Against "Democratizing AI"

Johannes Himmelreich, Syracuse University

This paper argues against the call to democratize artificial intelligence (AI). Several authors demand to reap purported benefits that rest in direct and broad participation: In the governance of AI, more people should be more involved in more decisions about AI—from development and design to deployment. This paper opposes this call. The paper presents five objections against broadening and deepening public participation in the governance of AI. The paper begins by reviewing the literature and carving out a set of claims that are associated with the call to "democratize AI". It then argues that such a democratization of AI (1) rests on weak grounds because it does not answer to a demand of legitimization, (2) is redundant in that it overlaps with existing governance structures, (3) is resource intensive, which leads to injustices, (4) is morally myopic and thereby creates popular oversights and moral problems of its own, and finally, (5) is neither theoretically nor practically the right kind of response to the injustices that animate the call. The paper concludes by suggesting that AI should be democratized not by broadening and deepening participation but by increasing the democratic quality of the administrative and executive elements of collective decision making. In a slogan: The question is not so much whether AI should be democratized but how.

Keywords: Artificial Intelligence; Governance; AI Ethics; Democracy; Participation; Public Administration.

1. Introduction

Anyone who values democracy likely welcomes the call to democratize the governance of artificial intelligence (AI). While the idea to democratize technology is by no means novel (Sclove 1995), this call to "democratize AI" could recently be heard and seen in various venues—from *Wired* magazine, to the *Boston Review*, to academic journals (e.g. Rahwan 2018; Gould 2019; Wong 2020; Zimmermann, Di Rosa, and Kim 2020; Cammaerts and Mansell 2020). There is a lot at stake. The issues range from performance standards for autonomous vehicles, equity standards for facial recognition, to decisions which research avenues companies should pursue, which datasets they may use in training and testing AI models, and which AI technologies public and private organizations should deploy. The call to democratize AI often responds to existing injustices of AI in hiring, criminal justice, policing, education, or social service administration. The call is hence animated by concerns of fairness, freedom, and equality.[1]

[1] These three ideas—fairness, freedom and equality—mean different things. By "freedom" I understand the capacity to see one's will carried out and, more generally, a robust congruence between one's actions and the conditions of one's life on the one hand and the authentic expression of one's values on the other. By "equality" I understand, first, the tenet that each individual has the same moral worth and, second, that this tenet finds its expression in how individuals relate to each

In short, the call to democratize AI sounds attractive. The animating concerns are important, a lot is at stake, and—more generally—it is just good to democratize collective practices. But although the call to democratize AI sounds attractive, I argue that what this call demands is not persuasive upon reflection. In this paper, I describe five objections against the demand to democratize AI.

The upshot of this paper may sound anti-democratic. It is not. I take myself to be playing for the same team as those who demand to democratize AI—team democracy. I am worried about our game plan. The demand to democratize AI seems to be driven by understanding "democracy" as "rule by the people." Or at least, the demand emphasizes the value and importance of broad and inclusive participation. Moreover, the demand is abstract. It concentrates on the *whether* of participation not on the *how*. This abstract emphasis could be a fatal mistake. We see that authoritarian figures misappropriate the idea of democracy to advance their kleptocratic or nationalist mission. These authoritarian figures not only co-opt the name "democracy", they also weaponize the idea's substance (Mayer 2001).[2] With the gesture and solicitude of a tribune, authoritarians, claiming to be at one with the masses, take the "rule by the people" element of democracy to turn democracy against itself. Our times therefore require that democratic theorists identify the pitfalls and limits of participation. To clarify these limits is to defend democracy—both the ideal and the practice. This paper is part of this effort.

I hope that this paper does its part in changing the game plan of team democracy. At least when it comes to the governance of AI, we might be looking to increase democracy in the wrong places (the value of participation; defending the democratic ideal) and we, thereby—perhaps—, even misconstrue what democracy ought to look like.[3] This paper, in its conclusion, offers an outlook to counterbalance the emphasis on participation: It is important to improve the democratic quality of the bureaucracy, that is, of the administrative and executive components of our collective decision-making. Democratizing the administrative state—although not necessarily furthering participation—allows to overcome existing and future injustices—not only of AI.

I start with the question: What is meant exactly with admonishments that "AI is political" or demands that "AI needs to be democratized"? Such resounding calls for democracy are often shrouded in an aura of radicalism. They articulate a defect of the *status quo* and voice a call for change. But what exactly ought this change to be? The first task in this paper will be to carve out a set of claims from the slogan that "AI should be governed democratically". I review existing arguments and contributions to the recent literature to develop such a set of claims. The central demand is that, to address issues of injustice in AI, novel institutional avenues of legislative participation should be established. These avenues should employ direct individual participation from a broad and inclusive constituency and

other—that they relate to each other as equals. By "fairness" I mean an impartial appraisal of the reasons that each individual could offer on matters of common concern.

[2] Be there no doubt: Those who emphasize the participatory and populist elements of democracy are in no way to blame or even complicit in this authoritarian misappropriation.

[3] I argue for this latter claim elsewhere: Proposals, which nominally aim to improve democracy, often hollow out its values (Himmelreich 2022).

exert binding practical authority over both the coarse-grained regulatory domain as well as over fine-grained decisions concerning the development and deployment of AI. It is towards this demand, this set of claims, against which I then formulate five objections.

2. "Democratizing AI"

"Artificial Intelligence", or "AI", here refers to a set of machine learning (ML) techniques and the software features that are built with these techniques. The kinds of things that AI can do are vast. AI can be used to estimate traffic patterns, house prices, or crime rates. AI is already being used to identify students at the risk of suicide, unemployment insurance claims that are fraudulent, or harmful digital content. AI is particularly good at recognizing objects in images, matching faces to persons, or generating music or text. Among the somewhat catchy examples, AI has been used to pen opening lines for flirts on *Tinder*, defeat world champions in games such as Go and StarCraft II, or compose a Nirvana song.

Several authors demand that AI should be governed democratically. A classic early statement of the more general idea that technology should be democratized is already found in Sclove (1995). He gives his "basic argument" with the following syllogism.

> (1) citizens ought to be empowered to participate in their society's basic circumstances and (2) technologies profoundly affect and partly constitute those circumstances, it follows that (3) technological design and practices should be democratized. (Sclove 1995, ix)

More recently, similar demands have been made for algorithms,[4] AI, and big data research more specifically (e.g. Rahwan 2018; Gould 2019; Wong 2020; Zimmermann, Di Rosa, and Kim 2020; Cammaerts and Mansell 2020).[5] In what follows, I review some of these demands and arguments in order to distill out how the idea to "democratize AI" should be understood.

One ambiguity of the call to "democratize AI" should be clarified upfront. The demand is, broadly, that AI should be subject to novel or different forms of democratic governance. In contrast, "democratize AI" may instead refer to the demand that AI be made widely available or useable or that individuals should be empowered to develop or use AI. Training data should be accessible, AI technologies should be affordable, and their use should be encouraged. But "availability of technology" or "empowerment of individuals", although a common sense of "democratization", is not the sense of "democratization" with which this paper is concerned. The subject here is instead the organizational form and the mode of governance over how AI is developed and used.

Zimmermann, Di Rose and Kim (2020) formulate a call to democratize AI in a recent piece in the *Boston Review*. Their animating concerns are "algorithmic injustice", "algorithmic fairness" and

[4] Strictly speaking, algorithms are abstract objects, like theorems and arithmetic operations. It is not obvious how this vast class of abstract objects—and not just their implementations—are supposed to give rise to ethical problems.
[5] A related set of claims is defended for the practice of scientific research across the board by Kitcher (2011).

"algorithmic bias". They survey examples, such as the use of AI in recidivism risk-scoring—in which the average African-American is twice as likely to be disadvantageously misclassified as the average white American (Angwin et al. 2016)— and they arrive at the following main claims. They speak of "weak AI", by which they mean what is meant with just "AI" in this paper.

