exemplified by the M-Pesa project in Kenya. [Mbiti and Weil, 2016]

Bitcoin [Nakamoto, 2008] and its many derivative cryptocurrencies represent another class of technologies often loosely called “democratizing” – this time usually in the sense of enabling people to perform cash-like transactions electronically without relying on trusted parties such as banks. While Bitcoin may *in principle* be usable by anyone without banks, however, to use it one must either buy Bitcoin from someone, or *mine* it yourself by competing to solve cryptographic puzzles. Because the nature Bitcoin mining confers huge advantages on those with access to cheap energy and the latest specialized hardware, however, mining is no longer economically viable to ordinary users – or to anyone but a few large entrenched specialists, in fact. [Vorick, 2018] Thus, to use Bitcoin the “have-nots” must buy or borrow it from the “haves.” Bitcoin and most cryptocurrencies thus merely replicate and digitally automate the inequality-increasing status quo, and cannot justifiably be described as “democratic” at least in a sense of equality of participation or inclusiveness.

An increasingly-popular idea is to replace, or augment, traditional social “safety net” programs with a *universal*

*basic income* or UBI – a regular income that citizens of some jurisdiction receive automatically regardless of whether or how much they work. [Parijs, 2017, Standing, 2017, Jackson and Victor, 2018, Bidadanure, 2019] UBI is an intriguing idea that has seen limited experiments. [Forget, 2011, Koistinen and Perkiö, 2014] There is at least one important downside of the usually-proposed approach of implementing UBI in a political jurisdiction such as a town, state, or country, [Wagner, 2017] however: it would create an incentive for anyone outside the relevant jurisdiction to move in – or try to gain residency fraudulently – thereby exacerbating already-inflamed xenophobic and protectionist tendencies.

An interesting alternative approach would be to build a cryptocurrency with a built-in UBI. [Ford, 2019c] Like Bitcoin, such a *crypto-UBI* currency would be usable by anyone “across borders” and not tied with existing geopolitical jurisdictions or currencies. Several cryptocurrency startups are already attempting such projects, in fact. [eternalgloom, 2018] Such crypto-UBI currencies could also conceivably be designed to offer collectively-financed economic rewards to the producers of information that the user community finds useful. This possibility suggests new ways to fund news, media, and open source technologies more democratically, rather than via traditional profit-motivated or philanthropic channels. [Cagé, 2018]

Many social, economic, and technical issues remain to determine whether and in what form the crypto-UBI idea is viable, however. And like many of the other potentially democratizing technology ideas discussed above, no crypto-UBI scheme can operate fairly or improve equality unless it can identify *real people* and distinguish them from fake identities of fraudsters, the fundamental challenge we focus on next.

## 5 Identity, Personhood, and Digital Citizenship

When humans interact in the real world, we use multiple senses tuned over millions of years of evolution to detect and distinguish other humans, creatures, inanimate objects, and unknown potential threats. We therefore take it for granted that we can easily and reliably distinguish *people* from *non-people*. Professions such as computer graphics and robotics that try to simulate human forms have discovered just how difficult it is to fool our senses, due to the widely-observed *uncanny valley* effect, in which almost-but-not-quite-perfect simulations can unintentionally trigger deep emotional reactions. [Mori, 2012, MacDorman, 2006] Our physical-world identity challenges thus tend to focus on classifying and differentiating *between* people: known or unknown, friend or enemy, attractive or unattractive, insider or outsider, member or non-member, citizen or foreigner.

But our intuitive assumption that distinguishing between real and fake people is easy completely fails to translate into the digital world – in part because our electronic devices do not have human senses with their millions of years of evolutionary tuning. By default, digital technologies know a “person” only as an electronic record or account someone

entered claiming, correctly or incorrectly, to represent a person. This inability to recognize *personhood* underlies one of the most fundamental unsolved challenges in our technology ecosystem: preventing abusers from creating several (or many) fake identities – whether for fun, for profit, or to undermine democracy. In distributed systems, this problem is termed the Sybil attack, [Douceur, 2002] after a famous psychiatric case of multiple-personality disorder. [Schreiber, 1973]

The reason the Sybil attack is so important is that it renders most of the common and intuitive defenses against abuse ineffective. Blocking a spammer’s E-mail address is useless because the spammer will just create many more fake identities automatically, leaving the spam problem unsolved after decades of attempted solutions. [Dwork and Naor, 1992, Cranor and LaMacchia, 1998, Koprowski, 2003, Templeton, 2003, Shaw, 2004, Chellapilla et al., 2005, Ramachandran et al., 2006, Mislove et al., 2008] As technology companies try to employ more sophisticated automated algorithms such as machine learning to detect fake identities, professional spammers and trolls adapt the same automation technologies to create ever-more-convincing fake identities, [Ferrara et al., 2016] with increasingly-serious consequences. [Bessi and Ferrara, 2016, Broniatowski et al., 2018] Automated Turing tests such as CAPTCHAs [von Ahn et al., 2003] fail to stem fake accounts in part because machine learning algorithms are getting better than real humans at solving such tests. [Chellapilla et al., 2005, May, 2019] Sybil attacks allow trolls to amplify their voices in collaborative forums by creating *sockpuppets* supporting their cause. [Bu et al., 2013, Solorio et al., 2014, Liu et al., 2016, Yamak et al., 2016] Because the Internet cannot distinguish between real and fake people, and ideologically-, politically-, or profit-motivated users can exploit this fundamental vulnerability at increasingly-massive scales without significant risk, our digital ecosystem is evolving into one in which a large fraction of the “people” – and their “discourse” – is fake. [Berger, 2018, Read, 2018]

This increasingly-true perception that so much of the Internet is fake, including a large portion of its supposed inhabitants, marginalizes *real* people online and fuels the growing technology backlash. [Doward, 2018] Further, the fact that such a high percentage of likes, upvotes, reviews, or any other online artifacts purportedly representing the opinions of “people” are well-known to be fake or readily forgeable, undermines any presumption that *anything* about the Internet can be justifiably or legitimately called “democratic.”

The scope, generality, and global consequences of the fake identity problem demand a correspondingly robust and general solution, but all of the currently-popular proposed solutions have significant flaws. Bitcoin’s attempt to address Sybil attacks via proof-of-work [Dwork and Naor, 1992] failed to ensure either equality or inclusiveness in participation, [Vorick, 2018] but also created an environmentally disastrous runaway competition to waste energy. [de Vries, 2018, Digiconomist, 2019] The cryptocurrency community has explored many variations such as memory-hard proof-of-work, [Boneh et al., 2016] proof-of-space, [Park et al., 2018] and proof-of-stake, [Kiayias et al., 2016,

[Gilad et al., 2017](#)]. All these variations reward participants in proportion to some form of *investment*, however, whether in terms of computation, memory, or purchasing and “staking” existing currency. All these investment-centric mechanisms, therefore, can be expected to retain “rich get richer” tendencies toward inequality in the power and influence of participants, and thus cannot hope to offer *person-centric* fairness or ensure equality of participation in a truly democratic sense.

