Cambridge Forum on AI: Law and Governance (2026), **2**, e4, 1–15
doi:10.1017/cfl.2026.10045

Cambridge Forum

RESEARCH ARTICLE

The boundary problem of democratising AI in the corporation

William Chan

Leverhulme Centre for the Future of Intelligence, University of Cambridge, Cambridge, UK
Email: waichuenchan.william@gmail.com

(Received 6 August 2025; revised 17 February 2026; accepted 25 February 2026)

Abstract

The boundary problem (BP) asks who ought to be included by democratic procedure. Initially articulated as a challenge for democratic citizenship, the BP is later reframed as a difficulty faced by the 'democratiing AI' (DAI) agenda. By mapping different senses of DAI, this article aims to provide a resolution to the BP in the context of DAI. In particular, it notes that DAI shifts our attention away from a crucial factor to make ethical decisions around AI: the power disparities within organizations that put employees advancing ethical AI in a vulnerable position. In this context, I show that workplace democracy provides a more promising normative aspiration than the DAI agenda, although some elements of the latter, such as its emphasis on subjecting tech firms to greater democratic control, serve to mitigate certain problems that workplace democracy is unable to address, for example, when members of tech firms unanimously prioritise ethically concerning AI practices.

Keywords: the boundary problem; democratising AI; workplace democracy; corporate leadership; AI ethics

The boundary problem (BP) is a long-standing concern confronting theorists of democratic politics: *who* should constitute the 'demos', the individuals who share equal rights and responsibilities as citizens of the relevant democratic society (Abizadeh, 2012; Miller, 2020; Owen, 2012; Simmons, 2013; Song, 2012; Whelan, 1983)? Notably, democratic citizens have an equal right and responsibility to participate in the decision-making procedure of their society, for instance, by voting, setting policy agenda, lobbying, and running for public office. However, because democratically made decisions often have far-reaching impact on both citizens and non-citizens, such as those related to climate change and immigration, there remains continuous controversy over the criteria for which one can rightfully claim democratic citizenship.

Like political decisions, corporates' decisions on AI are no less influential. A corporation, for example, can make profits from AI tools trained on the data of billions of users, who may not even be aware of whether and how their personal data is being processed by such tools. Moreover, as a number of corporates are invested in the development and deployment of AI systems integral to the everyday lives of individuals, what they do about AI is highly likely to affect numerous users.

Against this backdrop, the 'Democratiing AI' (DAI) agenda has gained traction in recent years. A key aim of DAI, among others, is to widen participation in the development, deployment, and governance of AI in an organizational setting. Such participation, moreover, should ideally be extended to the general public, not to the employees alone. The DAI agenda is usually celebrated on the basis that

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

it can promote several other AI ethical values, such as accessibility, safety, algorithmic justice, fairness, decision-making legitimacy, and trust in algorithmic decisions (Banifatemi et al., 2021; Castro et al., 2024; Lin, 2024; Okidegbe, 2023). However, to widen participation in AI, a DAI-specific boundary problem naturally arises: who exactly should be involved in making AI-related decisions, and on what basis? According to certain views, the BP is a pressing challenge for DAI's advocates, exactly because there are uncertainties around where the boundary should lie (Himmelreich, 2023; Murphy & Taylor, 2024). While considerable research has been done on the DAI agenda, how the BP should be resolved remains largely unappreciated. More generally, the BP would be relevant to anyone interested in who ought to be included in decision-making processes over AI design, deployment and/or governance.

This article has three aims. The first aim is to explore how the BP for the DAI agenda may be resolved, drawing on existing contributions to the BP in normative democratic theory and their implications for AI democratization. The second aim is to show that many DAI initiatives blur a crucial factor that demotivates corporates to make ethical decisions around AI: the power disparities within organizations that put employees advancing ethical AI in a vulnerable position. The third aim is to demonstrate that workplace democracy provides a more promising normative aspiration than the DAI agenda. But some interpretations of DAI, notably its emphasis on subjecting tech firms to greater democratic control, will complement certain problems that workplace democracy is unable to address. These three aims are intertwined, however, as they all intersect with the DP in certain ways.

Here is the structure of the article. Section 1 shows that DAI, as it is framed by existing scholarships, is an inflated agenda inclusive of almost all 'good' practices of AI. It also considers a few problems of associating the DAI agenda with 'equitable distribution of AI's costs and benefits' or 'making AI systems more accessible'. Specifically, I introduce three major initiatives in widening participation in corporations' AI decisions: *R&D Inclusion, Inclusive Intra-Organisational Oversight* and *Democratic Control of Tech Firms*. Section 2 turns to the BP, and examines four considerations underpinning one's response to the problem: the *all-affected interests principle, solidarity*, the *all-subjected principle* and the *effects of inclusion*. I argue that these considerations, taken together, would provide a case for *excluding* non-employees from DAI initiatives aimed at widening participation, for example, in the R&D of AI products and the intra-organisation oversight bodies of AI. However, Section 3 shows that, even if the BP can be resolved, DAI's impact is largely limited by the power disparities between employees at different levels of their organizational hierarchy. I also discuss how workplace democracy provides a more promising normative aspiration, compared to DAI, in response to the challenge that corporations lack incentives to take ethical AI seriously. Ultimately, however, workplace democracy must work together with *some* DAI initiatives to maximise their value for driving ethical AI practices by corporations. Section 4 concludes the article.

1. Democratising AI: an inflated agenda

To begin, this section reviews three major understandings of DAI, and their corresponding strategies. However, it is argued that two of these understandings not only presuppose a controversial outlook of democracy, but they also overload DAI with too many commitments, rendering DAI an empty rather than a realistically action-guiding agenda for individuals. This gives us reason to limit our discussion of DAI to *widening participation in corporations' AI decisions*, an interpretation of DAI to be examined further in the following parts.

The term 'democracy' is usually interpreted in a manner inclusive of numerous 'good' political values, institutions or practices, such as freedom, autonomy, equality, solidarity, active participation, checks and balances, and so on. The DAI agenda is no less ambiguous, since it is linked by AI scholars and practitioners to a range of trending values or practices for ethical and/or responsible AI. For instance, there are the claims that '[democratizing AI] may help mitigate algorithmic injustice' (Lin, 2024, p. 102), that it may 'help address the knowledge gap in the development and deployment of AI technologies' (Banifatemi et al., 2021, p. 231), and that it may serve to 'mediate the powerlessness

experienced by oppressed groups' (Okidegbe, 2023, p. 1688). Sometimes DAI is used by tech firms to underpin their technical or commercial ambitions around AI products. For example, IBM interprets DAI as 'providing AI access to a wider range of users beyond machine learning (ML) experts. Common means of improving access include reducing AI costs and incorporating AI into tools and platforms people are already using' (Gomstyn & Jonker, 2024). This interpretation of DAI is clearly underpinning the efforts of tech firms to create more user-friendly AI products.