> [T]he responsible development and deployment of weak AI will involve not just developers and designers, but the public at large. […]
>
> *[D]eveloping and deploying* weak AI involves making consequential choices — choices that demand *greater democratic oversight* not just from AI developers and designers, but *from all members of society*. […]
>
> Broaching questions of algorithmic justice via the democratic process would give members of communities most impacted by algorithmic bias *more direct democratic power* over crucial decisions concerning weak AI—not merely after its deployment, but *also at the design stage*. (Zimmermann, Di Rosa, and Kim 2020 my emphasis)

Several aspects of these claims are noteworthy.

First, there is the aspect of *what* should be governed: Not only the *deployment* of AI—whether your local police department uses facial recognition—but also its *development* and *design*. This this includes questions such as whether specific facial recognition systems should be developed at all, if so, what data should be used in their development, and which specific model configuration and training techniques should be employed. The domain of democratization includes hence not only regulation but also individual decisions about development, design and deployment of AI.

Second, the claims articulate an aspiration about *who* should do the governing: "all members of society". Of course, this raises what is known as the democratic boundary problem (Miller 2009)—who exactly is meant with "all members"? This problem is particularly pressing since AI is global. As typical for software and the digital economy generally, AI, data or access to APIs can easily be transferred and marketed across the globe. The society that lives under AI is the global society.

Third, there is the aspect of *how* AI should be democratized: The authors call for increased or novel direct democratic powers.[6] They first demand more democratic "oversight". Because there is little point in demanding what is already the case, this claim pragmatically conveys a demand for something

[6] Because of this condition to increase or introduce direct democratic powers, the so-called Moral Machine experiment is not a form of democratizing AI. Some proponents of such surveys—and they are usually just *surveys* and *not experiments*—argue that public attitudes about ethics and technology must be studied, identified, and articulated to be "cognizant of public morality" (Awad et al. 2018). The idea is that the hence elicited public attitudes are to limit policymaking, because, otherwise, "societal push-back will drastically slow down the adoption of intelligent machines" (Awad et al. 2020). This approach is flawed (Jaques 2019; Himmelreich 2020). The overall idea contrasts with the call to "democratize AI" because it 1) aims mainly to inform and 2) sees individuals as *subjects* in an investigation. By contrast, the demand to democratize AI seeks to empower individuals, to endow marginalized groups with novel ways to make their voices heard (although in the wrong way as I argue here), and to give citizens greater direct influence—collective power—over decisions.

novel. The authors contend that this need not involve "entirely new democratic institutions and procedures" (Zimmermann, Di Rosa, and Kim 2020)—but obviously we need to go beyond whatever oversight we currently have. An example of an existing oversight institution might be Facebook's global Oversight Board (albeit whether this institution is democratic is another matter). Another existing oversight institution might be citizen review boards of local police departments (albeit whether these boards have sufficient power is doubtful). These examples show already that *mere* oversight is too often inefficacious. Thus, even where it doesn't involve new institutions, democratic oversight must have new powers to intervene, veto, or act. It therefore makes sense for the authors to arrive at the claim later in their piece that there should be "more direct democratic power".

Although he is writing in an earlier era of technology, Sclove (1995) is sympathetic to these three claims. He proposes increased "worker participation in corporate RD&D [research development and design]", "citizen oversight or participation in large corporations' technology development efforts" (Sclove 1995, 209), so as to "[allow] communities, groups, and citizens—including those today least empowered—to help directly initiate some RD&D programs and design technologies responsive to their needs." (Sclove 1995, 207). The *what*, *who*, and *how* of democratization seem almost identical here.

In sum, the demand to democratize AI consists of up to *three claims*. First, AI should be governed by reformed existing or novel institutions. These institutions should have greater—even if still limited—practical authority, that is, greater decision-making power. Second, these institutions should operate on direct participation by individuals. Third, the polity of these institutions is conceived as being very broad and highly inclusive. This aspiration is translated into the claim that democratizing AI should involve "all members of society".

Other proposals in the literature make similar claims. Some of the proposals may make the claims, or at least some of them, even more forcefully. One proposal is that AI should be governed by councils. This proposal is inspired by the organizational form of governance in ancient Athens (Carugati 2020; McQuillan 2018). Like the Athenian *ekklesia*, "[p]eople's councils are *bottom-up, confederated structures* that act as *direct democratic* assemblies." The details on how these councils would operate are somewhat opaque. But aspirations about how they should function are stated upfront. Councils are purported to "restore collective subjectivity and agency" so that "machine learning can be ethically reclaimed" (McQuillan 2018, 6–7). More specifically, councils are purported to increase transparency about how AI is used, they are said to prevent that AI is used for nefarious purposes; and they promise to legitimize difficult trade-offs between competing statistical definitions of "fairness". Overall, these councils will help address issues of justice, since "ethical problems raised by machine learning are primarily issues of justice" (ibid.). In other words, a democratization of AI through councils—governance by direct democratic assemblies—is instrumental for attaining justice, or such is the hope.

Another proposal has it that a democratization of AI would not be instrumental for, but constitutive of justice. Sloane et al. (2020) argue to deepen the participation in technology design. Currently, technology design has participation that is mostly superficially: Individuals participate as workers—if they

produce or provide data—or as consultants. As workers or consultants, individuals are "stakeholders" in an "episodic, short-term" fashion. Proper participation, what Sloane et al. call "participation as justice", is instead an "ongoing [relationship] based on mutual benefit, reciprocity, equity and justice." Examples of such participation are resident associations—an instance of local self-governance. On this proposal, the governance of AI broadly should take inspiration from such associations and be organized around participation that is deeper, that is, more direct, more active, and more powerful.

A final proposal concerns data trusts (S. Mills 2019). A data trust pools data of its members and makes these data selectively available, for example, to social media companies or developers of AI products. Such a data trust "resembles a data cooperative or a data union" and "negotiates access to the pooled data" in a way that "can be democratically coordinated with members." The data trust proposal, again, advocates to democratize AI.[7] The governance mode of data trusts resembles that of tenants' associations. The picture of democracy that is invoked, again, emphasizes broad and direct participation, modelled after small group deliberation.

To structure these ideas: These proposals each take up two questions: *what* should be democratized and *how* should it be democratized. The first issue—the *what* question—is the question about the *domain* of democratizing AI, whereas the second issue—the *how* question (and, derivatively, the *who* question)—can be called a question about the *organizational form*.

On these two questions, the proposals surveyed above give answers that are often heterogeneous and sometime vague. Although I understand the call to "democratize AI", on the *how* question, as a call for direct participation, some of these proposals above—especially the proposal on data trusts—may be most effective when organized using representative participation instead. Insofar as there are general trends in the literature surveyed above, the following ideas can plausibly be discerned.

On the organizational form, the *how* question, all the proposals surveyed above seem to subscribe to the three claims of democratizing AI. They contend that democratizing AI requires or consists in: (1) novel (legislative) institutions or avenues with (2) practical authority that is rooted in (3) direct individual and broad participation. To be clear, the demand is, effectively, not just to reform existing democratic institutions but to establish new ones. That these new institutions have practical authority means that they can make decisions that others will have to abide by. Finally, the proposals often signal at least a preference for direct participation taking, for example, Athenian direct democracy as their guiding inspiration.

On the domain of democratization—the *what* question—the call to "democratize AI" includes both coarse-grained issues of regulation as well fine-grained issues of concrete decisions. Issues of regulation include standards of performance, such as accuracy standards for facial recognition used by the police. Moreover, issues of regulation include norms of practice. One example for such norms of

[7] This proposal by Mills (2019) can be distinguished into a proposal about organizational function (the data trust) and a proposal about the trust's mode of governance (deeper participation). My argument is only about the latter.

practice is the US fair credit reporting act, which defines, roughly, a standard of fairness for credit scores, namely, that fairness consists in not using certain information about protected characteristics.

Importantly, in addition to such coarse-grained issues of regulation, the call to "democratize AI" may refer to fine-grained, concrete decisions about the development and deployment of AI, such as: Should this particular AI that aims to predict criminality from voice data be developed into a prototype? What data should its model be trained on? Should a model to predict recidivism be trained only on data about violent crimes or all crimes? Finally, with whom should a specific company do business? Should it offer AI cloud services to military or intelligence agencies?—these questions are examples fine-grained issues of *what* should be democratized.

