

**WEB SCIENCE
INSTITUTE**



The Algorithmic Hand: Artificial Intelligence, Democracy, and Collective Action at Scale

Position Paper 2025-02
November 2025

Tiago C. Peixoto



About the Web Science Institute

The Web Science Institute (WSI) brings together the University of Southampton's world-class, interdisciplinary, socio-technical expertise in web science, data science and artificial intelligence (AI) to leverage the unique role of online technologies in tackling global challenges. We work to create collaborations within the University and with industry, governments and third sector organisations that bring interdisciplinary socio-technical insights and impacts to the world's most pressing problems.

<https://www.southampton.ac.uk/research/institutes-centres/web-science-institute>

About the Centre for Democratic Futures

The Centre for Democratic Futures (CDF) is an interdisciplinary research hub that brings together a diverse group of academics from across the social sciences, humanities, physical and life sciences, and beyond. Our team includes experts in democratic theory, governance, health science, computer science, public policy, history, sociology, environmental studies, new technologies, and more. The centre is dedicated to exploring the complex challenges facing democracies today and to developing innovative strategies to promote democratic resilience and sustainability. We foster a collaborative environment where research is informed by a rich array of theoretical perspectives and methodological approaches, spanning qualitative, quantitative, and mixed-methods research

<https://www.southampton.ac.uk/research/institutes-centres/centre-for-democratic-futures>

Copyright © the authors 2025

The opinions expressed in this publication are those of the authors and do not necessarily reflect the views of the University of Southampton, the Web Science Institute or its Board of Directors.



This work is licensed under a Creative Commons Attribution — Non-commercial — No Derivatives Licence. To view this licence, visit (www.creativecommons.org/licenses/by-nc-nd/3.0/). For reuse or distribution, please include this copyright notice.

Web Science Institute, Building 32, Highfield Campus, University of Southampton, SO17 1BJ ws@oton.ac.uk (DOI:10.5258/SOTON/WSI-WP014)

About the Author

Tiago C. Peixoto is an international civil servant and a visiting professor at the Centre for Democratic Futures at the University of Southampton.

He has been recognized as one of the 20 Most Innovative People in Democracy and one of the 100 Most Influential People in Digital Government, and is a recipient of the Louis Brownlow Award, which honours outstanding contributions on topics of wide contemporary interest to both practitioners and scholars in public administration.

His work has been featured in major media outlets and published in leading academic journals, including the *British Journal of Political Science*, *Public Administration Review*, *Journal of Information Technology and Politics*, and *UCLA Law Review*.

Democracy reformers have always had a technological imagination.

Since at least the 18th century, new technologies have prompted hopes that they might reinvent politics. When the optical telegraph was introduced in revolutionary France, one philosopher speculated that it could solve Rousseau's puzzle: how to achieve democratic self-rule at scale (Vandermonde 1994). If communication across distance could match the experience of a living-room conversation, why not extend it across the fledgling republic?

This belief that technologies could bring us to a new democratic era has reappeared throughout history. In 1919, the San Francisco Chronicle claimed that the telephone would revive the participatory spirit of ancient Athens. In the 1980s, civic futurists saw in cable television a platform for "teledemocracy." In the 1990s, Al Gore envisioned a "new Athenian Age" enabled by the Information Superhighway. And by the 2010s, open data and civic tech platforms were supposed to bring about an era of transparency, participation, and accountable governance (Yu & Robinson 2011, Peixoto & Steinberg 2019).

But history tells a more complex story. As sociologist Armand Mattelart (1999) observed, these cycles of enthusiasm often rely on what he called a "strange alchemy of cynicism, naïveté, and amnesia". The assumption that better technology can fundamentally alter democracy has proven resilient, and just as often wrong.

And yet, in the last three decades, few tribes within the world of democratic theory and practice have been more captivated by the promise of technology than the deliberative democrats. The democracy field, after all, is a landscape of tribes who believe passionately in the collective intelligence of everyone – except, it sometimes seems, those in other democratic tribes (e.g., participatory democracy, electoral democracy).

Within this ecosystem, deliberative democrats have long struggled with a particular constraint: scaling thoughtful dialogue beyond the small settings where it most clearly flourishes. Nowhere is this clearer than in the case of citizens' assemblies, the emblematic instrument of the "deliberative turn," whose promise and limits reveal both the achievements and the frustrations of the field. Artificial intelligence (AI) reopens that challenge, not only because it promises deliberation at scale, but because it may change the basic economics of collective action: the very bottleneck deliberative democracy has never overcome.

What becomes of deliberative democracy when it meets AI? This paper does not claim immunity from the historical pattern just described. It risks the same contradiction of diagnosing techno-solutionism while proposing technical means. Yet it proceeds from a

different premise. Current and past waves of optimism have treated new technologies as means to increase participation, with little regard for the depth of democratic commitment or the structures of power that determine whether participation is consequential.

Here, the difference lies in what technology is asked to do. AI cannot create democratic commitment where none exists, but it might alter coordination costs. By lowering barriers for citizens and raising the costs of institutional unresponsiveness, it could change the calculus of power rather than bypassing it. Whether civic actors can build and govern such systems responsibly remains uncertain. What follows explores that possibility.

The starting point is familiar: despite decades of innovation in deliberative design, most experiments have remained small in scale and have rarely translated into institutional responsiveness. This essay examines why scale matters for deliberative democracy, then argues that scaled deliberation must connect to collective action, exploring how AI-enabled coordination might lower the costs of organizing while raising the costs of institutional unresponsiveness. It concludes by confronting the structural obstacles of power and funding, and by identifying what deliberative processes must target to matter politically.

Deliberative Democracy and Scalability

At its core, deliberative democracy is the idea that democracy is not about just voting or voicing preferences. Above all, it's about exchanging reasons. It holds that public decisions should emerge from inclusive, informed, and reflective dialogue among citizens, not just aggregated opinions or partisan combat. The contemporary vision of what counts as deliberative democracy - also known as *mini-publics* - is as noble as it is demanding: a group of people (often randomly selected to promote diversity) learns about the issue at stake, considers diverse points of view, debates, takes part in facilitated dialogues, and ultimately decide through reasoned argument (Guttmann & Thompson 2004, Chambers 2018, Curato et al. 2021). How those decisions then affect the world outside the deliberations is often left unsaid.

Rooted in the work of German philosopher Jürgen Habermas, deliberative democrats embrace the belief that, as reasonable beings, we can arrive at better collective decisions through what he famously called “the unforced force of the better argument” (Habermas 1996, p. 306). In essence, Habermas argued that the legitimacy of political and moral decisions should rest on the quality of the reasoning behind them, and not on force, manipulation, or the influence of power. For him, persuasion is only valid when it arises from open, inclusive, and egalitarian dialogue in which participants aim not to win but to understand one another – and to be guided by the strength of the argument itself, not by status or pressure.

Yet this vision has long faced powerful critiques, particularly from agonistic democrats like Chantal Mouffe (1999), who argue that politics is fundamentally about managing irreconcilable conflicts rather than transcending them through reason. Deliberation, in this view, often privileges certain forms of speech and certain speakers, reproducing existing power relations under the guise of neutrality. The turn toward collective action later in this essay acknowledges this critique: democratic progress requires not just better arguments, but organized power capable of making those arguments politically unavoidable.

Citizens' assemblies represent deliberative democracy's most ambitious institutional bet: that purposefully designed processes can produce outputs that inform public decisions.¹ Like other minipublics, they are designed to embody the ideals of deliberation in practice. These assemblies bring together a cross-section of the public, usually selected by lot, to learn about a policy issue, engage with expert testimony and diverging points of view, deliberate with one another, and collectively issue recommendations (Smith 2009). What was once a theoretical aspiration has gained real institutional traction: governments in countries as diverse as Ireland, Canada, France, and Mongolia have embraced citizens' assemblies to, at least allegedly, grapple with complex or contested policy issues. International organizations such as the OECD, the UN, and the European Commission have also lent their support, promoting deliberative processes as a means of renewing democratic ideals.