Despite the many senses of DAI, there are three relatively popular understandings of it. The first understanding of DAI sees it as the endeavours to ensure equitable distributions, from both the public and private sectors, of the economic costs/benefits of AI (Clough & Otterbacher, 2023; Coeckelbergh, 2024; Gomstyn & Jonker, 2024; Randieri, 2024; Rubeis, Dubbala & Metzler, 2022). This echoes the familiar concerns over AI's potential to exacerbate job displacement, income inequality, and the power disparities between individuals and firms, as noted by Randieri (2024). Concerns have been raised over AI as an exploitative industry, the quick advancement of which is often predicated on tech firms' unfair extraction of the natural resources, environment and labour of the Global South (Crawford, 2021). To mitigate these issues, scholars have proposed various initiatives. One example is universal basic income (UBI), to shield individuals against the risks of job displacement driven by automation, or to leverage AI's power to administer the complex financial frameworks required to make UBI possible (Allegri & Foschi, 2021; James, 2020; Jozaghi, 2024; Marr, 2024). Another example is the Windfall Clause, according to which there ought to be some profits 'a firm could not earn without achieving fundamental, economically transformative breakthroughs in AI capabilities' (O'Keefe et al., 2020, p. 327). Moreover, those profits should be *ex ante* donated to initiatives aimed at benefiting humanity widely. There are also proposals to introduce a tax on the use or development of AI (e.g. robots), in order to mitigate the economic disadvantage that individuals may suffer from emerging AI technologies (Abbott & Bogenschneider, 2018; Mazur, 2018; Oberson, 2025). Some propose digital commons (i.e. publicly owned technologies and resources, such as data and software) as a way of widening the beneficiaries of AI (Coeckelbergh, 2024, 2025; Furendal, 2025; Muldoon, 2025). This sense of 'democratising' AI, in short, is about allocating the costs and benefits of AI in ways that avoid disadvantaging certain social groups unfairly.

The second understanding centres DAI on the *accessibility* of AI systems. It refers broadly to the measures aimed at making it easier for individuals to use AI (Clough & Otterbacher, 2023; Coeckelbergh, 2024; Gomstyn & Jonker, 2024; Lin, 2024; Sclove, 1995). An individual's access to AI technologies can be limited by a number of factors. One factor, as Gomstyn and Jonker (2024) pointed out, is people's general knowledge of the essential technologies (e.g. machine learning) in relation to AI. Without such knowledge, people may find it challenging to use AI. Another factor is the opacity of AI systems, which 'might deliver the right results, but they do not provide users or affected parties any insight as to how they came to produce those results' (Vaassen, 2022, p. 2). This opacity problem often contributes to people's hesitancy to trust AI-produced decisions, especially in contexts (e.g. healthcare, admissions, hiring) where allocation of responsibility is crucial (Lebovitz, Lifshitz-Assaf & Levina, 2022; Smith, 2021; Zeide, 2022). In addition, it was found that people have differential access to information communication technology, depending on their demographics, socio-economic circumstances, education, digital infrastructure and so on, whereas this phenomenon is often described as *the digital divide* (Lythreatis, Singh & El-Kassar, 2022). To address these issues, some have advocated for initiatives focused on establishing AI's *trustworthiness*, namely a set of properties or standards AI systems must exhibit to demonstrate 'that they are deserving of our trust' (Durán & Pozzi, 2025, p. 16). This approach is partly underscored by the emergence of AI ethical guidelines, developed by both public and private organisations, that specify the 'best practices' for AI development and deployment, such as the EU's Ethics Guidelines for Trustworthy Artificial Intelligence, IBM's and NVIDIA's principles for trustworthy AI, (Díaz-Rodríguez et al., 2023; EU, 2019; Gomstyn, Jonker & McGrath, 2024; Lu, Zhu, Whittle & Xu, 2023; Pope, 2024). In response

to the digital divide, international organisations, companies and scholars have also been calling for widening computer and internet access, for instance (Mathiesen, 2012; Reglitz, 2020, 2024).

While these two understandings of DAI highlight a range of goals and values that merit our attention, they will not be examined any further, for two reasons. Regarding 'equitable distribution of AI's cost and benefits' as an instance of 'AI democratisation' is potentially misleading, or at least rests on a controversial view of the defining purposes of democracy. Here, a distinction must be drawn between *distributive justice* and *democracy*. Although democracy is a form of rule partly characterised by its commitment to equalising the influence and participation of the relevant members, it rarely promises to warrant distributive justice, that is, the just allocation (by whatever standards) of resources, burdens and opportunities among individuals. To clarify, this does not mean that advocates for democracy or democratisation are never concerned about distributive justice, since certain issues in distributive justice (e.g. economic inequality) have implications for the extent to which democratic ideals (e.g. citizens' equal influence on politics) are realised. But distributive justice is *not* a defining ideal of democracy that is widely shared.[1] Likewise, the distributive justice of AI does matter for all sorts of reasons, but putting it under the umbrella of AI democratisation risks diluting the values that have traditionally been attached to the term 'democracy', and therefore blurring the boundaries between DAI and other normative agendas. Without establishing what makes DAI distinctive, the agenda is far less likely to provide any added value to other attempts to make AI more equitable or user-friendly.

In the rest of this article, therefore, I shall only focus on the third understanding DAI: widening participation (especially from the general public) in the decisions around the design, development and/or governance of AI systems (Coeckelbergh, 2024; Gomstyn & Jonker, 2024; Lin, 2024; Noorman & Swierstra, 2023; Zimmermann, Rosa & Kim, 2020). This is also the understanding of DAI to which the BP is more relevant. According to this understanding, there should be measures aimed at including, as many stakeholders as possible, in the making, application and regulation of AI technologies. As Noorman and Swierstra argued, 'all those impacted by these [AI] technologies should have a say over the values and political dispositions built into these technologies…democratic processes should also, self-reflectively, be directed at diagnosing and guiding the ways technologies affect democratic governance' (Noorman & Swierstra, 2023, p. 570). The idea is not just that more people, rather than a few economic and technical elites, should have an opportunity to shape how AI is built or used. It also expresses the thought that our political processes should govern AI systems in such a way that upholds democratic values.

There are three more widely discussed approaches aligned with this third understanding of DAI. The first approach, *R&D Inclusion*, is to widen participation in organisational efforts in researching, developing and designing new AI products. Participation, however, should not only come from members within an organisation. It should also, as Sclove phrased it, allows 'communities, groups, and citizens – including those today least empowered – to help directly initiate some R&D programs and design technologies responsive to their needs' (1995, p. 207). This echoes the calls for participatory design of AI, which place an emphasis on including a wider range of stakeholders in developing new AI systems (Birhane et al., 2022; Buddemeyer et al., 2022; Gill, 1989; Schmager, Pappas & Vassilakopoulou, 2025; Zytko, Wisniewski, Guha, Baumer & Lee, 2022).

The second approach, *Inclusive Intra-Organisational Oversight*, is to reform or establish intra-organisation bodies that review or oversee an organisation's AI practices. In particular, those bodies should represent, as much as possible, the voices of different stakeholders, including those outside of the organisation (Blackman, 2022). A range of tech firms have already attempted to set up internal teams to oversee their AI practices, such as IBM's AI Ethics Board, Google's short-lived Advanced Technology External Advisory Council (ATEAC), and OpenAI's Safety and Security Committee

[1] One exception might be the Marxist conception of democracy, according to which a defining aim of democracy should be to drive socialist policies that mitigate the distributive injustices that are characteristic of capitalism. See Marx (1844 [2016]) and Held (2006), for instance. Still, characterising 'democracy' in terms of 'distributive justice' presents a very specific outlook of democracy, which may be controversial.