Of course, not all the above proposals contend that fine-grained issues of AI should be democratized. Most of the proposals do not specify their governance domain clearly. But the proposals use such fine-grained issues as examples. Generally, the proposals above often respond to concrete cases of injustice that arise from fine-grained issues in the design, development and deployment of AI in a broad range of public and private contexts.

3. Five reasons against democratizing AI

Although I argue against democratizing AI, with those who call for this democratization, I largely agree on the problem: AI governance has a democracy gap. The ideal of democracy—as a system of social governance that gives everyone a fair and equal opportunity of influence—is far from realized today. In fact, this problem may be larger than AI governance. The broader system of scientific inquiry may have a democracy gap (Kitcher 2011).[8] The call to democratize AI rides on similar diagnoses and arguments as the call to democratize science. Similarly, in AI research as in scientific research, the democracy gap is deepened by the increasing privatization of scientific inquiry (Kitcher 2011, 126; Jurowetzki et al. 2021).[9]

But although we might agree on the problem, we disagree about the solution. More democracy does not mean more participation. Some of those who advocate to democratize science or technology, recognize that increasing participation is problematic. Kitcher (2011, 113), who advocates for a democratization of science, is opposed to increasing participation because of a tension between participation and expert judgment. Similarly Sclove (1995, 211), who advocates for a democratization of technology, concedes that "experts must continue to play a role" even though "lay members may be marginalized or intimidated" as a result. This tension between participation and expertise is one to which we, the community of researchers in AI and society, perhaps still need to wake up to.[10]

[8] Kitcher writes (2011, 127): "Current scientific research neglects the interests of a vast number of people, except insofar as their interests coincide with those of people in the affluent world."

[9] Kitcher also writes (2011, 126): "Privatization of scientific research will probably matters worse." Given that much research on AI is privatized and proprietary, the problems that animate Kitcher are amplified in the case of AI.

[10] Admittedly, some of the relevant experts in cases of AI injustice are those who suffer the injustice. It is their expertise that must find its way into our deliberation and collective decision-making. But I disagree that broadening and deepening participation is the right way of doing so.

Making the governance of AI more democratic is hard. It requires, on a theoretical level, to harmonize apparently discordant values (freedom and equality), and on a practical level, to balance countervailing design aspirations (deliberation and participation)—all while considering complex empirical realities.[11] The current proposals over-emphasize the aspirations and benefits of broad and deep participation. The proposals operate with a picture of democracy in which participation looms large. Participation, that is, that involves "everyone", that includes participants directly, as opposed to via representatives, and that endows participants with decision-making powers.

This picture of democracy as participation—as "rule by the people"—leads to problems. I describe five objections against the call to "democratize AI", which operates with this picture. The first two objections aim at the claim that in order to democratize AI, novel institutions ought to be established. The next two objections target the claim that democratizing AI should involve direct and broad participation. Finally, I raise a general and principled objection: Democracy—understood with this emphasis on participation—is just not the right kind of answer to the kind of problems that animate the call to democratize AI.

3.1. Weak grounds

Here are some generally accepted tenets in democratic theory. Democracy has both intrinsic—or perhaps better: *final*—as well as instrumental value.[12] Democracy uniquely aims at a certain ideal of fairness and equality and thereby leads to better decisions, or so many argue. For the purposes here, more important than the arguments themselves is what prompts them. Arguments for democracy respond to requirements of legitimization. Arguments in favor of democracy are offered, among other things, to argue that that democracy *legitimizes* the state. Some go even further and argue that democracy is *necessary* for legitimacy—or in a slogan: without democracy, no legitimacy.

I contend that the reverse of this slogan holds: without legitimacy, no democracy. More precisely, unless an institution needs to be legitimized, the grounds for democratizing this institution are much weaker. Triggers of legitimization requirements are the question to which democratic governance is the answer. At least one trigger of legitimization requirements needs to be in place for democratic governance to have a compelling case. When these triggers are absent, democracy may still have instrumental and final value—at least in theory (whether they are realized in practice is another matter). But absent the triggers of legitimization, demands for democracy are on markedly weak ground. This is what the slogan "without legitimacy, no democracy" stands for.

[11] On the tension between participation and deliberation see Cohen (2009, sec. 5).

[12] The distinction intrinsic vs. instrumental value conflates a distinction about values' location (intrinsic vs. extrinsic) with a distinction about their relations (final vs. instrumental). See Korsgaard (1983).

Democratic states pose two justificatory demands: first, the practical authority of the state generally; second, democracy as the system to govern the state.[13] Whereas democracy might well be justified as the governance system of a *state*, and while it may contribute to justifying *state* authority, democracy might *not* be justified as a governance system *beyond* the state. Democracy has a home in states because states trigger the appropriate legitimization requirements. Democracy responds to the demands of legitimization that states trigger. But democracy comes with moral "costs"; it can be burdensome, inefficient, or lead to suboptimal outcomes. In the case of states—because here democracy legitimizes a system that needs legitimization—these "costs" of democracy are justified. By contrast, the objects in the domain of claim to democratize AI do not trigger requirements for legitimization (or, at least, not to the same extent).

Three triggers of legitimization requirements can be distinguished: (1) subjection to coercion, (2) pervasive impact, and (3) involvement in a scheme of social cooperation (Abizadeh 2007; Miller 2009).

First, it is widely accepted that coercion requires legitimization. This is because coercion abridges natural liberties or restricts individual autonomy (esp. autonomy as sovereignty). Insofar as an institution subjects individuals with the threat of coercion, this institution needs legitimization.

Second, it is likewise widely accepted that perversive impact triggers a legitimization requirement. This requirement for legitimization again arises out of a concern of autonomy (here rather autonomy as non-alienation).[14] Insofar as individuals have a right to shape how their life goes and insofar as circumstances should comport with their desires or interests, a system that pervasively impacts individuals requires legitimization.

Finally, involvement in a cooperative system might require legitimization. A scheme of social cooperation is a form of social coordination that "is guided by publicly recognized rules and procedures which those cooperating accept as appropriate to regulate their conduct" (Rawls 2001, sec. 2.2). A system of social cooperation requires legitimization because it constitutes a form of collective autonomy, or collective agency, of which an individual is part. Legitimization here concerns the rules that govern the system and whether they give appropriate expression to some kind of collective interest.

These triggers of legitimization requirements can be found in the state, within the state, and beyond the state. A state itself is by definition coercive, its actions typically have pervasive impacts on many, and a state may often be co-extensive with a system of cooperation (or, absent cooperation on appropriate terms, with a system of oppression). Within the state, triggers of legitimization can be found in

[13] Questions about the legitimacy of the *state*: Why should you respect what the state asks you to do? Why can some demands of the state be enforced, even coercively? Questions about justifying *democracy*: Why should you value, and perhaps choose, democracy over alternative systems?

[14] The distinction between autonomy as sovereignty and autonomy as non-alienation in these terms is due to Enoch (Enoch 2017; 2020).

the workplace (Anderson 2017).[15] Finally, triggers of legitimization can also be found beyond the state in inter- or supra-governmental institutions, which animates theories of cosmopolitanism and investigations into "democracy's domain".

At the same time, democratic governance is costly in several ways. First, democracy is resource intensive (more below). Democratic participation, deliberation and association require time, cognitive and financial means. Even if democracy might produce good outcomes, it may do so in a rather inefficient way. Second, democracy may lead to outcomes that are sub-optimal. In part because democracy is resource intensive, the decisions that are made democratically might reflect incompetence, ignorance, psychological biases, or even malice and spite. Such concerns have nurtured the long tradition of elitist arguments against democracy (Plato 2008; Brennan 2016). Finally, democracy might conflict with equality or individual liberties. As such, not only might the outcomes of democracy be flawed, but its promise of freedom and equality might be hollow.

Because of these costs, democracy needs to be justified. Because democracy can legitimize systems that trigger requirements of legitimization, democracy is often justified by its legitimizing properties. But if the call to democratize AI does not concern such systems that trigger legitimization requirements, then there are insufficient grounds to justify democracy.[16]

The objects in the domain of the democratization of AI—the *what* that should be democratized—do not sufficiently trigger distinct legitimization requirements.

