"Can AI bring deliberative democracy to the masses?"

Hélène Landemore, Yale University
helene.landemore@yale.edu

Abstract: A core problem in deliberative democracy is the tension between two seemingly equally important conditions of democratic legitimacy: deliberation on the one hand and mass participation on the other. Might artificial intelligence help bring quality deliberation to the masses? The paper first examines the conundrum in deliberative democracy around the tradeoff between deliberation and mass participation by returning to the seminal debate between Joshua Cohen and Jürgen Habermas about the proper model of deliberative democracy. It then turns to an analysis of the 2019 French Great National Debate, a low-tech attempt to involve millions of French citizens in a structured exercise of collective deliberation over a two-month period. Building on the shortcomings of this empirical attempt, the paper then considers two different visions for an algorithm-powered scaled-up form of mass deliberation—Mass Online Deliberation on the one hand and a multiplicity of rotating randomly selected mini-publics on the other—theorizing various ways Artificial Intelligence could play a role in either of them.

A core problem in deliberative democracy is the tension between two seemingly equally important tenets of democratic legitimacy: deliberation on the one hand and mass participation on the other. Both ideas are indeed contained in the deliberative democracy conception of legitimacy, which says that laws and policies have legitimacy—that is, the moral authority to command obedience—only to the extent that they have been produced through a public and inclusive deliberation among free and equal individuals. Deliberation is defined as an exchange of reasons, arguments, and justifications. Democratic deliberation, specifically, is supposed to be inclusive of all members of the demos on equal grounds.

The problem, however, is that deliberation of the kind deliberative democrats have in mind only works well in small groups. Past a threshold, which is probably between a few dozen and a few hundred people, deliberation becomes impossible and needs to be delegated to a subset of citizens, usually elected representatives. As a result, most people end up excluded from the deliberations that shape laws and policies, creating a deficit of democratic legitimacy that needs to (but might not be able fully to) be compensated in other ways (consent of the governed cannot do all the work).

Do we need then to *choose* between deliberation and mass participation, as some have suggested (e.g., James Fishkin 2009), or must we find a way for deliberation to involve the entire population as Cristina Lafont (2018) has counter-argued, because as she puts it there can be "no democratization without mass participation"? If we choose deliberation over participation, then we risk falling into elitist and undemocratic territory; if we insist on keeping mass participation, it will be at the expense of a robust concept and practice of deliberation taking place between individuals.

Herein lies the possible relevance of technological change, including the rise of artificial intelligence. Might digital technology, which has connected millions of people across the globe over the last 20 years or less, provide a possible way to reconcile these two values? Could Artificial Intelligence (AI) in particular help us scale deliberation to the masses?

In this paper I agree that digital technologies and artificial intelligence in particular can certainly help connect more people's minds and thus expand the possibility of quality deliberation to a larger number of people. I also argue that we should let go of the ideal of all minds engaged in one common deliberation and instead settle for an approximation that can take at least two different forms: many minds working on different parts of a common question, with the mediation of algorithms handling the complexity of the data or many randomly selected deliberative assemblies exposing their members to the full range of views on a given question over time. Both versions of mass deliberation present advantages and drawbacks, and both can be considerably augmented by the use of AI.

The first section examines the conundrum in deliberative democracy around the tradeoff between deliberation and mass participation by returning to the seminal debate between Joshua Cohen and Jürgen Habermas about the proper model of deliberative democracy. The second section examines an attempt to scale face-to-face deliberation, the

French Great National Debate, and theorizes the nature of the participatory and deliberative "augmentation" to representative democracy that the process represents. I argue that the Habermasian two-track model of the public sphere got expanded in the French case into a three-track model. The third section speculates about the ways AI could have been used in such an experiment, and could be used in future processes, to help expand the number of people involved in the deliberations. Finally, I ask whether AI could in fact be used to simulate or predict the results of a fully inclusive mass deliberation and considers anew the benefits and risks of proxy mass deliberation. While the original conundrum returns, it does so in an attenuated form and with such potential instrumental (epistemic) benefits that they may outweigh, up to a point, the concerns around legitimacy that critics will inevitably raise.

In what follows, I take AI to mean the ability of a computer system to perform complex tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, translation between languages, and solving games.

1. *Deliberative democracy as a theory of political legitimacy*

Deliberative democracy is a theory of political legitimacy according to which fundamental laws are legitimate only to the extent that they are the product of a public, inclusive deliberation among equals. It was developed in part as an alternative to and improvement over purely aggregative theories of democracy, whereby mere preference aggregation and majoritarian outcomes are what defines legitimate law and policy.

According to a description by one of its first theoreticians, "[t]he notion of a deliberative democracy is rooted in the intuitive ideal of a democratic association in which the justification of the terms and conditions of association proceeds through public arguments and reasoning among equal citizens" (Cohen 1989). Deliberation as an exchange of public

arguments and a form of collective reasoning generates acceptable reasons to endorse the "terms and conditions of association"—i.e., laws and policies, for everyone.

Deliberative democracy is a demanding theory of legitimacy but one that respects individuals as autonomous beings endowed with both a capacity and a right to understand the reasons behind the rules that are meant to apply to them. Additionally, according to so-called "epistemic" deliberative democrats, there are instrumental benefits to setting the bar for legitimacy so high. On their view, inclusive deliberation among equals is also valuable for its knowledge-aggregating benefits, in that it can be expected to produce better laws and policies than less deliberative and less inclusive processes. Inclusive deliberation thus generates both process legitimacy and output legitimacy.

The belief in the intrinsic and instrumental benefits of deliberation is the reason why this paper will not entertain the use of AI as mere aggregator of raw opinions. Some authors have argued that an improvement over representative democracy would be to track the rich preferences of citizens by having them write up their political opinions on a piece of paper and feeding this input to an AI. Martin Hilbert (2009), for example, theorized what he called Deep Democratic Neural Network: a system whereby an AI would synthesize citizens' political views and preferences, as expressed through richly descriptive natural language written contributions, into an aggregated policy platform ready to be implemented by politicians, now unburdened of the task of poorly "representing" the people on the thin informational basis of a single vote per person cast every few years. This model of an AI building a consensus out of written inputs by individual citizens is also a model currently explored by researchers at Deepmind.[1]

[1] Personal conversation with Christopher Summerfield, June 20, 2022.

While the ability of AI to paint a detailed, accurate picture of the preferences of a given population of participants at any point in time is certainly a massively valuable tool, improving over cruder tools like polls, the result lacks a crucial deliberative component, which is where the legitimacy of laws and policies is supposed to come from (at least on the model of political legitimacy endorsed in this paper and presented in the first section). For a deliberative democrat, therefore, these kind of high-definition snapshots of the people would not substitute for mass deliberation, though it may offer the starting point of one.