(OpenAI, 2024; Rossi, 2020; Wakefield, 2019). However, these initiatives, are often criticised on the grounds that they are instances of 'ethics washing', that corporates' commercial interests often trump the priorities and decisions of such teams, that such teams often play an advisory role with no substantive influence on a corporation's AI decisions, and so on (Gilbert, 2020; Himmelreich, 2023; Lee & Getahun, 2024). Therefore, there is the argument that such AI oversight bodies within organisations should have greater power and include a wider range of representatives, such as employees at different levels of the organisational hierarchy, experts external to the relevant organisation, affected citizens, and the like (Himmelreich, 2023; Schuett, Reuel & Carlier, 2025).

The final approach, *Democratic Control of Tech Firms*, is to subject tech firms to greater control by the democratic state. Risse, for instance, argues that the power of big tech firms 'needs to be harnessed for democratic purposes. It is hard to see how that can be done unless the Big Tech companies are either dismantled entirely … or treated and thus regulated as public utilities alongside enterprises like phone companies' (2023, p. 47). A key implication of this view is that aligning AI companies with democratic values will require greater state control of such companies, although the state must be democratic itself. Similarly, Coeckelbergh maintained that 'one could consider establishing a new institution that better links technology-development contexts to democratically elected bodies and does so in ways that redistribute power towards the elected bodies' (2024, p. 75). In a like manner, Aytac has proposed the establishment of Citizen Boards of Governance (CBGs), which are randomly selected citizens 'authorized to govern the algorithmic infrastructure of Big Tech platforms' (2024, p. 1431). Overall, the aim of these proposals is to hold tech firms accountable to social or political bodies where the core members are selected by democratic means (e.g. elections or sortition).

2. Democratising AI and the boundary problem

The previous section has shown that DAI, understood as widening participation in the decisions around AI's design, development and governance, typically entails three approaches: (1) R&D Inclusion, (2) Inclusive Intra-Organisational Oversight and (3) Democratic Control of Tech Firms. This section turns to the implications of the BP for these approaches, by revisiting several major principles and considerations from which political theorists have attempted to diffuse the BP of democratic citizenship. More importantly, this section proposes a solution to the boundary problem of DAI, arguing that corporations have legitimate reasons to exclude *non-employees* from their efforts in fostering R&D Inclusion as well as Inclusive Intra-Organisational Oversight. I shall also highlight the immense complexities involved, as we seek to address the BP in the context of Democratic Control of Tech Firms.

Initially, the BP is a question more relevant to issues in democratic citizenship: who should constitute the 'demos', the citizens who share equal rights and responsibilities (notably, to vote) in a democratic society (Abizadeh, 2012; Miller, 2020; Owen, 2012; Simmons, 2013; Song, 2012; Whelan, 1983)? Here, an important point should be flagged: citizens within the same 'demos' are supposed to share *equal* rights and responsibilities. Citizens and non-citizens, even when they reside in the same territory and are subject to the authority of the same government, do not have equal rights and responsibilities. Expanding the list of rights or responsibilities of non-citizens does not immediately turn them into democratic citizens, the full members of the 'demos.' Non-citizens become full members of the 'demos' only when they have rights and responsibilities that are equal, or at least comparable to the citizens. This is a crucial feature of the BP to be observed, as we consider whether it is indeed challenging for DAI.

Who should be included in the decision-making processes on AI, a question in the spirit of the BP, has often been considered a potential challenge for the DAI agenda as well (Han, 2025; Himmelreich, 2023; Murphy & Taylor, 2024; Noorman & Swierstra, 2023). This concern is most clearly expressed by Himmelreich, who mentions how the global impact of AI troubles any attempts to 'include' more people in the decisions around AI: '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' (2023, p. 1335). Before specifying the senses in which the BP may challenge DAI initiatives, I shall consider several major principles aimed at addressing the BP, in normative democratic theory.

First, the *all-affected interests principle* is probably one of the most celebrated approaches to the BP (Dahl, 1990; Fung, 2013; Goodin, 2007; Owen, 2012; Rosenberg, 2020; Wilson, 2022). The principle, in its most basic formulation, states that 'All whose interests are actually affected by a decision should be able to participate as equals in the democratic decision-making process' (Owen, 2012, p. 131). The intuitive appeal of this principle is widely recognised, since a lot of citizens and even non-citizens (e.g. immigrants) are subject to the impact of the decisions made by their democratic government. Likewise, individuals, perhaps all of us around the world, are strongly impacted by corporates' decisions around AI, which essentially sustains a number of digital tools. However, if the all-affected interests principle holds, then DAI would seem to require including everyone affected in corporates' decision-making processes over AI. This is practically impossible.

Second, *solidary* has also been cited as a crucial consideration in resolving the BP (Miller, 1999, 2009, 2020; Song, 2009, 2012). In Song's terms, 'Solidarity might be forged through a shared history, shared culture, and/or shared values' (Song, 2012, p. 47). On one view, solidarity motivates people to 'listen to and understand one another' rather than make decisions based on self-interests (Song, 2012, p. 47). If the 'demos' lacks solidarity, its characteristic activity (i.e. democratic participation) will produce outcomes based on self-interests rather than mutual understanding. To ensure the solidarity of participants in any democratic initiatives (including those related to AI), therefore, there is good reason to include only people with a relatively solid collective identity, marked by shared values and culture. Clearly, this criterion has implications for whether participants in DAI endeavours (e.g. intra-organisation oversight bodies of AI) ought to include only people sharing a similar identity.

Third, the *all-subjected principle* has also been proposed as a solution to the BP (Abizadeh, 2021; Goodin, 2016; Goodin & Arrhenius, 2024; Honohan, 2018; Inoue, 2024). The essence of the principle is to include only 'subjects' within the demos. For individuals to be 'subjects' of a democratic body, they must be *coerced* into following the decisions of that body (Abizadeh, 2012; Goodin, 2016). The all-subjected and all-affected interest principles are different, however. The all-affected interest principle is much broader in scope: it implies that the relevant democratic decision-making body on AI should include individuals affected by its outcomes, even if those outcomes are enforced in *non-coercive* manner. For instance, the all-affected principle would imply that I have a say over a company's decisions over an AI tool it developed, even if I use that tool on an entirely voluntary basis. But the all-subjected principle will urge the body to include only those who will be forced or coerced into following its decisions. Imagine, for instance, the state requires citizens to record their health conditions through a government-designed AI tool every week during a pandemic; anyone failing to do so may be prosecuted. According to the all-subjected principle, this will be a particularly concerning scenario that calls for widening citizens' participation in the decisions around that tool, with the reason being that they are coerced into using that app.