First, AI is deployed in systems that are already coercive. AI may increase the effectiveness of coercion, for example because "law enforcement" by the state and compliance at the workplace can increasingly be automated. But I have yet to think of an example of a hitherto uncoercive system that becomes coercive because of the use of AI. The development and deployment of AI as such do not trigger novel or distinct demands for legitimization.

Second, AI by itself does not have pervasive impacts.[17] For example, when AI augments decisions about pre-trail detention, deportation, hiring, or scarce resource allocation, it is not the use of AI that has pervasive impacts, but the pervasive impacts originate from the fact that there are social systems that detain, deport, hire, or allocate scarce resources. These systems would have similarly pervasive impacts absent the use of AI. AI may scale the impact of the criminal justice system, the immigration system, or of hospitals and health care. AI may amplify the impact of these systems, or it may modulate their impact. But AI does not have a pervasive impact on lives in any way other than technological progress has always impacted lives.

---

[15] Many associations are governed democratically. Labor unions, recreational clubs, or church parish administrations are examples. In addition to exhibiting triggers of legitimatization requirements, these associations can also be seen as essential parts of a democratic society. In other words, they might be part of a state democracy and part of meeting legitimization requirements that are triggered by the state.

[16] This assumes, of course, that there are feasible alternatives that have fewer of the costs outlined above.

[17] I here argue against the second assumption of Sclove's argument presented earlier.

Third, AI does not extend or create involvement in cooperative systems. AI, on the way it is understood here, are tools and techniques that can be deployed in various ways where prediction, categorization, or analysis are needed. As such, AI is embedded in existing systems of social coordination, but it does not constitute a system of social coordination. Since every system of cooperation is a system of social coordination, AI does not extend or create involvement in systems of cooperation. AI may facilitate trade and expand the reach of global markets. It may lead to monopolies and increase the size of corporations (cf. Goldfarb and Tucker 2019). In these ways, AI raises issues of political justice that ought to be addressed. But AI does not raise sufficient, novel, and distinct triggers of requirements of legitimization because it does not extend or create a cooperative system.

In sum, insofar as democratization without legitimization is on weak grounds, and insofar as there is no need for legitimization because of AI, the case for the democratization of AI on these grounds is weak. I have argued that democracy needs to be justified. Because democracy is costly, it generally cannot be justified on strong grounds unless democracy responds to requirements of legitimization. But AI has none of three possible triggers of legitimization requirements. Therefore, one positive case in favor of democratizing AI falters. Moreover, in light of its costs, democracy may not be justified when there is no novel and distinct legitimizing work to be done.

3.2. Redundancy

The democratization of AI should be constrained. To the extent that the demand to democratize AI entails the claim that novel institutions should be established, there is a problem of redundancy. This leads to a constraint: The efforts to democratize AI via novel institutions should not be extended to places that are already governed, albeit imperfectly, by (statist) democratic institutions.

This argument appeals to what can be called the "no overlap" principle: In domains where the state governs legitimately through democratic institutions, no other institution ought to compete over practical authority. Because there exists already a basic structure that is governed by statist institutions, this raises the question of what there is left to democratize. AI is used widely in existing policy domains such as criminal justice, social policy, health policy, or education policy. Moreover, existing regulatory regimes such equal opportunity legislation, consumer protection, or investment protection and market oversight already cover most domains in which AI is used. Thus few, if any, regulatory domains are left in which to democratize AI.

This redundancy argument may limit the demand to democratize AI to the fine-grained *decision* domain. A proponent of the claim that AI should be democratized might accept the no overlap principle and that no novel institutions are needed insofar as many areas of AI applications are already subject to regulation and legislation. But the proponent might argue that the claim to democratize AI targets individual decisions that organizations make and that these decisions are, so far, not determined by regulation or legislation.

Democratizing AI in the fine-grained domain of decisions would change who makes decisions. Whereas regulation leaves agency over decisions unaffected, democratization in this fine-grained domain would change not only how decisions are made but by whom. Regulation may constrain and create options, yet regulation does not change matters of basic agency. The market for medical equipment is highly regulated. Decisions to develop, sell, and buy medical equipment are governed by the respective regulations. Still, the decisions to develop, sell, and buy medical equipment are made by individual organizations and not by regulators. Democratizing these decisions would mean changing corporate governance.

But this raises problems of its own. There may be strong reasons for why a state rarely dictates individual decisions within the policy areas it regulates. Those who make decisions about AI development and deployment may have a right to make these decisions (some liberty- or autonomy-based argument). Or those who make these decisions are more likely than others to make good decisions (perhaps because they have the relevant information or more effective incentives). In short, there might be good reasons for not democratizing the decision domain. Thus, answering the question of how plausible the demand to democratize AI is when it comes to the fine-grained domain of concrete decisions, needs to be left for a separate more context-specific investigation. Perhaps, in some contexts, reasons for democratization outweigh reasons that speak against it—perhaps not.

At any rate, it should be granted that especially corporate governance in the US has a democratic deficit. There might, that is, be decisive reasons for changing how and by whom decisions are made. Workers at Google and Facebook should have a say in more workplace-related matters including what products they develop and how and to whom they are marketed. Such a workplace democratization seems overdue, even if the grounds for democratization may have little to do with AI as such. The argument from redundancy, if plausible, thus at least constrains the claim to democratize AI to this decision domain.

3.3. Resource intensity

A traditional challenge for democracy is that practicing democracy is hard. It is somewhat of a puzzle why democracy "works" when it does. Democracy requires various resources. In result, democracy—as it is demanded by the call to democratize AI—can be infeasible, inefficient, or inequitable. But those who call to democratize AI rarely acknowledge this as problematic. This is unfortunate, to say the least, because the problems of resource intensity are severe.

First, democracy—especially direct participation—requires cognitive resources, financial resources, and time. This resource intensity is particularly acute for participation, especially when participation is intended to be broad and deep, that is, when at least a large group of individuals is to be directly involved in decision making. Democratizing AI hence might be inefficient or practically infeasible. To the extent that this feasibility problem for democracy is caused by the elements of direct democracy, it is this part of the demand to democratize AI that ought to be given up. However, this may negatively affect the desirability of the demand: Some of the value-grounding features of democracy may depend

on direct democratic participation. If one aim of democratizing AI is to "restore collective subjectivity and agency" (McQuillan 2018, 6), then it is unclear whether this aim can still be met by a model of representative instead of direct democracy.

Second, in addition to cognitive, financial, and time resources, participation requires social resources. Those who are expected to participate in democracy face a collective action problem. They need to organize and articulate their interests collectively.

To make matters worse, the resource intensity problem here is asymmetric. Industry players—capital—have more resources than individuals. By contrast, the most marginalized people in a society are likely also those who have the fewest resources to participate. They lack the time, the money, the relevant social capital, or simply the freedom to associate and organize effectively. They are thus least likely to organize successfully (Gilens and Page 2014; Shughart and Thomas 2019). This taints the democratic process. The (social) resource intensity results in "a tendency within the political system as a whole to generate the exploitation of the unorganized by the organized." (Heath 2020, 74). *The call to democratize AI may therefore have the opposite of its intended effect*. A democratization of AI that focusses on participation may compound the problem it was trying to solve.

Then there is a problem of who shows up. Democratic participation has a well-known representation problem (Irvin and Stansbury 2004, 59). Those "who show up to participate are those whose time has a very low opportunity cost (e.g., retirees), or else those with a significant economic stake in the decision (e.g., industry spokespeople)." (Heath 2020, 74). It has been known for a long time that "[l]ow-income and minority citizens in particular tend to be inadequately represented and unable to afford what is required to make a significant contribution." (Checkoway 1981, 569; see also Gilens and Page 2014; Weaver and Prowse 2020)

The problem of who shows up is, again, rooted in participation's resource intensity. The depth of this problem becomes clear when the variegated practical problems of participation are considered. For starters, participation requires collective cultural resources. Just to make "participation" sound remotely attractive, a general tendency of complacency and cynicism needs to be overcome (Irvin and Stansbury 2004, 58). Next, those who would indeed like to participate need to have informational resources, or pay search costs. Opportunities to participate—especially in "conventional participation" such as public hearings—are often announced in obscure places. This is a problem of pre-participation procedures. Next, participants need time. Participation typically happens in the middle of the workday. Even if participants can spare the time, absent appropriate place resources, participants pay a psychological price. Participation usually takes place in government offices. Many participants perceive these places as "foreign territory" (Checkoway 1981, 567). Even the layout of the physical spaces may induce the feeling of being out of place (Nabatchi and Leighninger 2015, chap. 8).