The most obvious solution to identifying *people* is simply to transplant traditional identity documentation and verification processes online. Today’s “know your customer” (KYC) regulations for anti-money-laundering (AML) in banking have made ID checking a critical step in online finance businesses, and have created a booming market in identity-verification startups. [[nanalyze, 2017](#), [Abhishek and Mandal, 2017](#)] These verification processes typically involve asking users to present a physical photo ID over a video-chat session, and use machine learning for automation [[Doughty, 2005](#), [nanalyze, 2017](#)] and liveness detection techniques such as eyeball tracking in attempt to detect spoofing attacks. [[Pan et al., 2008](#), [Bao et al., 2009](#), [Wen and Jain, 2015](#)]

Besides being privacy-invasive, however, this approach is not particularly secure either. Computer-generated imagery technology has already traversed the “uncanny valley” to produce *deep fakes*, or video-simulated people that look convincing to real people. [[Chesney and Citron, 2018](#), [Mack, 2018](#)] Since abusers can use this ever-improving simulation technology to counter advances in detection, this approach offers at best an endless arms race that will likely end in deep fakes eventually becoming *more* reliably convincing than real people to detection algorithms.

A second problem with ID verification processes is that the physical IDs that they verify are not difficult or costly to fake. Digital passport scans sell for \$15 on the black market, for example, with forged passports sufficient for online ID verification but not to cross borders selling for \$1,000, and *genuine* passports usable to cross borders available for around \$15,000. [[Durden, 2015](#), [Bischoff, 2018](#), [Havocscope, 2019](#)] Since these are *retail* prices of black market IDs sold individually, the vendors and their corrupt sources can doubtless perform wholesale ID forgery much more cheaply. In effect, ID verification appears inevitably destined to become little more than security theater [[Kline, 2008](#), [Sethi et al., 2017](#)] and a legal compliance checkbox while offering no real protection against determined identity forgery.

Another approach to Sybil attack protection utilizes automated graph analysis of social trust networks. These algorithms typically rely on the assumption that a Sybil attacker can easily and cheaply create *nodes* (fake identities) in the social graph, but has a harder time creating *edges* (trust relationships) connecting them to real people. [[Mislove et al., 2008](#), [Yu et al., 2008](#), [Tran et al., 2009](#), [Viswanath et al., 2012](#)] As mentioned above, however, it is doubtful that popular online social networks actually constitute trust networks satisfying this assumed property. [[Mislove et al., 2007](#), [Viswanath and Post, 2010](#)] For example, many Twitter users intentionally engage in *link farming*, or following

a large number of other accounts on the (well-placed) hope that a significant fraction will reciprocate. [Ghosh et al., 2012, Messias et al., 2013] The presence of a significant number of link-promiscuous real users makes it easy for Sybil accounts to hide in that group and defeat graph algorithms that assume attackers are link-limited. In more sophisticated infiltration attacks, social bots interact with other users by forwarding or synthesizing content. [Freitas et al., 2015, Ferrara et al., 2016, Bessi and Ferrara, 2016, Broniatowski et al., 2018]

Even if real users could build a well-disciplined trust network, these graph algorithms would detect only *egregious* Sybil attacks, such as the case of one attacker creating a large number of fake identities. Graph algorithms could not and would not prevent *many* users from each cheating *a little bit*, by coordinating with their friends to create a few fake identities for example. Finally, Sybil resistance via trust networks would also be exclusionary against real people or groups who are poorly-connected socially, who would likely be falsely eliminated as Sybil identities.

Biometrics present another approach to identity and Sybil attack protection, as exemplified in India's Aadhaar digital identity project, which has issued biometric identities to over a billion people. [Bhatia and Bhabha, 2017, Chaudhuri and König, 2017] While attractive in terms of usability, the use of biometrics for identification is problematic in numerous ways. First, protecting against Sybil attacks and ensuring that each person registers only once requires *de-duplication* checks during enrollment, or comparing the new enrollee's biometrics against all existing ones, *i.e.*, over a billion in the case of Aadhaar. [Abraham et al., 2017] This de-duplication requires all users' biometrics to be collected in a massive searchable database, creating huge privacy and surveillance concerns, [Dixon, 2017, Srinivasan et al., 2018] in part because biometrics are effectively "passwords you can't change." [Schneier, 2009, Chanthadavong, 2015] Second, since biometric matching is inherently imprecise, it can both falsely accept duplicate enrollments and falsely reject legitimate new users. The Aadhaar program estimated a 0.035% false accept rate in 2017, [Abraham et al., 2017] but a different method produced an estimate of 0.1% the following year, implying that hundreds of thousands of Aadhaar records might be duplicates. [Abraham et al., 2018] There are signs that false rejections may be an increasing problem, as well, leading to another form of digital exclusion. [Venkatanarayanan, 2017, Mathews, 2016] Biometric exclusion threatens not just the participation opportunities, but even the very lives of unlucky or marginalized people who fall through the inevitable gaps left by biometric technologies. [Ratcliffe, 2019]

Having exhausted the commonly-proposed but uniformly-flawed solutions to distinguishing real people from fake Sybil accounts, what else is left? One idea is to create digital *proof-of-personhood* tokens via *pseudonym parties*. [Ford and Strauss, 2008, Ford, 2015, Borge et al., 2017] This idea builds on a "back to basics" security foundation, by relying on a person's *physical presence* at some time and place. For now, real people still have only one body each, and thus can be in only one place at a time. Expendable clones are still science fiction, [Brin, 2005] and robots have yet to follow Hollywood across the uncanny valley. [Scott, 2017]

Leveraging this property, a few times per year we might organize concurrent *personhood parties* at various locations, wherever a suitable group of organizers is available to run one. Before a certain critical moment, synchronized across a set of coordinated events, anyone is allowed to enter an enclosed or cordoned-off space. After the critical moment, people may only leave, getting one anonymous credential scanned on the way out, such as a QR code displayed on a smart phone or printed on paper. If properly run, witnessed, and recorded for public transparency, such a process could ensure that any participant can get one *and only one* “verified real person” credential valid for a given time period. Because pseudonym parties rely only on physical presence for their security, they avoid requiring any privacy-invasive identity checks or biometrics, or problematic security assumptions about social trust networks.

There is ample precedent for people participating in events requiring physical presence. The billions of members of the world’s largest religious traditions often attend in-person ceremonies in churches or temples several times a year, once a week, or more. Two Swiss cantons, Glarus and Appenzell Interior, have used open-air assemblies or *Landsgemeinde* for direct democracy for hundreds of years. [Dürst, 2004, Reinisch, 2007, Schaub, 2012] Political protests play a regular role in many democracies despite producing only rough media estimates of numbers present, and to uncertain and non-obvious concrete political effect. [Madestam et al., 2013, Acemoglu et al., 2018, Enikolopov et al., 2018] Scheduling and organizing such events to double as personhood parties could provide both the organizers and the public more precise statistics on attendance, and could give the attendees themselves verifiably Sybil-resistant anonymous credentials that might eventually become useful for many purposes.