Fourth, the *effects* of inclusion also matter for what and how people should be included by the 'demos.' For instance, in liberal political theory as well as epistemic theories of democracy, a democratic procedure that includes more voices is valuable, because the increased diversity of inputs will ultimately enhance the quality of the procedure's outcomes (Anderson, 2006; Cohen, 2002; Estlund, 2008; Goodin & Spiekermann, 2018; Landemore, 2017; Mill, 1859 [2015], 1861 [1861 [2015]]). There may be reason to include more individuals, whoever they are, insofar as this improves corporations' decisions around AI design, development and governance. Meanwhile, there is the view that one's understanding of politics is enhanced through participation in democratic activities (Aristotle, 1995; Landemore, 2017; Mill, 1859 [2015], 1861 [1861 [2015]]). To understand how to vote well, one must have the right to vote; likewise, one may need an opportunity to participate in the decision-making activities on AI, to gradually build the literacy essential for informed decisions over AI-related issues.

The upshot is that DAI initiatives can be justified on *educational* grounds, building individuals' ethical and AI literacy. Finally, solidarity is not necessarily a precondition for democratic membership; rather, the latter might be a crucial condition for solidarity instead. In democratic theory, for example, some have stressed how equal democratic participation serves to foster the collective identity and reciprocal relationships between participants (Christiano, 2004, 2008; Rawls, 1999, 2001, 2005). Seen in this light, certain DAI proposals can be conceived as endeavours to build solidary among those who are included.

Having clarified multiple responses to the BP, let us now turn to their implications for the following three DAI initiatives: (1) R&D Inclusion, widening participation in organisational efforts in researching, developing and designing new tech products; (2) Inclusive Intra-Organisational Oversight, reforming or establishing intra-organisation bodies that oversee an organisation's AI practices, whereas those bodies should represent as far as possible the voices of different stakeholders; and (3) Democratic Control of Tech Firms, subjecting tech forms to greater control by the democratic state. Clearly, each of these approaches invoke a different sense of 'demos'.

For instance, the 'demos' of (1) refers to the participants in corporations' R&D processes; the 'demos' of (2) refers to participants in the relevant AI oversight bodies within an organisation; the 'demos' of (3) refers much more broadly to the citizenry of a democratic society. Such differential constitution of the 'demos' makes the BP challenging for these approaches in different ways.

I shall first comment on (1) and (2). In my view, while there may be some value of widening participation, there is no compelling reason to *equalise* participation in the R&D of AI products, as well as their intra-organisational AI oversight bodies, *to those outside of the same organisational ecosystem*. To clarify, the view does not say that the R&D and oversight teams of an organisation's AI practices should never invite the inputs of non-members (e.g. communities, citizens) of the organisation. The BP may not be concerning at all, if what is being demanded is merely that inputs of non-members should have *some* role to play. The reason, as I noted earlier, is that the BP is concerned with the criteria for which someone should be granted rights and responsibilities, *on a par* with other members within an organisation or society. For example, it is common for corporates to invite non-employees to beta-test their AI products before rolling them out on the market, to improve the user experience of such products. But this does not mean that the inputs of employees or non-employees share an equal degree of authority over how AI products ought to be developed and governed. Yet, if what DAI is demanding is merely that non-employees get to shape a corporation's strategies for AI development or governance in some moderate ways, regardless of how much influence these non-employees can effectively exercise, it is not obvious what makes the DAI agenda ambitious enough to differentiate itself from other proposals of participatory AI design and/or governance.

Furthermore, there is a compelling case for the employees within a corporation to remain the key decision-makers of R&D and oversight of AI, while corporations are justified in assigning to non-members less influence over these matters. The main ground for this view is that non-members of organisations do not plausibly fulfil three out of four criteria for democratic inclusion above. Although the interests of individuals are often affected by organisations' AI practices, of which they have no membership, such individuals lack the kind of solidary and subjection to an organisation's authority that supports granting those individuals equal rights and responsibilities, within the relevant decision-making processes. But employees of an organisation are different. Compared to non-employees, employees are marked by a higher level of shared interests, culture and values, that constitute their collective identity. Because they belong to the hierarchical structure of the organisation, employees are also often coerced into prioritising the commands and interests of those from senior levels of management; employees have relatively little leeway to disobey, unless they are prepared to be marginalised or even leave their organisation. The mere fact that individuals' interests are being affected by companies invested in AI development and application may only justify that their voices be included by organisations to a modest degree, but this does not warrant *equal* participation of members and non-members of an organisation.

In particular, it is unclear how the positive effects of inclusion must be achieved by providing *both* employees and non-employees with an equal say over R&D and organisational AI governance. As mentioned earlier, if collective decision-making is constituted by participants with a relatively loose shared identity, they are likely to deliberate and make decisions based on self-interests rather than mutual understanding. Also, citizens' skills and knowledge of AI research, development and oversight are not acquired exclusively through participation, on an equal footing, with intra-organisational professionals working on these domains. It has become a lot more common today for individuals to have access to the tools or software to develop AI systems, outside of an organisational setting. Given the availability of resources and opportunities for individuals to educate themselves about AI-related issues, it is also hard to justify giving non-employees substantive power, *comparable to employees*, to contribute to decisions around the R&D and governance of AI. Finally, the view that inclusive decision-making enhances the quality of decisions being made, as proposed by epistemic theorists of democracy, often assumes that the participants, on average, should pass some modest thresholds of competency, or that there are *expert moderators* who serve to organise the wisdom of everyone in ways contributing to high-quality decision-making (Brennan, 2011, 2016; Somin, 2014). Yet, if the participants in the R&D and governance of AI are supposed to meet these requirements, this does not seem to deviate meaningfully from existing common practices by corporations to include outsiders or laymen to improve their research, development and oversight of AI, while allowing expert employees to moderate the inputs of such individuals and transform them into actionable items for the corporate.

In short, these facts reveal a broader challenge of widening participation in the development and/or governance of AI products. If widening participation merely means non-employees have *some* say over AI development and/or governance of corporations, this does not seem highly controversial and effectively trivialises the DAI agenda. But if widening participation means equalising the say and responsibilities of employees and non-employees with respect to AI development and/or governance, this requires a much stronger justification as it involves revisiting the criteria for membership within the 'demos', which refers to the corporation in this context.

Let us turn to Democratic Control of Tech Firms. This proposal is by nature different from the previous two, since it is not primarily concerned with providing employees and non-employees with an equal or at least comparable say over AI-related corporate decisions. It is essentially about expanding the state's control of (tech) firms, in such a way that their AI practices take democratic values more seriously. Here, the 'demos' refers broadly to the entire citizenry of democratic society. However, the tendency of tech firms to disregard democratic values in their AI practices does not seem clearly owed to the fact that our 'demos' is not sufficiently inclusive. Rather, that issue is more obviously linked to two general facts: (a) tech firms have tremendous power, compared to other citizens, in shaping the agenda and outcomes of democratic processes, and (b) existing democratic institutions do not have adequate mechanisms to hold corporations accountable to AI practices that undermine democratic politics. These facts, however, are not clearly addressable by granting more people with the rights and responsibilities of citizens, at least not without some further justifications. Also, resolving the BP of who should quality as full members of a democratic *society* is a much broader question, an informed response to which requires more careful consideration of the moral rights, interests and responsibilities of individuals, across a wide spectrum of political issues beyond AI governance. The BP, therefore, can hardly be framed as a suitably specific threat to Democratic Control of Tech Firms.