In sum, democracy is highly resource intensive. Individuals need to organize—indeed, organize better than capital interests—they need to find the information, the time, the confidence, and the material

resources to participate effectively. And all this presupposes that they are interested in participating and will be disposed in a way that averts myopic or partisan discussions (see section 3.4).

When we move from the practical to the theoretical level, there is a conceptual problem. Because of its resource intensity, democracy has an internal tension between participation and deliberation (J. Cohen 2009). An increase in participation usually leads to a decrease in (the quality of) deliberation. A broad polity can participate directly by voting, but "giving a voice" to each individual and incorporating each individual contribution into deliberations is technically hard, inefficient, and likely practically infeasible. To the extent that the quality of deliberations grounds some of the value of democracy, there is a dilemma between deliberation and the value of deep and universal participation. For the demand to democratize AI, this cautions against its emphasis on participation.

There are, in sum, two ways in which democracy is resource intensive and in which this is problematic. First, democracy is simply expensive in terms of material, cognitive, cultural, informational, and social resources. Second, this leads to a conceptual conflict between participation and deliberation. At least the practical problems of resource intensity can be overcome at least some of the time (Nabatchi and Leighninger 2015). A more deliberative democracy is possible and has much to recommend it (Curato et al. 2017; Dryzek et al. 2019). But insofar as "democratizing AI" involves broadening and deepening participation, these practical and conceptual problems speak against democratizing AI.

3.4. Popular oversights

A further problem for democracy—especially when a high emphasis is placed on its participatory components—is that democracies extend serious moral concern only to present citizens. When participation is broadened and deepened, this increases the risk of popular oversights—errors in collective perception, reasoning, and agency—, which leads to harm, wrongdoing, or morally bad outcomes. My point here has psychological roots and a moral upshot.

Human reasoning is bedeviled by a myriad of biases and flaws of rationality (Heap et al. 1992, chap. 3). For starters, humans are well-known to make inconsistent or otherwise irrational decisions under risk, in particular when the probability of an event's occurrence is small, even if the magnitude of the effect is great (Tversky and Kahneman 1974; Kahneman 2011, pt. 2). Additionally, human reasoning is bias-prone when things are political. Our reasoning is motivated: We see the world through a partisan lens (L. M. Bartels 2002; Westen 2008). We make choices based on group affiliation not on policies' content (G. L. Cohen 2003; Ditto, Pizarro, and Tannenbaum 2009; Lodge and Taber 2013). The problem is not close-mindedness but that we use our cognitive capacities to affirm our loyalty to affinity groups (Kahan 2012). Such motivated reasoning may increase polarization (Lord, Ross, and Lepper 1979; Stanley et al. 2020). Add to this the fact that voters have often not even a command of simplest relevant facts (Brennan 2016, 24–27), and that, when deliberation the occurs in enclaves—think of echo chambers and filter bubbles—,the resulting collective attitudes move even further to the extremes (Cass R. Sunstein 2002). These different phenomena are summarized in various popular books (Haidt 2012; Kahneman 2011; Lakoff 2008; Lenz 2013; Cass R. Sunstein 2006).

This is the psychological side of the story. It is important to emphasize that there is hope in this dark picture. Structured and careful deliberation can help overcome many of these problems (Dryzek et al. 2019). But this deliberative aspect is precisely one that the call to democratize AI does *not* emphasize. The call to democratize AI emphasizes instead participation. Given the tension between participation and deliberation (see section 3.3), the call to democratize AI risks to weaken deliberation and its potentially salutary effects in favor of the psychological perils of participation.

This psychological story may connect to a moral oversight: Democracies fail in their moral obligations. Democracies—and we as citizens of democracies—do far from enough to help the poor (Hulme 2016). We are a moral failure (Singer 1972; Unger 1996, chaps. 1–6; Shue 1980; O'Neill 2016, chap. 2; Gabriel 2018). We are, perhaps, even responsible—causally or morally—for poverty and injustice (Pogge 2005; 2008; Ronzoni 2009). This suspected moral failure is the animating contention for the field of global justice (entry points to this literature: Brock 2021; Blake 2008; Brooks 2020). Considering climate change, democracies not only fail to extend moral concern to those who are geographically distant but also to future generations, who are distant in time (Broome 2012; Moellendorf 2015; Shue 2020, chap. 8).

This popular oversight materializes not only in problems of global justice, climate change, or domestic inequality, but also in wars. Some argue that there is a distinct form of "democratic torture". That is, democracies get away with torture by inflicting it, more or less explicitly, only on non-citizens and doing so in a way that leaves few if any traces of injuries. Whereas "[i]n ancient republics, torture […] was inflicted exclusively against slaves, foreigners, and 'barbarians'", "[t]he victims of torture in democracies today are not spoken of as slaves, but as 'street children, vagrants, loiterers, and illegal immigrants . . . [who] fall into a class of quasi- citizens that is perceived as vicious.'" (Rejali 2009; cited in Bell 2016, 46–47).

Of course, democracy might not be worse in its moral oversights and flaws than non-democratic systems. But the label "democracy" can be an ideological trojan horse. It allows to smuggle in self-interest in the name of a good thing—democracy, participation, the power of the people. Participation may bestow a falsely understood virtue of respectability onto (morally) problematic decisions (Irvin and Stansbury 2004, 59–60).

These various shortcomings—from biases and irrationality to lack of moral concern for distant others—illustrate that democracy may suffer from popular oversights. On the assumption that such oversights generally occur in all democratic practices that emphasize participation, we arrive at a further reason against democratizing AI: democratizing AI may cause or sustain moral problems.

3.5. Theoretical and practical inefficacy

A final reason against the call to democratize AI is not that there is something wrong with it, but that the call to democratize AI fails to speak to the concerns that animate it. The call to democratize AI is

incongruent to, and ineffective for addressing, the problems that motivate it. The call to democratize AI is typically made in response to AI-related injustice and oppression. Recall how Zimmerman et al. (2020) motivated their demand to democratize AI from problems of "algorithmic injustice". For example, the problem that AI leads to wrongful discrimination in hiring decisions, the problem that AI entrenches harmful stereotypes or unfair practices, or the problem that powerful AI tools are used to enforce inhumane policies, are often what motivates the call to democratize AI.

But democracy is generally not normatively well-equipped to afford a proper response to problems of injustice and oppression. It is far from clear, for example, that democracy—as an idea—has the normative resources to condemn and address problems of structural injustice.[18] We should not expect democracy to solve problems that it is not meant to solve. What is needed is not a theory of democracy but, more substantively, a theory of justice that targets "the institutional conditions necessary for the development and exercise of individual capacities and collective communication and cooperation." (Young 1990, 39). Democracy is not necessarily such a theory of justice. This is the theoretical problem. On a practical level, not a call for more democracy but direct action might be the better theory of change. In this spirit, Crawford writes (2021, 223): "To suggest that we democratize AI to reduce asymmetries of power is a little like arguing for democratizing weapons manufacturing in the service of peace. As Audre Lorde reminds us, the master's tools will never dismantle the master's house."

Democracy is normatively thin. That is, the theory itself contains only relatively few or relatively undemanding normative claims. This is true in particular for minimalist and aggregative conceptions of democracy. Minimalist conceptions see democracy as consisting in competitive elections and, perhaps, the rule of elected elites. Aggregative conceptions see democracy, roughly, as a process of fair aggregation. Democracy pools information and attitudes and then yields collective decisions. As a proponent of this conception, Dahl (1989, chap. 8) characterizes the ideal of democracy as meeting criteria of inclusiveness, enlightened understanding, agenda control and effective participation. There is little to nothing in minimalist or aggregative conceptions of democracy to address oppression and injustice.