Deliberative democracy arguably takes as its models the philosophy seminar room, the jury, or other small-scale deliberative spaces, only expanded to society as a whole. No mention is made, in the early version by Joshua Cohen, of the need for representatives or any restriction of the public exchange of arguments to a smaller subset of the citizenry.

Jürgen Habermas took issue with some aspects of this earlier model of deliberative democracy, not just as too demanding and practically unfeasible for large societies, but indeed as normatively undesirable. He argues that public deliberation as a legitimizing procedure cannot be something that structures society as a whole:

> In contrast to Cohen, I would like to understand it [i.e., public deliberation] as the core structure in a **separate, constitutionally organized political system, but not as a model for all social institution (and not even for all government institutions)** (1997, chapter 7).

Why is Habermas eager to confine the legitimizing procedure of public deliberation to a "separate, constitutionally organized political system"? It is because otherwise, without this restriction, deliberative politics would have "to be inflated into a structure shaping the totality of society." And this is "*impossible,*" Habermas explains, "for the simple reason that democratic procedure *must* be embedded in contexts it *cannot* itself regulate."

There is some ambiguity in this passage as to whether Habermas thinks of the impossibility as a feasibility constraint, perhaps even a logical impossibility, or a normative injunction. Presumably it is all of the above. Deliberation must take place against the

backdrop of a context that it is itself not deliberated about. And it would be bad to try and regulate that backdrop context anyway. The "must" injunction can read as normative. It is not that the democratic procedure could not regulate the totality of society (though it would be presumably very costly and impractical). It is that such political regulation, even of a democratic kind, would be a bad thing, presumably because it would encroach on the freedom of individuals in the larger society to associate freely and come up with their own agenda, independently from political pressures. This is arguably a worry that animated Hannah Arendt too about the possibility of the social being entirely subsumed by the political (*On Revolution*?).

As a result of his disagreement with Cohen, Habermas puts forward instead a model of the public sphere divided into two tracks, only one of which is regulated by public deliberation and includes formal political institutions like the Parliament, the Courts, and administrative agencies. The other corresponds to the anarchical space of the larger society where a thousand ideas can freely bloom. It is a space for "deliberation in the wild" as Habermas poetically calls it. Habermas had in mind for it the historical example of 18th century French public sphere of coffee shops, newspapers, and political associations but offers this two-track public sphere as a normative reconstruction based more around the ideals underlying our institutions than around practical limitations. In Habermas' vision track 1 is the space of formal decision-making and provides a context of justification for laws and policies whereas track 2, which one could argue comes first chronologically, sets an agenda for track 1 and forms the context of discovery for the various ideas behind the laws and policies then formalized in track 1.

Track One in many ways is the implementation of Cohen's ideal. It is a "strong" public sphere characterized by, among other things, the fact that processes of deliberation take place in argumentative form, that is, through the regulated exchange of information and

reasons among parties who introduce and critically test proposals. This is also a space where deliberations are inclusive, public, and free of any external as well as internal coercion that could detract from the equality of the participants. Each individual has an equal opportunity to be heard, to introduce topics, to make contributions, to suggest and criticize proposals. The taking of yes/no positions is motivated solely by the unforced force of the better argument. In this track, deliberations aim at rationally motivated agreement but must be concluded by majority decisions. Finally, political deliberations extend to any matter that can be regulated in the equal interest of all.

By contrast, the "weak" public sphere of Track Two is characterized by a lack of structure. It is "[e]ffected in an open and inclusive network of overlapping, subcultural publics having fluid temporal, social, and substantive boundaries." Second, whatever structures exist "develop more or less spontaneously." They are not imposed or even constructed. Third, the various currents of public communication generated in this track "are channeled by mass media and flow through different publics that develop informally inside associations." Taken together, these publics "form a 'wild' complex that resists organization as a whole." One way to describe Track Two is as the "cosmos" to Track One's "taxis" (to use famous Hayekian's categories).

The two-tracks are thus distinguished by both their function and their degree of formalization. Track One is where public deliberation operates as a guiding norm oriented towards the production of "will-formation" (the production of decisions) on the basis of the agenda defined by Track Two. Track One is also described by Habermas as a "context of justification." By contrast, Track Two is the realm of ideas, a space where "opinion-formation" takes place. The opinion-formation process results in the generation of an agenda for the formal sphere. Habermas describes this second track at the "context of discovery."

By what mechanisms are these two tracks connected? Here Habermas resorts to the metaphor of the "sluice"—a system of water channels controlled at their heads by a gate—as capturing the relation between the two communicative tracks of the public sphere. Those sluices correspond to the role of intermediaries such as the media, political parties, and the pressure of an informal public opinion formed in civil associations of all kinds. It is through these "sluices" that ideas from the larger public sphere filter up to Track One and it is also through them that ideas from the formal public sphere are disseminated back to Track Two.

In Habermas' usage, the metaphorical gate is meant to both ensure transmission of information from the outer periphery of diffuse public opinion to the center where decision-making takes place, and to properly filter that information.[2] The metaphor is meant to capture the ways in which the two tracks—ordinary citizens on the one hand and their representatives on the other—are connected in constructive ways. But it also emphasizes the ways in which these tracks are meant to be kept separate. It may help to visualize this model as two concentric circles connected by arrows (the arrows symbolizing the reciprocal interactions between the two tracks via the sluices).

There is a lot to admire in this model, which has rightfully been very influential. However, it should also be obvious that it runs into a series of limitations. It is certainly too idealized as a normative sociological reconstruction of the past. But even as a pure normative model it has problems of its own. For one thing, the metaphors on which it relies are problematic. The two tracks suggest a separation between the world of representatives and other officials and the people themselves, with the former at the center of the system and the

[2] See Patberg 2016 for the double meaning of a sluice (or "lock") and corresponding dual functionality (gate and filter) suggested by the metaphor.

people at the periphery. The sluice's connotations are also mechanical, hierarchical, rigid and slow.