3. Insights from workplace democracy

Section 2 has examined the senses in which the BP may or may not be challenging for DAI. Specifically, it has shown that either the 'widening participation' proposals associated with DAI are not ambitious enough to render themselves vulnerable to the BP, or such proposals disregard the legitimate interests of corporations in prioritising employees over non-employees, when it comes to AI development and governance. But the BP is no more threatening to Democratic Control of Tech

Firms, as the problem does not respond aptly to the distinctive concerns such expanded control is meant to address.

So far, I have defended the thought that the BP for DAI, understood as widening participation in AI development and governance, is largely resolvable. However, underneath the BP-motivated struggles between different options to broaden participation in AI-relevant decision-making processes, is a potentially concerning shift of people's attention away from a major factor disincentivising corporate leadership to prioritise ethical AI practices: the power disparities among members within the same corporate hierarchy. In other words, the main problem is not whether and how inclusive the process of AI development or governance should be, but the power dynamics that suppress voices around ethical AI. The aim of this section, then, is to explain why *workplace democracy*, rather than DAI, presents a more promising normative aspiration in response to such power dynamics, although the former needs to be supplemented by Democratic Control of Tech Firms to have a more significant effect.

Workplace democracy refers broadly to the institutional proposals aimed at infusing democratic values into a corporate setting (Foley & Polanyi, 2006; Frega, 2021; Frega, Herzog & Neuhäuser, 2019; Landemore & Ferreras, 2016; Mayer, 2000; Vrousalis, 2019; Yeoman, 2014). The most notable feature of workplace democracy is that it seeks to turn workplace into 'systems for the exercise of power by workers or their representatives' (Frega et al., 2019, p. 1). There are a variety of arguments for workplace democracy, appealing to the similarities between the state and firms, the importance of equality, democratic education, meaningful work and so on (Frega et al., 2019). Unlike the DAI agenda, however, expanding organisational participation from *non-employees* is not often one of the characteristic arguments for workplace democracy. Rather, their focus is usually placed on creating less hierarchical organisational structures, or any other ways to mitigate the power disparities between employees at different organisational levels.

The DAI agenda is often pitched by ethicists and AI practitioners as a way of accelerating responsible or ethical AI decisions. Because many DAI initiatives focus on widening participation in AI design and governance, who should be included in the relevant participatory proposals – motivated by the BP – naturally emerges as a matter of concern. However, an over-emphasis on broadening participation unfortunately shifts our attention away from a main factor driving corporate leadership to make unethical decisions on AI: that is, employees advancing ethical AI decisions are in a fragile position within their organisation. As is observed by a recent report of Ada Lovelace Institute, 'While calling for democratising AI, many of the biggest technology companies have also made mass redundancies to internal teams that are focused on 'responsible AI' and ethics initiatives – the precise teams that would presumably undertake public participation work' (Groves, 2023, p. 11). If corporate leaders already had limited incentives to enforce the advice around ethical AI from their employees – who, compared to non-employees, can be reasonably expected to know better the organisational contexts that produce AI decisions – then why would they prioritise the opinions of the outsiders, who may even be less familiar with the culture, interests and priorities of the corporation?

The problem, in short, is not that corporate leaders are unaware of what decisions should be made to ensure ethical AI development and governance, but rather that ethical AI practices do not weigh as much as profitability in the eyes of many corporate leaders. Voices around ethical AI, whether they come from the team responsible for advancing it or from other employees, are usually subordinate to the interests of corporate leadership. This often results from the relationship of *domination* between corporate leaders and employees. This relationship inevitably exists, however, insofar as employees need to stay in their current organisations for an income, without substantive exit options. But it is possible to democratise the workplace, to moderate this relationship of domination. Ways of democratising the workplace 'range from rather bland organizational reforms of hierarchical structures to delegation systems based on workers' representation, such as the German codetermination model, to more involved forms of workers' direct participation through works councils…[as well as] ownership of the firms by workers' (Frega et al., 2019, p. 1). It is beyond the scope

of this article to defend a specific proposal of workplace democracy. However, the normative aspiration of it (i.e. addressing domination in the workplace) responds more directly to a more crucial factor impeding corporates' interest in taking ethical AI practices seriously, rather than as a symbolic gesture, while this factor rarely invites serious inquiry in the scholarships on AI democratisation.

It is beyond the scope of this article to develop a full defence of workplace democracy, and there are several legitimate concerns over my view that should be acknowledged. My responses to these concerns, though, will shed light on several ways in which workplace democracy and Democratic Control of Tech Firms, as a paradigmatic initiative of the DAI agenda, complement each other. The first concern is that, if some corporations object to DAI on the basis that this impedes innovation and efficiency, why would they opt for a more radical proposal to democratise *themselves*? Arguably, corporations are unlikely to use workplace democracy to elevate those responsible for ethical AI, at least voluntarily. Corporate leadership is, in many cases, concerned with profitability and generally views ethical AI as a cost that interferes with this goal.[2] To be sure, it will be challenging to enforce any proposal to empower workers in the workplace, whether this takes the form of widening participation in AI governance and development, or limiting the power disparities between employers and employees. Realistically, perhaps the most viable way to incentivise corporations to take ethical AI seriously is to impose greater political and economic costs on corporate leaders when they fail to hold AI practices up to sound ethical standards. This is when Democratic Control of Tech Firms becomes complementary, or even essential. Public institutions authorised through democratic means serve as external forces imposing pressures on corporations, to hold them accountable not just for their ethically concerning decisions on AI, but also for the ways in which corporations moderate the relationship of domination between employers and employees.

The second concern is that workplace democracy may erode efficiency and corporate freedom, as two drivers for AI innovations. For example, some argued that monitoring from higher levels of corporate leadership promote employees' performance at work (Holmstrom & Milgrom, 1994). But such monitoring may not be possible in a democratic workplace. Moreover, capital owners have a legitimate interest in entrepreneurship, which implies creating organisational structure and culture in accordance with their own values (Tomasi, 2012). In response, two points should be flagged. On the one hand, efficiency and corporate freedom are not the only decisive factors in evaluating the moral promise of different governance strategies of AI. An emphasis on efficiency and corporate freedom, in fact, echoes the familiar but contested utilitarian assumption that good AI policies depend enormously, if not fully on the benefits they can produce. But this is not a comfortable assumption to be granted, as the pursuit of AI's benefits must be weighed against the moral rights and interests of other individuals as well. Efficiency and corporate freedom are desirable, but not always decisive. On the other hand, these concerns are not characteristic of workplace democracy, but of various DAI initiatives as well. Similar arguments, for instance, can be made of employees' interest in having more ownership of the strategies they impose on their own AI products or governance structures as well, due to employees' shared culture, interests or identities that make such products or structure possible. It is thus reasonable for corporations to assign less weight to the inputs of non-employees, when making decisions around AI development or governance, and prioritise the voices of their own employees. While this is an implication of my earlier comments, it does not follow that a similar point could be made of workplace democracy. In essence, workplace democracy aims to equalise the power of the individuals *within*, not those outside of a single organisational structure. Employees have a stronger claim to be given rights and responsibilities comparable to those of their leaders, precisely because their corporate membership is a plausible basis for their shared culture, interests and identity. Thus, even when capital owners have a legitimate interest in entrepreneurship, it should be competed against the interests of employees in having meaningful control of their common workplace.