Both the minimalist and aggregative conceptions of democracy are largely procedural.[19] They incorporate only a merely formal and not a substantive and material demand of equality. In result, minimalist and aggregative conceptions are positively agnostic about moral views concerning, for example, obligations to non-citizens, the justifications of military interventions, the permissibility of affirmative action, or the adequate rectification of historic injustices. Moreover, minimalist and aggregative conceptions of democracy might stand in the way of justice. These conceptions might hold dearly the idea of formal equality of opportunity, that is, the demand that everyone has the same opportunities, enjoys the same rights, and is treated the same. This formal equality contrasts with material equality, which recognizes that structural causes subtly undermine equality of opportunities. Even if everyone

[18] Structural injustice are systematic violations of particularly important moral claims or liberties, the maintenance of which is explained by non-individual entities such as cultures, norms, or practices.
[19] Although some argue that the distinction between procedure and substance collapses (J. Cohen 1993; 1997). Thus, there might be no such thing as a purely proceduralist conception of democracy.

has the same formal rights, these rights have not the same value to everyone (cf. Rawls 1971, sec. 36). Material equality aims to equalize the value of rights. It permits, or even demands, to violate formal equality of opportunity, for example, by giving individuals advantages to make up for structural disadvantages. Currently, this conflict between formal and material equality goes also by the name of a conflict between "equality" and "equity". Insofar as minimalist and aggregative conceptions of democracy are only committed to formal but not to material equality, they have little, if anything, to offer in response to precisely the problems that motivate the claim to democratize AI.

Being so normatively thin is a feature and not a bug of democracy. Democracy is intended to be part of what Rawls calls an "overlapping consensus" (Rawls 1993). Individuals who disagree profoundly about value and morality still have reason to recognize a broadly democratic system as common ground. The fact that democracy is normatively thin is, in a way, what makes this overlapping consensus possible. Democracy—as a theoretical idea—should not be loaded with thicker values of material equality or equity. Insofar as one can reasonably reject these values of material equality or equity, not everyone would then have reasons to accept democracy as part of the overlapping consensus. This takes away legitimacy of the system and makes it unstable.[20] The current literature on ethics and technology does not seem to take sufficiently seriously the fact of (reasonable) pluralism, that is, that societies can be highly pluralistic for (good reasons) (Himmelreich 2020).

One might say: Even if democracy—as a theoretical idea—does not incorporate thick values, perhaps a democratization of AI unleashes, emboldens, or channels the popular powers in the right way. There might be a contingent case that, as a matter of mere actual (as opposed to necessary) fact, the democratization of AI—even on the minimalist or aggregative conceptions of democracy—will solve problems of bias, unfairness, wrongful discrimination, and historic injustice.

But even if democracy were the vehicle that takes us towards justice in AI, it is at least questionable whether this vehicle can take us all the way and whether it is the best means of getting there. Instead of defending the procedures and values of democracy, it would be better to make the case against injustice and oppression in the fora of public reason: that it is wrong to be complicit in drone strikes by supplying AI tools to the Pentagon, that it is wrong that black women have a greater chance of not being recognized in pictures, and that it is wrong to not give preference to minority applicants in hiring decisions. The existing democratic structures, domestically and globally—through institutions such as the WTO or the OECD—provide avenues to articulate such moral demands. This would be a call for justice, however, and not a call for more democracy or more participation. Calls to democratize AI seem to confuse the two.

Democracy need not be normatively thin. The deliberative conception of democracy opposes both the minimalist as well as the aggregative conception of democracy in this respect. The deliberative conception emphasizes fairness, that is, that a democracy should offer not just formally but materially equal opportunities of participation. Moreover, as the name suggests, the deliberative conception

[20] The question, of course, is whether a rejection of material equality is reasonable.

emphasizes the importance of *public reasoning* (J. Cohen 1989), that is, that collective decisions are, or could, be justified to everyone.[21] Proper public reasoning requires a democratic society that honors not just rights to free speech but free expression,[22] that elevates marginalized voices, that sees to it that material needs are satisfied, and that offers, in short, material equality of opportunity. As such, a deliberative conception of democracy may have the resources to address at least some of the problems that motivate the claim to democratize AI.

Yet, even deliberative conceptions of democracy might be insufficiently thick. Although perhaps not positively agnostic, deliberative conceptions are still rather ambiguous about obligations to non-citizens, the justifications of military interventions, the permissibility of affirmative action, or the adequate rectification of historic injustices. Famously, it is rather unclear to what extent Rawls' theory of justice and political liberalism is consistent with affirmative action (C. W. Mills 2013; 2015; 2018; Shelby 2003; Taylor 2009). Thus, even on a deliberative conception of democracy, the call to democratize AI might still be incongruent to the problems that motivate it.

And finally—it is worth mentioning again—the call to democratize AI emphasizes participation over deliberation. Those who call to democratize AI, put front and center that more people should be more involved in more decisions about AI governance. Even if deliberative democracy were sufficiently normatively thick, those who call for AI to be democratized show too little concern that their emphasis on participation is in tension—theoretically and practically—with deliberation.

4. Conclusion

One upshot of this paper is that the substance behind the call to "democratize AI" is somewhat underdeveloped. It is often not fully clear what it would mean to "democratize AI" and what arguments speak in favor of doing so. The objections described in this paper highlight that three questions should be asked.

The first question is: the democratization of *what*? This question is driven by the redundancy concern. Proposals to democratize AI should be clear how novel democratic institutions mesh with existing ones. A next question to ask is: *Why* should AI be democratized? This question is driven by the argument of insufficient grounds and the argument of insufficient substantive import. This *why* question matters because the democratization of AI should be congruent with the reasons that motivate it. Are there unmet requirements of legitimacy? Is the call to democratize AI motivated by concerns about oppression and structural injustice? There is a risk that the call to democratize AI is motivated by a concern for justice—for which, as I suggested, there is only limited room within the idea of democracy. The call to democratize AI may thus get things the wrong way around: It is not the call for democracy that sets us on a path to justice, instead it might be the call for justice that will set us on a

[21] Some deliberative democrats want *public reasons*. That is, they demand that this justification should be based on reasons that everyone can accept. Roughly, the same justification should be offered to everyone. By contrast, others argue that each individual can be offered a different justification as long as each can be offered some reasons. Roughly, they contend that different reasons can be offered to different people.

[22] Free expression is the broader category, it includes, for example, artistic expression.

path of democracy. Finally, one should ask: *How* should AI be democratized? Even if there is no positive general answer to this question—answers should differ from one context to another—there is a general negative one: In most context, direct democracy is not the model to emulate. Concerns about popular oversights and resource intensity militate against doing so.

The call to democratize AI emphasizes participation: More people should be more involved in more decisions about the governance of AI. Another upshot of this paper is that such an emphasis on participation is problematic.

In addition to the problems of resource intensity and popular oversights, the idea of democracy as participation and majority rule is ideologically exploited from different sides. Authoritarians co-opt it (see introduction); but so do powerful domestic interests. Participation might be merely a chimera that masks a consolidation and centralization of power—capital power or bureaucratic power. After all, it is a puzzle—if there is so much value in participation—why there is so little of it. Since the problems of formalized participation are so well-known, demanding more participation would only help to "contain citizen demands, … to channel citizens in ways acceptable to officials" (Checkoway 1981, 574).

But "democracy" means more than participation, voting, and majority rule. The main upshot of this paper is hence the need to clarify what is meant by "democracy". The call to "democratize AI" neglects other elements of democracy, such as deliberation, equality, or fairness. Democracy is a system of social governance that gives everyone a fair and equal opportunity of influence. What does it mean for everyone to have "a fair and equal opportunity to influence"? This question cannot have an easy answer. But a clear answer is needed to correctly diagnose how AI falls democratically short (cf. Himmelreich 2022). This paper gives you a partial answer: Whatever the ideal of democracy might be, participation plays a relatively small part in it.

In addition to emphasizing participation, it is also noteworthy that the proposals to democratize AI often seem to concentrate on what seems to be a legislative angle. But one important—and too often overlooked—arena of the democratic process is public administration. Bureaucracy is powerful. Bureaucrats can not only lead the charge for reform proposals and draft legislative acts, they have discretion in preparing international agreements and cooperation between domestic institutions. Bureaucrats make regulation. Bureaucrats evaluate AI technology, such as the performance of facial demographic classification. What makes the workings of bureaucracy more democratic? What should public administration do about AI technologies? How can the state support society in public reasoning about AI? How else—if participation is not the right way—can the democratic values of fairness and equality find their way into bureaucratic decision making?