Most problematically however, why should we trust that a self-organizing but fundamentally anarchic sphere could generate a political agenda of normative import? Why does deliberation in the wild have any redeeming normative virtues just from the fact that it is spontaneous and "unregulated"? There is a strange, almost neo-liberal faith in Habermas' assumption of a self-organizing civil society that can yield agendas which do not merely reproduce existing power imbalances. It is true that he additionally assumes background economic equality and a perfectly sound media ecosystem, which lies far from the highly inegalitarian, skewed, and polarized media world we actually live in. But even assuming an idealized background does not take care of all potential problems. In fact, Habermas himself acknowledges an inevitable trade-off between freedom of communication and vulnerability to distortion:

> On account of its anarchic structure, the general public sphere is, on the one hand, **more vulnerable to the repressive and exclusionary effects of unequally distributed social power, structural violence, and systemically distorted communication** than are the institutionalized public spheres of parliamentary bodies. On the other hand, it has **the advantage of a medium of unrestricted communication**." (Habermas 1996, 307, my emphasis)

Habermas seems to suggest that on balance there is more to gain from unrestricted communication in the larger public sphere and that this gain is worth the risks of distortion. But he does not give much reassurance as to why we should believe him.

Beyond the asymmetries of power and inevitable distortions, there is also downright exclusion in "wild" deliberations. Some people are silenced and never given a chance to have their voices heard. Additionally, the idea that the decentralized deliberations of the citizens who do deliberate "in the wild" add up to a meaningful way of setting the agenda is not convincing. There are many reasons to think that the larger public sphere is itself shaped by the formal deliberative track in a way that is not fully reciprocal. The collective action

problems faced by "the public" are enormous compared to those faced the by the smaller number of agents at the center of formal decision structures. Moreover, even in the best-case scenario of a functional public sphere, why should we expect a series of haphazard, unregulated, and decentralized deliberations among groups of different sizes and compositions, which are not intentionally oriented towards this outcome, to be the proper way of setting the agenda for the formal deliberative track? Does such "deliberation in the wild" even amount to proper deliberation?

Absent the proper background conditions assumed by Habermas, we already know what happens in practice: a lack of representativeness of whatever deliberation and resulting decision process take place in the formal public sphere. The problem presumably lies both in the difficulty encountered by civil society to produce an agenda that is actually representative of the needs of the larger population, and in the distortions of the agenda as it passes through the "sluices" and is appropriated by the formal institutions. The discrepancy between the public's actual preferences in Track Two and the policies generated by Track One can be enormous. This will eventually result in outbursts of anger and frustration, expressed through the Trump vote in the United States, the Brexit vote in the United Kingdom, or the Yellow Vest movement in France. These outbursts force the system to realign its priorities somewhat, but in a reactive and suboptimal way.

Habermas' attempt to solve the problem of deliberation at scale thus ends up giving up on an actual exchange of reasons among equals in favor of a decentralized and dehumanized network of communication fluxes. So we are back to the original dilemma: either we can have quality deliberation between equal individuals at the small scale of the seminar room (the scale that inspired early conceptions of deliberative democracy) or we have mass participation that is not truly deliberative. Given the recent focus on mini-publics and small-scale deliberation in the empirical literature, Simone Chambers deplores that deliberative

democracy has all but abandoned mass democracy (2009). Meanwhile, Cristina Lafont similarly criticizes the gap on a normative level. She argues, contra Fishkin who accepts the trade-off, that there can be "no democratization without improved mass deliberation" (2015: 45). Both Chambers and Lafont call for a more appealing vision of deliberative democracy that would be less elitist and more accessible to more people. Their respective solutions to this conundrum however remains quite elusive. For Chambers, a certain kind of rhetoric—deliberative rhetoric as opposed to plebiscitary rhetoric--is supposed to ensure that the exchanges that take place in "deliberation in the wild" are truly discursive, even if they don't strictly resemble the more structured deliberations within highly curated small-scale venues. Lafont helpfully offers the concept of deliberative rhetoric as a yardstick by which to assess the actual deliberative quality of the public sphere. But she does not provide guidance as to how to implement this ideal. As to Lafont, who makes the powerful argument that deliberative democrats cannot abandon mass participation in the name of quality deliberation, her solution is surprisingly elitist and metaphorical. She envisions the Supreme Court in the United States, and the possibility for ordinary citizens to take their grievances to this type of institution, as a model of quality deliberation available to the masses. But this solution only allows for very indirect participation of the masses, as the "deliberation" is actually performed by an appointed few. The call for bridging the gap between micro and macro deliberation is right but the solutions remain quite elusive.

What if instead we returned to Cohen's initial vision to expand Track 2, but in a way that is kept safe from some of the problems identified by Habermas? It is perhaps possible to imagine a way in which deliberation would be used as the structuring norm but would involve a lot more people than the few hundreds involved in Track 1 and with the function of agenda-setting rather than decision-making. In the next section I turn to what I see as just such a (limited and imperfect) effort to square that circle, namely the French Great National

Debate. I will argue that it can be read as an effort to create an intermediate layer between Habermas' Track One and Track Two, namely a Third Track, where the public is given structured deliberative opportunities to speak and build an agenda for Track One.

2. A low-tech attempt to expand/scale deliberation

In this section I turn to what seems to me to amount to a first, imperfect attempt to construct an intermediary third track between Habermas' two tracks: the French 2019 Great National Debate.

In December 2018, President Macron decided to address the crisis of the Yellow Vests, which erupted over a gas tax increase, via a deliberative democratic strategy: the "Great National Debate," a two-month process that attempted to involve the whole population into a large-scale deliberation about 4 broad themes, including taxation, state services and organization, ecological transition, and democracy and participation. The process ultimately led to some policy changes and was crowned by an additional deliberative process known as the Citizens' Convention for Climate (which is fascinating in its own right but which I will not explore here because while it took place at the national level it only directly involved 150 participants and does not qualify as large-scale).[3]

What did this experiment in large-scale deliberation look like in a diverse country of 67 million people? Part of the challenge the French government set for itself, besides the rushed 2-month timeline, was the sheer scale and scope of the event. It is one thing to organize mass referendums, in which all people have to do is cast a vote on a predetermined question. It is another thing to ask millions of citizens to deliberate with each other on the scale of a large nation, even on a restricted agenda of four but nevertheless large and

[3] See instead my other paper: "In Defense of Citizen-Legislators" (working manuscript).

important questions. Past democratic experiments of the kind were still far and few at the times. The most promising ones had taken place in smaller and/or homogenous countries—Iceland, Canada, Ireland…—and often on single issues. The most recent precedent was barely a few months old, with about 1,000 townhall meetings across France on the topic of Europe, but those had been conducted largely under the radar of the media and public perception. Unsurprisingly, many observers could thus confidently predict that the Great Debate could not be successfully conducted within the announced timeframe and was going to end in disaster.