After a pseudonym party, for example, attendees could later use their merely to “prove they were there,” or to form anonymous but abuse-resistant online forums for followup discussion or deliberation with attendance restricted to the in-person attendees. More broadly, attendees could use personhood badges to obtain “verified real person” status similar to verified account (“blue checkmark”) status on websites and social networks. Attendees could use their credentials as voting tokens in online polls or deliberative forums. They could use credentials to represent a one-person notion of stake to build “proof-of-personhood” decentralized blockchains and crypto-UBI currencies. [Borge et al., 2017]

Regular attendance of personhood parties could eventually become part of a social contract that offers a kind of Sybil-resistant *formal digital citizenship* with various rights and abilities unlocked by personhood credentials. These rights include secure, private, and democratically equitable online participation, together with the necessary protection from abuse, trolling, and ballot stuffing by fake identities. Because proof-of-personhood tokens have limited value and validity period, they are inherently “renewable” simply by showing up at a future pseudonym party anywhere. Digital citizenship rights attached to such time-limited but renewable tokens may therefore prove both more democratically equitable (fair) and more inalienable (inclusive) than offline or online identity-based approaches can achieve. The

main cost to citizens imposed by this social contract is simply to show up and “prove personhood” periodically.

While promising, many social, process, and implementation challenges remain to develop and test the viability of proof-of-personhood. Addressing these challenges remains an ongoing research project. [Ford, 2019e]

## 6 An Architecture for Digital Democracy

We have explored several levels of societal functionality that appear to be critical to effective democracy: deliberation and choice, information curation, inclusion and economic empowerment, personhood and digital citizenship. We have explored ways technology attempts to address these levels of functionality, ways it fails to do so, and potential ways we might improve our technology to address some of those flaws.

We now attempt to stitch these functionality levels together and look at them as a sketch for a potential *architecture for digital democracy*. This architectural perspective is directly inspired by classic layered network architectures such as the OSI model, [Day and Zimmermann, 1983] which attempts to decompose functionality into layers so that higher-level layers providing more sophisticated functionality depends only on that of the simpler services of lower layers. Taking this inspiration, we might arrange the functional layers of digital democracy described in the above sections as follows:

Democratic Deliberation and Choice
Information Filtering and Curation
Inclusion and Economic Empowerment
Personhood and Digital Citizenship

Although this is only one of no doubt many potential architectural perspectives and is most likely not complete or perfect, we can at least briefly justify this particular layering as follows, from bottom to top.

At the base level we need personhood and digital citizenship – specifically, some technological mechanism to different real people from fake accounts, whether or not that means identifying them in a traditional sense – in order to enable all the layers above to function securely and provide inclusion and democratic equality. Without a secure personhood foundation, financial inclusion technologies such as cryptocurrencies cannot allocate stake or resources (*e.g.*, crypto-UBI) fairly among real people, information filtering and curation technologies are vulnerable to sock puppetry and content reputation manipulation attacks, and deliberation and choice mechanisms are vulnerable to trolling and ballot stuffing. Online democracy can never be legitimate, either in fact or in public perception, without a legitimate *demos* comprised of real people.

At the next level up, citizens of democracies need a stable social and economic “floor” to stand on before they



can be expected to take time for or prioritize enlightened participation in democracy. This is simply the inevitable principle of “survival first.” A UBI or crypto-UBI might or might not be the right economic mechanism to help ensure such an economic floor and the assurance of personal independence of dignity it is intended to provide. However, it seems that every conceivable such mechanism, if *democratic*, will need to rely on some notion of personhood to allocate resources and services of all kinds equitably, and thus must be built atop some form of personhood and digital citizenship layer.

Given sufficient social and economic freedom to participate in democracy, citizens then need access to good information with which to make decisions, whose provision in whatever form is the function of the information filtering and curation layer. Again, we have explored some potential ways current abuse-ridden social information filtering and reputation systems might be improved and made more democratic, for example by relying on representative sample populations, with or without delegation capability, to prioritize topics, evaluate and filter information, and choose experts in a democratically egalitarian fashion. While we have not much discussed models for the funding and compensation of news and information, leaving that topic to other chapters in this volume, [Cagé, 2018] we might envision such funding and reward mechanisms building on the economic empowerment mechanisms of the layer below, such as cryptocurrencies supporting collective rewards or micropayments for information content. Regardless, since almost all realistic filtering and curation mechanisms become vulnerable if abusers can use Sybil attacks to inject fake upvotes/downvotes or reviews, this layer depends like the others on the personhood and citizenship foundation.

Finally, at the top level, we feel ready to envision more solid digital mechanisms for democratic participation, deliberation, and choice, building on all the functionality of the lower layers. It may not be too far off the mark to consider this layer the “mind” of the democratic digital collective: the decentralized organ at which the deliberative body hopefully achieves awareness and well-informed collective decision-making capability. We can hope for this collective “mind” to make truly democratic decisions, reflecting the interests of the entire population, only if it has the critical lower architectural layers to build on: layers ensuring that people have good information with which to make decisions, that guarantee a universal baseline of access to the time and economic resources to do so, and that protect participants’ rights as people from both individual exclusion and collective manipulation through digital fakery.

Again, we offer this perspective only as a likely-incomplete and imperfect sketch of a potential reference model fitting together a few of the critical support functions for digital democracy. The hope is merely that it provide be a useful starting point to think about and build on.

## 7 Conclusion

Inspired by Robert Dahl’s analysis of critical elements of effective democracy, [Dahl, 1961, Dahl, 1989] we have attempted a high-level exploration of key areas of functionality where digital technologies seem relevant to the mechanisms of democracy, but are currently failing to fill these roles reliably or securely. In this exploration, we have attempted to fit together these functionality areas into a layered architectural perspective designed around the principle of ensuring that higher layers depend only on lower layers, but derive from those lower layers all the functional services they need in order to operate in a reliable, secure, abuse-resistant, and democratically egalitarian fashion. All elements of digital democratic functionality seem to depend fundamentally on a currently-missing personhood or digital citizenship foundation to distinguish real people from fake Sybil accounts. The inclusion and economic empowerment layer depends on the personhood layer to build a “floor” of economic freedom and financial empowerment for all digital citizens to stand on and be able to have the time for real democratic participation. The information filtering and curation layer ensures that citizens have access to good information, depending on the economic layer to fund the production of information and the personhood layer to ensure that information filtering and curation is broad, representative, and objectively unbiased. Finally, the deliberation and choice layer builds on all lower layers – personhood to ensure “one-person-one-vote” equality in participation, the economic layer to assure the time and economic freedom to participate, and the information layer to support enlightened understanding. While only the barest sketch, this architectural perspective might help us break down and think about the complex problems of digital democracy in a more modular, systematic fashion than has been typical, and hopefully will provide a baseline for more detailed future architectural models for digital democracy to build from.

## References

- [lib, 2003] (2003). ‘Democratising’ expertise, ‘expertising’ democracy: what does this mean, and why bother? *Science and Public Policy*, 30(3):146–150.
- [Abdul-Rahman, 1997] Abdul-Rahman, A. (1997). The PGP Trust Model. *Journal of Electronic Commerce*, 10(3):27–31.
- [Abhishek and Mandal, 2017] Abhishek, K. and Mandal, D. (2017). Digital ID Verification: Competitive Analysis of Key Players. MEDICI.
- [Abraham et al., 2018] Abraham, R., Bennett, E. S., Bhusal, R., Dubey, S., Li, Q. S., Pattanayak, A., and Shah, N. B. (2018). State of Aadhaar Report 2017-18. Technical report, IDinsight.