[2] This is a crucial point made by one of the reviewers.

Relatedly, Democratic Control of Tech Firms is valuable, because it honours institutional mechanisms through which citizens express their views on how, if at all, workplace democracy, or other AI democratisation initiatives, should be enforced.

Finally, it might be argued that, although reducing the power inequality between employers and employees provides the latter with greater opportunity to advance ethical AI decisions, this does not necessarily incentivise employees to prioritise ethical AI over the profit. Arguably, employees could have even stronger incentives to prioritise their firms' profitability over everything else. The reason is that, in a relatively democratic workplace where the firms' power and benefits are distributed more equitably, employees can be expected to have a greater sense of ownership of their work and solidarity. If, for instance, there is a moderately democratic organisation that is commercially successful mainly because it prioritise innovations above everything else, challenging existing paradigms of what technology should and cannot do, then members of the organisation are likely to share that particular emphasis on innovation rather than safety. In fact, a democratic workplace in which employees unanimously support ethically concerning AI practices, would provide such practices with even greater legitimacy. This is, in my view, the most challenging objection to workplace democracy being the panacea for corporations' limited incentives to enforce ethical AI practices. After all, organisations known for their flat hierarchy may still actively endorse AI strategies underpinning democratic values. This points out that why accountability burdens on corporations coming from the wider democratic society are crucial. In cases where an organisation is marked by its members' unanimous support of ethically controversial AI practices, external forces from other democratic citizens must be in place to constrain them.

4. Conclusion

In this article, I have offered a resolution to the boundary problem (BP) as a challenge for the 'democratising AI' (DAI) agenda. While the BP is less challenging than it first appears, however, many DAI initiatives misleadingly shift our attention away, from the main reason why corporations are demotivated to make ethical decisions around AI – that is, the power disparities within organisations that put employees advancing ethical AI in a vulnerable position – to the inclusion of non-employees in corporations' processes of the R&D and governance of AI. I argued that workplace democracy provides a more promising normative aspiration than the DAI agenda, although some elements of the latter will complement certain problems that workplace democracy is unable to address, for example, when members of tech firms unanimously prioritise ethically concerning AI practices.

Acknowledgements. I am grateful to the colleagues and students who attended the CFI departmental seminar, where I presented the earlier version of this article.

Funding statement. The author has not received any funding dedicated to the production of this article.

Conflict of interests. The author declares that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

Abbott, R., & Bogenschneider, B. (2018). Should robots pay taxes: Tax policy in the age of automation. *Harvard Law & Policy Review*, *12*(1), 145–176.

Abizadeh, A. (2012). On the demos and its Kin: Nationalism, democracy, and the boundary problem. *American Political Science Review*, *106*(4), 867–882. https://doi.org/10.1017/S0003055412000421

Abizadeh, A. (2021). The scope of the all-subjected principle: On the logical structure of coercive laws. *Analysis*, *81*(4), 603–610. https://doi.org/10.1093/analys/anab041

Allegri, G., & Foschi, R. (2021). Universal basic income as a promoter of real freedom in a digital future. *World Futures*, *77*(1), 1–22. https://doi.org/10.1080/02604027.2020.1792600

Anderson, E. (2006). The epistemology of democracy. *Episteme*, *3*(1-2), 8–22. https://doi.org/10.3366/epi.2006.3.1-2.8

Aristotle (1995) *Politics* (E. Barker, Trans. R. F. Stalley Ed.). Oxford University Press.

Aytac, U. (2024). Big tech, algorithmic power, and democratic control. *The Journal of Politics*, *86*(4), 1431–1445. https://doi.org/10.1086/729938

Banifatemi, A., Miailhe, N., Buse Çetin, R., Cadain, A., Lannquist, Y., & Hodes, C. (2021). Democratizing AI for humanity: A common goal. In B. Braunschweig & M. Ghallab (Eds.), *Reflections on artificial intelligence for humanity* (pp. 228–236). Springer International Publishing.

Birhane, A., Isaac, W., Prabhakaran, V., Diaz, M., Elish, M. C., Gabriel, I., & Mohamed, S. (2022). *Power to the people? Opportunities and challenges for participatory AI*. Paper presented at the Proceedings of the 2nd ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, Arlington, VA, USA. https://doi.org/10.1145/3551624.3555290

Blackman, R. (2022). Why you need an AI ethics committee. Retrieved August 1, 2025, from, https://hbr.org/2022/07/why-you-need-an-ai-ethics-committee.

Brennan, J. (2011). The right to a competent electorate. *The Philosophical Quarterly*, *61*(245), 700–724.

Brennan, J. (2016). *Against democracy*. Princeton University Press.

Buddemeyer, A., Nwogu, J., Solyst, J., Walker, E., Nkrumah, T., Ogan, A., & Stewart, A. (2022). *Unwritten Magic: Participatory Design of AI Dialogue to Empower Marginalized Voices*. Paper presented at the Proceedings of the 2022 ACM Conference on Information Technology for Social Good, Limassol, Cyprus. https://doi.org/10.1145/3524458.3547119

Castro, K. A., Siwady, J. A., Castillo, E., Alonzo, A., Cardona, M., & Perdomo, M. E. (2024). *Artificial Intelligence for All: Challenges and Harnessing Opportunities in AI Democratization*. Paper presented at the 2024 IEEE International Conference on Machine Learning and Applied Network Technologies (ICMLANT).

Christiano, T. (2004). The authority of democracy. *Journal of Political Philosophy*, *12*(3), 266–290.

Christiano, T. (2008). *The constitution of equality: Democratic authority and its limits*. Oxford University Press.

Clough, P. D., & Otterbacher, J. (2023). Democratizing AI: From theory to practice. In E. Carayannis & E. Grigoroudis (Eds.), *Handbook of research on artificial intelligence, innovation and entrepreneurship* (pp. 402–418). Edward Elgar Publishing.

Coeckelbergh, M. (2024). *Why AI undermines democracy and what to do about it*. John Wiley & Sons.

Coeckelbergh, M. (2025). Artificial intelligence, the common good, and the democratic deficit in AI governance. *AI and Ethics*, *5*(2), 1491–1497. https://doi.org/10.1007/s43681-024-00492-9

Cohen, J. (2002). Deliberation and democratic legitimacy. In D. Estlund (Ed.), *Democracy* (pp. 87–106). Blackwell Publishers.

Crawford, K. (2021). *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.

Dahl, R. A. (1990). *After the revolution? Authority in a Good Society*. Yale University Press.

Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López de Prado, M., Herrera-Viedma, E., & Herrera, F. (2023). Connecting the dots in trustworthy artificial intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. *Information Fusion*, *99*, 101896. https://doi.org/10.1016/j.inffus.2023.101896

Durán, J. M., & Pozzi, G. (2025). Trust and trustworthiness in AI. *Philosophy & Technology*, *38*(1), 16. https://doi.org/10.1007/s13347-025-00843-2

Estlund, D. (2008). *Democratic authority: A philosophical framework*. Princeton University Press.

EU (2019). *Ethics guidelines for trustworthy AI*. Author.