In the end, around 10, 000 town-hall meetings involving anywhere between 10 and 300 people took place all over the territory; 18, 847 grievance books were written in over 16,000 municipalities; 21 randomly selected assemblies of 100 citizens or so were organized[4]; and 4 thematic conferences (gathering intermediary bodies' representatives) were organized at the national level on each of the 4 themes delineated by President Macron. Meanwhile the online governmental platform gathered 1.9 million contributions; 16,874 emails and mails were received; and 5,400 contributions were gathered at "proximity stands" located in train station and post offices across the country. All in all, around 500,000 people were involved in the local meetings. Between 500,000 and 1.5 million people contributed online. Around 1400 people participated in the regional assemblies. In addition, the Great National Debate triggered conversations among roughly 45,000 participants in the so-called "True Debate" organized in parallel by opponents and critics of the Great National Debate. In some fashion, one can conservatively estimate that around 1.5% of the population directly participated in the debates.

[4] Including 13 in the 13 regions of France, 7 in the 5 French overseas territories and 1 among the youth (35 year-old adults or less) at the national level

1.5% is a lot of people. The French case thus suggests that deliberation at the large scale rather than the scale of small groups is possible, even if it will never take the form of deliberation involving the entire population at once. At the same time 1.5% is still a small fraction of the total population, even as it seems an upper bound figure compared to similar experiments elsewhere.

Additionally, as participation was mostly on the basis of self-selection (except in the 21 randomly selected regional assemblies), the people involved were highly unrepresentative of the larger population. In the local assemblies, for example, which gathered a lot more people (500,000 people or so) in groups ranging between 12 and 300 individuals, self-selection led to an overrepresentation of retired people, men, and Macron sympathizers, while the demographic whose actions can be said to have launched the Great National Debate, namely the peri-urban lower working-class people who started the Yellow Vest movement, were massively underrepresented.

Also problematically, most of the engagement was not truly deliberative in nature, consisting instead of online posts on the government's website, which did not allow for opportunities for discursive exchanges, either between participants or between participants and the government members meant to read the input (if any were looking, which is unknown).

Finally, even when the engagement was actually deliberative, the deliberation was not always structured according to the highest standards and best practices of deliberative democracy. The chairs of the self-selected meetings were also usually self-appointed and there were no facilitators to ensure equal speaking opportunities. In my observation, this led to classic hegemonies of age, gender, and race to assert themselves unimpeded, leaving many perspectives and voices either out of the room or silenced.

Nevertheless, even with such caveats, a very large number of people engaged in some form of however imperfect deliberation, that is, an inclusive exchange of views and arguments among individuals on the assumed basis of their equality of political status. Additionally, the deliberations taking place in the Great National Debate had ripple effects outside of it, spurring what might be termed "shadow deliberations" among the Yellow Vests in their previously mentioned counter debate (the so-called "True Debate") as well as among the larger population, who did not necessarily participate but read about it in the papers and discussed it at home and among friends, family, and colleagues. In that sense, the (semi-) structured deliberation of the Great National Debate produced and stimulated more unstructured "deliberation in the wild."

The French process was surprisingly low-tech for something so ambitious. Technologies were not usefully mobilized to facilitate online deliberation (if anything the official government website made it impossible). They were not used to connect the local meetings or the randomly selected assemblies or the thematic assemblies (though this could have been done and was briefly considered). Where it happened, deliberation took the form of regular face-to-face, small-scale deliberation but de-multiplied over the whole territory and with a (weak) effort at connecting the participants engaged in various steps.

Digital technologies and big data, however (though not AI) proved essential and indeed absolutely necessary to analyze and process the enormous amount of data generated by the Great Debate. 6 million online contributions were analyzed by two companies (Opinion Way and Kwam) using automated text analysis. Everything else (emails, grievances books, the output of regional assemblies, …) was analyzed by three companies (Roland Berger, Bluenove and Cognito) using visualization methods dating back to the 60s, centrally the "knowledge trees" of Michel Serres that help display consensus areas as the trunk of a

tree, polarized or different ideas as separate branches, and dissenting or unique ideas as "leaves" springing from the trunk or branches.

While all these techniques help map out the main areas of consensus among an enormous amount of data, the way the process generated this data had been structured made it very hard to derive any conclusions from it. How could one indeed meaningfully aggregate the output of a myriad self-selected meetings, that of 21 randomly selected assemblies, the content of grievance books, and that of online comments? The government tried to articulate local assemblies and randomly selected regional assemblies by having organizers present an hour-long PowerPoint about the provisional results of the Great National Debate up to that point on the first night of the regional assemblies. The point was presumably to have the random samples at the regional level deliberate on the basis of the input of previously held self-organized local meetings but the articulation of these two sequences was very poorly thought-through and it is unclear that the presentation had any influence on the deliberations of the participants in the regional assemblies.

The Great National Debate was thus a flawed but interesting attempt to create a third deliberative track between the formal track of decision-making and the informal track of "deliberation in the wild," which at the very least stimulated more action in the latter. Did this semi-structured third track produce what could qualify as "mass deliberation"? And if not, what more could have been done to do so?

It would be hard to justify calling the participation of 1.5% of the population in the Great National Debate an example of mass deliberation, for all the reasons mentioned above. The question is then: could more people have been brought in and been engaged in more inclusive, at the very least more representative, and in more deliberative ways?

This third track surely could have been structured in a more coherent way. The National Center for Public Deliberation, which had initially been entrusted with the

organization of the GND but was ultimately sidelined for political reasons in favor of the more pliable CESE, had proposed a design that hierarchized more logically various forms of participation and types of contributions. Another question worth pondering here is whether more could have been done, using new technologies, to help scale deliberation, and specifically quality deliberation, to a larger number of people. In what ways could AI, specifically, help scale deliberation to the masses in a way that resolves the tension between mass participation and quality deliberation? I now turn to these two questions.

3. Scaling deliberation with AI?

Mass deliberation is, at the limit, all-inclusive deliberation of all minds with all minds taking place in one space and at the same time. Let us first admit that this ideal is probably unreachable. But can we get close enough to it in the sense that vast numbers of people could engage in synchronous deliberation together, and how could AI help bring this approximation about? In this section I consider in turn two different solutions. One is the vision of "mass online deliberation" proposed by a Russian Engineer named Cyril Velikanov, which makes use of algorithms to distribute clustering, facilitation and evaluation tasks among participants. In the second section, I then turn to the more decentralized vision—inspired by certain aspects of the French Great National Debate—of a multiplicity of rotating citizens' assemblies (either physical or virtual). I propose four ways in which AI could help either model scale further.