- [Abraham et al., 2017] Abraham, R., Bennett, E. S., Sen, N., and Shah, N. B. (2017). State of Aadhaar Report 2016-17. Technical report, IDinsight.
- [Acemoglu et al., 2018] Acemoglu, D., Hassan, T. A., and Tahoun, A. (2018). The Power of the Street: Evidence from Egypt’s Arab Spring. *The Review of Financial Studies*, 31(1):1–42.
- [Agarwal, 2018] Agarwal, A. (2018). On-chain Liquid Democracy. *Medium*.
- [Aker and Mbiti, 2010] Aker, J. C. and Mbiti, I. M. (2010). Mobile Phones and Economic Development in Africa. *Journal of Economic Perspectives*, 24(3):207–232.
- [Alvarez et al., 2009] Alvarez, R. M., Hall, T. E., and Trechsel, A. H. (2009). Internet voting in comparative perspective: The case of estonia. *Political Science & Politics*, 42(3):497–505.
- [Ananny, 2018] Ananny, M. (2018). Presence of absence:exploring the democratic significance of silence. Digital Technology and Democratic Theory Workshop.
- [Armingeon and Schädel, 2015] Armingeon, K. and Schädel, L. (2015). Social Inequality in Political Participation: The Dark Sides of Individualisation. *West European Politics*, 38(3).
- [Asongu, 2013a] Asongu, S. A. (2013a). The impact of mobile phone penetration on African inequality. Technical Report AGDI Working Paper, No. WP/13/021, African Governance and Development Institute (AGDI).
- [Asongu, 2013b] Asongu, S. A. (2013b). How has Mobile Phone Penetration Stimulated Financial Development in Africa? *Journal of African Business*, 14(1):7–18.
- [Bao et al., 2009] Bao, W., Li, H., Li, N., and Jiang, W. (2009). A Liveness Detection Method for Face Recognition Based on Optical Flow Field. In *International Conference on Image Analysis and Signal Processing*.
- [Barberá et al., 2015] Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., and Bonneau, R. (2015). Tweeting From Left to Right: Is Online Political Communication More Than an Echo Chamber? *Psychological Science*, 26(10):1531–1542.
- [Becker, 2012] Becker, S. (2012). Liquid Democracy: Web Platform Makes Professor Most Powerful Pirate. *Spiegel Online*.
- [Behrens, 2014] Behrens, J. (2014). The evolution of proportional representation in LiquidFeedback. *Liquid Democracy Journal*, 1.

- [Berger, 2018] Berger, A. (2018). Bot vs. Bot: Will the Internet Soon Be a Place Without Humans? Singularity Hub.
- [Bernholz, 2018] Bernholz, L. (2018). Toward a new nonprofit: Civil society institutions in the digital age. Digital Technology and Democratic Theory Workshop.
- [Bessi and Ferrara, 2016] Bessi, A. and Ferrara, E. (2016). Social bots distort the 2016 U.S. Presidential election online discussion. *First Monday*, 21(11).
- [Bhatia and Bhabha, 2017] Bhatia, A. and Bhabha, J. (2017). India’s Aadhaar scheme and the promise of inclusive social protection. *Oxford Development Studies*, 45(1):64–79.
- [Bidadanure, 2019] Bidadanure, J. U. (2019). The Political Theory of Universal Basic Income. *Annual Review of Political Science*, 22.
- [Bischoff, 2018] Bischoff, P. (2018). Passports on the dark web: how much is yours worth? *Comparitech*.
- [Blinder, 2019] Blinder, A. (2019). New Election Ordered in North Carolina Race at Center of Fraud Inquiry. *The New York Times*.
- [Blum and Zuber, 2016] Blum, C. and Zuber, C. I. (2016). Liquid democracy: Potentials, problems, and perspectives. *The Journal of Political Philosophy*, 24(2):162–182.
- [Blumler and Gurevitch, 2001] Blumler, J. G. and Gurevitch, M. (2001). The New Media and Our Political Communication Discontents: Democratizing Cyberspace. 4(1):1–13.
- [Boldi et al., 2011] Boldi, P., Bonchi, F., Castillo, C., and Vigna, S. (2011). Viscous democracy for social networks. *Communications of the ACM*, 54(6).
- [Boneh et al., 2016] Boneh, D., Corrigan-Gibbs, H., and Schechter, S. (2016). Balloon Hashing: A Memory-Hard Function Providing Provable Protection Against Sequential Attacks. In *Asiacrypt*.
- [Borge et al., 2017] Borge, M., Kokoris-Kogias, E., Jovanovic, P., Gailly, N., Gasser, L., and Ford, B. (2017). Proof-of-Personhood: Redemocratizing Permissionless Cryptocurrencies. In *1st IEEE Security and Privacy on the Blockchain*.
- [Brewster, 2003] Brewster, M. (2003). *Unaccountable: How the Accounting Profession Forfeited a Public Trust*. Wiley.
- [Brin, 2005] Brin, D. (2005). *Kiln People*. Tor Books.

- [Broniatowski et al., 2018] Broniatowski, D. A., Jamison, A. M., Qi, S., AlKulaib, L., Chen, T., Benton, A., and Dredze, S. C. Q. M. (2018). Weaponized Health Communication: Twitter Bots and Russian Trolls Amplify the Vaccine Debate. *American Journal of Public Health*.
- [Browning, 2018] Browning, C. R. (2018). The Suffocation of Democracy. *New York Review of Books*.
- [Bu et al., 2013] Bu, Z., Xia, Z., and Wang, J. (2013). A sock puppet detection algorithm on virtual spaces. *Knowledge-Based Systems*, 37:366–377.
- [Budak et al., 2016] Budak, C., Goel, S., and Rao, J. M. (2016). *Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis*, 80(S1):250–271.
- [Cagé, 2018] Cagé, J. (2018). From philanthropy to democracy: Rethinking governance and funding of high-quality news in the digital age. Digital Technology and Democratic Theory Workshop.
- [Caplan, 2018] Caplan, R. (2018). Content standards and the (private) governance of speech. Digital Technology and Democratic Theory Workshop.
- [Chanthadavong, 2015] Chanthadavong, A. (2015). Biometrics: The password you cannot change. *ZDNet*.
- [Chaudhuri and König, 2017] Chaudhuri, B. and König, L. (2017). The Aadhaar scheme: a cornerstone of a new citizenship regime in India? *Contemporary South Asia*, 26(2):127–142.
- [Chellapilla et al., 2005] Chellapilla, K., Larson, K., Simard, P., and Czerwinski, M. (2005). Computers beat humans at single character recognition in reading based human interaction proofs (HIPs). In *2nd Conference on E-mail and Anti-Spam*.
- [Chesney and Citron, 2018] Chesney, R. and Citron, D. K. (2018). Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security. *California Law Review*.
- [Cho, 2015] Cho, Y. (2015). How Well are Global Citizenries Informed about Democracy? Ascertain the Breadth and Distribution of Their Democratic Enlightenment and Its Sources. *Political Studies*, 63(1):240–258.
- [Cohen and Fung, 2018] Cohen, J. and Fung, A. (2018). Democracy, design, and the digital public sphere. Digital Technology and Democratic Theory Workshop.
- [Collins and Evans, 2002] Collins, H. and Evans, R. (2002). The Third Wave of Science Studies: Studies of Expertise and Experience. 32(2):235–296.