A landmark example is the Irish Citizens' Assembly, where 99 randomly selected citizens were tasked with deliberating on issues such as same-sex marriage and abortion law (Suiter et al. 2016). After weeks of learning and dialogue, the Assembly issued recommendations that laid the groundwork for constitutional referendum. In the case of abortion, the assembly's proposals directly informed the ballot, ultimately helping to overturn one of the world's most restrictive abortion bans, an outcome widely seen as a triumph of deliberative process over polarization.

But that vision has always run up against the limits of logistics: time, space, and resources. In short, *scale*. By the mid-1990s, some deliberative democrats began to look to the Internet as a possible way to transcend those constraints (Schneider 1996, White 1997, Coleman & Gotze 2001). Three decades later, it is hard to engage with the fields of democracy or artificial intelligence without encountering proposals to use AI to enable deliberation at scale.

What is as striking now as it was then, is how rarely the push for scale is grounded in a clear account of what scaling is meant to solve. The assumption that expanding the

¹ Citizens assemblies and other types of mini-publics are by no means the only forms of deliberative action – many of which have achieved mass participation without resorting to technology, such as National Public Policy Conferences in Brazil (Pogrebinschi & Samuels 2014) and Gram Sabhas in India (Sanyal and Rhao 2019).

number of participants will inherently deepen democratic outcomes is more often implicit than justified (e.g., Klein et al. 2024, Small et al. 2021, McKinney 2024, Pentland and Tsai 2024, Tessler et al. 2024, Tsai et al. 2024). The normative case for deliberation at scale has too often remained implicit, as if the ability to scale were a self-evident good. And this conceptual leap, from technical feasibility to democratic desirability, deserves closer scrutiny.

If scale is the solution, then what is the problem?

Deliberative Democracy Without Scale

To begin answering that question, it's worth identifying the kinds of democratic goods that deliberation at scale might advance. While rarely spelled out, the promise of scaling up deliberative practices often gestures toward deeper learning, broader inclusion, and stronger legitimacy. In what follows, I outline seven areas where scale could make a difference.

1. Civic competence: One of the core tenets of deliberative democracy is that participation doesn't just serve decision-making: it builds civic capacity (Fishkin 1997, Nabatchi 2010). When citizens engage in structured dialogue that requires listening, reason-giving, and respect for opposing views, they tend to develop an understanding of complex issues, sharpen their reasoning skills, and gain a more nuanced appreciation of opposing perspectives. Deliberation, in this sense, is not merely a mechanism for producing outcomes; it is also a school for democratic citizenship. The more people are exposed to this kind of engagement, the more widespread these civic benefits become. Scaling deliberation, then, is not just about quantity. It is about extending the formative experience of political learning to more members of the public, transforming participation into a vehicle for civic education at scale.²

2. Depolarization: Another compelling argument for scaling deliberation is its potential to counteract political polarization. One of the most persistent challenges in contemporary democracies is the entrenchment of partisan identities. Deliberative democrats have long argued that reasoned dialogue with diverse others can mitigate such polarization by fostering mutual understanding and shifting the frame from tribal allegiance to shared reasoning (Dryzek 2005, Fishkin 2009). Until recently, however, this remained more of a normative hope than an

² Enhancing civic competence at scale may also generate positive externalities such as greater tax morale. Evidence increasingly suggests that participatory experiences can raise citizens' propensity to comply voluntarily with taxation. For governments, this can translate into higher revenues and improved fiscal legitimacy, offering both a political and administrative incentive to invest in participatory and deliberative processes (see Touchton, Wampler, and Peixoto 2021; Sjöberg et al. 2019; Peixoto 2012).

empirically grounded claim. That is beginning to change. Evidence from larger scale deliberative experiments, including those using AI-assisted platforms, suggests that participation in deliberations can produce depolarizing effects, particularly among individuals who begin with the most extreme positions (Fishkin et al. 2025, Tessler et al. 2024). When citizens are given space to consider competing arguments seriously, alongside others from different political backgrounds, partisan gaps narrow – not because people are persuaded to switch sides, but because they come to see issues through a more complex, less binary lens. If these effects could be extended to larger segments of the population, then deliberation would become, at least hypothetically, not just an educational exercise, but a democratic antidote to the social fragmentation that plagues public life.

3. Inclusiveness: Despite its egalitarian aspirations, deliberative democracy as typically practiced is exclusionary by design. Citizens' assemblies and other similar methods rely on randomized, stratified selection to ensure demographic diversity and representativeness, but this mechanism limits participation to the few (Smith 2009). The underlying principle is not that every citizen has an equal *right* to participate, but that everyone has an equal *chance* of being selected (Sintomer 2023). As a result, most members of the public are structurally excluded from it (Spada & Peixoto 2025). This exclusion may be justifiable in a world of resource constraints. But in this case, limiting deliberation to its typical sizes begins to look less like a design feature and more like a constraint in search of a workaround. Scaling deliberation offers one such workaround. By expanding the number of people who can engage in structured, reasoned dialogue, we move closer to fulfilling the inclusive promise that deliberative democracy so often claims but rarely realizes. Scale, in this sense, is not just a logistical challenge; it is a matter of democratic justice.

These first three benefits, civic competence, depolarization, and inclusiveness, are on their own compelling reasons to pursue deliberation at scale. Both speak to what deliberative democracy already promises: a more thoughtful, more connected democratic citizenry. In that sense, they strengthen the case for *deliberation itself*, with scale serving mainly to extend its reach and deepen its effects.

By contrast, the remaining justifications put forward here are of a different kind. They reflect limitations that even the most well-designed citizens' assemblies, in their typical size, fall short of overcoming. Precisely because they are not intrinsic to traditional deliberative practice, and refer to its external shortcomings, they are less frequently addressed by deliberative democracy scholars and advocates. Yet they may prove just

as decisive in determining whether deliberation can, and should, shape democratic governance.

4. Statistical representation: One of the rationales behind minipublics is that, by selecting a diverse cross-section of the population through random sampling, they can serve as a kind of democratic microcosm: a window into what the public might think if given the chance to deliberate under good conditions. This idea has long underpinned claims that such bodies are “broadly representative of the population.”

Yet the silent truth is that most deliberative arrangements, especially citizens’ assemblies, fall short of any statistical representativeness, let alone intersectional dimensions such as race, gender, class, education, or disability. The field has only recently begun to confront this limitation, as the insistence on earlier assumptions has become untenable in light of well-established evidence (Spada & Peixoto 2025).³

But the structural tension remains: if representation is the goal, then small sample sizes, however well stratified, will always be a limiting factor. Scaling deliberation offers a potential corrective. By increasing the number of participants, we also increase the chances of capturing the full heterogeneity of a population, not just across broad demographic categories, but within them. If deliberation is to be both inclusive and statistically representative, scale may be the only viable path.

5. Collective Intelligence: Another argument for scaling deliberation lies in the recognition that wisdom depends not just on how decisions are made, but on who gets to contribute to them. Some advocates of AI see its promise in the ability to “listen at scale,” to extract preferences from vast populations and optimize for collective outcomes. But from a deliberative democratic perspective, this raises a deeper question: what if the answer isn’t just out there waiting to be aggregated, but instead waiting to be formulated through public reasoning? Small deliberative bodies, however well selected, risk excluding segments of the population and the knowledge, values, and lived experiences they hold (Klein 2021, Peixoto and Sifry 2017).

If participation is exclusionary by design, so are the epistemic resources available within the room. Insights, especially those grounded in local, tacit, or marginalized forms of knowledge, may never be surfaced. Scaling deliberation, then, is not simply about adding more voices. It is about unlocking collective intelligence that

³ Recent recognition of these challenges is evident, for instance, in Yale ISPS’s September 2025 conference ‘Designing Participation: Recruitment and Technology in Deliberative Democracy,’ which brought together researchers and practitioners to address selection bias, sampling, and recruitment in citizens’ assemblies.

is otherwise out of reach. If deliberative democrats truly believe that the many are smarter than the few (Hong & Page 2012, Landemore 2012, Noveck 2019, Neves & Peixoto 2020), then expanding who gets to deliberate is not just a democratic imperative: it's a precondition for better decisions.