1. Mass Online Deliberation

The concept of "Mass Online Deliberation" has been developed by a retired Russian engineer named Cyril Velikanov. His vision combines human judgment and the capacity of AI for clustering ideas and proposals into one possibly attractive vision of deliberative democracy

for the masses. In a paper written with co-author Alexander Prossove, they theorize Mass Online Deliberation (MOD) as a process whereby thousands, perhaps hundreds of thousands of people could be brought into a common virtual space, where they are able to engage in a multi-stage deliberative process of ideation, commenting and exchanging arguments, evaluation, and decision-making with the help of algorithms. They emphasize that for them "mass" means many people deliberating *together* in one common "room." This is by contrast with the common method of having several small groups deliberating *separately* in several "rooms"" (2022: 3 and section 6). They distinguish their model from Deliberative Polls (and by implication Citizens' Assemblies), which typically break up large groups into smaller deliberative units and then develop various strategies to integrate the outputs of these separate units at the collective level afterwards. By contrast, they insist that:

> our MOD model defines an *integrated deliberation space*, where every participant deliberatively addresses the whole community, and gets back deliberation data, somehow integrated, from the whole community. "Somehow integrated" means a set of backstage procedures that work permanently, or at regular intervals, on the whole space of deliberation data (proposals, comments, appraisals etc.) produced by individual participants (Velikanov and Prossove 2022: 12, my emphasis).

Integration is the key idea here, offering the vision of 'all minds in one room' even at the scale of hundreds of thousands. In the end, as the authors summarize it, MOD is intended as "as a mode of integrated communication of the whole community, regardless of its size, in one common "room", with everybody addressing the whole community and having an integrated vision of the deliberation results of the whole community" (2022: 15).

In this vision of mass online deliberation, algorithms play an important albeit limited role. It sorts out and clusters proposals to offer "a "bird's eye view onto the whole sea of participants' contributions" so as to "make it easy for any participant to navigate across it" (Velikanov and Prossove 2022: 32). The exchanges among self-selected participants in the process are supposed to be content-moderated, facilitated, partially structured and organized,

and ultimately judged and evaluated by random draws of humans themselves, on the basis of prodding by the algorithm. Velikanov and Prossove also entertain the possibility for AI to play a role as a translator in multilinguistic MODs, but again not so much to provide the translation as to distribute the function of translators to volunteers with the right skills.

This vision is fascinating and could be, to my mind, made even more plausible still by the use of smart algorithms unburdening humans of the moderating, organizing and clustering, editing and translating tasks that Velikanov still wants them to perform. If AI can automate such tasks, it would leave humans with the sole evaluation task of judging the quality of a proposal and how much they agree with it according to their own internal standards. Minimizing the time constraints and cognitive burdens on humans might even make participation in such MODs more attractive to an even greater number of people.

Nonetheless they are remaining difficulties with the model of Mass Online Deliberation. One, obviously, is that much of the vision for MOD remains to be implemented and tried out empirically. The evaluation task has already been piloted in a Finnish crowdsourced process on offroad traffic law reform, with an algorithm capable of eliciting various forms of ranking and rating from random draws of participants (Lee, Goel, Aitamurto, Landemore 2014). But the rest of the model is still in need of a proof of concept. The big unknown, in particular, is whether the model can really accommodate as many people as it claims and whether it could accommodate the entire population of any given country. Even a clustered bird's eye view of the deliberative landscape among millions might take too long or be too difficult to produce at that scale, and it would only cover a small fraction of a large country like the United States.

There are also conceptual difficulties in terms of the original ideal of all-inclusive deliberation.

Note, first, that even in this integrated model, not everyone reads all the arguments and comments, let alone reply to them all, and we need the intermediary steps of an algorithm clustering ideas, arguments, and proposals to expose participants to a much reduced and more manageable number of them. This division of labor, to be fair, is a structural necessity for mass deliberation in a way that it isn't for mass voting. Unlike mass voting, mass deliberation needs to be decomposed in smaller tasks and to take place into sufficiently smaller groups of people. So people are in one integrated room but they do not talk to everyone in it nor do they see every detail of the whole picture.

Second, the number of people who are able to participate, though massive in absolute terms may still be a very small percentage of the target population. If we could force participation of all or at the very least a sufficient percentage of the population (high enough to satisfy legitimacy requirements if those can fall short of absolute inclusion), then we run into the feasibility constraint again. Can MOD be run for millions of people? And at that scale are people really engaging in deliberation or a more superficial form of participation?

Third, it is not clear, even at the intermediary scale for which it is theorized, whether MOD really allows for the quality deliberation among all we are after or whether it simply elicits a superficial engagement with the question, where most people end up passively scrolling through the comments and at best editing proposals at the margin rather than engaging in a constructive back and forth about the reasons for such proposals. With a sample of participants so large, the incentives to free ride are enormous and one sense of agency quite reduced. Would people feel sufficiently seen, valued, and safe to try and persuade others of their views? Would, in particular, traditionally underrepresented or vulnerable people show up and, if they show up, actively fight for minority or misunderstood views rather than passively go with the emerging consensus?

Finally, and on a related note, to the extent that enrollment is voluntary and participation once enrolled on the sole basis of self-selection, the participants may be very unrepresentative of the target population and thus the resulting deliberation highly biased. Velikanov does not seem to think it is such a problem, essentially on the basis of an observation made by Jim Fishkin and Robert Luskin that the differences between the people from a large random draw of the population who chose to enroll in a deliberative poll and those who did not were not all that significant empirically (Luskin and Fishkin 2005: 40). Fishkin and Luskin credit the lack of distinctiveness between the self-selected sample of participants and the others to the peculiar incentives for participation baked in the design of their deliberative poll.

That self-selected participants representatively track the main features of a given population, however, is typically not the case in other processes based on self-selection, such as Participatory Budgeting or crowdsourced processes, where the participating samples are usually very biased. Whether financial and other incentives fundamentally change the matter sounds dubious. It is true that even when the samples are biased, the proposals coming out of them might still be quite good and useful with respect to the preferences of the larger public. This would explain why Wikipedia, which is almost entirely based on self-selection and thus like almost every other online crowdsourced process overrepresents educated white males, is remarkably accurate not just in tracking facts but also the moral consensuses ultimately underlying the interpretation of many of those facts (this is not to say that editorial wars over controversial issues like the Iraq war or the January 6 Capital invasion are not real). Similarly, the online platform pol.is used by the Taiwanese government on issues ranging from Airbnb or Uber regulations or, more recently, pandemic management has been able to yield reasonable consensuses (built by the algorithm on the basis of the crowd's written online inputs) that the larger population could get behind. Nonetheless the self-selection bias

should still give us pause from a conceptual and democratic point of view and make us wary of using a self-selected sample, even a large one, as a substitute for a randomly selected one, despite the attractiveness of a non-gated process over a gated, albeit egalitarian one like a sortition-based assembly.