- [Comminos, 2011] Comminos, A. (2011). Twitter revolutions and cyber crackdowns: User-generated content and social networking in the Arab spring and beyond.
- [Cost, 2015] Cost, J. (2015). *A Republic No More: Big Government and the Rise of American Political Corruption*. Encounter Books.
- [Cranor and LaMacchia, 1998] Cranor, L. F. and LaMacchia, B. A. (1998). Spam! *Communications of the ACM*, 41(8).
- [Crichton, 2018] Crichton, D. (2018). Liquid democracy uses blockchain to fix politics, and now you can vote for it. *TechCrunch*.
- [Dahl, 1961] Dahl, R. A. (1961). *Who Governs? Democracy and Power in an American City*. Yale University Press.
- [Dahl, 1989] Dahl, R. A. (1989). *Democracy and Its Critics*. Yale University Press.
- [Daian et al., 2018] Daian, P., Kell, T., Miers, I., and Juels, A. (2018). On-Chain Vote Buying and the Rise of Dark DAOs. Hacking, Distributed.
- [Day and Zimmermann, 1983] Day, J. D. and Zimmermann, H. (1983). The OSI Reference Model. *Proceedings of the IEEE*, 71(12).
- [de Vries, 2018] de Vries, A. (2018). Bitcoin’s Growing Energy Problem. *Joule*, 2(5):801–805.
- [Digiconomist, 2019] Digiconomist (2019). Bitcoin Energy Consumption Index.
- [Dixon, 2017] Dixon, P. (2017). A Failure to “Do No Harm” – India’s Aadhaar biometric ID program and its inability to protect privacy in relation to measures in Europe and the U.S. *Health and Technology*, 7(4):539–567.
- [Douceur, 2002] Douceur, J. R. (2002). The Sybil Attack. In *1st International Workshop on Peer-to-Peer Systems (IPTPS)*.
- [Doughty, 2005] Doughty, C. (2005). Know your customer: Automation is key to comply with legislation. *Business Information Review*, 22(4):248–252.
- [Doward, 2018] Doward, J. (2018). The big tech backlash: Tech giants are drawing political fire over fake news and Russian meddling. *The Guardian*.
- [Dubois and Blank, 2018] Dubois, E. and Blank, G. (2018). The echo chamber is overstated: the moderating effect of political interest and diverse media. 21(5):729–745.

- [Durdun, 2015] Durdun, T. (2015). From \$1,300 Tiger Penis To \$800K Snipers: The Complete Black Market Price Guide. *ZeroHedge*.
- [Dürst, 2004] Dürst, H. (2004). The “landsgemeinde”: the cantonal assembly of Glarus (Switzerland), history, present and future. In *IX Congreso Internacional del CLAD sobre la Reforma del Estado y de la Administración Pública*.
- [Dwork and Naor, 1992] Dwork, C. and Naor, M. (1992). Pricing via Processing or Combatting Junk Mail. In *12th Advances in Cryptology (CRYPTO)*.
- [Eberl et al., 2015] Eberl, J.-M., Boomgaarden, H. G., and Wagner, M. (2015). One Bias Fits All? Three Types of Media Bias and Their Effects on Party Preferences. *Communication Research*, 44(8):1125–1148.
- [Edens et al., 2012] Edens, J. F., Smith, S. T., Magyar, M. S., Mullen, K., Pitta, A., and Petrila, J. (2012). “Hired Guns,” “Charlatans,” and Their “Voodoo Psychobabble”: Case Law References to Various Forms of Perceived Bias Among Mental Health Expert Witnesses. *Psychological Services*, 9(3):259–271.
- [Enikolopov et al., 2018] Enikolopov, R., Makarin, A., and Petrova, M. (2018). Social Media and Protest Participation: Evidence from Russia.
- [Esau et al., 2017] Esau, K., Friess, D., and Eilders, C. (2017). Design Matters! An Empirical Analysis of Online Deliberation on Different News Platforms. *Policy and Internet*, 9(3):321–342.
- [eternalgloom, 2018] eternalgloom (2018). Overview of Universal Basic Income Crypto Projects. Bitcoin Forum.
- [Farrell and Schwartzberg, 2018] Farrell, H. and Schwartzberg, M. (2018). The democratic consequences of the new public sphere. Digital Technology and Democratic Theory Workshop.
- [Feldman, 2006] Feldman, L. C. (2006). *Citizens without Shelter: Homelessness, Democracy, and Political Exclusion*. Cornell University Press.
- [Ferrara et al., 2016] Ferrara, E., Varol, O., Davis, C., Menczer, F., and Flammini, A. (2016). The Rise of Social Bots. *Communications of the ACM*, 59(7).
- [Filetti, 2016] Filetti, A. (2016). Participating Unequally? Assessing the Macro-Micro Relationship Between Income Inequality and Political Engagement in Europe. *Partecipazione e Conflitto*, 9(1).
- [Fishkin, 1993] Fishkin, J. S. (1993). *Democracy and Deliberation: New Directions for Democratic Reform*. Yale University Press.

- [Flavin, 2015] Flavin, P. (2015). Campaign Finance Laws, Policy Outcomes, and Political Equality in the American States. *Political Research Quarterly*, 68(1).
- [Ford, 2002] Ford, B. (2002). Delegative Democracy.
- [Ford, 2014] Ford, B. (2014). Delegative Democracy Revisited.
- [Ford, 2015] Ford, B. (2015). Let's verify real people, not real names.
- [Ford, 2019a] Ford, B. (2019a). The Remote Voting Minefield: from North Carolina to Switzerland.
- [Ford, 2019b] Ford, B. (2019b). A Liquid Perspective on Democratic Choice. Draft (forthcoming).
- [Ford, 2019c] Ford, B. (2019c). Democratic Value and Money for Decentralized Digital Society. Draft (forthcoming).
- [Ford, 2019d] Ford, B. (2019d). Experts and Charlatans, Breakthroughs and Bandwagons: Collectively Distinguishing Signal from Noise Under Attack. Draft (forthcoming).
- [Ford, 2019e] Ford, B. (2019e). Privacy-preserving foundation for online personal identity. ONR grant N000141912361.
- [Ford and Strauss, 2008] Ford, B. and Strauss, J. (2008). An offline foundation for online accountable pseudonyms. In *1st International Workshop on Social Network Systems (SocialNets)*.
- [Forget, 2011] Forget, E. L. (2011). The Town with No Poverty: The Health Effects of a Canadian Guaranteed Annual Income Field Experiment. 37(3):283–305.
- [Freitas et al., 2015] Freitas, C. A., Benevenuto, F., Ghosh, S., and Veloso, A. (2015). Reverse Engineering Socialbot Infiltration Strategies in Twitter. In *Advances in Social Networks Analysis and Mining (ASONAM)*, pages 25–32.
- [Gangadharan, 2018] Gangadharan, S. (2018). Digital exclusion as a political resource. Digital Technology and Democratic Theory Workshop.
- [Gemberling and Cramer, 2014] Gemberling, T. M. and Cramer, R. J. (2014). Expert Testimony on Sensitive Myth-Ridden Topics: Ethics and Recommendations for Psychological Professionals. *Professional Psychology: Research and Practice*, 45(2):120–127.
- [Gerlach and Gasser, 2009] Gerlach, J. and Gasser, U. (2009). Three Case Studies from Switzerland: E-Voting. Technical Report Research Publication No. 2009-03.1, Berkman Center.