6. *Legitimacy*: Legitimacy is not a property that deliberative processes can declare for themselves; it must be granted by others. It emerges through a relationship between those who claim authority and those from whom that authority is sought. In this sense, legitimacy is not an internal characteristic of a process, but a social construct: it depends on whether those outside the deliberation accept its outcomes as valid and binding. Or, as in the words of Mansbridge et al. (2007:12) “for those excluded, no deliberative democratic legitimacy is generated”. This is where the limits of small-scale deliberation become especially stark (Landemore 2022). Even when participants are carefully selected, those left out of the process may not recognize themselves in the results. If deliberation changes minds, it also risks alienating those who were never part of the conversation (Lafont 2023).

And yet, many deliberative democrats remain uncomfortable with the idea that legitimacy might need to be *earned*: through referenda, public uptake, or broader democratic engagement. The much-celebrated Ireland's Citizens' Assembly on abortion is a case in point: its political force came only after the public ratified its recommendations through a referendum. The deliberative process was critical at the agenda-setting stage, whereas the policy impact was enabled via non-deliberative, direct democratic voting. For deliberation to have democratic weight, it must be reasoned and inclusive, but also recognized. And recognition, in a democracy, is hard to achieve without scale.

7. *Responsiveness*: Of all the reasons to scale deliberation, perhaps the most important and least discussed is political responsiveness. Deliberative processes typically suffer from a voice–responsiveness gap (Geissel 2023, Giraudet et al. 2022, Jacquet & van der Does 2020). They give participants the space to speak, to reason, even to agree; but not necessarily the power to be heard by those who govern. This is not a minor flaw; it is a limitation that haunts the deliberative model, and should haunt it even more as the number of deliberative initiatives multiplies while actual government responsiveness remains low or unchanged. A proliferation of citizens' assemblies without a corresponding increase in institutional response risks creating the appearance of democratic renewal without its substance, further eroding democracy.

Scale changes that equation. All else being equal, the more people who participate in a deliberative process, the higher the political cost of ignoring its outcomes

should be. But it is important to clarify what *responsiveness* means. It is not simply about offering a reply, nor about selectively embracing recommendations that align with preexisting agendas. Governments may “respond” by issuing polite acknowledgments or justifying inaction on procedural or legal grounds. Others may cherry-pick the recommendations that are easiest to implement or least politically risky, while quietly discarding the rest (Font et al. 2018). These forms of response may check boxes, but they do not live up to the democratic ideal.

True responsiveness means treating the outcomes of public deliberation as substantively relevant, even when they are politically inconvenient. It means demonstrating, in good faith, how citizen input has altered the trajectory of state action. It also means creating institutional pathways for follow-up, feedback, and iteration, so that deliberation is not a one-off performance, but part of an ongoing democratic conversation (Fox et al. 2024). Without that, scaling deliberation risks scaling only frustration: more participation, same disregard.

Deliberative democrats have often sought legitimacy through the *quality* of the process (Curato 2015, Parkinson 2003, Raelin 2012). But legitimacy without responsiveness is not democracy: it is simulation. Public deliberation earns its democratic meaning by being consequential, not merely thoughtful. That means embedding deliberation within democratic systems, linking it to mechanisms of accountability and, critically, building it at a scale where it cannot be ignored. Responsiveness is not just another item on the list of deliberative goods. It is the condition that determines whether any of the others matter.

Taken together, these seven areas begin to clarify why scale matters. Not simply as a technical aspiration, but as a democratic necessity. Deliberation is often praised for what it cultivates in small spaces: reflection, understanding, mutual respect. But if it is to shape public life, it must also contend with the realities of political systems; systems that respond to scale, to numbers, to organized voice. The goods of deliberation, such as learning, inclusion, legitimacy, and intelligence, can be meaningful in isolation. But without responsiveness, they risk becoming ornamental. And without scale, they risk remaining marginal.

In the next section, I turn to a different but closely related question: not what deliberation offers citizens, but what citizens can do together, and how scaling deliberation may open new pathways for *collective action* itself.

Bringing Together Deliberative Democracy and Collective Action

If scaled deliberation is to matter, it must connect to the larger project of democratization. Democratization is, at its core, a struggle over enfranchisement: not just the right to vote

but the ability to shape collective decisions and their implementation. That struggle has never been won through dialogue alone. Democratic gains, whether the right to vote, to organize, or to be heard, come not from the quiet force of better arguments but from the noisy work of collective action. Democratic reformers must engage seriously with power: how it is mobilized, challenged, and redistributed.

This does not mean abandoning deliberation. It means recognizing its limits. Politics cannot be cleansed of contestation. Reformers who imagine that democratic progress flows naturally from inclusive deliberation mistake a procedural means for a political end. Rosa Parks' worst use of her time would have been sitting in a randomly selected panel to discuss segregation. Her power came not from deliberating about injustice but from refusing to comply with it, from sparking a movement that made inaction politically untenable. While deliberation clarifies what ought to be done, only collective action forces whether it gets done.

Framing deliberation and collective action as opposing forces misses the point. The productive question is: How might deliberation at scale fuel action? How can reflective dialogue become infrastructure for organizing rather than a substitute for it? What if scaled deliberation were treated not as a dialogical tool but as a means of movement-building?

The civil rights movement offers a reminder. Ella Baker, one of its most influential organizers, believed deeply in the value of grassroots discussion, but not for its own sake. Her countless community meetings were sites of political education, strategic alignment, and collective resolve (Tutashinda 2010; Anthony and Elpers 2023). They transformed diffuse frustration into coordinated disruption. These were deliberative spaces that mattered because they were embedded in power-building projects. That is what AI-enabled deliberation must emulate if it is to translate reflection into action rather than approximate a philosophical ideal where the best argument gently prevails.

Recent work on AI-enabled deliberation has explored how digital public squares and collective dialogue systems can improve the quality and scale of civic discussion (Goldberg et al. 2024; Revel & Penigaud 2025; Summerfield et al. 2024). These innovations represent important progress in scaling public reasoning. Yet they remain oriented toward communicative exchange rather than political leverage. They help publics reason together but stop short of helping publics act together: mobilizing pressure, coordinating advocacy, converting shared preferences into institutional responsiveness.

The relationship between deliberation and coordination determines whether democracy gains or loses agency. AI coordination tools could operate without deliberative inputs, mobilizing publics around pre-existing preferences, organizational agendas, or algorithmically identified opportunities. But coordination without deliberation lacks the

justifying force of public reasoning, offers no mechanism for preference refinement or conflict mediation, and becomes nearly indistinguishable from manipulation when citizens cannot trace how mobilization targets are chosen. Conversely, when scaled deliberation produces broad consensus, say on campaign finance or pharmaceutical pricing, coordination infrastructure can convert that political currency into institutional pressure. Deliberation provides direction, justification, and intelligence: what to demand, why it matters, where alignment exists across divides.

The coordination hypothesis is that each strengthens the other. Deliberation becomes consequential when it leads to coordinated action; coordination gains legitimacy when it flows from public reasoning. When citizens see deliberative consensus translate into actual pressure, when their reasoned judgment moves institutions, they invest more deeply. When institutions face coordinated publics armed with deliberated positions, dismissal becomes costlier. This is not a one-time conversion but a reinforcing cycle that transforms scattered agreement into sustained political force.