Because of these question marks about the feasibility and representativeness of mass online deliberation, let me turn instead to an alternative vision of AI-powered mass deliberation, one that would take as its starting point the only physical approximation of an all-inclusive deliberation that I think is conceptually possible, even as it presents some of the problems of lack of integration that MOD is supposed to overcome.

2. *A multiplicity of rotating randomly selected assemblies*

Let me return to the Great National Debate and what seems to me the most promising example of quality deliberations in it, namely the deliberations that took place in the 21 randomly selected regional assemblies.

The beauty of random selection (or rather stratified random sampling in most cases) is that if the group is large enough, participation either mandatory or sufficiently well incentivized to draw most people out, and their deliberations are truly inclusive of all the participants in the process, this process offers a good simulation of what the millions of people forming the larger public would think if they were given a chance to deliberate for a day and a half on the same issues as the randomly selected citizens were. How do we know that? It is a prediction on the basis of the formal properties of random selection with large numbers (see also Fishkin 2018). In practice citizens' assemblies have often delivered recommendations tracking relatively closely what the larger public is ready to endorse when exposed to their proposals. For example, 2/3 of the voting population in the 2018 referendum on abortion in Ireland supported decriminalization of abortion, a striking match with the

proportion by which the Irish citizens'assembly had itself voted two years before. In France, the Citizens' Convention for Climate produced after 9 months of intense deliberations 149 proposals. While many critics and the participants themselves feared a disconnect between the outcome of the 9-month process and the rest of the public, pollsters checked the popularity of the 14 major proposals among the larger public and found that except for a single proposal, all found majoritarian approval in the larger public (by a admittedly lesser margin than in the Citizens' Convention itself) (REF).

Critics complain that this more or less rough match is not a guaranteed outcome because the outcome of citizens' assemblies is bound to vary considerably depending on all kinds of factors. Dominique Schnapper thus argues that deliberation would in fact produce different results each time, even among a strictly identical group of randomly selected people, if, for example, the first person to start the conversation tilted the agenda in a different direction (REF).

One way to counter this argument would be to show that citizens' assemblies drawn at random from the same pool and run along similar lines would, in fact, systematically converge on similar recommendations or conclusions. Empirically there is not much evidence in that direction since deliberative assemblies tend to be unique events. So far there has been no way of verifying or falsifying the claim that deliberative assemblies run in parallel would result in the same overall recommendations, since all the known cases of deliberative mini-publics were singular events. Most of them do not even include a control group, that is a group of citizens drawn at random from the same original pool, which can serve as the counterfactual for the deliberative treatment given to the other group. Only James Fishkin's deliberative polls do include such a control group, which allows him and his team to measure more precisely the effect of deliberation on pre-deliberative preferences.

The French Great National Debate came closest to offering simultaneous citizens' assemblies, with its multiple randomized regional assemblies run within roughly the same two weeks under very similar conditions. Unfortunately, the participants in each were not random draws from the entire French population, as each draw was from a different region of France. And yet it is striking how converging some of their ultimate judgments and recommendations were, perhaps because on big picture questions the regional differences ultimately did not matter that much. On the issue of the environment, for example, 12 out of 18 assemblies converged on the specific idea that a new form of democratic governance needed to be put in place for the specific purpose of managing the green transition. This convergence is apparently what convinced President Macron to launch a Citizens' Convention for Climate the following fall.[5] So, the Great National Debate suggests that if we could run at the same time and under sufficiently similar conditions enough true random draws of the population separately and observe them converge in their deliberations, we would have a pretty good sense of the positions most people enrolled in such a deliberative process would land on any given topic.

Let us run for a minute with this thought experiment. Imagine, in fact, that we could get 67 million French people speaking to each other in a structured way, namely by being enrolled in as many mini-publics as necessary to map out the entire population. This could take place physically or online. We can imagine 670,000 assemblies of 100 people each, or 134,000 assemblies of 500 people each. Within these assemblies, smaller groups would deliberate (in addition to attending plenaries and engaging in other interactions) and then people in them would be randomly rotated until they talked to all of the other members.

[5] Personal communication from one of the organizers of the regional assemblies who presented their results to President Macron in April 209.

Similarly, once each process within individual assemblies has been completed, we would rotate the members of the assemblies and reconfigure new assemblies at random, so that a second batch of assemblies can deliberate further about the same issue. We would iterate the process until everyone has talked to everyone else at least once.

That is conceptually one way to envisage mass deliberation as deliberation of all with all. Nonetheless such a system would still be too time-consuming and presumably very redundant (since in order for everyone to speak with everyone else at least once, many people would often end up having very similar conversations with sufficiently similar people). There are clearly diminishing returns to seeking exhaustivity of the process in terms of having absolutely everyone talk to absolutely everyone else.

But are we allowed to compromise the original ideal of deliberative democracy as involving deliberation among all? If the goal is mass participation in the sense of all-inclusive deliberation, then we should rotate all the way to the end until everyone has talked to everyone and thus ensure a deliberation of everyone with everyone. That is the original ideal of deliberative democracy in Joshua Cohen's version for example. At the same time, at scale, we do make compromises with similar ideals all the time. Consider the ideal of voting. In theory the process should involve absolutely everyone. In practice referenda are deemed legitimate when at least half the voting population shows up. In countries like France, where abstention has become a structural factor, we maintain the fiction of legitimacy even as participation drop below 50%. Yet under that threshold many referenda are often deemed illegitimate, or at least much less legitimate. This is so because the outcome seems decided by a minority and, additionally, a minority that is likely not to represent all existing subgroups within the population.

The same way that mass voting does not imply that everyone votes, but only that a critical mass of people vote, mass deliberation should probably mean that not all but a

sufficient number of people deliberate. Now, if 50% is the floor for voting, what should be the right floor for deliberation among randomly selected participants? Presumably it will be less than half the population, since deliberation is a much more demanding task than voting and because since we appoint people at random (preferably on the basis on mandatory participation or with strong incentives to ensure a high take-up rate), the sample will be sufficiently demographically representative as long as it is sufficiently large. So the cut-off point would probably be somewhere between the 1.5% participation rate of the Great National Debate, which seems too little, and 50% of the population, which seems too large to be realistic.

I could see a case made for a threshold as low as 10 to 15%. In other words, as long as 10 to 15% of the population ends up being enrolled in randomly selected assemblies rotated a sufficient number of times, we might have met the legitimacy threshold for mass deliberation. Of course, one might object that a big difference between mass voting and mass deliberation thus conceived is that in mass voting even if not everyone shows up everyone had an *opportunity* to vote. If they had wanted to vote, they could have. That cannot be the case for participation in citizens' assemblies since, by construction, not everyone who wants to join them can. So here the bar for legitimacy should be something else, like the fact of being entered by default in the pool from whom the participants in randomly selected assemblies are selected.