Foley, J. R., & Polanyi, M. (2006). Workplace democracy: Why bother? *Economic and Industrial Democracy*, *27*(1), 173–191. https://doi.org/10.1177/0143831x06060595

Frega, R. (2021). Employee involvement and workplace democracy. *Business Ethics Quarterly*, *31*(3), 360–385. https://doi.org/10.1017/beq.2020.30

Frega, R., Herzog, L., & Neuhäuser, C. (2019). Workplace democracy—the recent debate. *Philosophy Compass*, *14*(4), e12574.

Fung, A. (2013). The principle of affected interests: An interpretation and defense. In J. Nagel & R. Smith (Eds.), *Representation: Elections and beyond* (pp. 236–268). University of Pennsylvania Press.

Furendal, M. (2025). Collective ownership of AI. *A Companion to Applied Philosophy of AI*, 372–386.

Gilbert, S. (2020). Big tech and data ethics. Retrieved August 1, 2025, from, https://www.bennettschool.cam.ac.uk/blog/big-tech-and-data-ethics/?utm_source=chatgpt.com.

Gill, K. S. (1989). Reflections on participatory design. *AI & Society*, *3*(4), 297–314. https://doi.org/10.1007/BF01908620

Gomstyn, A., & Jonker, A. (2024). Democratizing AI: What does it mean and how does it work? Retrieved August 1, 2025, from, https://www.ibm.com/think/insights/democratizing-ai.

Gomstyn, A., Jonker, A., & McGrath, A. (2024). What is trustworthy AI? Retrieved August 1, 2025, from, https://www.ibm.com/think/topics/trustworthy-ai.

Goodin, R. E. (2007) Enfranchising all affected interests, and its alternatives. *Philosophy and Public Affairs*, *35*(1), 40–68. Retrieved from http://www.jstor.org/stable/4623780

Goodin, R. E. (2016). Enfranchising all subjected, worldwide. *International Theory*, *8*(3), 365–389. https://doi.org/10.1017/S1752971916000105

Goodin, R. E., & Arrhenius, G. (2024). Enfranchising all subjected: A reconstruction and problematization. *Politics, Philosophy & Economics*, *23*(2), 125–153. https://doi.org/10.1177/1470594x241232023

Goodin, R. E., & Spiekermann, K. (2018). *An epistemic theory of democracy*. Oxford University Press.

Groves, L. (2023). *Going Public: Exploring Public Participation in Commercial AI Labs*. Retrieved August 1, 2025, from, https://www.adalovelaceinstitute.org/report/going-public-participation-ai/.

Han, S. J. (2025). The question of AI and democracy: Four categories of AI governance. *Philosophy & Technology*, *38*(2), 83. https://doi.org/10.1007/s13347-025-00904-6

Held, D. (2006). *Models of Democracy*. Cambridge University Press.

Himmelreich, J. (2023). Against "Democratizing AI". *AI & Society*, *38*(4), 1333–1346. https://doi.org/10.1007/s00146-021-01357-z

Holmstrom, B., & Milgrom, P. (1994) The firm as an incentive system. *The American Economic Review*, *84*(4), 972–991. Retrieved from http://www.jstor.org/stable/2118041.

Honohan, I. (2018). Republicanism and the all subjected principle as the basis of democratic membership. In R. Bauböck (Ed.), *Democratic inclusion* (pp. 143–159). Manchester University Press.

Inoue, A. (2024). The proper scope of the all-subjected principle. *Political Studies Review*, *22*(2), 427–435. https://doi.org/10.1177/14789299231160513

James, A. (2020). Planning for mass unemployment: Precautionary basic income. In S. M. Liao (Ed.), *Ethics of artificial intelligence* (pp. 0). Oxford University Press.

Jozaghi, E. (2024). The initial stage of the artificial intelligence revolution: Access to basic income is a human rights issue. *Journal of Health Care for the Poor and Underserved*, *35*(2), xv–xvi.

Landemore, H. (2017). *Democratic reason: politics, Collective intelligence, and the rule of the many*. Princeton University Press.

Landemore, H., & Ferreras, I. (2016). In defense of workplace democracy:Towards a justification of the firm–state analogy. *Political Theory*, *44*(1), 53–81. https://doi.org/10.1177/0090591715600035

Lebovitz, S., Lifshitz-Assaf, H., & Levina, N. (2022). To engage or not to engage with AI for critical judgments: How professionals deal with opacity when using ai for medical diagnosis. *Organization Science*, *33*(1), 126–148. https://doi.org/10.1287/orsc.2021.1549

Lee, L., & Getahun, H. (2024). Former OpenAI board members say the company can't be trusted to govern itself. Retrieved August 1, 2025, from, https://www.businessinsider.com/former-openai-board-members-challenge-self-governance-sam-altman-2024-5.

Lin, T.-A. (2024). "Democratizing AI" and the concern of algorithmic injustice. *Philosophy & Technology*, *37*(3), 103. https://doi.org/10.1007/s13347-024-00792-2

Lu, Q., Zhu, L., Whittle, J., & Xu, X. (2023). *Responsible AI: Best practices for creating trustworthy AI systems*. Addison-Wesley Professional.

Lythreatis, S., Singh, S. K., & El-Kassar, A.-N. (2022). The digital divide: A review and future research agenda. *Technological Forecasting and Social Change*, *175*, 121359. https://doi.org/10.1016/j.techfore.2021.121359

Marr, B. (2024). Will AI make universal basic income inevitable? Retrieved August 1, 2025, from, https://www.forbes.com/sites/bernardmarr/2024/12/12/will-ai-make-universal-basic-income-inevitable/.

Marx, K. (1844 [2016]). On the Jewish question. In R. Blaug, & J. Schwarzmantel (Eds.), *Democracy: A reader* (pp. 268–273). Columbia University Press.

Mathiesen, K. (2012). The human right to Internet access: A philosophical defense. *The International Review of Information Ethics*, *18*, 9–22.

Mayer, R. (2000) Is there a moral right to workplace democracy? *Social Theory and Practice*, *26*(2), 301–325. Retrieved from http://www.jstor.org/stable/23559088

Mazur, O. (2018). Taxing the Robots. *Pepperdine Law Review*, *46*(2), 277–330.

Mill, J. S. (1859 [2015]). On liberty. In M. Philp, & F. Rosen (Eds.), *On liberty, utilitarianism and other essays* (pp. 5–114). Oxford University Press.

Mill, J. S. (1861 [2015]). Considerations on representative government. In M. Philp, & F. Rosen (Eds.), *On liberty, utilitarianism and other essays* (pp. 115–180). Oxford University Press.

Miller, D. (1999). Bounded Citizenship. In K. Hutchings, & R. Dannreuther (Eds.), *Cosmopolitan citizenship* (pp. 60–80). Palgrave Macmillan UK.