Yet this cycle depends on foundations that algorithms alone cannot provide. AI-enabled coordination tools are not seed capital for movements that do not yet exist; they are leverage for movements that already do. Deliberative democracy has rarely intersected with organizing traditions, emphasizing procedural inclusiveness while remaining detached from the networks and trust infrastructures that sustain collective action. Coordination infrastructure functions like any commons: it produces value only when those who use it can define boundaries, enforce norms, and exercise shared control (Ostrom 1990). Without such foundations, AI tools risk amplifying noise, or worse, enabling manipulation disguised as mobilization.

The relationship between deliberation and coordination also reveals a governance challenge. If coordination tools select targets and tactics opaquely, agenda-setting power concentrates in whoever controls the algorithms. But when coordination operates downstream from deliberative processes – where citizens have shaped priorities through structured dialogue – power remains distributed. The deliberative stage functions as a democratic checkpoint: those who coordinate must justify their priorities through public reasoning, not algorithmic black boxes. This does not eliminate the risk of manipulation, but it raises the visibility and political cost of attempting it.

Social movements often deliberate within struggle, not before it. The relationship between reflection and action varies by context. No single sequence fits all democratic work. Yet when the goal is making institutions responsive to broad publics rather than narrow interests, linking deliberation to coordination offers strategic advantages that neither provides alone. Deliberation clarifies what publics want and why. Coordination converts that clarity into sustained pressure. Together, they shift the burden: instead of citizens

hoping institutions will listen, institutions must explain why they refuse to respond. This is the democratic wager of the algorithmic hand: that infrastructure designed to lower coordination costs can make public voice harder to dismiss and costlier to ignore.

The Algorithmic Hand: Re-engineering Democratic Action at Scale

Most thinking about AI within deliberative democracy has been unimaginative, fixated on one objective: enabling deliberation at scale. This scaling up promises to address inclusiveness, civic learning, and shared understanding, all democratic goods deliberative advocates already champion. Yet the field has ignored what comes after scale is achieved. Deliberative democrats have asked how AI can make mass deliberation possible while neglecting how it might make deliberation consequential. The question is not whether thousands can reason together, but whether each additional thousand makes ignoring them costlier. Whether scale builds power or merely amplifies a consensus institutions can still safely dismiss.

The answer is already visible in democracies where consensus exists but responsiveness does not. Finding common ground at scale, a much-celebrated virtue of AI-enabled deliberation (Konya et al. 2025, Summerfield et al. 2024, Tessler et al. 2024), may foster mutual understanding, but it does not guarantee political consequence. The United States offers stark evidence: policies such as universal background checks for gun ownership or Medicaid expansion command broad popular support across the political spectrum, yet political action remains stalled. The bottleneck is not the absence of public agreement but that representative institutions respond to wealthy interests rather than popular majorities (Carnes and Lupu 2015; Gilens 2005; Gilens & Page 2014; Peters and Ensink 2015; Soontjens and Persson 2024). The central challenge for deliberative democracy is not how preferences are formed but whether they matter.

Thomas Schelling's (2006) insight is apt here: individual choices, innocuous in isolation, can aggregate into unexpected social patterns. Deliberation may shift opinions, but without mechanisms that synchronize and mobilize those shifts, the macro-behavior on which democracy depends never materializes. This is where AI enters not as a deliberative facilitator but as a coordination infrastructure: the algorithmic hand. Unlike the invisible hand of markets that aggregates individual self-interest, or the visible hand of state authority that commands compliance, the algorithmic hand coordinates voluntary collective action at scale. It can reshape the information environment and the incentives that govern mobilization, translating individual shifts in opinion into synchronized pressure on institutions.

If large-scale deliberation generates broad consensus, the question becomes: what now? AI coordination infrastructure could match deliberative outputs with the institutional levers

most likely to yield responsiveness, such as identifying which legislators are vulnerable, when budget cycles create opportunities, where media attention amplifies pressure. The goal is no longer just mass deliberation but collective leverage: converting reasoned agreement into coordinated action that institutions cannot safely ignore.

The vision that follows reimagines political influence itself, shifting it from organized special interests toward coordinated publics. It draws on the consensus and shared priorities that deliberation produces, translating collective reflection into strategic coordination. Whereas traditional channels of influence concentrate access and resources in the hands of a few, AI-enabled coordination could extend comparable strategic capacity to deliberative publics. The result is not the replacement of advocacy but its democratization: distributed, data-driven, and anchored in public reasoning rather than concentrated power.

The applications that follow are speculative. None exists at scale today. But each addresses real coordination problems that have long constrained collective action: the challenges Olson (1971) and Ostrom (2010) identified around free-riding, monitoring, and collective choice, alongside contemporary concerns about momentum loss, skill mismatches, and reciprocal mistrust. Rather than digitizing traditional organizing tactics, these applications imagine how AI systems and human actors might collaborate in new configurations – identifying opportunities, synchronizing timing, tracking contributions, matching skills, and maintaining momentum. They illustrate how deliberative consensus might translate into coordinated pressure: not through a single intervention but through persistent infrastructure that makes collective action less costly and institutional non-response more visible.

1. Dynamic Focal-Point Generators with Predictive Feedback

Large language model–powered coordination engines could continuously mine policy calendars, legislative agendas, protest-permit databases, transit feeds, and even meteorological data to identify high-leverage moments for convergence: a committee markup, a budget vote, or a narrow 48-hour window when public pressure might decisively shift outcomes. These systems would not only detect opportunities but also estimate their expected impact.⁴ By integrating legislative scraping, media sentiment analysis, lobbying activity, and fundraising flows, they could quantify the marginal payoff of additional effort: "Five thousand more constituent calls by Wednesday raises the probability of amendment passage from 32% to 47%." Such predictive feedback helps movements recognize when to

⁴ For instance, the UK government recently adopted Parlex, which predicts how members of parliament are likely to respond to proposed policies by analysing their past contributions, giving policy teams early estimates of potential support or opposition.

escalate, pivot, or consolidate partial wins, transforming diffuse desire into precisely timed, coordinated action.

2. Verifiable Contribution Ledgers for Conditional Cooperation

Collective action falters when effort is invisible or easily free-ridden. Verifiable ledgers (paired with AI-based entity resolution to link contributions across data sources) could create a real-time map of civic effort, tracking actions such as door knocks, donations, petition signatures, or expert briefings. Contributors might receive micro-credits or public recognition, while dynamic thresholds (e.g., "200 more calls by midnight unlock a field organizer") transform abstract goals into concrete milestones. Contributors could receive tailored prompts, data, or micro-trainings matched to emerging needs and their demonstrated skills. This visibility enables conditional cooperation: citizens act because they can see others doing the same, anchoring participation in reciprocal trust.

3. Adaptive Coalition Architectures

AI systems trained on policy priorities, voting behavior, and media narratives could detect latent alignments among actors who rarely collaborate. They might reveal, for example, that food-justice advocates and anti-waste environmentalists share overlapping goals on agricultural subsidy reform, suggesting a temporary joint campaign. These systems could recommend shared frames, coordinate tactics, and even propose exit conditions, allowing alliances to form rapidly around moments of opportunity and dissolve once strategic alignment fades. By expanding coalition politics from ideological affinity to pragmatic synergy, AI could enable movements to build power across traditional boundaries, forming alliances that would never emerge through conventional organizing but that multiply pressure on specific institutional decisions.

These three capabilities illustrate a shift from digital tools that enable participation to algorithmic infrastructures that amplify its consequences. The mechanism is sequential: deliberation produces consensus, coordination infrastructure enables mobilization, and mobilization raises the costs of institutional non-response. Unlike isolated platforms, these infrastructures integrate coordination, feedback, and contribution tracking into persistent architecture for influence. For example, a predictive system that ties outreach activity to real-time legislative outcomes does not merely support participation; it makes visible the relationship between collective effort and institutional movement, transforming how citizens understand and exercise power.

Yet this vision confronts questions of feasibility and power. First, these capabilities do not currently exist at scale, and their feasibility remains unproven. Second, even if buildable,

each could be weaponized: the same systems that coordinate democratic publics could enable micro-targeted manipulation or real-time counter-mobilization by adversaries with greater resources. Third, the asymmetry in technical capacity means that those threatened by democratic coordination will likely deploy comparable tools first and more effectively. The question is not whether AI can theoretically serve democratic ends, but whether those who control AI development have any incentive to build systems that threaten their own power.