One may want to refine the threshold of mass deliberation beyond the somewhat arbitrary percentage of the population just mentioned. A better argument might be to look at a cut-off point corresponding to the point at which a sufficiently large range of views on the given topic has been explored, to minimize the risk of excluding good arguments and ideas. How much of the population would need to be enrolled and rotated into deliberative

assemblies for the relevant coverage of arguments and ideas to be sufficient? The answer would probably vary depending on the topic but there is surely a way to answer this question, if only approximatively.

An objector might ask: If what we are interested in is capturing a diversity of views rather than include absolutely everyone, then why not settle for a large Citizens' Assembly like the 150-citizens strong Citizens' Convention for Climate , perhaps combined with crowdsourced open consultations in the vein of Taiwan's pol.is or Velikanov's concept of Mass Online Deliberation?

The answer is that having a single citizens' assembly at the national level, even combined with participatory forms of crowdsourcing as a supplement, does not solve the legitimacy deficit that critics claim emerge from limiting deliberation to a few hundred people while millions watch on or even contribute from afar. If we are serious about radically expanding the numbers of people involved in deliberation, thousands of rotated assemblies are closer to the mark (because the ceiling is not as low or rigid). These assemblies might end up involving millions of people and the only point of cutting the rotation short would be to save time and money while still meeting the bar for "mass participation" deemed critical for democratic legitimacy. As to the epistemic performance of either solution, it is unclear whether a large central CA would produce the same results as multiple rotating ones. With the singular large assembly, one would need to choose the size beforehand (which may end up not capturing all perspectives), whereas the rotation structure allows for ever-increasing amounts of citizen interactions, ensuring that all perspectives will always be represented.[6]

[6] I thank Andrew Sorota for that point.

3. *Role for AI in mass deliberation*

We have just encountered two visions of mass deliberation, both making use of AI in some fashion. What other roles than the ones already mentioned could AI play in either of them? There are at least four obvious ways in which AI could help scale and even improve deliberation in these two scenarios: facilitation, translation, fact-checking, and data clustering and aggregation.

3.1 Facilitation

Facilitators are often argued to be essential to quality deliberation, especially in the early stages of the implementation of deliberative processes, when participants may lack the awareness, knowledge, or practice of certain norms aiming to ensure equality of voice. Facilitators may also prove essential even when those norms have been sufficiently internalized—such as in the heat of the discussion, when people often tend to revert to tendencies such as talking too much or interrupting women and shy people—and some discipline needs to be constantly reintroduced. The problem with human facilitators is that they are expensive and often biased, even with proper training. Additionally, if the goal for participants is ultimately some kind of autonomous decision-making, it should probably happen without the micro-management of external facilitators and be replaced, in the best case scenario, by internalized norms, self-monitoring, and peer pressure among the participants for best practices as well as procedures and the equivalent of Robert's Rules of Order for deliberative assemblies of citizens.

If we deem facilitators necessary, however, the first use of AI could thus be to serve as a cheap, impartial moderator and facilitator of group discussions, provided those are conducted online (for now, as perhaps an embodied AI could one day conduct physical meetings as well). Stanford University has already developed an algorithm called "Alice"

(after Alice Siu, the person who gave it her voice) that has facilitated online conversations among groups of 7 to 8 people. AI "Alice" currently works more as a smart clock than an intelligent person (or perhaps even an intelligent algorithm, on a demanding notion of intelligence that involves learning over time), but it is likely that these minimal functions of keeping time and distributing speaking rights equitably could be upgraded to a fuller range of human-like abilities to maintain a good and productive flow of conversations among a larger group of people. Content moderation—flagging insults, advertisements, and gibberish—and either removing them or signaling them to a human supervisor—is already a routine task assigned to algorithms on social media and could simply be extended to deliberative platforms.

3.2 Translation

Second, an AI could serve as an instant translator among participants in a multi-lingual context. Currently the standard alternative is professional translators, and the costs of this are rather prohibitive, as evidenced in the context of the deliberative poll "Tomorrow's Europe" conducted by Fishkin in the European Parliament in Brussels in 2009, where an army of translators helped communication among people speaking in no fewer than 23 languages. Translations is one area where AI has exceeded expectations and would truly facilitate cross-linguistic dialogues. This is not to say, however, that AI translators could necessarily ever amount to "political translators" (Doerr 2018), a complex cognitive and ethical task that involves being able to pick up and explain the background inequalities between speakers from different regions of the world and advocating on behalf of the weaker groups and individuals. From that perspective the full automation of translation tasks might not be desirable and hybrid systems might prove preferable.

3.3 Fact-checking

Third, an AI could also help in providing accurate information in real time for the human deliberators and serving generally as a basic fact-checker, in lieu of the actual humans sometimes used for that function in deliberative settings. Think of an "Alexa" for deliberation.

3.4 Clustering and organizing arguments

Fourth, an AI could help cluster ideas and arguments in an organized way so that the intersecting contributions of large amounts of people can actually be reduced to a manageable amount of content. This could be done in real time to help move the deliberation forward or ex post facto to make sense of a large amount of data produced in processes like the Great National Debate.

These four functions are currently performed by humans, and while relying on AI to perform these functions may certainly present issues (of privacy, impartiality, and bias, among others), these are not fundamentally different than the issues raised by human facilitators, translators, fact-checkers, and data processors. There is to my knowledge no existing deliberative platform that makes use of all of these capabilities, thus suggesting the real constructive impact that an intervention with AI could have in these settings.

We should also consider, however, the potential responsibilities of AI that *cannot* be analogously performed by humans in deliberative spaces.

Among the tasks that an AI would be expected to perform in something like this one could list the following:

3.5 Track all the exchanges anyone's had with anyone else and measure the degree of overlap in content

If we seek to rotate assemblies until the point of diminishing return in terms of exposure to diverse ideas and arguments, we need a sense of when that point has been reached. The first, low-tech solution would simply be to ask people regularly whether they feel they have

reached a point at which they are not learning anything anymore or the conversations they have are no longer that productive. The feeling of going in circles presumably tracks the reality that enough people have been exposed to enough other people. So, when a majority of people, or perhaps a super majority, decides that it is time to move on, the deliberation could come to an end and the process could move on to, say, a vote. Another option, however, would be to let an AI algorithm track the arch of the conversations, analyzing the similarities in content and counting up the number of new ideas and measuring the point at which repetitions multiply (or whatever other way is used to measure the fact of "going in circles") and diminishing returns are reached. The AI could then inform the participants that such a point has been reached, which could be before they themselves identify it or perhaps quite some time after they feel ready to move on.