- [Germann and Serdült, 2017] Germann, M. and Serdült, U. (2017). Internet voting and turnout: Evidence from Switzerland. *Electoral Studies*, 47:1–12.
- [Ghosh et al., 2012] Ghosh, S., Viswanath, B., Kooti, F., Sharma, N. K., Korlam, G., Benevenuto, F., Ganguly, N., and Gummadi, K. P. (2012). Understanding and Combating Link Farming in the Twitter Social Network. In *21st International Conference on World Wide Web (WWW)*.
- [Gilad et al., 2017] Gilad, Y., Hemo, R., Micali, S., Vlachos, G., and Zeldovich, N. (2017). Algorand: Scaling Byzantine Agreements for Cryptocurrencies.
- [Gilens and Page, 2014] Gilens, M. and Page, B. I. (2014). Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens. *Perspectives on Politics*, 12(3):564–581.
- [Gölz et al., 2018] Gölz, P., Kahng, A., Mackenzie, S., and Procaccia, A. D. (2018). The Fluid Mechanics of Liquid Democracy. arXiv preprint 1808.01906v1.
- [Green-Armytage, 2014] Green-Armytage, J. (2014). Direct Voting and Proxy Voting. *Constitutional Political Economy*, 26(2):190–220.
- [Grönlund et al., 2009] Grönlund, K., Strandberg, K., and Himmelroos, S. (2009). The challenge of deliberative democracy online – a comparison of face-to-face and virtual experiments in citizen deliberation. *Information Polity*, 14(3):187–201.
- [Hajnal et al., 2017] Hajnal, Z., Lajevardi, N., and Nielson, L. (2017). Voter Identification Laws and the Suppression of Minority Votes. *The Journal of Politics*, 79(2):363–379.
- [Hardt and Lopes, 2015] Hardt, S. and Lopes, L. C. R. (2015). Google Votes: A Liquid Democracy Experiment on a Corporate Social Network.
- [Hauben and Hauben, 1997] Hauben, M. and Hauben, R. (1997). *Netizens: On the History and Impact of Usenet and the Internet*. IEEE Computer Society Press.
- [Havocscope, 2019] Havocscope (2019). Fake ID Cards, Driver Licenses, and Stolen Passports.
- [Hawkins et al., 2018] Hawkins, S., Yudkin, D., Juan-Torres, M., and Dixon, T. (2018). Hidden Tribes: A Study of America’s Polarized Landscape. More in Common report.
- [Heinlein, 1966] Heinlein, R. A. (1966). *The Moon is a Harsh Mistress*. G. P. Putnam’s Sons.

- [Hicks et al., 2015] Hicks, W. D., McKee, S. C., Sellers, M. D., and Smith, D. A. (2015). A Principle or a Strategy? Voter Identification Laws and Partisan Competition in the American States. *Political Research Quarterly*, 68(1):18–33.
- [Highton, 2017] Highton, B. (2017). Voter Identification Laws and Turnout in the United States. 20:149–167.
- [Hill and Hughes, 2007] Hill, K. A. and Hughes, J. E. (2007). Is the Internet an Instrument of Global Democratization? *Democratization*, 6(2):99–127.
- [Horowitz, 2018] Horowitz, J. (2018). The Mystery Man Who Runs Italy’s ‘Five Star’ From the Shadows. *The New York Times*.
- [Howard and Hussain, 2013] Howard, P. N. and Hussain, M. M. (2013). *Democracy’s Fourth Wave? Digital Media and the Arab Spring*. Oxford University Press.
- [Iyengar et al., 2003] Iyengar, S., Luskin, R. C., and Fishkin, J. S. (2003). Facilitating Informed Public Opinion: Evidence from Face-to-face and Online Deliberative Polls. In *Annual Meeting of the American Political Science Association*.
- [Iyengar and Westwood, 2015] Iyengar, S. and Westwood, S. J. (2015). Fear and Loathing across Party Lines: New Evidence on Group Polarization. *American Journal of Political Science*, 59(3):690–707.
- [Jackson and Victor, 2018] Jackson, T. and Victor, P. (2018). Confronting inequality in a post-growth world – Basic income, factor substitution and the future of work. Technical Report Working Paper No 11, Centre for the Understanding of Sustainable Prosperity.
- [Kalla and Broockman, 2016] Kalla, J. L. and Broockman, D. E. (2016). Campaign Contributions Facilitate Access to Congressional Officials: A Randomized Field Experiment. *American Journal of Political Science*, 60(3):545–558.
- [Kamvar et al., 2003] Kamvar, S. D., Schlosser, M. T., and Garcia-Molina, H. (2003). The EigenTrust Algorithm for Reputation Management in P2P Networks. In *12th International Conference on World Wide Web (WWW)*, pages 640–651.
- [Kara, 2017] Kara, S. (2017). *Modern Slavery: A Global Perspective*. Columbia University Press.
- [Kaye and Johnson, 2016] Kaye, B. K. and Johnson, T. J. (2016). Across the Great Divide: How Partisanship and Perceptions of Media Bias Influence Changes in Time Spent with Media. *Journal of Broadcasting & Electronic Media*, 60(4):604–623.

- [Keller and Kelly, 2015] Keller, E. and Kelly, N. J. (2015). Partisan Politics, Financial Deregulation, and the New Gilded Age. *Political Research Quarterly*, 68(3):428–442.
- [Kiayias et al., 2016] Kiayias, A., Russell, A., David, B., and Oliynykov, R. (2016). Ouroboros: A Provably Secure Proof-of-Stake Blockchain Protocol. Cryptology ePrint Archive, Report 2016/889.
- [Kline, 2008] Kline, C. L. (2008). Security Theater and Database-Driven Information Markets: a Case for an Omnibus U.S. Data Privacy Statute. *The University of Toledo Law Review*, 39(1).
- [Koistinen and Perkiö, 2014] Koistinen, P. and Perkiö, J. (2014). Good and Bad Times of Social Innovations: The Case of Universal Basic Income in Finland. *Basic Income Studies*, 9(1–2).
- [Kokoris-Kogias et al., 2018] Kokoris-Kogias, E., Jovanovic, P., Gasser, L., Gailly, N., Syta, E., and Ford, B. (2018). OmniLedger: A Secure, Scale-Out, Decentralized Ledger via Sharding. In *39th IEEE Symposium on Security and Privacy (SP)*, pages 19–34. IEEE.
- [Koprowski, 2003] Koprowski, G. J. (2003). Spam filtering and the plague of false positives. *TechNewsWorld*.
- [Kotzee, 2012] Kotzee, B. (2012). Expertise, fluency and social realism about professional knowledge. *Journal of Education and Work*, 27(2):161–178.
- [Landemore, 2018] Landemore, H. (2018). Open democracy and digital technologies. Digital Technology and Democratic Theory Workshop.
- [Lee et al., 2018] Lee, D., Levi, M., and Brown, J. S. (2018). Democratic societal collaboration in a whitewater world. Digital Technology and Democratic Theory Workshop.
- [Litvinenko, 2012] Litvinenko, A. (2012). Social media and perspectives of liquid democracy: The example of political communication in the pirate party in germany. In *12th European Conference on e-Government*, pages 403–408.
- [Liu et al., 2016] Liu, D., Wu, Q., Han, W., and Zhou, B. (2016). Sockpuppet gang detection on social media sites. *Frontiers of Computer Science*, 10(1):124–135.
- [Loucaides, 2019] Loucaides, D. (2019). What Happens When Techno-Utopians Actually Run a Country. *Wired*.
- [Luechinger et al., 2007] Luechinger, S., Rosinger, M., and Stutzer, A. (2007). The impact of postal voting on participation: Evidence for switzerland. *Swiss Political Science Review*, 13(2):167–202.