The governance question is not hypothetical. The civil rights movement's use of Wide Area Telephone Service (WATS) lines illustrates how infrastructure becomes either empowering or constraining depending on who controls it.⁵ Organizations like the Student Nonviolent Coordinating Committee paid for shared lines that enabled activists across the South of the US to report violence, coordinate responses, and bypass hostile white switchboard operators who routinely blocked or monitored their calls (Cammaerts 2018). Organizers knew the lines were likely surveilled; later investigations confirmed FBI wiretapping. Yet the lines remained invaluable because movements governed their use: deciding who could access them, how information was shared, and which incidents merited escalation. When adversaries gained access, as with the FBI's surveillance of Martin Luther King Jr., the same infrastructure became an instrument of control (Garrow 1986). The outcome was determined not by the technology but by who governed it. AI coordination infrastructure faces the same test: whether it amplifies democratic action or enables sophisticated control will be decided by who builds it, who owns it, and whose interests it serves.

Technology amplifies political capacity, but it does not create it. Nor can it substitute for the learning that comes from building it. Democratic capacity develops through organizing practice: through the friction of coalition-building, the judgment required to identify pressure points, the resilience built by navigating failure. The Civil Rights Movement strengthened American democracy not merely through policy victories but because millions of citizens learned to organize, to sustain pressure through setbacks, and to make maintaining the status quo costlier than changing it. These are not skills that algorithms can teach, and they were learned by people who lacked formal accountability mechanisms like voting rights.

⁵ The precursor to modern toll-free numbers, WATS lines allowed organizations to make unlimited long-distance calls within a defined region for a fixed monthly rate. Throughout much of the rural South in the early 1960s, direct-dial service was unavailable, requiring callers to route connections through telephone operators - positions Southern telecommunications companies reserved exclusively for white workers, many aligned with law enforcement or white supremacist organizations who routinely monitored or refused civil rights calls. WATS lines circumvented this barrier by enabling direct calls to movement offices without operator assistance, with charges billed to the organization (Veterans of the Civil Rights Movement, n.d.).

This creates a design constraint for coordination infrastructure. If AI systems simply tell citizens when and where to act (optimizing targets, timing, and tactics without requiring strategic judgment) they risk producing mobilization without movement-building. Efficient bursts of coordinated action that fade when the algorithmic prompts stop, leaving no organizational residue, no enhanced capacity, no relationships of trust that persist beyond individual campaigns. Democracy's calculator dependency: effective at generating outputs while eroding the capability it claims to serve.

The distinction between enhancing and replacing is therefore foundational. Coordination tools must lower transaction costs without eliminating the learning costs that build democratic capacity.⁶ This means systems that surface opportunities but leave tactical decisions to organizers. Platforms that track contributions but require human interpretation of what patterns mean. Infrastructure that makes the work of raising institutional costs more effective without making it effortless. The goal is not frictionless coordination but strategic leverage that preserves the educative function of collective action. Without this principle, AI coordination risks accelerating the very problem it aims to solve: a democracy where citizens have voice but lack the organized capacity to make ignoring that voice politically expensive.

Despite these obstacles, it is worth examining what such systems would enable if they could be built and democratically governed. In this emerging vision, mass deliberative publics are not just larger: they are more purposeful and agentic. That is, they possess a capacity to initiate and coordinate action. AI-powered orchestration allows citizens, civil society actors, and autonomous agents to operate in tandem – deliberating, mobilizing, fundraising, monitoring, adapting. Each agent, human or machine, becomes part of a larger repertoire of democratic action. For example, an AI agent might continuously surface windows of legislative vulnerability while volunteers mobilize outreach during those windows. This coordination creates a hybrid form of leverage that balances human strengths (social trust, physical presence) with machine capabilities (pattern recognition, timing).

This choreography will depend on more than mobilization. It will depend on insight. If deliberative outputs are to influence decisions, we must understand what drives responsiveness. AI systems can help uncover the behavioral patterns and threshold dynamics that make institutions shift: when media exposure matters, which districts tip votes, what frames resonate, and under what conditions they break through. Over time, this would produce a predictive model of responsiveness – a continuously updated civic

⁶ An example that offers a glimpse of this is Peter Levine's *Civic Helpdesk*, an AI tool that generates agendas, recruitment messages, and other basic organizational scaffolds to lower coordination barriers for grassroots groups. See Peter Levine, "Civic Helpdesk" (2025), available online.

intelligence that replaces guesswork with evidence about what makes decision-makers move.

Yet the data and analytic capacity needed to identify where and when institutions are most responsive remain fragmented across jurisdictions, while the capability to build such systems currently lies with actors whose incentives are commercial rather than civic. Bridging that gap will demand investment in public-interest research and democratically governed prototypes, ensuring that the tools designed to coordinate influence do not simply reproduce existing asymmetries of power.

The challenge for deliberative democracy is no longer just how to scale mutual understanding, but how to translate it into consequence. The speculative tools outlined here address that gap: not by making deliberation bigger, but by connecting it to the coordination infrastructure that makes institutional non-response costly. They provide citizens with strategic intelligence about where pressure matters, real-time feedback about whether it's working, and mechanisms to sustain collective action across the time horizons that institutional change requires. By lowering coordination costs while preserving the learning function of organizing, AI could shift the balance between concentrated interests and coordinated publics. Whether it does so depends on who builds these systems, who governs them, and whether democratic movements can claim them before those threatened by democratic coordination do.

Democratic Safeguards: The Visible Algorithmic Hand

The coordination infrastructure described above carries evident risks. Technologies designed to empower can become instruments of control through surveillance, manipulation, or capture by those they were meant to constrain. The question is not whether to build such systems but how to govern them so they preserve citizen agency and resist co-optation.

Governing the algorithmic hand requires institutional safeguards embedded from the outset. Democratic coordination tools must not be governed by the same centralized interests they are meant to challenge. Unbridled corporate ownership invites commercialization; state control risks authoritarian drift; philanthropic stewardship concentrates agenda-setting power in elite hands. Platform cooperatives, multi-stakeholder oversight boards, or public utility governance could ensure distributed control, but only if democratic movements build and claim these systems before others do.

Transparency is necessary but insufficient. Open-source models, traceable data inputs, and funding source disclosures allow scrutiny of how coordination engines identify opportunities or prioritize action. Yet as Mendonça, Filgueiras, and Almeida (2023) argue, algorithms now function as institutional arrangements that shape decision-making, yet

often without the democratic pillars of authorization and accountability. Treating AI systems as civic infrastructure requires institutional democratization, since governance problems cannot be solved through disclosure and audit alone.

Recent work demonstrates how collective governance produces different outcomes than technocratic evaluation. Even in domains such as healthcare, where questions appear purely technical, Hamna et al. (2025) show that when Indian civil-society organizations and community members controlled the evaluation criteria for chatbots, their assessments diverged markedly from automated metrics. If technical experts overlook what communities value in health information, a relatively bounded domain, the gap is likely far wider in politically contested arenas such as mobilization and responsiveness. Recent experiments also show that participatory alignment can be scaled. Huang et al. (2024) demonstrate that Anthropic's language model, when fine-tuned on principles derived from collective public deliberation rather than developer-defined norms, can reduce bias and maintain performance, illustrating the feasibility of institutional democratization in AI design. The question, then, is not whether coordination tools can be built, but who builds them and for whom.

Privacy is another cornerstone. For citizens to participate in organizing efforts, especially those critical of powerful institutions, they must be protected from surveillance. This is especially urgent for marginalized groups, who often face disproportionate risks. Privacy-preserving design should be the default: opt-in participation, anonymization, and data minimization must be built into the architecture. At the same time, systems like contribution ledgers must balance transparency with protection. Emerging techniques, such as zero-knowledge proofs or differential privacy, can offer ways to show collective momentum without exposing individuals, enabling conditional cooperation without surveillance.