3.6 Measuring the quality of deliberation

An AI could also measure and plot the quality of deliberation in real time, for example using a properly adapted index like the Discourse Quality Index (Bächtiger et al. REF). If quality drops, AI may need to step up its facilitator or fact-checking functions or bring in a human supervisor.

3.7 Taking cognitive pictures of the group

AI could allow the group to take a "cognitive group selfie"—by producing an instant visualization, perhaps projected on a screen, of where the group stands on a given question at any point in time, including in terms of the various dissenting minorities—so that individual members know where they stand at all times in terms of emerging consensus, degrees of polarization, dissenting points of views, quality of deliberation etc.

3.8 Sharing consensus across groups

AI could perhaps share an emerging consensus in one group with other groups to see if it helps speed up convergence (e.g., all 18 assemblies might have converged in the French

GND). Of course, this probably depends on what we are trying to achieve. Perhaps it is epistemically better to preserve the independence of assemblies up to a point. But recent literature indicates that ultimately deliberation beats pure aggregation when it comes to generating superior predictions and solutions (Navajas et al REF) and so this points to the need to have the members of each assembly be put in dialogue with the members of other assemblies, once each assembly has reached a sufficient consensus or some natural endpoint of the conversations.

3.9 Seeding assemblies with high-potential ideas

AI could pick up on highly contagious ideas and once they are fact-checked and otherwise vetted by the proper authorities make sure they get transferred to other groups even if deliberation ends before they are properly discussed or become consensual (for example because they emerge too late in the deliberation) If the ideas spread quickly in other assemblies once planted there early, they might help the conversations or become consensus themselves much faster.

How would we know which ideas are the most "promising"? One method is the one my colleagues and I have used in the Finnish experiment (Aitamurto and Landemore 2016, Aitamurto, Landemore and Salle 2018), and also that considered by Velikanov. It consists in delegating to a random draw of participants the ranking, rating, and general evaluation of a random draw of the ideas or proposals. Another option is simply for the AI to identify the ideas that quickly garner attention within a group—ideas with a high degree of potential virality so to speak—and make sure to transplant them in other assemblies before they die out in their assembly of origin, which might just not be a fertile ground for them regardless of the quality of the idea.

All these proposed uses of AI raise a number of issues, none of which the scope of this paper allows me to tackle but which deserve mention. What if the AI picks up on terrible

viral ideas, like raw anti-immigration sentiments and anti-minority views (or views that are not fact-checkable or are even true but toxic) and help spread populist wildfires rather than reasonable consensus? Among the safeguards that could be put in place are an ethics committee, which would monitor the AI's findings and to which the AI would be made to refer any proposals with potentially problematic content. Participants themselves would be allowed to report issues and trigger some procedural investigation into the breaches of established norms of good deliberation.

Conclusion:

AI can help scale deliberative democracy in multiple capacities. But in the end, we have to let go of the original ideal of full inclusive, simultaneous deliberation of all with all. The best we can probably have is either millions of self-selected people processing arguments and ideas together on an AI-augmented online platform (Mass Online Deliberation) or a substantial percentage of the population engaged in deliberative mini-publics being rotated a number of times (the regional assemblies of the Great National Debate on steroids). Even these approximations of all-inclusive deliberation, however, are incredibly demanding, time-consuming, and presumably costly. But they may be the bridges we need to bridge the ideal of deliberative democracy and that of mass participation.

References

Ackerman, Bruce, and James Fishkin. 2005. *Deliberation Day.* New Haven, CT: Yale University Press.

Aitamurto, Tanja and Hélène Landemore. 2016. "Crowdsourced Deliberation: The Case of an Off-Traffic Law Reform in Finland." *Policy & Internet* May 8(2): 174-196.

Aitamurto, Tanja, Hélène Landemore and Jorge S. Galli. 2017."Unmasking the Crowd: Participants' Motivation Factors, Expectations, and Profile in a Crowdsourced Law Reform", *Information, Communication, and Society* 20(8): 1239-1260.

Chambers, Simone. 2009. "Rhetoric and the Public Sphere: Has Deliberative Democracy Abandoned Mass Democracy?" *Political Theory* 37(3): .

Cohen, Joshua. 1989. "Deliberation and Democratic Legitimacy." In A. Hamlin and P. Pettit, eds., *The Good Polity*. New York: Basil Blackwell, 17–34.

Doerr, Nicole. 2018. *Political Translation: How social movement democracies survive.* Cambridge; Cambridge University Press.

Fishkin, James. 2009. *When the People Speak: Deliberative Democracy and Public Consultation*. Oxford: Oxford University Press.

Fishkin, James. 2018. *Democracy When the People Are Thinking: Revitalizing Our Politics Through Public Deliberation*. Oxford: Oxford University Press.

Fishkin, James and Robert Luskin. 2005. "Experimenting with a Democratic Ideal: Deliberative Polling and Public Opinion." *Acta Politica 2005.*

Habermas, Jürgen. 1996. *Between Facts and Norms*. Translated by William Rehg. Cambridge, MA: MIT Press.

Hilbert, Martin. 2009. "The Maturing Concept of e-democracy. From e-voting and Online Consultations, to Democratic Value Out of Jumbled Online Chatter." *Journal of Information Technology and Politics*.

Iandoli Luca, Mark Klein and Giuseppe Zollo 2009. "Can We Exploit Collective Intelligence for Collaborative Deliberation? The Case of the Climate Change Collaboratorium." *MIT Sloan Research Paper No. 4675-08*.

Lafont, Cristina. 2020. *Democracy Without Shortcuts*. Oxford: Oxford University Press.

Lafont, Cristina. 2015. "Deliberation, Participation, and Democratic Legitimacy: Should Deliberative Mini-publics Shape Public Policy?." *Journal of Political Philosophy* 23(1): 40–63.

Landemore, Hélène. 2015. "Inclusive Constitution-Making: The Icelandic Experiment." *Journal of Political Philosophy* 23(2): 166-191.

Lee, David, Ashish Goel, Tanja Aitamurto, and Hélène Landemore. 2014. "Crowdsourcing for Participatory Democracies: Efficient Elicitation of Social Choice Functions." *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing* 2.

Navajas, Joaquin, Tamara Niella, Gerry Garbulsky, Bahador Bahrami and Mariano Sigman. "Aggregated knowledge from a small number of debates outperforms the wisdom of large crowds." https://arxiv.org/ftp/arxiv/papers/1703/1703.00045.pdf

Schnapper, Dominique.

Velikanov, Cyril and Alexander Prosser. "Mass Online Deliberation in Participatory Policy-Making." Working paper.