- [Luu et al., 2016] Luu, L., Narayanan, V., Zheng, C., Baweja, K., Gilbert, S., and Saxena, P. (2016). A Secure Sharding Protocol For Open Blockchains. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS '16*, pages 17–30, New York, NY, USA. ACM.
- [MacDorman, 2006] MacDorman, K. F. (2006). Subjective ratings of robot video clips for human likeness, familiarity, and eeriness: An exploration of the uncanny valley. In *ICCS/CogSci-2006 Symposium: Toward Social Mechanisms of Android Science*, page 26–29.
- [Mack, 2018] Mack, D. (2018). This PSA About Fake News From Barack Obama Is Not What It Appears. *Buzzfeed News*.
- [Madestam et al., 2013] Madestam, A., Shoag, D., Veuger, S., and Yanagizawa-Drott, D. (2013). Do Political Protests Matter? Evidence from the Tea Party Movement. *The Quarterly Journal of Economics*, 128(4):1633–1685.
- [Manza and Uggen, 2008] Manza, J. and Uggen, C. (2008). *Locked Out: Felon Disenfranchisement and American Democracy*. Oxford University Press.
- [Mathews, 2016] Mathews, H. V. (2016). Flaws in the UIDAI Process. *Economic & Political Weekly*, 51(9).
- [May, 2019] May, M. (2019). Inaccessibility of CAPTCHA: Alternatives to Visual Turing Tests on the Web. W3C Working Draft 14.
- [Mbiti and Weil, 2016] Mbiti, I. and Weil, D. N. (2016). Mobile Banking: The Impact of M-Pesa in Kenya. In Edwards, S., Johnson, S., and Weil, D. N., editors, *African Successes, Volume III: Modernization and Development*, pages 247–293. University of Chicago Press.
- [Mendez and Serdült, 2017] Mendez, F. and Serdült, U. (2017). What drives fidelity to internet voting? Evidence from the roll-out of internet voting in Switzerland. *Government Information Quarterly*, 34(3):511–523.
- [Messias et al., 2013] Messias, J., Schmidt, L., Oliveira, R., and Benevenuto, F. (2013). You followed my bot! Transforming robots into influential users in Twitter. 18(7).
- [Miller III, 1969] Miller III, J. C. (1969). A program for direct and proxy voting in the legislative process. *Public Choice*, 7:107–113.
- [Mislove et al., 2007] Mislove, A., Marcon, M., Gummadi, K. P., Druschel, P., and Bhattacharjee, B. (2007). Measurement and analysis of online social networks. In *Internet Measurement Conference (IMC)*.

- [Mislove et al., 2008] Mislove, A., Post, A., Druschel, P., and Gummadi, K. P. (2008). Ostra: Leveraging trust to thwart unwanted communication. In *5th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 15–30.
- [Mori, 2012] Mori, M. (2012). The Uncanny Valley. *IEEE Robotics & Automation Magazine*. Translated by Karl F. MacDorman and Norri Kageki.
- [Moynihan, 2004] Moynihan, D. P. (2004). Building secure elections: E-voting, security, and systems theory. *Public Administration Review*, 64(5):515–528.
- [Nakamoto, 2008] Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System.
- [nanalyze, 2017] nanalyze (2017). 6 Digital Identity Verification Startups to Check Out.
- [Packer, 2018] Packer, G. (2018). A New Report Offers Insights Into Tribalism in the Age of Trump. *The New Yorker*.
- [Pan et al., 2008] Pan, G., Wu, Z., and Sun, L. (2008). Liveness Detection for Face Recognition. In *Recent Advances in Face Recognition*. IntechOpen.
- [Parijs, 2017] Parijs, P. V. (2017). *Basic Income: A Radical Proposal for a Free Society and a Sane Economy*. Harvard University Press.
- [Park et al., 2018] Park, S., Kwon, A., Fuchsbauer, G., Gaži, P., Alwen, J., and Pietrzak, K. (2018). SpaceMint: A Cryptocurrency Based on Proofs of Space.
- [Piketty, 2017] Piketty, T. (2017). *Capital in the Twenty-First Century*. Belknap Press.
- [Prior, 2013] Prior, M. (2013). Media and political polarization. *Annual Review of Political Science*, 16:101–127.
- [Quinlan, 2015] Quinlan, S. (2015). Facilitating the Electorate: A Multilevel Analysis of Election Timing, Registration Procedures, and Turnout. *Irish Political Studies*, 30(4):482–509.
- [Ramachandran et al., 2006] Ramachandran, A., Dagon, D., and Feamster, N. (2006). Can DNS-based blacklists keep up with bots? In *3rd Conference on Email and Anti-Spam*.
- [Ratcliffe, 2019] Ratcliffe, R. (2019). How a glitch in India’s biometric welfare system can be lethal. *The Guardian*.
- [Read, 2018] Read, M. (2018). How Much of the Internet Is Fake? Turns Out, a Lot of It, Actually. *New York Magazine*.