Finally, no democratic system is invulnerable to failure or capture. AI-enabled coordination tools are likely to be attractive targets for elite manipulation through subtle algorithmic bias, false data inputs, or strategic funding. Resilience requires error detection, rapid response protocols, and democratic override mechanisms. When failures occur, whether technical or political, communities must have the power to intervene, adapt, or abandon compromised platforms. Resilience here is less about perfect security than about democratic adaptability. Embedding accountability, transparency, and participatory governance at every layer would increase the odds that the algorithmic hand becomes a tool of strategic empowerment rather than a machine for elite control.

But even the best technical safeguards cannot by themselves make deliberation matter politically. Governance design can protect integrity, but it cannot create commitment or sustain participation without resources and power. The deeper constraint lies in the

political economy that determines which forms of public voice are funded, amplified, and absorbed into decision-making, and which are excluded from the institutional circuits through which power flows.

The Funding Problem and the Power Agenda: Why Deliberation Hasn't Scaled

This vision of AI-enabled deliberation as infrastructure for collective action raises two questions. First, if the benefits of scaling deliberation are as compelling as claimed, why hasn't it happened already? And second, assuming those barriers can be overcome, what should scaled deliberation actually focus on to maximize democratic impact?

The answer to both questions converges on the same structural problem: deliberative democracy hasn't scaled not because the technology is underdeveloped or the costs are prohibitive, but because the political incentives to fund it meaningfully are weak, especially when it comes to issues that challenge existing concentrations of power. The paradox is this: those most capable of funding deliberative processes at scale are also those least interested in the outcomes such processes might generate, particularly when those outcomes threaten their influence.

Consider the kinds of actors who have funded most existing deliberative initiatives: philanthropic foundations, public institutions, and occasionally civic-minded local governments. These efforts are often modest in scale not because deliberation is prohibitively expensive, but because scaling it up would require investment from actors with a vested interest in limiting broad-based democratic voice.

Corporations spend billions annually on lobbying, marketing, and public relations. Yet it is hard to imagine them investing in citizen assemblies about pharmaceutical pricing, banking regulation, or data privacy, issues on which broad citizen input might produce recommendations in direct conflict with corporate interests. And indeed, they don't. This reflects a pattern: there is often an inverse relationship between how readily institutions embrace citizen deliberation and how much that deliberation threatens redistributing power.

This gives rise to a strategic dilemma for deliberative democracy advocates. It is not just that powerful interests avoid funding deliberation on threatening topics; they may actively support deliberation on unthreatening ones. A corporate-sponsored citizen assembly on "innovation policy" or a foundation-funded process on "community resilience" may offer useful insights or contribute to procedural learning. But they can also serve to channel democratic energy away from more contentious, consequential questions, domesticating

the "deliberation in the wild". They risk giving the appearance of inclusion while reinforcing the arrangements that render public voice ineffective.

More bluntly: better to have fewer deliberative processes that confront structural inequalities than many that leave power untouched. Deliberative democracy faces a persistent risk of co-optation: many of its intermediaries depend for their legitimacy and funding on the very institutions they seek to influence. That dependence breeds caution, potentially steering deliberation toward safe, consensual issues and away from the political fault lines where responsiveness is most at stake: let citizens deliberate about bike lanes while decisions about public transportation concessions are made in boardrooms and lobbying offices.

The test for whether a deliberative initiative is worth pursuing should be clear: does it engage the institutional levers that determine responsiveness, or is it a distraction from them?

What Scaled Deliberation Should Target: Constitutional and Meta-Political Questions

To matter, scaled deliberation cannot stop at mirroring public preferences once deliberation has taken place. It must reshape the institutions that determine whether those preferences influence decisions. This means targeting "constitutional" or "meta-political" questions: who decides, how decisions are made, and whose interests are structurally advantaged or excluded.

Money in politics sits at the top of this agenda. Deliberative democracy cannot afford indifference to how economic power distorts political responsiveness. Deliberation on lobbying, political advertising, and corporate donations isn't peripheral: it's foundational. Without it, even the best-designed processes risk becoming performative.

The same logic applies to other pillars of democratic design: voting systems, judicial selection, regulatory capture, and media ownership. These are not just technical issues: they are the plumbing of democracy, shaping participation, voice, agenda-setting, and decision-making. Citizen deliberation holds particular value here: there's often no expert consensus, lived experience carries weight, and public preferences are more open to compromise than elite narratives suggest.

Deliberation on these institutional questions directly strengthens democratic responsiveness. Reforms that curb undue influence or boost electoral competitiveness help ensure citizen voice resonates across policy domains. In systems with referenda or recall, large-scale deliberation can amplify these tools; raising issue visibility, coordinating action, and increasing the political cost of inaction. When elected officials ignore widely

supported deliberation outcomes, they risk political backlash. Deliberation, then, starts to become part of an accountability infrastructure.

None of this will be easy. Deliberation aimed at structural reform will encounter stronger resistance than deliberation focused on less contentious policy areas. Entrenched interests have little incentive to legitimize or implement recommendations that threaten their dominance. This is the "bootstrap problem" of deliberative democracy: to redistribute power, one often needs power.

This reveals a deeper limitation of the argument presented here. If organized power to force responsiveness is the bottleneck, then imagining AI coordination tools may displace the harder work of actual organizing. The honest assessment: deliberative democrats, including this author, find it easier to design technical systems than to build political movements. That doesn't make the technical work worthless, but it clarifies where the real constraints lie.

Confronting this bootstrapping problem makes strategic timing and coalition-building essential. Moments of institutional crisis, such as corruption scandals, electoral deadlocks, regulatory failures, may open political space for deliberative interventions. But capitalizing on such windows requires infrastructure: the capacity to rapidly convene credible deliberative processes, the networks to mobilize public attention, and the alliances to pressure for implementation.

This, in turn, calls for a shift in the deliberative democracy field, from a focus on designing optimal procedures to a focus on political strategy. Advocates will need to behave more like organizers: building long-term coalitions, raising sustained funding, and developing the tactical awareness to align deliberative opportunities with moments of political vulnerability. The skill set required is not simply deliberation design, but strategic political engagement.

This shift also means letting go of the field's habitual suspicion toward aggregative methods, an attitude that sometimes verges on deliberative fundamentalism. Large-scale deliberation is only as powerful as its ability to yield binding outcomes. That means combining it with mechanisms like ballot initiatives, referendums, recall, and even elections, tools that translate public will into political effect. Aggregation is not the enemy of deliberation; it is the means through which deliberation stands the greatest chance of enforceability.

These binding mechanisms also address a practical constraint: funding. Rather than depend on philanthropic or elite backing, scaled deliberation could leverage democratic mechanisms like citizen initiatives, ballot measures, or constitutional conventions to institutionalize deliberative processes' public financing. If broad consensus on issues like

campaign finance reform becomes binding law, the result could both limit concentrated wealth and secure long-term funding for deliberation. The goal, in short, is not just to enable deliberation, but to embed it in democratic governance.

This strategic reframing carries a warning: not all deliberative opportunities are worth pursuing (Wilson & Mellier 2023). Deliberation that avoids confronting power structures, regardless of how well-designed, reinforces the problem it claims to address. In some cases, deliberation on safe topics can even preempt pressure for more consequential reform.

Ultimately, the challenge of scaling deliberation isn't technological or procedural: it's political. If deliberative democracy is to be more than a democratic sideshow, it needs strategic clarity and the alliances to confront the forces that limit responsiveness. The question is not whether deliberation can scale, but whether its advocates are prepared to do what consequential scaling requires. No algorithm can substitute for that work.

Concluding Remarks

Each wave of technological innovation, from the optical telegraph to the Internet, has stirred hopes of democratic renewal, only to falter against entrenched power structures. Today's enthusiasm for AI-enabled deliberation risks repeating that cycle unless it confronts the institutional barriers that have historically blocked democratic reform.