Miller, D. (2009). Democracy's Domain. *Philosophy and Public Affairs*, *37*(3), 201–228. https://doi.org/10.1111/j.1088-4963.2009.01158.x

Miller, D. (2020). Reconceiving the democratic boundary problem. *Philosophy Compass*, *15*(11), e12707. https://doi.org/10.1111/phc3.12707

Muldoon, J. (2025). Data-owning democracy or digital socialism? *Critical review of international social and political philosophy*, *28*(4), 570–591. https://doi.org/10.1080/13698230.2022.2120737

Murphy, J. W., & Taylor, R. R. (2024). To democratize or not to democratize AI? That is the question. *AI and Ethics*, *4*(4), 1357–1363. https://doi.org/10.1007/s43681-023-00313-5

Noorman, M., & Swierstra, T. (2023). Democratizing AI from a sociotechnical perspective. *Minds and Machines*, *33*(4), 563–586. https://doi.org/10.1007/s11023-023-09651-z

O'Keefe, C., Cihon, P., Garfinkel, B., Flynn, C., Leung, J., & Dafoe, A. (2020). *The Windfall Clause: Distributing the Benefits of AI for the Common Good*. Paper presented at the Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society, New York, NY, USA. https://doi.org/10.1145/3375627.3375842

Oberson, X. (2025). Taxation of artificial intelligence. In Woodrow Barfield (Ed.), *Research handbook on the law of artificial intelligence* (pp. 636–651). Edward Elgar Publishing.

Okidegbe, N. (2023) To democratize algorithms toward an abolitionist future. *UCLA Law Review*, *69*(6), 1688–1743. Retrieved from https://heinonline.org/HOL/P?h=hein.journals/uclalr69&i=1726.

OpenAI (2024). OpenAI board forms safety and security committee. Retrieved August 1, 2025, from, https://openai.com/index/openai-board-forms-safety-and-security-committee/.

Owen, D. (2012). Constituting the polity, constituting the demos: On the place of the all affected interests principle in democratic theory and in resolving the democratic boundary problem. *Ethics & Global Politics*, *5*(3), 129–152. https://doi.org/10.3402/egp.v5i3.18617

Pope, N. (2024). What is trustworthy AI? Retrieved August 1, 2025, from, https://blogs.nvidia.com/blog/what-is-trustworthy-ai/.

Randieri, C. (2024). The democratization of AI: Bridging the gap between monopolization and personal empowerment. Retrieved August 1, 2025, from, https://www.forbes.com/councils/forbestechcouncil/2024/03/25/the-democratization-of-ai-bridging-the-gap-between-monopolization-and-personal-empowerment/.

Rawls, J. (1999). *A theory of justice*. The Belknap Press of Harvard University Press.

Rawls, J. (2001). *Justice as fairness: A restatement*. The Belknap Press of Harvard University Press.

Rawls, J. (2005). *Political liberalism*. Columbia University Press.

Reglitz, M. (2020). The human right to free internet access. *Journal of Applied Philosophy*, *37*(2), 314–331.

Reglitz, M. (2024). *Free internet access as a human right*. Cambridge University Press.

Risse, M. (2023). *Political theory of the digital age: Where artificial intelligence might take us*. Cambridge University Press.

Rosenberg, J. H. (2020) The all-affected principle reconsidered. *Social Theory and Practice*, *46*(4), 847–867. Retrieved from http://www.jstor.org/stable/45302468.

Rossi, F. (2020). How IBM is working toward a fairer AI. Retrieved August 1, 2025, from, https://hbsp.harvard.edu/product/H05Z6V-PDF-ENG.

Rubeis, G., Dubbala, K., & Metzler, I. (2022). "Democratizing" artificial intelligence in medicine and healthcare: Mapping the uses of an elusive term. *Frontiers in Genetics*, *13*, –2022. https://doi.org/10.3389/fgene.2022.902542

Schmager, S., Pappas, I. O., & Vassilakopoulou, P. (2025). Understanding human-centred AI: A review of its defining elements and a research agenda. *Behaviour & Information Technology*, 1–40. https://doi.org/10.1080/0144929X.2024.2448719

Schuett, J., Reuel, A.-K., & Carlier, A. (2025). How to design an AI ethics board. *AI and Ethics*, *5*(2), 863–881. https://doi.org/10.1007/s43681-023-00409-y

Sclove, R. (1995). *Democracy and technology*. Guilford Press.

Simmons, A. J. (2013). Democratic authority and the boundary problem. *Ratio Juris*, *26*(3), 326–357. https://doi.org/10.1111/raju.12017

Smith, H. (2021). Clinical AI: Opacity, accountability, responsibility and liability. *AI & Society*, *36*(2), 535–545. https://doi.org/10.1007/s00146-020-01019-6

Somin, I. (2014). Why political ignorance undermines the wisdom of the many. *Critical Review*, *26*(1-2), 151–169. https://doi.org/10.1080/08913811.2014.907047

Song, S. (2009). Democracy and noncitizen voting rights. *Citizenship Studies*, *13*(6), 607–620. https://doi.org/10.1080/13621020903309607

Song, S. (2012). The boundary problem in democratic theory: Why the demos should be bounded by the state. *International Theory*, *4*(1), 39–68. https://doi.org/10.1017/S1752971911000248

Tomasi, J. (2012). *Free market fairness*. Princeton University Press.

Vaassen, B. (2022). AI, opacity, and personal autonomy. *Philosophy & Technology*, *35*(4), 88. https://doi.org/10.1007/s13347-022-00577-5

Vrousalis, N. (2019). Workplace democracy implies economic democracy. *Journal of Social Philosophy*, *50*(3), 259–279. https://doi.org/10.1111/josp.12275

Wakefield, J. (2019). Google's ethics board shut down. Retrieved August 1, 2025, from, https://www.bbc.co.uk/news/technology-47825833.

Whelan, F. G. (1983). Prologue: Democratic theory and the boundary problem. *Nomos*, *25*, 13–47. Retrieved from http://www.jstor.org/stable/24219358.

Wilson, J. L. (2022) Making the all-affected principle safe for democracy. *Philosophy and Public Affairs*, *50*(2), 169–201. Retrieved from https://heinonline.org/HOL/P?h=hein.journals/philadp50&i=171.

Yeoman, R. (2014). *Meaningful work and workplace democracy: A philosophy of work and a politics of meaningfulness*. Springer.

Zeide, E. (2022). The silicon ceiling: How artificial intelligence constructs an invisible barrier to opportunity. *UMKC Law Review*, *91*(2), 403–436.

Zimmermann, A., Rosa, E. D., & Kim, H. (2020). Technology can't fix algorithmic injustice. Retrieved August 1, 2025, from, https://www.bostonreview.net/articles/annette-zimmermann-algorithmic-political/.

Zytko, D., Wisniewski, P. J., Guha, S., Baumer, E. P. S., & Lee, M. K. (2022). *Participatory Design of AI Systems: Opportunities and Challenges Across Diverse Users, Relationships, and Application Domains*. Paper presented at the Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems, New Orleans, LA, USA. https://doi.org/10.1145/3491101.3516506

Dr William Chan is a Teaching Fellow at the Leverhulme Centre for the Future of Intelligence. His research focuses on the intersection of normative political theory and AI ethics, especially in the domains of democratic politics, education and healthcare.

Cite this article: Chan, W. (2026). The boundary problem of democratising AI in the corporation. *Cambridge Forum on AI: Law and Governance*, *2*, e4, 1–15. https://doi.org/10.1017/cfl.2026.10045