- [Reich, 2018] Reich, R. (2018). *Just Giving Why Philanthropy Is Failing Democracy and How It Can Do Better*. Princeton University Press.
- [Reinisch, 2007] Reinisch, C. (2007). Swiss *Landsgemeinden*: a deliberative democratic evaluation of two outdoor parliaments. In *ECPR Joint Sessions*.
- [Ribeiro et al., 2018] Ribeiro, F. N., Henriqueo, L., Benevenuto, F., Chakraborty, A., Kulshrestha, J., Babaei, M., and Gummadi, K. P. (2018). Media bias monitor: Quantifying biases of social media news outlets at large-scale. In *12th International AAAI Conference on Web and Social Media (ICWSM)*.
- [Saez-Trumper et al., 2013] Saez-Trumper, D., Castillo, C., and Lalmas, M. (2013). Social Media News Communities: Gatekeeping, Coverage, and Statement Bias. In *22nd ACM International Conference on Information & Knowledge Management (CIKM)*, pages 1679–1684.
- [Samuel, 2018] Samuel, I. (2018). Rigging the vote: how the American right is on the way to permanent minority rule. *The Guardian*.
- [Sayke, 2003] Sayke (2003). Liquid democracy. <https://web.archive.org/web/20040726071737/http://twistedmatrix.com/wiki/python/LiquidDemocracy>.
- [Schaub, 2012] Schaub, H.-P. (2012). Maximising Direct Democracy – by Popular Assemblies or by Ballot Votes? *Swiss Political Science Review*, 18(3):305–331.
- [Schneier, 2009] Schneier, B. (2009). Tigers use scent, birds use calls – biometrics are just animal instinct. *The Guardian*.
- [Schreiber, 1973] Schreiber, F. R. (1973). *Sybil: the true story of a woman possessed by sixteen separate personalities*. Warner Books.
- [Schryen and Rich, 2009] Schryen, G. and Rich, E. (2009). Security in Large-Scale Internet Elections: A Retrospective Analysis of Elections in Estonia, The Netherlands, and Switzerland. *IEEE Transactions on Information Forensics and Security*, 4(4):729–744.
- [Scott, 2017] Scott, G. L. (2017). 5 Lifelike Robots That Take You Straight Into the Uncanny Valley. Inverse.
- [Serdült et al., 2015] Serdült, U., Germann, M., Mendez, F., Portenier, A., and Wellig, C. (2015). Fifteen Years of Internet Voting in Switzerland: History, Governance and Use. In *2nd International Conference on eDemocracy & eGovernment (ICEDEG)*.

- [Sethi et al., 2017] Sethi, T. S., Kantardzic, M., and Ryu, J. W. (2017). ‘Security Theater’: On the Vulnerability of Classifiers to Exploratory Attacks. In *12th Pacific Asia Workshop on Intelligence and Security Informatics (PAISI)*.
- [Shao et al., 2018] Shao, C., Ciampaglia, G. L., Varol, O., Yang, K.-C., Flammini, A., and Menczer, F. (2018). The spread of low-credibility content by social bots. *Nature Communications*, 9.
- [Shaw, 2004] Shaw, R. (2004). Avoid the Spam Filter. *iMedia Connection*.
- [Shlapentokh and Woods, 2011] Shlapentokh, V. and Woods, J. (2011). *Feudal America: Elements of the Middle Ages in Contemporary Society*. Penn State University Press.
- [Smith, 2014] Smith, W. (2014). Political donations corrupt democracy in ways you might not realise.
- [Solorio et al., 2014] Solorio, T., Hasan, R., and Mizan, M. (2014). Sockpuppet Detection in Wikipedia: A Corpus of Real-World Deceptive Writing for Linking Identities. In *9th International Conference on Language Resources and Evaluation (LREC)*.
- [Srinivasan et al., 2018] Srinivasan, J., Bailur, S., Schoemaker, E., and Seshagiri, S. (2018). The Poverty of Privacy: Understanding Privacy Trade-Offs From Identity Infrastructure Users in India. *International Journal of Communication*, 12:1228–1247.
- [Stallings, 1995] Stallings, W. (1995). The PGP Web of Trust. *BYTE Magazine*, 20(2):161–164.
- [Standing, 2011] Standing, G. (2011). *The Precariat: The New Dangerous Class*. Bloomsbury Academic.
- [Standing, 2017] Standing, G. (2017). *Basic Income: A Guide for the Open-Minded*. Yale University Press.
- [Swierczek, 2014] Swierczek, B. (2014). 5 years of Liquid Democracy in Germany. *The Liquid Democracy Journal*, 1.
- [Templeton, 2001a] Templeton, B. (2001a). I Remember USENET. O’Reilly Network.
- [Templeton, 2001b] Templeton, B. (2001b). Origin of the term “spam” to mean net abuse. O’Reilly Network.
- [Templeton, 2003] Templeton, B. (2003). Reflections on the 25th Anniversary of Spam. O’Reilly Network.
- [Tisdall, 2018] Tisdall, S. (2018). American democracy is in crisis, and not just because of Trump.
- [Tran et al., 2009] Tran, N., Min, B., Li, J., and Submaranian, L. (2009). Sybil-resilient online content voting. In *6th Symposium on Networked System Design and Implementation (NSDI)*, pages 15–28.

- [Tullock, 1967] Tullock, G. (1967). *Toward a Mathematics of Politics*. The University of Michigan Press.
- [Venkatanarayanan, 2017] Venkatanarayanan, A. (2017). Enrolment rejections are accelerating. Medium.
- [Viswanath et al., 2012] Viswanath, B., Mondal, M., Gummadi, K. P., Mislove, A., and Post, A. (2012). Canal: Scaling social network-based sybil tolerance schemes. In *EuroSys Workshop on Social Network Systems (SNS)*.
- [Viswanath and Post, 2010] Viswanath, B. and Post, A. (2010). An Analysis of Social Network-Based Sybil Defenses. In *ACM SIGCOMM (SIGCOMM)*.
- [von Ahn et al., 2003] von Ahn, L., Blum, M., Hopper, N. J., and Langford, J. (2003). CAPTCHA: using hard AI problems for security. In *Eurocrypt*.
- [Vorick, 2018] Vorick, D. (2018). The State of Cryptocurrency Mining.
- [Wagner, 2017] Wagner, C. (2017). The Swiss Universal Basic Income Vote 2016: What’s Next?
- [Weitzer, 2015] Weitzer, R. (2015). Human Trafficking and Contemporary Slavery. *Annual Review of Sociology*, 41:223–242.
- [Wen and Jain, 2015] Wen, D. and Jain, H. H. A. K. (2015). Face Spoof Detection With Image Distortion Analysis. *IEEE Transactions on Information Forensics and Security*, 10(4).
- [Woolley, 2016] Woolley, S. C. (2016). Automating power: Social bot interference in global politics. *First Monday*, 21(4).
- [Woolley and Guilbeault, 2017] Woolley, S. C. and Guilbeault, D. R. (2017). Computational Propaganda in the United States of America: Manufacturing Consensus Online. Working Paper No. 2017.5.
- [Yamak et al., 2016] Yamak, Z., Saunier, J., and Vercouter, L. (2016). Detection of Multiple Identity Manipulation in Collaborative Projects. In *25th International Conference Companion on World Wide Web (WWW Companion)*, pages 955–960.
- [Yu et al., 2008] Yu, H., Gibbons, P. B., Kaminsky, M., and Xiao, F. (2008). SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks. In *29th IEEE Symposium on Security and Privacy (S&P)*.
- [Yu et al., 2009] Yu, H., Shi, C., Kaminsky, M., Gibbons, P. B., and Xiao, F. (2009). DSybil: Optimal sybil-resistance for recommendation systems. In *30th IEEE Symposium on Security and Privacy (S&P)*.



[Zetter, 2019] Zetter, K. (2019). Experts Find Serious Problems With Switzerland's Online Voting System Before Public Penetration Test Even Begins. *Motherboard*.

[Zhang and Zhou, 2019] Zhang, B. and Zhou, H.-S. (2019). Statement Voting. In *23rd International Conference on Financial Cryptography and Data Security 2019 (FC)*.

[Zinn, 2005] Zinn, H. (2005). *A People's History of the United States*. Harper Perennial Modern Classics.