Deliberative democracy has not failed to scale because of technical or financial constraints, but because those with the power to support it have had little incentive to do so, particularly when its outcomes might challenge their interests. This misalignment between democratic aspiration and institutional support cannot be corrected through design tweaks or smarter algorithms alone.

If deliberation is to achieve impact at scale, it must target the institutional arrangements that shape responsiveness, such as lobbying regimes, electoral systems, media ownership, and judicial appointments. These meta-political arenas determine how public voice is translated into policy or systematically ignored.

This reframing also redefines the role of AI. Rather than treating it as an instrument for scaling conversation, we must examine how AI can enable strategic coordination, amplify citizen agency, and sustain organizing that challenges institutional inertia. The speculative applications outlined in this essay, from coordination engines to predictive feedback systems, matter not because they perfect deliberative procedure, but because they might help convert consensus into political leverage.

No system, however sophisticated, can replace the political labor required for democratic change. Algorithms do not build coalitions, identify strategic openings, or sustain engagement through adversity. These are human tasks. If deliberative democracy is to move beyond symbolic participation, its advocates must be willing to take them on.

Above all, deliberative democracy must shift from process optimization to power building. It must align with other democratic movements, not only those seeking inclusion in policymaking but also those contesting the rules that structure it. As long as deliberation remains confined to domains that leave power arrangements untouched, it risks becoming democratic theater: symbolic engagement without consequence.

The promise of deliberative democracy lies not in transcending conflict but in redistributing the power to shape it. AI may help, but only if it serves a strategy of structural change. The question is not whether technology can perfect deliberation, but whether it can help citizens make their deliberation matter.

The honest answer is that we do not yet know. The information highway, too, is littered with good democratic intentions. Yet refusing to test this coordination hypothesis is equivalent to accepting unresponsiveness as inevitable. If publics are often aligned but unable to convert agreement into leverage, tools that lower coordination costs are worth building and testing, provided their governance is democratic. The goal is not better algorithms. The goal is redistributed power.

List of References

- Anthony, C. S., & Elpers, N. (2023). "Give Light and The People Will Find a Way": Democratic Deliberation and Public Achievement at Colorado College. In *Creating Space for Democracy* (pp. 249-256). Routledge.
- Barandiaran, X. E., Calleja-López, A., Monterde, A., & Romero, C. (2024). *Decidim, a technopolitical network for participatory democracy: philosophy, practice and autonomy of a collective platform in the age of digital intelligence* (p. 133). Springer Nature.
- Cammaerts, B. (2019). A genealogy of communicative affordances and activist self-mediation practices. In *Citizen Media and Practice* (pp. 98-112). Routledge.
- Carnes, N., & Lupu, N. (2015). Rethinking the comparative perspective on class and representation: Evidence from Latin America. *American Journal of Political Science*, 59(1), 1-18.
- Chambers S (2018) The philosophic origins of deliberative ideals. In: Bächtiger A, Dryzek J, Mansbridge J, et al. (eds) *The Oxford Handbook of Deliberative Democracy*. Oxford Handbooks. Oxford: Oxford University Press, pp. 55–69.
- Cohen, J., & Fung, A. (2023). Democratic responsibility in the digital public sphere. *Constellations: An International Journal of Critical & Democratic Theory*, 30(1).
- Coleman, S., & Gøtze, J. (2001). *Bowling together: Online public engagement in policy deliberation* (pp. 39-50). London: Hansard Society.
- Curato, N. (2015). Deliberative capacity as an indicator of democratic quality: The case of the Philippines. *International Political Science Review*, 36(1), 99-116.
- Curato, N., Farrell, D. & Geissel, B. (2021) *Deliberative Mini-Publics*, Bristol: Policy Press.
- Hong, L., & Page, S. E. (2012). Some microfoundations of collective wisdom. *Collective wisdom*, 56-71.
- Fishkin, J. S. (1997). *The voice of the people: Public opinion and democracy*. Yale university press.
- Fishkin, J. (2009). *When the people speak: Deliberative democracy and public consultation*. Oup Oxford.

Fishkin, J., Siu, A., Diamond, L., & Bradburn, N. (2021). Is deliberation an antidote to extreme partisan polarization? Reflections on “America in one room”. *American Political Science Review*, 115(4), 1464-1481.

Fishkin, J., Bolotnyy, V., Lerner, J., Siu, A., & Bradburn, N. (2025). Scaling Dialogue for Democracy: Can Automated Deliberation Create More Deliberative Voters?. *Perspectives on Politics*, 1-18.

Font, J., Smith, G., Galais, C., & Alarcon, P. (2018). Cherry-picking participation: Explaining the fate of proposals from participatory processes. *European Journal of Political Research* 57 (3), 615 – 636.

Fox, J., Halloran, B., Fölscher, A., & McGee, R. (2024). Disentangling Government Responses: How Do We Know When Accountability Work Is Gaining Traction?.

Garrow, D. J. (1986). *Bearing the cross: Martin Luther King, Jr., and the southern Christian leadership conference*. Open Road Media.

Geissel, B. (2023). 5 Linking citizens’ assemblies to policymaking: Real-life and visionary connections. *De Gruyter Handbook of Citizens’ Assemblies*, 1, 59.

Giraudet, L. G., Apouey, B., Arab, H., Baeckelandt, S., Bégout, P., Berghmans, N., ... & Tournus, S. (2022). “Co-construction” in deliberative democracy: lessons from the French Citizens’ Convention for Climate. *Humanities and Social Sciences Communications*, 9(1), 1-16.

Gilens, M. (2005). Inequality and democratic responsiveness. *Public Opinion Quarterly*, 69(5), 778-796.

Gilens, M., & Page, B. I. (2014). Testing theories of American politics: Elites, interest groups, and average citizens. *Perspectives on politics*, 12(3), 564-581.

Goldberg, B., Acosta-Navas, D., Bakker, M., Beacock, I., Botvinick, M., Buch, P., ... & Zhang, A. (2024). AI and the Future of Digital Public Squares. *arXiv preprint arXiv:2412.09988*.

Gutmann, A., & Thompson, D. F. (2004). *Why deliberative democracy?*. Princeton University Press.

Habermas, J. (1996). *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy* (Trans. William Rehg). Cambridge, MA: The MIT Press.

Hamna, Bhat, G., Mukherjee, S., Lalani, F., Hadfield, E., Siddarth, D., Bali, K., & Sitaram, S. (2025). Building benchmarks from the ground up: Community-centered evaluation of LLMs in healthcare chatbot settings. arXiv preprint arXiv:2509.24506.

Huang, S., Siddarth, D., Lovitt, L., Liao, T. I., Durmus, E., Tamkin, A., & Ganguli, D. (2024, June). Collective constitutional ai: Aligning a language model with public input. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1395-1417).

Jacquet V, van der Does R (2020) The consequences of deliberative minipublics: systematic overview, conceptual gaps, and new directions. *Representation* 0:1–11.

Klein, M. (2021). Crowd-scale deliberation for group decision-making. *Handbook of group decision and negotiation*, 355-369.

Klein, M., Spada, P., & Paulson, L. (2023). From shouting matches to argument maps: An online deliberation experiment in Italy. In *The Routledge handbook of collective intelligence for democracy and governance* (pp. 190-201). Routledge.

Konya, A., Thorburn, L., Almasri, W., Leshem, O. A., Procaccia, A., Schirch, L., & Bakker, M. (2025, June). Using collective dialogues and AI to find common ground between Israeli and Palestinian peacebuilders. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency* (pp. 312-333).

Lafont, C. (2023). Democracy Without Shortcuts: An Institutional Approach to Democratic Legitimacy. In *Another Universalism: Seyla Benhabib and the Future of Critical Theory* (pp. 151-165). Columbia University Press.

Landemore, H. (2012). Democratic reason: Politics, collective intelligence, and the rule of the many.