To democratize AI, we should think about democratizing the administrative state without having to rely on broad and direct participation (e.g. Tutt 2017). Here the task may include, not to create new institutions, but to improve existing ones. Some opportunities for participation already exist. Think,

for example, of citizen review boards of local police departments. These are likely lacking. In the case of police oversight, the problems include the power of police unions, the vast legal protections of police officers (the qualified immunity doctrine in US law), the strategic alignment between police departments and prosecutors, the cultural valorization of police forces and the resulting epistemic injustice—that is, differences in whose testimony counts—between police officers and victims of police misconduct. None of these are problems of participation, but each of them is a problem of democracy. Each of these problems is a lack of democracy. What this lack exactly consists in, however, is not easy to identify—a lack of accountability, of rule of law, of the value civil liberties? The practical-political task of rectifying this lack is even harder still.

The main upshot of this paper is hence a call for a certain reorientation: The issue is perhaps not so much *whether* AI should be democratized but *how*. The call to "democratize AI" should not emphasize only participation. What a democratic governance of AI would even look like—both on a theoretical level of values and on an organizational level of institutions and processes—and how the shortcomings of existing democratic institutions can be overcome are the question to ask.

Acknowledgements
In writing this paper I benefited greatly from early discussions with Iason Gabriel who suggested objections, their names, and relevant literature. I thank Tina Nabatchi for helping me with references to literature on participation and collaborative governance. I also am grateful for the discussions and the feedback I received at the exploratory seminar on The Ethics of Technology: Beyond Privacy and Safety at the Radcliffe Institute for Advanced Study at Harvard University, the Embedding AI in Society Symposium at NC State University, the brownbag seminar of the Information Society Project at Yale Law School, and the workshop Business Ethics in the 6ix.

References

Abizadeh, Arash. 2007. "Cooperation, Pervasive Impact, and Coercion: On the Scope (Not Site) of Distributive Justice." *Philosophy & Public Affairs* 35 (4): 318–58.

Anderson, Elizabeth. 2017. *Private Government: How Employers Rule Our Lives (and Why We Don't Talk about It)*. Princeton, N.J.: Princeton University Press.

Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. "Machine Bias." ProPublica. May 23, 2016. https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing.

Awad, Edmond, Sohan Dsouza, Jean-François Bonnefon, Azim Shariff, and Iyad Rahwan. 2020. "Crowdsourcing Moral Machines." *Communications of the ACM* 63 (3): 48–55. https://doi.org/10.1145/3339904.

Awad, Edmond, Sohan Dsouza, Richard Kim, Jonathan Schulz, Joseph Henrich, Azim Shariff, Jean-François Bonnefon, and Iyad Rahwan. 2018. "The Moral Machine Experiment." *Nature*, October. https://doi.org/10.1038/s41586-018-0637-6.

Bartels, Larry M. 2002. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24 (2): 117–50. https://doi.org/10.1023/A:1021226224601.

Bell, Daniel A. 2016. *The China Model: Political Meritocracy and the Limits of Democracy*. Princeton, N.J.: Princeton University Press.

Blake, Michael. 2008. "International Justice." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2008. http://plato.stanford.edu/archives/win2008/entries/international-justice/.

Brennan, Jason. 2016. *Against Democracy*. Princeton, N.J.: Princeton University Press.

Brock, Gillian. 2021. "Global Justice." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2021. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2021/entries/justice-global/.

Brooks, Thom. 2020. *The Oxford Handbook of Global Justice*. Oxford: Oxford University Press.

Broome, John. 2012. *Climate Matters: Ethics in a Warming World*. New York: W. W. Norton & Company.

Cammaerts, Bart, and Robin Mansell. 2020. "Digital Platform Policy and Regulation: Toward a Radical Democratic Turn." *International Journal of Communication* 14 (22): 135–54.

Carugati, Federica. 2020. "A Council of Citizens Should Regulate Algorithms." *Wired*, December 6, 2020. https://www.wired.com/story/opinion-a-council-of-citizens-should-regulate-algorithms/.

Checkoway, Barry. 1981. "The Politics of Public Hearings." *The Journal of Applied Behavioral Science* 17 (4): 566–82. https://doi.org/10.1177/002188638101700411.

Cohen, Geoffrey L. 2003. "Party Over Policy: The Dominating Impact of Group Influence on Political Beliefs." *Journal of Personality and Social Psychology* 85 (5): 808–22. https://doi.org/10.1037/0022-3514.85.5.808.

Cohen, Joshua. 1989. "Deliberation and Democratic Legitimacy." In *The Good Polity: Normative Analysis of the State*, edited by Alan P. Hamlin and Philip Pettit, 17–34. Oxford: Blackwell.

———. 1993. "Pluralism and Proceduralism." *Chicago-Kent Law Review*, no. 3 (1994): 589–618.

———. 1997. "Procedure and Substance in Deliberative Democracy." In *Deliberative Democracy: Essays on Reason and Politics*, edited by James Bohman and William Rehg, 407–37. Cambridge Mass.: The MIT Press.

———. 2009. "Reflections on Deliberative Democracy." In *Contemporary Debates in Political Philosophy*, edited by Thomas Christiano and John Christman, 247–63. Chichester: John Wiley & Sons. http://onlinelibrary.wiley.com/book/10.1002/9781444310399.

Crawford, Kate. 2021. *The Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven and London: Yale University Press.

Curato, Nicole, John S. Dryzek, Selen A. Ercan, Carolyn M. Hendriks, and Simon Niemeyer. 2017. "Twelve Key Findings in Deliberative Democracy Research." *Daedalus* 146 (3): 28–38. https://doi.org/10.1162/DAED_a_00444.

Dahl, Robert Alan. 1989. *Democracy and Its Critics*. New Haven and London: Yale University Press.

Ditto, Peter H., David A. Pizarro, and David Tannenbaum. 2009. "Motivated Moral Reasoning." In *Psychology of Learning and Motivation: Moral Judgment and Decision Making*, edited by Daniel M. Bartels, Christopher W. Bauman, Linda J. Skitka, and Douglas L. Medin, 307–38. San Diego: Academic Press.

Dryzek, John S., André Bächtiger, Simone Chambers, Joshua Cohen, James N. Druckman, Andrea Felicetti, James S. Fishkin, et al. 2019. "The Crisis of Democracy and the Science of Deliberation." *Science* 363 (6432): 1144–46. https://doi.org/10.1126/science.aaw2694.

Enoch, David. 2017. "Hypothetical Consent and the Value(s) of Autonomy." *Ethics* 128 (1): 6–36. https://doi.org/10.1086/692939.

———. 2020. "False Consciousness for Liberals, Part I: Consent, Autonomy, and Adaptive Preferences." *The Philosophical Review* 129 (2): 159–210. https://doi.org/10.1215/00318108-8012836.

Gabriel, Iason. 2018. "The Problem with Yuppie Ethics." *Utilitas* 30 (1): 32–53. https://doi.org/10.1017/S0953820817000024.

Gilens, Martin, and Benjamin I. Page. 2014. "Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens." *Perspectives on Politics* 12 (3): 564–81. https://doi.org/10.1017/S1537592714001595.

Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3–43. https://doi.org/10.1257/jel.20171452.

Gould, Carol C. 2019. "How Democracy Can Inform Consent: Cases of the Internet and Bioethics." *Journal of Applied Philosophy* 36 (2): 173–91. https://doi.org/10.1111/japp.12360.

Haidt, Jonathan. 2012. *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York: Penguin.

Heap, Shaun Hargreaves, Martin Hollis, Bruce Lyons, Robert Sugden, and Albert Weale. 1992. *The Theory of Choice: A Critical Guide*. Oxford: Blackwell.

Heath, Joseph. 2020. *The Machinery of Government: Public Administration and the Liberal State*. Oxford: Oxford University Press.

Himmelreich, Johannes. 2020. "Ethics of Technology Needs More Political Philosophy." *Communications of the ACM* 63 (1): 33–35. https://doi.org/10.1145/3339905.