Landemore, H. (2022, October). Can AI bring deliberative democracy to the masses. In *Human Centered Artificial Intelligence Seminar, Stanford University*.

Mansbridge, J. (2007). “Deliberative democracy” or “democratic deliberation”? In *Deliberation, participation and democracy: Can the people govern?* (pp. 251-271). London: Palgrave Macmillan UK.

Mansbridge, J., Bohman, J., Chambers, S., Christiano, T., Fung, A., Parkinson, J., ... & Warren, M. E. (2012). A systemic approach to deliberative democracy. *Deliberative systems: Deliberative democracy at the large scale*, 1-26.

Mattelart, A. (1999). La communication et la promesse de rédemption. *Quaderni*, 40(1), 69-78.

McKinney, S. (2024). Integrating artificial intelligence into citizens' assemblies: Benefits, concerns and future pathways. *Journal of Deliberative Democracy*, 20(1).

Mellon, J., Sjöberg, F. M., Peixoto, T., & Lueders, J. (2025). The haves and the have nots: Civic technologies and the pathways to government responsiveness. *Government Information Quarterly*, 42(1), 102007.

Mendonça, R. F., Almeida, V., & Filgueiras, F. (2024). *Algorithmic Institutionalism: the changing rules of social and political life*. Oxford University Press.

Mouffe, C. (1999). Deliberative democracy or agonistic pluralism?. *Social research*, 745-758.

Nabatchi, T. (2010). Deliberative democracy and citizenship: In search of the efficacy effect. *Journal of Deliberative Democracy*, 6(2).

Neves, B., & Peixoto, T. (2020). O sorteio na democracia: uma alternativa viável. *Insight Inteligência. Edição*, 89.

Noveck, B. S. (2019). Why Dewey Was Wrong. *Digital media and democratic futures*, 101-120.

Olson Jr, M. (1971). *The Logic of Collective Action: Public Goods and the Theory of Groups, with a new preface and appendix* (Vol. 124). Harvard University Press.

Ostrom, E. (2010). Beyond markets and states: polycentric governance of complex economic systems. *American economic review*, 100(3), 641-672.

Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge university press.

Parkinson, J. (2003). Legitimacy problems in deliberative democracy. *Political studies*, 51(1), 180-196.

Peixoto, T. (2012). The benefits of citizen engagement: A (brief) review of the evidence. Available at <http://democracyspot.net/2012/11/24/the-benefits-of-citizen-engagement-a-brief-review-of-the-evidence/>

Peixoto, T. (2012). The Uncertain Relationship between Open Data and Accountability: A Response to Yu and Robinson's the New Ambiguity of Open Government. *UCLA I. Rev. Discourse*, 60, 200.

Peixoto, T., & Fox, J. (2016). When does ICT-enabled citizen voice lead to government responsiveness? *IDS Bulletin*, 47(1), 23–40.

Peixoto, T., & Sifry, M. L. (2017). Civic tech in the global south. *World Bank Publications-Books*.

Peixoto, T. C., & Steinberg, T. (2019). Citizen engagement: Emerging digital technologies create new risks and value. The World Bank.

Pogrebinschi, T., & Samuels, D. (2014). The impact of participatory democracy: evidence from Brazil's National Public Policy Conferences. *Comparative Politics*, 46(3), 313-332.

Pentland, A., & Tsai, L. (2024). Toward building deliberative digital media: from subversion to consensus. *PNAS nexus*, 3(10), pgae407.

Peters, Y. and Ensink, S.J., (2015). Differential responsiveness in Europe: The effects of preference difference and electoral participation. *West European Politics*, 38(3), pp.577-600.

Raelin, J. A. (2012). Dialogue and deliberation as expressions of democratic leadership in participatory organizational change. *Journal of Organizational Change Management*, 25(1), 7-23.

Revel, M., & Pénigaud, T. (2025). AI-Facilitated Collective Judgements. *arXiv preprint arXiv:2503.05830*.

Sanyal, P., & Rao, V. (2019). *Oral democracy: Deliberation in Indian village assemblies*. Cambridge University Press.

Spada, P., & Peixoto, T. C. (2025). The limits of representativeness in citizens assemblies: A critical analysis of democratic minipublics. *Journal of Sortition*, 1(1), 137-159.

Schelling, T. C. (2006). *Micromotives and macrobehavior*. WW Norton & Company.

Schneider, S. M. (1996). Creating a democratic public sphere through political discussion: A case study of abortion conversation on the Internet. *Social Science Computer Review*, 14(4), 373-393.

Sjoberg, F. M., Mellon, J., Peixoto, T. C., Hemker, J. Z., & Tsai, L. L. (2019). Voice and punishment: A global survey experiment on tax morale. *World Bank policy research working paper*, (8855).

Small, C., Bjorkegren, M., Erkkilä, T., Shaw, L., & Megill, C. (2021). Polis: Scaling deliberation by mapping high dimensional opinion spaces. *Recerca: revista de pensament i anàlisi*, 26(2).

Summerfield, C., Argyle, L., Bakker, M., Collins, T., Durmus, E., Eloundou, T., ... & Botvinick, M. (2024). How will advanced AI systems impact democracy?. *arXiv preprint arXiv:2409.06729*.

Tessler, M. H., Bakker, M. A., Jarrett, D., Sheahan, H., Chadwick, M. J., Koster, R., ... & Summerfield, C. (2024). AI can help humans find common ground in democratic deliberation. *Science*, 386(6719), eadq2852.

Tsai, L. L., Pentland, A., Braley, A., Chen, N., Enríquez, J. R., & Reuel, A. (2024). Generative AI for Pro-Democracy Platforms. MIT GovLab.

Tutashinda, K. (2010). The Grassroots Political Philosophy of Ella Baker: Oakland, California Applicability. *Journal of Pan African Studies*, 3(9).

Sjoberg, F. M., Mellon, J., Peixoto, T. C., Hemker, J. Z., & Tsai, L. L. (2019). Voice and punishment: A global survey experiment on tax morale. *World Bank policy research working paper*, (8855).

Smith, G. (2009). *Democratic innovations: Designing institutions for citizen participation*. Cambridge University Press.

Soontjens, K. and Persson, M., (2024). Lacking Incentives, Not Information. Why Politicians tend to be Less Responsive to Lower-Income Citizens. *Legislative Studies Quarterly*.

Sintomer, Y. (2023). *The government of chance: Sortition and democracy from Athens to the present*. Cambridge University Press.

Suiter, J., Farrell, D. M., & O'Malley, E. (2016). When do deliberative citizens change their opinions? Evidence from the Irish Citizens' Assembly. *International Political Science Review*, 37(2), 198-212.

Touchton, M., Wampler, B., & Peixoto, T. (2021). Of democratic governance and revenue: Participatory institutions and tax generation in Brazil. *Governance*, 34(4), 1193-1212.

Vandermonde, A. (1994). Quatrième leçon d'économie politique, 23 ventôse/13 mars. *L'École normale de l'an III. Leçons d'histoire, de géographie, d'économie politique*, 385.

Veterans of the Civil Rights Movement. (n.d.). WATS line & phone reports. Civil Rights Movement Archive. <https://www.crmvet.org/docs/wats/watshome.htm>

White, C. S. (1997). Citizen participation and the Internet: Prospects for civic deliberation in the information age. *The Social Studies*, 88(1), 23-28.

Wilson, R. & Mellier, C. (2023) Getting Real About Citizens' Assemblies: A New Theory of Change for Citizens' Assemblies, [Online], <https://europeandemocracyhub.epd.eu/gettingreal-about-citizens-assemblies-a-new-theory-of-change-for-citizens-assemblies/>.

Wojciechowska, M. (2019). Towards intersectional democratic innovations. *Political Studies*, 67(4), 895-911.

Yu, H., & Robinson, D. G. (2011). The new ambiguity of open government. *UCLA L. Rev. Discourse*, 59, 178.