———. 2022. "Should We Automate Democracy?" In *The Oxford Handbook of Digital Ethics*, edited by Carissa Véliz. Oxford: Oxford University Press.

Hulme, David. 2016. *Should Rich Nations Help the Poor?* Cambridge: Polity.

Irvin, Renée A., and John Stansbury. 2004. "Citizen Participation in Decision Making: Is It Worth the Effort?" *Public Administration Review* 64 (1): 55–65. https://doi.org/10.1111/j.1540-6210.2004.00346.x.

Jaques, Abby Everett. 2019. "Why the Moral Machine Is a Monster." In , 10. Miami School of Law. https://robots.law.miami.edu/2019/wp-content/uploads/2019/03/MoralMachineMonster.pdf.

Jurowetzki, Roman, Daniel Hain, Juan Mateos-Garcia, and Konstantinos Stathoulopoulos. 2021. "The Privatization of AI Research(-Ers): Causes and Potential Consequences -- From University-Industry Interaction to Public Research Brain-Drain?" *ArXiv:2102.01648 [Cs]*, February. http://arxiv.org/abs/2102.01648.

Kahan, Dan M. 2012. "Ideology, Motivated Reasoning, and Cognitive Reflection: An Experimental Study." SSRN Scholarly Paper ID 2182588. Rochester, NY: Social Science Research Network. https://doi.org/10.2139/ssrn.2182588.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kitcher, Philip. 2011. *Science in a Democratic Society*. Amherst, NY: Prometheus Books.

Korsgaard, Christine M. 1983. "Two Distinctions in Goodness." *The Philosophical Review* 92:2 (April): 169–95. https://doi.org/10.2307/2184924.

Lakoff, George. 2008. *The Political Mind: A Cognitive Scientist's Guide to Your Brain and Its Politics*. New York: Penguin.

Lenz, Gabriel S. 2013. *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance*. University of Chicago Press.

Lodge, Milton, and Charles S. Taber. 2013. *The Rationalizing Voter*. Cambridge: Cambridge University Press.

Lord, Charles G., Lee Ross, and Mark R. Lepper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology* 37 (11): 2098–2109. https://doi.org/10.1037/0022-3514.37.11.2098.

Mayer, Robert. 2001. "Strategies of Justification in Authoritarian Ideology." *Journal of Political Ideologies* 6 (2): 147–68. https://doi.org/10.1080/13569310120053830.

McQuillan, Dan. 2018. "People's Councils for Ethical Machine Learning." *Social Media + Society* 4 (2): 2056305118768303. https://doi.org/10.1177/2056305118768303.
Miller, David. 2009. "Democracy's Domain." *Philosophy & Public Affairs* 37 (3): 201–28. https://doi.org/10.1111/j.1088-4963.2009.01158.x.
Mills, Charles W. 2013. "Retrieving Rawls for Racial Justice?: A Critique of Tommie Shelby." *Critical Philosophy of Race* 1 (1): 1–27. https://doi.org/10.5325/critphilrace.1.1.0001.
———. 2015. "Racial Equality." In *The Equal Society: Essays on Equality in Theory and Practice*, edited by George Hull, 43–72. Lanham M.D.: Lexington Books.
———. 2018. "I—Racial Justice." *Aristotelian Society Supplementary Volume* 92 (1): 69–89. https://doi.org/10.1093/arisup/aky002.
Mills, Stuart. 2019. "Who Owns the Future? Data Trusts, Data Commons, and the Future of Data Ownership." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3437936.
Moellendorf, Darrel. 2015. "Climate Change Justice: Climate Change Justice." *Philosophy Compass* 10 (3): 173–86. https://doi.org/10.1111/phc3.12201.
Nabatchi, Tina, and Matt Leighninger. 2015. *Public Participation for 21st Century Democracy*. Hoboken, N.J: John Wiley & Sons.
O'Neill, Onora. 2016. *Justice across Boundaries: Whose Obligations?* Cambridge: Cambridge University Press.
Plato. 2008. *Republic*. Translated by Robin Waterfield. Oxford: Oxford University Press.
Pogge, Thomas. 2005. "World Poverty and Human Rights." *Ethics & International Affairs* 19 (1): 1–7. https://doi.org/10.1111/j.1747-7093.2005.tb00484.x.
———. 2008. *World Poverty and Human Rights: Cosmopolitan Responsibilities and Reforms*. Cambridge: Polity.
Rahwan, Iyad. 2018. "Society-in-the-Loop: Programming the Algorithmic Social Contract." *Ethics and Information Technology* 20 (1): 5–14. https://doi.org/10.1007/s10676-017-9430-8.
Rawls, John. 1971. *A Theory of Justice*. 1999 Revised edition. Cambridge Mass: Harvard University Press.
———. 1993. *Political Liberalism*. New York: Columbia University Press.
———. 2001. *Justice as Fairness: A Restatement*. Cambridge Mass.: Harvard University Press.
Rejali, Darius. 2009. *Torture and Democracy*. Princeton, N.J.: Princeton University Press.
Ronzoni, Miriam. 2009. "The Global Order: A Case of Background Injustice?A Practice-Dependent Account." *Philosophy & Public Affairs* 37 (3): 229–56. https://doi.org/10.1111/j.1088-4963.2009.01159.x.
Sclove, Richard. 1995. *Democracy and Technology*. New York: Guilford Press.
Shelby, Tommie. 2003. "Race and Social Justice: Rawlsian Considerations." *Fordham L. Rev.* 72: 1697.
Shue, Henry. 1980. *Basic Rights: Subsistence, Affluence, and U.S. Foreign Policy*. Princeton, N.J.: Princeton University Press.
———. 2020. *Basic Rights: Subsistence, Affluence, and U.S. Foreign Policy: 40th Anniversary Edition*. Princeton, N.J.: Princeton University Press.
Shughart, William F., and Diana W. Thomas. 2019. "Interest Groups and Regulatory Capture." In *The Oxford Handbook of Public Choice, Volume 1*, edited by Roger D. Congleton, Bernard Grofman, and Stefan Voigt, 584–603. Oxford: Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190469733.013.29.
Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1 (3): 229–43.
Sloane, Mona, Emanuel Moss, Olaitan Awomolo, and Laura Forlano. 2020. "Participation Is Not a Design Fix for Machine Learning." *ArXiv:2007.02423 [Cs]*, August. http://arxiv.org/abs/2007.02423.

Stanley, Matthew L., Paul Henne, Brenda W. Yang, and Felipe De Brigard. 2020. "Resistance to Position Change, Motivated Reasoning, and Polarization." *Political Behavior* 42 (3): 891–913. https://doi.org/10.1007/s11109-019-09526-z.

Sunstein, Cass R. 2002. "The Law of Group Polarization." *Journal of Political Philosophy* 10 (2): 175–95.

———. 2006. *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press, USA.

Taylor, Robert S. 2009. "Rawlsian Affirmative Action." *Ethics* 119 (3): 476–506. https://doi.org/10.1086/598170.

Tutt, Andrew. 2017. "An FDA for Algorithms." *Administrative Law Review* 69 (1): 83–124.

Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science*, New Series, 185 (4157): 1124–31.

Unger, Peter. 1996. *Living High and Letting Die: Our Illusion of Innocence*. Oxford: Oxford University Press.

Weaver, Vesla M., and Gwen Prowse. 2020. "Racial Authoritarianism in U.S. Democracy." *Science* 369 (6508): 1176–78. https://doi.org/10.1126/science.abd7669.

Westen, Drew. 2008. *The Political Brain: The Role of Emotion in Deciding the Fate of the Nation*. New York: PublicAffairs.

Wong, Pak-Hang. 2020. "Democratizing Algorithmic Fairness." *Philosophy & Technology* 33: 225–44. https://doi.org/10.1007/s13347-019-00355-w.

Young, Iris Marion. 1990. *Justice and the Politics of Difference*. Princeton, N.J.: Princeton University Press.

Zimmermann, Annette, Elena Di Rosa, and Hochan Kim. 2020. "Technology Can't Fix Algorithmic Injustice." *Boston Review*, January 9, 2020. http://bostonreview.net/science-nature-politics/annette-zimmermann-elena-di-rosa-hochan-kim-technology-cant-fix-algorithmic.